

Amazon.Premium.SAA-C02.818q

Number: SAA-C02
Passing Score: 800
Time Limit: 120 min
File Version: 22.2



Exam Name: AWS Certified Solutions Architect - Associate SAA-C02
Certification Provider: Amazon
Corresponding Certification: AWS Certified Solutions Architect - Associate
Website: <https://VCEup.com/>
Free Exam: <https://vceup.com/exam-aws-certified-solutions-architect-associate-saa-c02/>



Exam A

QUESTION 1

A solutions architect is designing a solution where users will be directed to a backup static error page if the primary website is unavailable. The primary website's DNS records are hosted in Amazon Route 53 where their domain is pointing to an Application Load Balancer (ALB).

Which configuration should the solutions architect use to meet the company's needs while minimizing changes and infrastructure overhead?

- A. Point a Route 53 alias record to an Amazon CloudFront distribution with the ALB as one of its origins. Then, create custom error pages for the distribution.
- B. Set up a Route 53 active-passive failover configuration. Direct traffic to a static error page hosted within an Amazon S3 bucket when Route 53 health checks determine that the ALB endpoint is unhealthy.
- C. Update the Route 53 record to use a latency-based routing policy. Add the backup static error page hosted within an Amazon S3 bucket to the record so the traffic is sent to the most responsive endpoints.
- D. Set up a Route 53 active-active configuration with the ALB and an Amazon EC2 instance hosting a static error page as endpoints. Route 53 will only send requests to the instance if the health checks fail for the ALB.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Active-passive failover

Use an active-passive failover configuration when you want a primary resource or group of resources to be available the majority of the time and you want a secondary resource or group of resources to be on standby in case all the primary resources become unavailable. When responding to queries, Route 53 includes only the healthy primary resources. If all the primary resources are unhealthy, Route 53 begins to include only the healthy secondary resources in response to DNS queries.

To create an active-passive failover configuration with one primary record and one secondary record, you just create the records and specify Failover for the routing policy. When the primary resource is healthy, Route 53 responds to DNS queries using the primary record. When the primary resource is unhealthy, Route 53 responds to DNS queries using the secondary record.

How Amazon Route 53 averts cascading failures

As a first defense against cascading failures, each request routing algorithm (such as weighted and failover) has a mode of last resort. In this special mode, when all records are considered unhealthy, the Route 53 algorithm reverts to considering all records healthy.

For example, if all instances of an application, on several hosts, are rejecting health check requests, Route 53 DNS servers will choose an answer anyway and return it rather than returning no DNS answer or returning an NXDOMAIN (non-existent domain) response. An application can respond to users but still fail health checks, so this provides some protection against misconfiguration.

Similarly, if an application is overloaded, and one out of three endpoints fails its health checks, so that it's excluded from Route 53 DNS responses, Route 53 distributes responses between the two remaining endpoints. If the remaining endpoints are unable to handle the additional load and they fail, Route 53 reverts to distributing requests to all three endpoints.

Reference:

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

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

QUESTION 2

A solutions architect is designing a high performance computing (HPC) workload on Amazon EC2.

The EC2 instances need to communicate to each other frequently and require network performance with low latency and high throughput.

Which EC2 configuration meets these requirements?

- A. Launch the EC2 instances in a cluster placement group in one Availability Zone.
- B. Launch the EC2 instances in a spread placement group in one Availability Zone.
- C. Launch the EC2 instances in an Auto Scaling group in two Regions and peer the VPCs.
- D. Launch the EC2 instances in an Auto Scaling group spanning multiple Availability Zones.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Placement groups

When you launch a new EC2 instance, the EC2 service attempts to place the instance in such a way that all of your instances are spread out across underlying hardware to minimize correlated failures. You can use placement groups to influence the placement of a group of interdependent instances to meet the needs of your workload. Depending on the type of workload.

Cluster – packs instances close together inside an Availability Zone. This strategy enables workloads to achieve the low-latency network performance necessary for tightly-coupled node-to-node communication that is typical of HPC applications.

Reference:

<https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/placement-groups.html>

QUESTION 3

A company wants to host a scalable web application on AWS. The application will be accessed by users from different geographic regions of the world. Application users will be able to download and upload unique data up to gigabytes in

size. The development team wants a cost-effective solution to minimize upload and download latency and maximize performance. What should a solutions architect do to accomplish this?

- A. Use Amazon S3 with Transfer Acceleration to host the application.
- B. Use Amazon S3 with CacheControl headers to host the application.
- C. Use Amazon EC2 with Auto Scaling and Amazon CloudFront to host the application.
- D. Use Amazon EC2 with Auto Scaling and Amazon ElastiCache to host the application.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Reference:

<https://aws.amazon.com/ec2/autoscaling/>

QUESTION 4

A company is migrating from an on-premises infrastructure to the AWS Cloud. One of the company's applications stores files on a Windows file server farm that uses Distributed File System Replication (DFSR) to keep data in sync. A solutions architect needs to replace the file server farm. Which service should the solutions architect use?

- A. Amazon Elastic File System (Amazon EFS)
- B. Amazon FSx
- C. Amazon S3
- D. AWS Storage Gateway

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Migrating Existing Files to Amazon FSx for Windows File Server Using AWS DataSync We recommend using AWS DataSync to transfer data between Amazon FSx for Windows File Server file systems. DataSync is a data transfer service that simplifies, automates, and accelerates moving and replicating data between on-premises storage systems and other AWS storage services over the internet or AWS Direct Connect. DataSync can transfer your file system data and metadata, such as ownership, time stamps, and access permissions.

Reference:

<https://docs.aws.amazon.com/fsx/latest/WindowsGuide/migrate-files-to-fsx-datasync.html>

QUESTION 5

A company has a legacy application that processes data in two parts. The second part of the process takes longer than the first, so the company has decided to rewrite the application as two microservices running on Amazon ECS that can scale independently.

How should a solutions architect integrate the microservices?

- A. Implement code in microservice 1 to send data to an Amazon S3 bucket. Use S3 event notifications to invoke microservice 2.
- B. Implement code in microservice 1 to publish data to an Amazon SNS topic. Implement code in microservice 2 to subscribe to this topic.
- C. Implement code in microservice 1 to send data to Amazon Kinesis Data Firehose. Implement code in microservice 2 to read from Kinesis Data Firehose.
- D. Implement code in microservice 1 to send data to an Amazon SQS queue. Implement code in microservice 2 to process messages from the queue.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 6

A company captures clickstream data from multiple websites and analyzes it using batch processing. The data is loaded nightly into Amazon Redshift and is consumed by business analysts. The company wants to move towards near-real-time data processing for timely insights.

The solution should process the streaming data with minimal effort and operational overhead.

Which combination of AWS services are MOST cost-effective for this solution? (Choose two.)



- A. Amazon EC2
- B. AWS Lambda
- C. Amazon Kinesis Data Streams
- D. Amazon Kinesis Data Firehose
- E. Amazon Kinesis Data Analytics

Correct Answer: BD

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Kinesis Data Streams and Kinesis Client Library (KCL) – Data from the data source can be continuously captured and streamed in near real-time using Kinesis Data Streams. With the Kinesis Client Library (KCL), you can build your own application that can preprocess the streaming data as they arrive and emit the data for generating incremental views and downstream analysis.

Kinesis Data Analytics – This service provides the easiest way to process the data that is streaming through Kinesis Data Stream or Kinesis Data Firehose using SQL. This enables customers to gain actionable insight in near real-time from the incremental stream before storing it in Amazon S3.

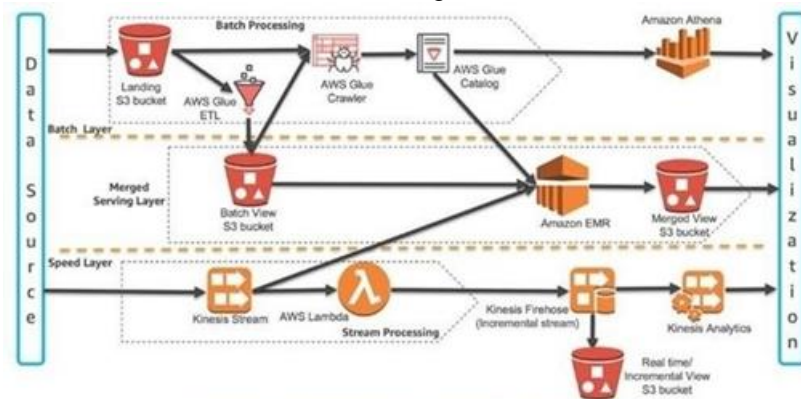


Figure 2: Lambda Architecture Building Blocks on AWS



Reference:

<https://d1.awsstatic.com/whitepapers/lambda-architecture-on-for-batch-aws.pdf>

QUESTION 7

A company's application runs on Amazon EC2 instances behind an Application Load Balancer (ALB). The instances run in an Amazon EC2 Auto Scaling group across multiple Availability Zones. On the first day of every month at midnight, the application becomes much slower when the month-end financial calculation batch executes. This causes the CPU utilization of the EC2 instances to immediately peak to 100%, which disrupts the application. What should a solutions architect recommend to ensure the application is able to handle the workload and avoid downtime?

- A. Configure an Amazon CloudFront distribution in front of the ALB.
- B. Configure an EC2 Auto Scaling simple scaling policy based on CPU utilization.
- C. Configure an EC2 Auto Scaling scheduled scaling policy based on the monthly schedule.
- D. Configure Amazon ElastiCache to remove some of the workload from the EC2 instances.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Scheduled Scaling for Amazon EC2 Auto Scaling

Scheduled scaling allows you to set your own scaling schedule. For example, let's say that every week the traffic to your web application starts to increase on Wednesday, remains high on Thursday, and starts to decrease on Friday. You can plan your scaling actions based on the predictable traffic patterns of your web application. Scaling actions are performed automatically as a function of time and date.

Reference:

https://docs.aws.amazon.com/autoscaling/ec2/userguide/schedule_time.html

QUESTION 8

A company runs a multi-tier web application that hosts news content. The application runs on Amazon EC2 instances behind an Application Load Balancer. The instances run in an EC2 Auto Scaling group across multiple Availability Zones and use an Amazon Aurora database. A solutions architect needs to make the application more resilient to periodic increases in request rates.

Which architecture should the solutions architect implement? (Choose two.)

- A. Add AWS Shield.
- B. Add Aurora Replica.
- C. Add AWS Direct Connect.
- D. Add AWS Global Accelerator.
- E. Add an Amazon CloudFront distribution in front of the Application Load Balancer.

Correct Answer: DE

Section: (none)

Explanation

Explanation/Reference:

Explanation:

AWS Global Accelerator

Acceleration for latency-sensitive applications

Many applications, especially in areas such as gaming, media, mobile apps, and financials, require very low latency for a great user experience. To improve the user experience, Global Accelerator directs user traffic to the application endpoint that is nearest to the client, which reduces internet latency and jitter. Global Accelerator routes traffic to the closest edge location by using Anycast, and then routes it to the closest regional endpoint over the AWS global network. Global Accelerator quickly reacts to changes in network performance to improve your users' application performance.

Amazon CloudFront

Amazon CloudFront is a fast content delivery network (CDN) service that securely delivers data, videos, applications, and APIs to customers globally with low latency, high transfer speeds, all within a developer-friendly environment.

Reference:

<https://docs.aws.amazon.com/global-accelerator/latest/dg/introduction-benefits-of-migrating.html>

QUESTION 9

An application running on AWS uses an Amazon Aurora Multi-AZ deployment for its database.

When evaluating performance metrics, a solutions architect discovered that the database reads are causing high I/O and adding latency to the write requests against the database.

What should the solutions architect do to separate the read requests from the write requests?

- A. Enable read-through caching on the Amazon Aurora database.
- B. Update the application to read from the Multi-AZ standby instance.
- C. Create a read replica and modify the application to use the appropriate endpoint.
- D. Create a second Amazon Aurora database and link it to the primary database as a read replica.



Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Amazon RDS Read Replicas

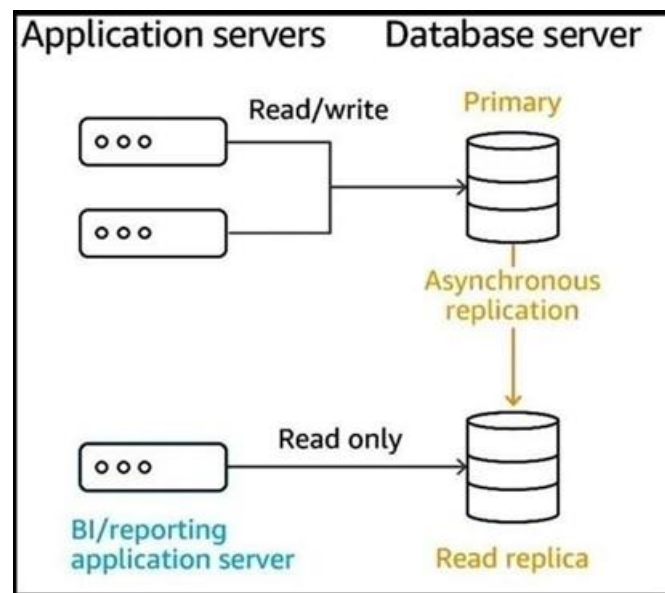
Amazon RDS Read Replicas provide enhanced performance and durability for RDS database

(DB) instances. They make it easy to elastically scale out beyond the capacity constraints of a single DB instance for read-heavy database workloads. You can create one or more replicas of a given source DB Instance and serve high-volume application read traffic from multiple copies of your data, thereby increasing aggregate read throughput. Read replicas can also be promoted when needed to become standalone DB instances. Read replicas are available in Amazon RDS for MySQL, MariaDB, PostgreSQL, Oracle, and SQL Server as well as Amazon Aurora.

For the MySQL, MariaDB, PostgreSQL, Oracle, and SQL Server database engines, Amazon RDS creates a second DB instance using a snapshot of the source DB instance. It then uses the engines' native asynchronous replication to update the read replica whenever there is a change to the source DB instance. The read replica operates as a DB instance that allows only read-only connections; applications can connect to a read replica just as they would to any DB instance.

Amazon RDS replicates all databases in the source DB instance.

Amazon Aurora further extends the benefits of read replicas by employing an SSD-backed virtualized storage layer purpose-built for database workloads. Amazon Aurora replicas share the same underlying storage as the source instance, lowering costs and avoiding the need to copy data to the replica nodes. For more information about replication with Amazon Aurora, see the online documentation.



Reference:
https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_ReadRepl.html
<https://aws.amazon.com/rds/features/read-replicas/>

QUESTION 10

A recently acquired company is required to build its own infrastructure on AWS and migrate multiple applications to the cloud within a month. Each application has approximately 50 TB of data to be transferred. After the migration is complete, this company and its parent company will both require secure network connectivity with consistent throughput from their data centers to the applications. A solutions architect must ensure one-time data migration and ongoing network connectivity. Which solution will meet these requirements?

- A. AWS Direct Connect for both the initial transfer and ongoing connectivity.
- B. AWS Site-to-Site VPN for both the initial transfer and ongoing connectivity.
- C. AWS Snowball for the initial transfer and AWS Direct Connect for ongoing connectivity.
- D. AWS Snowball for the initial transfer and AWS Site-to-Site VPN for ongoing connectivity.



Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
 Reference:
https://docs.aws.amazon.com/dms/latest/userguide/CHAP_LargeDBs.html
<https://aws.amazon.com/directconnect/>

QUESTION 11

A company serves content to its subscribers across the world using an application running on AWS. The application has several Amazon EC2 instances in a private subnet behind an Application Load Balancer (ALB). Due to a recent change in copyright restrictions, the chief information officer (CIO) wants to block access for certain countries. Which action will meet these requirements?

- A. Modify the ALB security group to deny incoming traffic from blocked countries.
- B. Modify the security group for EC2 instances to deny incoming traffic from blocked countries.
- C. Use Amazon CloudFront to serve the application and deny access to blocked countries.
- D. Use ALB listener rules to return access denied responses to incoming traffic from blocked countries.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
 Explanation:
 "block access for certain countries." You can use geo restriction, also known as geo blocking, to prevent users in specific geographic locations from accessing content that you're distributing through a CloudFront web distribution.
 Reference:

<https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/georestrictions.html>

QUESTION 12

A company is creating a new application that will store a large amount of data. The data will be analyzed hourly and modified by several Amazon EC2 Linux instances that are deployed across multiple Availability Zones. The application team believes the amount of space needed will continue to grow for the next 6 months.

Which set of actions should a solutions architect take to support these needs?

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

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Amazon Elastic File System

Amazon Elastic File System (Amazon EFS) provides a simple, scalable, fully managed elastic NFS file system for use with AWS Cloud services and on-premises resources. It is built to scale on demand to petabytes without disrupting applications, growing and shrinking automatically as you add and remove files, eliminating the need to provision and manage capacity to accommodate growth.

Amazon EFS is designed to provide massively parallel shared access to thousands of Amazon EC2 instances, enabling your applications to achieve high levels of aggregate throughput and IOPS with consistent low latencies.

Amazon EFS is well suited to support a broad spectrum of use cases from home directories to business-critical applications. Customers can use EFS to lift-and-shift existing enterprise applications to the AWS Cloud. Other use cases include: big data analytics, web serving and content management, application development and testing, media and entertainment workflows, database backups, and container storage.

Amazon EFS is a regional service storing data within and across multiple Availability Zones (AZs) for high availability and durability. Amazon EC2 instances can access your file system across AZs, regions, and VPCs, while on-premises servers can access using AWS Direct Connect or AWS VPN.

Reference:

<https://aws.amazon.com/efs/>

QUESTION 13

A company is migrating a three-tier application to AWS. The application requires a MySQL database. In the past, the application users reported poor application performance when creating new entries. These performance issues were caused by users generating different real-time reports from the application during working hours.

Which solution will improve the performance of the application when it is moved to AWS?

- A. Import the data into an Amazon DynamoDB table with provisioned capacity. Refactor the application to use DynamoDB for reports.
- B. Create the database on a compute optimized Amazon EC2 instance. Ensure compute resources exceed the on-premises database.
- C. Create an Amazon Aurora MySQL Multi-AZ DB cluster with multiple read replicas. Configure the application to use the reader endpoint for reports.
- D. Create an Amazon Aurora MySQL Multi-AZ DB cluster. Configure the application to use the backup instance of the cluster as an endpoint for the reports.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Amazon RDS Read Replicas Now Support Multi-AZ Deployments

Starting today, Amazon RDS Read Replicas for MySQL and MariaDB now support Multi-AZ deployments. Combining Read Replicas with Multi-AZ enables you to build a resilient disaster recovery strategy and simplify your database engine upgrade process.

Amazon RDS Read Replicas enable you to create one or more read-only copies of your database instance within the same AWS Region or in a different AWS Region. Updates made to the source database are then asynchronously copied to your Read Replicas. In addition to providing scalability for read-heavy workloads, Read Replicas can be promoted to become a standalone database instance when needed.

Amazon RDS Multi-AZ deployments provide enhanced availability for database instances within a single AWS Region. With Multi-AZ, your data is synchronously replicated to a standby in a different Availability Zone (AZ). In the event of an infrastructure failure, Amazon RDS performs an automatic failover to the standby, minimizing disruption to your applications.

You can now use Read Replicas with Multi-AZ as part of a disaster recovery (DR) strategy for your production databases. A well-designed and tested DR plan is critical for maintaining business continuity after a disaster. A Read Replica in a different region than the source database can be used as a standby database and promoted to become the new production database in case of a regional disruption.

You can also combine Read Replicas with Multi-AZ for your database engine upgrade process.

You can create a Read Replica of your production database instance and upgrade it to a new database engine version. When the upgrade is complete, you can stop applications, promote the Read Replica to a standalone database instance, and switch over your applications. Since the database instance is already a Multi-AZ deployment, no additional steps are needed.

Overview of Amazon RDS Read Replicas

Deploying one or more read replicas for a given source DB instance might make sense in a variety of scenarios, including the following:

Scaling beyond the compute or I/O capacity of a single DB instance for read-heavy database workloads. You can direct this excess read traffic to one or more read replicas.

Serving read traffic while the source DB instance is unavailable. In some cases, your source DB instance might not be able to take I/O requests, for example due to I/O suspension for backups or scheduled maintenance. In these cases, you

can direct read traffic to your read replicas. For this use case, keep in mind that the data on the read replica might be "stale" because the source DB instance is unavailable. Business reporting or data warehousing scenarios where you might want business reporting queries to run against a read replica, rather than your primary, production DB instance. Implementing disaster recovery. You can promote a read replica to a standalone instance as a disaster recovery solution if the source DB instance fails.

Reference:

<https://aws.amazon.com/about-aws/whats-new/2018/01/amazon-rds-read-replicas-now-support-multi-az-deployments/>

https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_ReadRepl.html

QUESTION 14

A solutions architect is deploying a distributed database on multiple Amazon EC2 instances. The database stores all data on multiple instances so it can withstand the loss of an instance. The database requires block storage with latency and throughput to support several million transactions per second per server.

Which storage solution should the solutions architect use?

- A. EBS Amazon Elastic Block Store (Amazon EBS)
- B. Amazon EC2 instance store
- C. Amazon Elastic File System (Amazon EFS)
- D. Amazon S3

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 15

Organizers for a global event want to put daily reports online as static HTML pages. The pages are expected to generate millions of views from users around the world. The files are stored in an Amazon S3 bucket. A solutions architect has been asked to design an efficient and effective solution.

Which action should the solutions architect take to accomplish this?

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



Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Using Amazon S3 Origins, MediaPackage Channels, and Custom Origins for Web Distributions Using Amazon S3 Buckets for Your Origin When you use Amazon S3 as an origin for your distribution, you place any objects that you want CloudFront to deliver in an Amazon S3 bucket. You can use any method that is supported by Amazon S3 to get your objects into Amazon S3, for example, the Amazon S3 console or API, or a third-party tool. You can create a hierarchy in your bucket to store the objects, just as you would with any other Amazon S3 bucket.

Using an existing Amazon S3 bucket as your CloudFront origin server doesn't change the bucket in any way; you can still use it as you normally would to store and access Amazon S3 objects at the standard Amazon S3 price. You incur regular Amazon S3 charges for storing the objects in the bucket.

Using Amazon S3 Buckets Configured as Website Endpoints for Your Origin You can set up an Amazon S3 bucket that is configured as a website endpoint as custom origin with CloudFront.

When you configure your CloudFront distribution, for the origin, enter the Amazon S3 static website hosting endpoint for your bucket. This value appears in the Amazon S3 console, on the Properties tab, in the Static website hosting pane.

For example: <http://bucket-name.s3-website-region.amazonaws.com> For more information about specifying Amazon S3 static website endpoints, see Website endpoints in the Amazon Simple Storage Service Developer Guide.

When you specify the bucket name in this format as your origin, you can use Amazon S3 redirects and Amazon S3 custom error documents. For more information about Amazon S3 features, see the Amazon S3 documentation.

Using an Amazon S3 bucket as your CloudFront origin server doesn't change it in any way. You can still use it as you normally would and you incur regular Amazon S3 charges.

Reference:

<https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/DownloadDistS3AndCustomOrigins.html>

QUESTION 16

A solutions architect is designing a new service behind Amazon API Gateway. The request patterns for the service will be unpredictable and can change suddenly from 0 requests to over 500 per second. The total size of the data that needs to be persisted in a backend database is currently less than 1 GB with unpredictable future growth. Data can be queried using simple keyvalue requests.

Which combination of AWS services would meet these requirements? (Choose two.)

- A. AWS Fargate
- B. AWS Lambda

- C. Amazon DynamoDB
- D. Amazon EC2 Auto Scaling
- E. MySQL-compatible Amazon Aurora

Correct Answer: BC

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Reference:

<https://aws.amazon.com/about-aws/whats-new/2017/11/amazon-api-gateway-supports-endpoint-integrations-with-private-vpcs>

QUESTION 17

A start-up company has a web application based in the us-east-1 Region with multiple Amazon EC2 instances running behind an Application Load Balancer across multiple Availability Zones. As the company's user base grows in the us-west-1 Region, it needs a solution with low latency and high availability. What should a solutions architect do to accomplish this?

- A. Provision EC2 instances in us-west-1. Switch the Application Load Balancer to a Network Load Balancer to achieve cross-Region load balancing.
- B. Provision EC2 instances and an Application Load Balancer in us-west-1. Make the load balancer distribute the traffic based on the location of the request.
- C. Provision EC2 instances and configure an Application Load Balancer in us-west-1. Create an accelerator in AWS Global Accelerator that uses an endpoint group that includes the load balancer endpoints in both Regions.
- D. Provision EC2 instances and configure an Application Load Balancer in us-west-1. Configure Amazon Route 53 with a weighted routing policy. Create alias records in Route 53 that point to the Application Load Balancer.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Register endpoints for endpoint groups: You register one or more regional resources, such as Application Load Balancers, Network Load Balancers, EC2 Instances, or Elastic IP addresses, in each endpoint group. Then you can set weights to choose how much traffic is routed to each endpoint.

Endpoints in AWS Global Accelerator

Endpoints in AWS Global Accelerator can be Network Load Balancers, Application Load Balancers, Amazon EC2 instances, or Elastic IP addresses. A static IP address serves as a single point of contact for clients, and Global Accelerator then distributes incoming traffic across healthy endpoints. Global Accelerator directs traffic to endpoints by using the port (or port range) that you specify for the listener that the endpoint group for the endpoint belongs to.

Each endpoint group can have multiple endpoints. You can add each endpoint to multiple endpoint groups, but the endpoint groups must be associated with different listeners.

Global Accelerator continually monitors the health of all endpoints that are included in an endpoint group. It routes traffic only to the active endpoints that are healthy. If Global Accelerator doesn't have any healthy endpoints to route traffic to, it routes traffic to all endpoints.

Reference:

<https://docs.aws.amazon.com/global-accelerator/latest/dg/about-endpoints.html>

<https://aws.amazon.com/global-accelerator/faqs/>

QUESTION 18

A solutions architect is designing a solution to access a catalog of images and provide users with the ability to submit requests to customize images. Image customization parameters will be in any request sent to an AWS API Gateway API. The customized image will be generated on demand, and users will receive a link they can click to view or download their customized image. The solution must be highly available for viewing and customizing images. What is the MOST cost-effective solution to meet these requirements?

- A. Use Amazon EC2 instances to manipulate the original image into the requested customizations.
Store the original and manipulated images in Amazon S3. Configure an Elastic Load Balancer in front of the EC2 instances.
- B. Use AWS Lambda to manipulate the original image to the requested customizations. Store the original and manipulated images in Amazon S3. Configure an Amazon CloudFront distribution with the S3 bucket as the origin.
- C. Use AWS Lambda to manipulate the original image to the requested customizations. Store the original images in Amazon S3 and the manipulated images in Amazon DynamoDB. Configure an Elastic Load Balancer in front of the Amazon EC2 instances.
- D. Use Amazon EC2 instances to manipulate the original image into the requested customizations.
Store the original images in Amazon S3 and the manipulated images in Amazon DynamoDB.
Configure an Amazon CloudFront distribution with the S3 bucket as the origin.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

AWS Lambda is a compute service that lets you run code without provisioning or managing servers. AWS Lambda executes your code only when needed and scales automatically, from a few requests per day to thousands per second. You pay only for the compute time you consume – there is no charge when your code is not running. With AWS Lambda, you can run code for virtually any type of application or backend service – all with zero administration. AWS Lambda runs your code on a high-availability compute infrastructure and performs all of the administration of the compute resources, including server and operating system maintenance, capacity provisioning and automatic scaling, code monitoring and logging. All you need to do is supply your code in one of the languages that AWS Lambda supports.

Storing your static content with S3 provides a lot of advantages. But to help optimize your application's performance and security while effectively managing cost, we recommend that you also set up Amazon CloudFront to work with your S3 bucket to serve and protect the content.

CloudFront is a content delivery network (CDN) service that delivers static and dynamic web content, video streams, and APIs around the world, securely and at scale. By design, delivering data out of CloudFront can be more cost effective than delivering it from S3 directly to your users.

CloudFront serves content through a worldwide network of data centers called Edge Locations.

Using edge servers to cache and serve content improves performance by providing content closer to where viewers are located. CloudFront has edge servers in locations all around the world.

Reference:

<https://docs.aws.amazon.com/lambda/latest/dg/welcome.html>

<https://aws.amazon.com/blogs/networking-and-content-delivery/amazon-s3-amazon-cloudfront-amatch-made-in-the-cloud/>

QUESTION 19

A company is planning to migrate a business-critical dataset to Amazon S3. The current solution design uses a single S3 bucket in the us-east-1 Region with versioning enabled to store the dataset. The company's disaster recovery policy states that all data multiple AWS Regions.

How should a solutions architect design the S3 solution?

- A. Create an additional S3 bucket in another Region and configure cross-Region replication.
- B. Create an additional S3 bucket in another Region and configure cross-origin resource sharing (CORS).
- C. Create an additional S3 bucket with versioning in another Region and configure cross-Region replication.
- D. Create an additional S3 bucket with versioning in another Region and configure cross-origin resource (CORS).

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Reference:

<https://medium.com/@KerrySheldon/s3-exercise-2-4-adding-objects-to-an-s3-bucket-with-crossregion-replication-a78b332b7697>



QUESTION 20

A company has application running on Amazon EC2 instances in a VPC. One of the applications needs to call an Amazon S3 API to store and read objects. The company's security policies restrict any internet-bound traffic from the applications.

Which action will fulfill these requirements and maintain security?

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

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 21

A company's web application uses an Amazon RDS PostgreSQL DB instance to store its application data. During the financial closing period at the start of every month, Accountants run large queries that impact the database's performance due to high usage. The company wants to minimize the impact that the reporting activity has on the web application.

What should a solutions architect do to reduce the impact on the database with the LEAST amount of effort?

- A. Create a read replica and direct reporting traffic to the replica.
- B. Create a Multi-AZ database and direct reporting traffic to the standby.
- C. Create a cross-Region read replica and direct reporting traffic to the replica.
- D. Create an Amazon Redshift database and direct reporting traffic to the Amazon Redshift database.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:

Explanation:
Amazon RDS uses the MariaDB, MySQL, Oracle, PostgreSQL, and Microsoft SQL Server DB engines' built-in replication functionality to create a special type of DB instance called a read replica from a source DB instance. Updates made to the source DB instance are asynchronously copied to the read replica. You can reduce the load on your source DB instance by routing read queries from your applications to the read replica.
When you create a read replica, you first specify an existing DB instance as the source. Then Amazon RDS takes a snapshot of the source instance and creates a read-only instance from the snapshot. Amazon RDS then uses the asynchronous replication method for the DB engine to update the read replica whenever there is a change to the source DB instance. The read replica operates as a DB instance that allows only read-only connections. Applications connect to a read replica the same way they do to any DB instance. Amazon RDS replicates all databases in the source DB instance.
Reference:
https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_ReadRepl.html

QUESTION 22

A company wants to migrate a high performance computing (HPC) application and data from onpremises to the AWS Cloud. The company uses tiered storage on premises with hot highperformance parallel storage to support the application during periodic runs of the application, and more economical cold storage to hold the data when the application is not actively running.
Which combination of solutions should a solutions architect recommend to support the storage needs of the application? (Choose two.)

- A. Amazon S3 for cold data storage
- B. Amazon Elastic File System (Amazon EFS) for cold data storage
- C. Amazon S3 for high-performance parallel storage
- D. Amazon FSx for Lustre for high-performance parallel storage
- E. Amazon FSx for Windows for high-performance parallel storage

Correct Answer: AD
Section: (none)
Explanation

Explanation/Reference:

Explanation:



QUESTION 23

A company's application is running on Amazon EC2 instances in a single Region. In the event of a disaster, a solutions architect needs to ensure that the resources can also be deployed to a second Region.
Which combination of actions should the solutions architect take to accomplish this? (Choose two.)

- A. Detach a volume on an EC2 instance and copy it to Amazon S3.
- B. Launch a new EC2 instance from an Amazon Machine Image (AMI) in a new Region.
- C. Launch a new EC2 instance in a new Region and copy a volume from Amazon S3 to the new instance.
- D. Copy an Amazon Machine Image (AMI) of an EC2 instance and specify a different Region for the destination.
- E. Copy an Amazon Elastic Block Store (Amazon EBS) volume from Amazon S3 and launch an EC2 instance in the destination Region using that EBS volume.

Correct Answer: BD
Section: (none)
Explanation

Explanation/Reference:

Explanation:
Cross Region EC2 AMI Copy
We know that you want to build applications that span AWS Regions and we're working to provide you with the services and features needed to do so. We started out by launching the EBS Snapshot Copy feature late last year. This feature gave you the ability to copy a snapshot from Region to Region with just a couple of clicks. In addition, last month we made a significant reduction (26% to 83%) in the cost of transferring data between AWS Regions, making it less expensive to operate in more than one AWS region.
Today we are introducing a new feature: Amazon Machine Image (AMI) Copy. AMI Copy enables you to easily copy your Amazon Machine Images between AWS Regions. AMI Copy helps enable several key scenarios including:
Simple and Consistent Multi-Region Deployment – You can copy an AMI from one region to another, enabling you to easily launch consistent instances based on the same AMI into different regions.
Scalability – You can more easily design and build world-scale applications that meet the needs of your users, regardless of their location.
Performance – You can increase performance by distributing your application and locating critical components of your application in closer proximity to your users. You can also take advantage of region-specific features such as instance types or other AWS services.
Even Higher Availability – You can design and deploy applications across AWS regions, to increase availability.
Once the new AMI is in an Available state the copy is complete.

Reference:

<https://aws.amazon.com/blogs/aws/ec2-ami-copy-between-regions/>

QUESTION 24

A solutions architect needs to ensure that API calls to Amazon DynamoDB from Amazon EC2 instances in a VPC do not traverse the internet. What should the solutions architect do to accomplish this? (Choose two.)

- A. Create a route table entry for the endpoint.
- B. Create a gateway endpoint for DynamoDB.
- C. Create a new DynamoDB table that uses the endpoint.
- D. Create an ENI for the endpoint in each of the subnets of the VPC.
- E. Create a security group entry in the default security group to provide access.

Correct Answer: AB

Section: (none)

Explanation

Explanation/Reference:

Explanation:

A VPC endpoint enables you to privately connect your VPC to supported AWS services and VPC endpoint services powered by AWS PrivateLink without requiring an internet gateway, NAT device, VPN connection, or AWS Direct Connect connection. Instances in your VPC do not require public IP addresses to communicate with resources in the service. Traffic between your VPC and the other service does not leave the Amazon network.

Gateway endpoints

A gateway endpoint is a gateway that you specify as a target for a route in your route table for traffic destined to a supported AWS service. The following AWS services are supported:

Amazon S3

DynamoDB

Reference:

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

QUESTION 25

A company's legacy application is currently relying on a single-instance Amazon RDS MySQL database without encryption. Due to new compliance requirements, all existing and new data in this database must be encrypted. How should this be accomplished?

- A. Create an Amazon S3 bucket with server-side encryption enabled. Move all the data to Amazon S3. Delete the RDS instance.
- B. Enable RDS Multi-AZ mode with encryption at rest enabled. Perform a failover to the standby instance to delete the original instance.
- C. Take a Snapshot of the RDS instance. Create an encrypted copy of the snapshot. Restore the RDS instance from the encrypted snapshot.
- D. Create an RDS read replica with encryption at rest enabled. Promote the read replica to master and switch the application over to the new master. Delete the old RDS instance.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

How do I encrypt Amazon RDS snapshots?

The following steps are applicable to Amazon RDS for MySQL, Oracle, SQL Server, PostgreSQL, or MariaDB.

Important: If you use Amazon Aurora, you can restore an unencrypted Aurora DB cluster snapshot to an encrypted Aurora DB cluster if you specify an AWS Key Management Service (AWS KMS) encryption key when you restore from the unencrypted DB cluster snapshot. For more information, see Limitations of Amazon RDS Encrypted DB Instances.

Open the Amazon RDS console, and then choose Snapshots from the navigation pane.

Select the snapshot that you want to encrypt.

Under Snapshot Actions, choose Copy Snapshot.

Choose your Destination Region, and then enter your New DB Snapshot Identifier.

Change Enable Encryption to Yes.

Select your Master Key from the list, and then choose Copy Snapshot.

After the snapshot status is available, the Encrypted field will be True to indicate that the snapshot is encrypted.

You now have an encrypted snapshot of your DB. You can use this encrypted DB snapshot to restore the DB instance from the DB snapshot.

Reference:

<https://aws.amazon.com/premiumsupport/knowledge-center/encrypt-rds-snapshots/>

QUESTION 26

A manufacturing company wants to implement predictive maintenance on its machinery equipment. The company will install thousands of IoT sensors that will send data to AWS in real time. A solutions architect is tasked with implementing a solution that will receive events in an ordered manner for each machinery asset and ensure that data is saved for further processing at a later time.

Which solution would be MOST efficient?

- A. Use Amazon Kinesis Data Streams for real-time events with a partition for each equipment asset. Use Amazon Kinesis Data Firehose to save data to Amazon S3.
- B. Use Amazon Kinesis Data Streams for real-time events with a shard for each equipment asset. Use Amazon Kinesis Data Firehose to save data to Amazon Elastic Block Store (Amazon EBS).
- C. Use an Amazon SQS FIFO queue for real-time events with one queue for each equipment asset. Trigger an AWS Lambda function for the SQS queue to save data to Amazon Elastic File System (Amazon EFS).
- D. Use an Amazon SQS standard queue for real-time events with one queue for each equipment asset. Trigger an AWS Lambda function from the SQS queue to save data to Amazon S3.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 27

A company's website runs on Amazon EC2 instances behind an Application Load Balancer (ALB). The website has a mix of dynamic and static content. Users around the globe are reporting that the website is slow. Which set of actions will improve website performance for users worldwide?

- A. Create an Amazon CloudFront distribution and configure the ALB as an origin. Then update the Amazon Route 53 record to point to the CloudFront distribution.
- B. Create a latency-based Amazon Route 53 record for the ALB. Then launch new EC2 instances with larger instance sizes and register the instances with the ALB.
- C. Launch new EC2 instances hosting the same web application in different Regions closer to the users. Then register instances with the same ALB using cross-Region VPC peering.
- D. Host the website in an Amazon S3 bucket in the Regions closest to the users and delete the ALB and EC2 instances. Then update an Amazon Route 53 record to point to the S3 buckets.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

What Is Amazon CloudFront?

Amazon CloudFront is a web service that speeds up distribution of your static and dynamic web content, such as .html, .css, .js, and image files, to your users. CloudFront delivers your content through a worldwide network of data centers called edge locations. When a user requests content that you're serving with CloudFront, the user is routed to the edge location that provides the lowest latency (time delay), so that content is delivered with the best possible performance. Routing traffic to an Amazon CloudFront web distribution by using your domain name.

If you want to speed up delivery of your web content, you can use Amazon CloudFront, the AWS content delivery network (CDN). CloudFront can deliver your entire website – including dynamic, static, streaming, and interactive content – by using a global network of edge locations. Requests for your content are automatically routed to the edge location that gives your users the lowest latency.

To use CloudFront to distribute your content, you create a web distribution and specify settings such as the Amazon S3 bucket or HTTP server that you want CloudFront to get your content from, whether you want only selected users to have access to your content, and whether you want to require users to use HTTPS.

When you create a web distribution, CloudFront assigns a domain name to the distribution, such asd111111abcdef8.cloudfront.net. You can use this domain name in the URLs for your content, for example: http://d111111abcdef8.cloudfront.net/logo.jpg Alternatively, you might prefer to use your own domain name in URLs, for example: http://example.com/logo.jpg If you want to use your own domain name, use Amazon Route 53 to create an alias record that points to your CloudFront distribution. An alias record is a Route 53 extension to DNS. It's similar to a CNAME record, but you can create an alias record both for the root domain, such as example.com, and for subdomains, such aswww.example.com. (You can create CNAME records only for subdomains.) When Route 53 receives a DNS query that matches the name and type of an alias record, Route 53 responds with the domain name that is associated with your distribution.

Reference:

<https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/routing-to-cloudfrontdistribution.html>

<https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/Introduction.html>

QUESTION 28

A company has been storing analytics data in an Amazon RDS instance for the past few years.

The company asked a solutions architect to find a solution that allows users to access this data using an API. The expectation is that the application will experience periods of inactivity but could receive bursts of traffic within seconds.

Which solution should the solutions architect suggest?

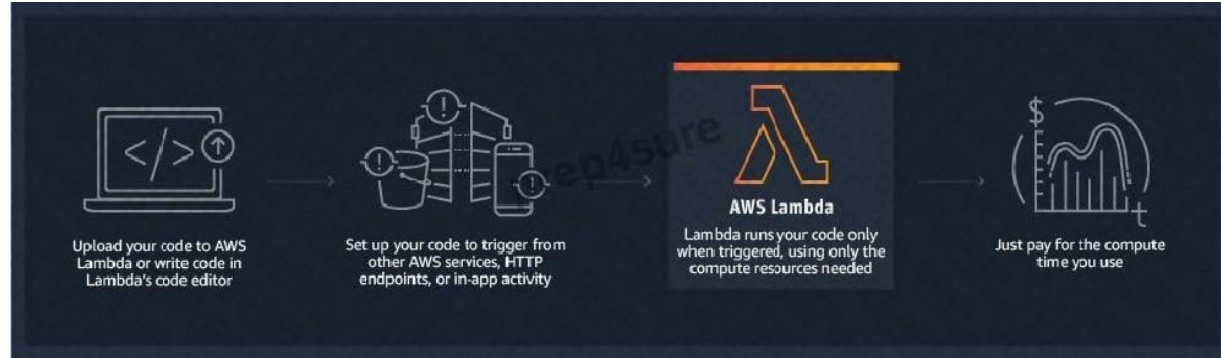
- A. Set up an Amazon API Gateway and use Amazon ECS.
- B. Set up an Amazon API Gateway and use AWS Elastic Beanstalk.
- C. Set up an Amazon API Gateway and use AWS Lambda functions.
- D. Set up an Amazon API Gateway and use Amazon EC2 with Auto Scaling.



Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:

Explanation:
AWS Lambda
With Lambda, you can run code for virtually any type of application or backend service – all with zero administration. Just upload your code and Lambda takes care of everything required to run and scale your code with high availability. You can set up your code to automatically trigger from other AWS services or call it directly from any web or mobile app.
How it works



Amazon API Gateway

Amazon API Gateway is a fully managed service that makes it easy for developers to create, publish, maintain, monitor, and secure APIs at any scale. APIs act as the "front door" for applications to access data, business logic, or functionality from your backend services. Using API Gateway, you can create RESTful APIs and WebSocket APIs that enable real-time two-way communication applications. API Gateway supports containerized and serverless workloads, as well as web applications.

API Gateway handles all the tasks involved in accepting and processing up to hundreds of thousands of concurrent API calls, including traffic management, CORS support, authorization and access control, throttling, monitoring, and API version management. API Gateway has no minimum fees or startup costs. You pay for the API calls you receive and the amount of data transferred out and, with the API Gateway tiered pricing model, you can reduce your cost as your API usage scales.

Reference:

<https://aws.amazon.com/lambda/>
<https://aws.amazon.com/api-gateway/>



QUESTION 29

A company must generate sales reports at the beginning of every month. The reporting process launches 20 Amazon EC2 instances on the first of the month. The process runs for 7 days and cannot be interrupted. The company wants to minimize costs.

Which pricing model should the company choose?

- A. Reserved Instances
- B. Spot Block Instances
- C. On-Demand Instances
- D. Scheduled Reserved Instances

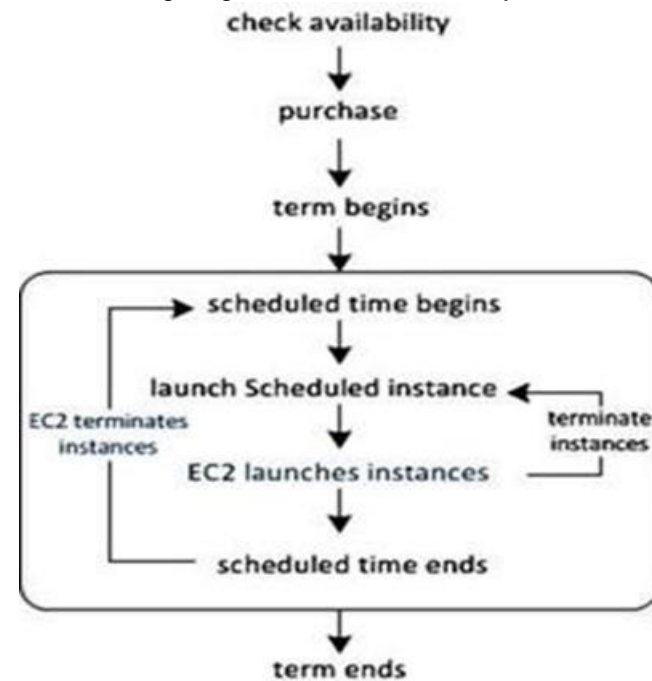
Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:

Explanation:
Scheduled Reserved Instances
Scheduled Reserved Instances (Scheduled Instances) enable you to purchase capacity reservations that recur on a daily, weekly, or monthly basis, with a specified start time and duration, for a one-year term. You reserve the capacity in advance, so that you know it is available when you need it. You pay for the time that the instances are scheduled, even if you do not use them.
Scheduled Instances are a good choice for workloads that do not run continuously, but do run on a regular schedule. For example, you can use Scheduled Instances for an application that runs during business hours or for batch processing that runs at the end of the week.
If you require a capacity reservation on a continuous basis, Reserved Instances might meet your needs and decrease costs.
How Scheduled Instances Work
Amazon EC2 sets aside pools of EC2 instances in each Availability Zone for use as Scheduled Instances. Each pool supports a specific combination of instance type, operating system, and network.
To get started, you must search for an available schedule. You can search across multiple pools or a single pool. After you locate a suitable schedule, purchase it.
You must launch your Scheduled Instances during their scheduled time periods, using a launch configuration that matches the following attributes of the schedule that you purchased: instance type, Availability Zone, network, and platform.
When you do so, Amazon EC2 launches EC2 instances on your behalf, based on the specified launch specification. Amazon EC2 must ensure that the EC2 instances have terminated by the end of the current scheduled time period so that the capacity is available for any other Scheduled Instances it is reserved for. Therefore, Amazon EC2 terminates the EC2 instances three minutes before the end of the current scheduled time period.

You can't stop or reboot Scheduled Instances, but you can terminate them manually as needed. If you terminate a Scheduled Instance before its current scheduled time period ends, you can launch it again after a few minutes. Otherwise, you must wait until the next scheduled time period.

The following diagram illustrates the lifecycle of a Scheduled Instance.



Reference:

<https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ec2-scheduled-instances.html>

QUESTION 30

A gaming company has multiple Amazon EC2 instances in a single Availability Zone for its multiplayer game that communicates with users on Layer 4. The chief technology officer (CTO) wants to make the architecture highly available and cost-effective.

What should a solutions architect do to meet these requirements? (Choose two.)?

- A. Increase the number of EC2 instances.
- B. Decrease the number of EC2 instances.
- C. Configure a Network Load Balancer in front of the EC2 instances.
- D. Configure an Application Load Balancer in front of the EC2 instances.
- E. Configure an Auto Scaling group to add or remove instances in multiple Availability Zones automatically.

Correct Answer: CE

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Network Load Balancer overview

A Network Load Balancer functions at the fourth layer of the Open Systems Interconnection (OSI) model. It can handle millions of requests per second. After the load balancer receives a connection request, it selects a target from the target group for the default rule. It attempts to open a TCP connection to the selected target on the port specified in the listener configuration.

When you enable an Availability Zone for the load balancer, Elastic Load Balancing creates a load balancer node in the Availability Zone. By default, each load balancer node distributes traffic across the registered targets in its Availability Zone only. If you enable cross-zone load balancing, each load balancer node distributes traffic across the registered targets in all enabled Availability Zones. For more information, see Availability Zones.

If you enable multiple Availability Zones for your load balancer and ensure that each target group has at least one target in each enabled Availability Zone, this increases the fault tolerance of your applications. For example, if one or more target groups does not have a healthy target in an Availability Zone, we remove the IP address for the corresponding subnet from DNS, but the load balancer nodes in the other Availability Zones are still available to route traffic. If a client doesn't honor the time-to-live (TTL) and sends requests to the IP address after it is removed from DNS, the requests fail.

For TCP traffic, the load balancer selects a target using a flow hash algorithm based on the protocol, source IP address, source port, destination IP address, destination port, and TCP sequence number. The TCP connections from a client have different source ports and sequence numbers, and can be routed to different targets. Each individual TCP connection is routed to a single target for the life of the connection.

For UDP traffic, the load balancer selects a target using a flow hash algorithm based on the protocol, source IP address, source port, destination IP address, and destination port. A UDP flow has the same source and destination, so it is consistently routed to a single target throughout its lifetime. Different UDP flows have different source IP addresses and ports, so they can be routed to different targets.

An Auto Scaling group contains a collection of Amazon EC2 instances that are treated as a logical grouping for the purposes of automatic scaling and management. An Auto Scaling group also enables you to use Amazon EC2 Auto Scaling features such as health check replacements and scaling policies. Both maintaining the number of instances in an Auto Scaling group and automatic scaling are the core functionality of the Amazon EC2 Auto Scaling service.

The size of an Auto Scaling group depends on the number of instances that you set as the desired capacity. You can adjust its size to meet demand, either manually or by using automatic scaling.

An Auto Scaling group starts by launching enough instances to meet its desired capacity. It maintains this number of instances by performing periodic health checks on the instances in the group. The Auto Scaling group continues to

maintain a fixed number of instances even if an instance becomes unhealthy. If an instance becomes unhealthy, the group terminates the unhealthy instance and launches another instance to replace it.

Reference:

<https://docs.aws.amazon.com/elasticloadbalancing/latest/network/introduction.html>

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

QUESTION 31

A company currently operates a web application backed by an Amazon RDS MySQL database. It has automated backups that are run daily and are not encrypted. A security audit requires future backups to be encrypted and the unencrypted backups to be destroyed. The company will make at least one encrypted backup before destroying the old backups.

What should be done to enable encryption for future backups?

- A. Enable default encryption for the Amazon S3 bucket where backups are stored.
- B. Modify the backup section of the database configuration to toggle the Enable encryption check box.
- C. Create a snapshot of the database. Copy it to an encrypted snapshot. Restore the database from the encrypted snapshot.
- D. Enable an encrypted read replica on RDS for MySQL. Promote the encrypted read replica to primary. Remove the original database instance.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

However, because you can encrypt a copy of an unencrypted DB snapshot, you can effectively add encryption to an unencrypted DB instance. That is, you can create a snapshot of your DB instance, and then create an encrypted copy of that snapshot. You can then restore a DB instance from the encrypted snapshot, and thus you have an encrypted copy of your original DB instance.

DB instances that are encrypted can't be modified to disable encryption.

You can't have an encrypted read replica of an unencrypted DB instance or an unencrypted read replica of an encrypted DB instance.

Encrypted read replicas must be encrypted with the same key as the source DB instance when both are in the same AWS Region.

You can't restore an unencrypted backup or snapshot to an encrypted DB instance.

To copy an encrypted snapshot from one AWS Region to another, you must specify the KMS key identifier of the destination AWS Region. This is because KMS encryption keys are specific to the AWS Region that they are created in.

Reference:

<https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/Overview.Encryption.html>



QUESTION 32

A company is hosting a website behind multiple Application Load Balancers. The company has different distribution rights for its content around the world. A solutions architect needs to ensure that users are served the correct content without violating distribution rights.

Which configuration should the solutions architect choose to meet these requirements?

- A. Configure Amazon CloudFront with AWS WAF.
- B. Configure Application Load Balancers with AWS WAF.
- C. Configure Amazon Route 53 with a geolocation policy.
- D. Configure Amazon Route 53 with a geoproximity routing policy.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Reference:

<https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/routing-policy.html> (geolocationrouting)

QUESTION 33

A solutions architect has created a new AWS account and must secure AWS account root user access.

Which combination of actions will accomplish this? (Choose two.)

- A. Ensure the root user uses a strong password.
- B. Enable multi-factor authentication to the root user.
- C. Store root user access keys in an encrypted Amazon S3 bucket.
- D. Add the root user to a group containing administrative permissions.
- E. Apply the required permissions to the root user with an inline policy document.

Correct Answer: AB

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 34

A solutions architect at an ecommerce company wants to back up application log data to Amazon S3. The solutions architect is unsure how frequently the logs will be accessed or which logs will be accessed the most. The company wants to keep costs as low as possible by using the appropriate S3 storage class.

Which S3 storage class should be implemented to meet these requirements?

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

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

S3 Intelligent-Tiering

S3 Intelligent-Tiering is a new Amazon S3 storage class designed for customers who want to optimize storage costs automatically when data access patterns change, without performance impact or operational overhead. S3 Intelligent-Tiering is the first cloud object storage class that delivers automatic cost savings by moving data between two access tiers – frequent access and infrequent access – when access patterns change, and is ideal for data with unknown or changing access patterns.

S3 Intelligent-Tiering stores objects in two access tiers: one tier that is optimized for frequent access and another lower-cost tier that is optimized for infrequent access. For a small monthly monitoring and automation fee per object, S3 Intelligent-Tiering monitors access patterns and moves objects that have not been accessed for 30 consecutive days to the infrequent access tier.

There are no retrieval fees in S3 Intelligent-Tiering. If an object in the infrequent access tier is accessed later, it is automatically moved back to the frequent access tier. No additional tiering fees apply when objects are moved between access tiers within the S3 Intelligent-Tiering storage class. S3 Intelligent-Tiering is designed for 99.9% availability and 99.99999999% durability, and offers the same low latency and high throughput performance of S3 Standard.

Reference:

<https://aws.amazon.com/about-aws/whats-new/2018/11/s3-intelligent-tiering/>

QUESTION 35

A company's website is used to sell products to the public. The site runs on Amazon EC2 instances in an Auto Scaling group behind an Application Load Balancer (ALB). There is also an Amazon CloudFront distribution, and AWS WAF is being used to protect against SQL injection attacks. The ALB is the origin for the CloudFront distribution. A recent review of security logs revealed an external malicious IP that needs to be blocked from accessing the website.

What should a solutions architect do to protect the application?

- A. Modify the network ACL on the CloudFront distribution to add a deny rule for the malicious IP address.
- B. Modify the configuration of AWS WAF to add an IP match condition to block the malicious IP address.
- C. Modify the network ACL for the EC2 instances in the target groups behind the ALB to deny the malicious IP address.
- D. Modify the security groups for the EC2 instances in the target groups behind the ALB to deny the malicious IP address.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

If you want to allow or block web requests based on the IP addresses that the requests originate from, create one or more IP match conditions. An IP match condition lists up to 10,000 IP addresses or IP address ranges that your requests originate from. Later in the process, when you create a web ACL, you specify whether to allow or block requests from those IP addresses.

AWS Web Application Firewall (WAF) – Helps to protect your web applications from common application-layer exploits that can affect availability or consume excessive resources. As you can see in my post (New – AWS WAF), WAF allows you to use access control lists (ACLs), rules, and conditions that define acceptable or unacceptable requests or IP addresses. You can selectively allow or deny access to specific parts of your web application and you can also guard against various SQL injection attacks. We launched WAF with support for Amazon CloudFront.

Reference:

<https://aws.amazon.com/blogs/aws/aws-web-application-firewall-waf-for-applicationloadbalancers/>

<https://docs.aws.amazon.com/waf/latest/developerguide/classic-web-acl-ip-conditions.html>

<https://docs.aws.amazon.com/waf/latest/developerguide/classic-web-acl-ip-conditions.html>

<https://aws.amazon.com/blogs/aws/aws-web-application-firewall-waf-for-application-loadbalancers/>

QUESTION 36

A solutions architect is designing an application for a two-step order process. The first step is synchronous and must return to the user with little latency. The second step takes longer, so it will be implemented in a separate component. Orders must be processed exactly once and in the order in which they are received. How should the solutions architect integrate these components?

- A. Use Amazon SQS FIFO queues.
- B. Use an AWS Lambda function along with Amazon SQS standard queues.
- C. Create an SNS topic and subscribe an Amazon SQS FIFO queue to that topic.
- D. Create an SNS topic and subscribe an Amazon SQS Standard queue to that topic.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Reference:

<https://docs.aws.amazon.com/AWSSimpleQueueService/latest/SQSDeveloperGuide/FIFOqueues.html>

QUESTION 37

A web application is deployed in the AWS Cloud. It consists of a two-tier architecture that includes a web layer and a database layer. The web server is vulnerable to cross-site scripting (XSS) attacks. What should a solutions architect do to remediate the vulnerability?

- A. Create a Classic Load Balancer. Put the web layer behind the load balancer and enable AWS WAF.
- B. Create a Network Load Balancer. Put the web layer behind the load balancer and enable AWS WAF.
- C. Create an Application Load Balancer. Put the web layer behind the load balancer and enable AWS WAF.
- D. Create an Application Load Balancer. Put the web layer behind the load balancer and use AWS Shield Standard.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Working with cross-site scripting match conditions

Attackers sometimes insert scripts into web requests in an effort to exploit vulnerabilities in web applications. You can create one or more cross-site scripting match conditions to identify the parts of web requests, such as the URI or the query string, that you want AWS WAF Classic to inspect for possible malicious scripts. Later in the process, when you create a web ACL, you specify whether to allow or block requests that appear to contain malicious scripts.

Web Application Firewall

You can now use AWS WAF to protect your web applications on your Application Load Balancers.

AWS WAF is a web application firewall that helps protect your web applications from common web exploits that could affect application availability, compromise security, or consume excessive resources.

Reference:

<https://docs.aws.amazon.com/waf/latest/developerguide/classic-web-acl-xss-conditions.html>

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

QUESTION 38

A company's website is using an Amazon RDS MySQL Multi-AZ DB instance for its transactional data storage. There are other internal systems that query this DB instance to fetch data for internal batch processing. The RDS DB instance slows down significantly when the internal systems fetch data. This impacts the website's read and write performance, and the users experience slow response times.

Which solution will improve the website's performance?

- A. Use an RDS PostgreSQL DB instance instead of a MySQL database.
- B. Use Amazon ElastiCache to cache the query responses for the website.
- C. Add an additional Availability Zone to the current RDS MySQL Multi-AZ DB instance.
- D. Add a read replica to the RDS DB instance and configure the internal systems to query the read replica.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Amazon RDS Read Replicas

Enhanced performance



You can reduce the load on your source DB instance by routing read queries from your applications to the read replica. Read replicas allow you to elastically scale out beyond the capacity constraints of a single DB instance for read-heavy database workloads. Because read replicas can be promoted to master status, they are useful as part of a sharding implementation. To further maximize read performance, Amazon RDS for MySQL allows you to add table indexes directly to Read Replicas, without those indexes being present on the master.

Reference:

<https://aws.amazon.com/rds/features/read-replicas>

QUESTION 39

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

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

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

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

"With target tracking scaling policies, you select a scaling metric and set a target value. Amazon EC2 AutoScaling creates and manages the CloudWatch alarms that trigger the scaling policy and calculates the scaling adjustment based on the metric and the target value. The scaling policy adds or removes capacity as required to keep the metric at, or close to, the specified target value. In addition to keeping the metric close to the target value, a target tracking scaling policy also adjusts to changes in the metric due to a changing load pattern. For example, you can use target tracking scaling to: Configure a target tracking scaling policy to keep the average aggregate CPU utilization of your Auto Scaling group at 40 percent. Configure a target tracking scaling policy to keep the request count per target of your Application Load Balancer target group at 1000 for your AutoScaling group."

Reference:

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

QUESTION 40

A company runs an internal browser-based application. The application runs on Amazon EC2 instances behind an Application Load Balancer. The instances run in an Amazon EC2 Auto Scaling group across multiple Availability Zones. The Auto Scaling group scales up to 20 instances during work hours, but scales down to 2 instances overnight. Staff are complaining that the application is very slow when the day begins, although it runs well by mid-morning.

How should the scaling be changed to address the staff complaints and keep costs to a minimum?

- A. Implement a scheduled action that sets the desired capacity to 20 shortly before the office opens.
- B. Implement a step scaling action triggered at a lower CPU threshold, and decrease the cooldown period.
- C. Implement a target tracking action triggered at a lower CPU threshold, and decrease the cooldown period.
- D. Implement a scheduled action that sets the minimum and maximum capacity to 20 shortly before the office opens.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Reference: <https://docs.aws.amazon.com/autoscaling/ec2/userguide/as-scaling-simple-step.html>

QUESTION 41

A financial services company has a web application that serves users in the United States and Europe. The application consists of a database tier and a web server tier. The database tier consists of a MySQL database hosted in us-east-1. Amazon Route 53 geoproximity routing is used to direct traffic to instances in the closest Region. A performance review of the system reveals that European users are not receiving the same level of query performance as those in the United States.

Which changes should be made to the database tier to improve performance?

- A. Migrate the database to Amazon RDS for MySQL. Configure Multi-AZ in one of the European Regions.
- B. Migrate the database to Amazon DynamoDB. Use DynamoDB global tables to enable replication to additional Regions.
- C. Deploy MySQL instances in each Region. Deploy an Application Load Balancer in front of MySQL to reduce the load on the primary instance.
- D. Migrate the database to an Amazon Aurora global database in MySQL compatibility mode.
Configure read replicas in one of the European Regions.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 42

A company hosts a static website on-premises and wants to migrate the website to AWS. The website should load as quickly as possible for users around the world. The company also wants the most cost-effective solution. What should a solutions architect do to accomplish this?

- A. Copy the website content to an Amazon S3 bucket. Configure the bucket to serve static webpage content. Replicate the S3 bucket to multiple AWS Regions.
- B. Copy the website content to an Amazon S3 bucket. Configure the bucket to serve static webpage content. Configure Amazon CloudFront with the S3 bucket as the origin.
- C. Copy the website content to an Amazon EBS-backed Amazon EC2 instance running Apache HTTP Server. Configure Amazon Route 53 geolocation routing policies to select the closest origin.
- D. Copy the website content to multiple Amazon EBS-backed Amazon EC2 instances running Apache HTTP Server in multiple AWS Regions. Configure Amazon CloudFront geolocation routing policies to select the closest origin.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

What Is Amazon CloudFront?

Amazon CloudFront is a web service that speeds up distribution of your static and dynamic web content, such as .html, .css, .js, and image files, to your users. CloudFront delivers your content through a worldwide network of data centers called edge locations. When a user requests content that you're serving with CloudFront, the user is routed to the edge location that provides the lowest latency (time delay), so that content is delivered with the best possible performance.

Using Amazon S3 Buckets for Your Origin

When you use Amazon S3 as an origin for your distribution, you place any objects that you want CloudFront to deliver in an Amazon S3 bucket. You can use any method that is supported by Amazon S3 to get your objects into Amazon S3, for example, the Amazon S3 console or API, or a third-party tool. You can create a hierarchy in your bucket to store the objects, just as you would with any other Amazon S3 bucket.

Using an existing Amazon S3 bucket as your CloudFront origin server doesn't change the bucket in any way; you can still use it as you normally would to store and access Amazon S3 objects at the standard Amazon S3 price. You incur regular Amazon S3 charges for storing the objects in the bucket.

Reference:

<https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/Introduction.html>

<https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/DownloadDistS3AndCustomOrigins.html>



QUESTION 43

A solutions architect is designing storage for a high performance computing (HPC) environment based on Amazon Linux. The workload stores and processes a large amount of engineering drawings that require shared storage and heavy computing.

Which storage option would be the optimal solution?

- A. Amazon Elastic File System (Amazon EFS)
- B. Amazon FSx for Lustre
- C. Amazon EC2 instance store
- D. Amazon Elastic Block Store (Amazon EBS) Provisioned IOPS SSD (io1)

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Amazon FSx for Lustre

Amazon FSx for Lustre is a new, fully managed service provided by AWS based on the Lustre file system. Amazon FSx for Lustre provides a high-performance file system optimized for fast processing of workloads such as machine learning, high performance computing (HPC), video processing, financial modeling, and electronic design automation (EDA).

FSx for Lustre allows customers to create a Lustre filesystem on demand and associate it to an Amazon S3 bucket. As part of the filesystem creation, Lustre reads the objects in the buckets and adds that to the file system metadata. Any Lustre client in your VPC is then able to access the data, which gets cached on the high-speed Lustre filesystem. This is ideal for HPC workloads, because you can get the speed of an optimized Lustre file system without having to manage the complexity of deploying, optimizing, and managing the Lustre cluster.

Additionally, having the filesystem work natively with Amazon S3 means you can shut down the Lustre filesystem when you don't need it but still access objects in Amazon S3 via other AWS Services. FSx for Lustre also allows you to also write the output of your HPC job back to Amazon S3.

Reference:

https://d1.awsstatic.com/whitepapers/AWS%20Partner%20Network_HPC%20Storage%20Options_2019_FINAL.pdf (p.8)

QUESTION 44

A company is performing an AWS Well-Architected Framework review of an existing workload deployed on AWS. The review identified a public-facing website running on the same Amazon EC2 instance as a Microsoft Active Directory

domain controller that was installed recently to support other AWS services. A solutions architect needs to recommend a new design that would improve the security of the architecture and minimize the administrative demand on IT staff. What should the solutions architect recommend?

- A. Use AWS Directory Service to create a managed Active Directory. Uninstall Active Directory on the current EC2 instance.
- B. Create another EC2 instance in the same subnet and reinstall Active Directory on it. Uninstall Active Directory.
- C. Use AWS Directory Service to create an Active Directory connector. Proxy Active Directory requests to the Active domain controller running on the current EC2 instance.
- D. Enable AWS Single Sign-On (AWS SSO) with Security Assertion Markup Language (SAML) 2.0 federation with the current Active Directory controller. Modify the EC2 instance's security group to deny public access to Active Directory.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

AWS Managed Microsoft AD

AWS Directory Service lets you run Microsoft Active Directory (AD) as a managed service. AWS Directory Service for Microsoft Active Directory, also referred to as AWS Managed Microsoft AD, is powered by Windows Server 2012 R2.

When you select and launch this directory type, it is created as a highly available pair of domain controllers connected to your virtual private cloud

(VPC). The domain controllers run in different Availability Zones in a region of your choice. Host monitoring and recovery, data replication, snapshots, and software updates are automatically configured and managed for you.

Reference:

https://docs.aws.amazon.com/directoryservice/latest/admin-guide/directory_microsoft_ad.html

QUESTION 45

A company hosts a static website within an Amazon S3 bucket. A solutions architect needs to ensure that data can be recovered in case of accidental deletion.

Which action will accomplish this?

- A. Enable Amazon S3 versioning.
- B. Enable Amazon S3 Intelligent-Tiering.
- C. Enable an Amazon S3 lifecycle policy.
- D. Enable Amazon S3 cross-Region replication.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Data can be recovered if versioning is enabled, also it provides extra protection like file delete, MFA delete. MFA Delete only works for CLI or API interaction, not in the AWS Management Console.

Also, you cannot make version DELETE actions with MFA using IAM user credentials. You must use your root AWS account.

Object Versioning

Use Amazon S3 Versioning to keep multiple versions of an object in one bucket. For example, you could store my-image.jpg (version 111111) and my-image.jpg (version 222222) in a single bucket.

S3 Versioning protects you from the consequences of unintended overwrites and deletions. You can also use it to archive objects so that you have access to previous versions.

You must explicitly enable S3 Versioning on your bucket. By default, S3 Versioning is disabled.

Regardless of whether you have enabled Versioning, each object in your bucket has a version ID.

If you have not enabled Versioning, Amazon S3 sets the value of the version ID to null. If S3 Versioning is enabled, Amazon S3 assigns a version ID value for the object. This value distinguishes it from other versions of the same key.

Reference:

<https://books.google.com.sg/books?id=ww45DQAAQBAJ&pg=PA39&lpg=PA39&dq=hosts+a+static+website+within+an+Amazon+S3+bucket.+A+solutions+architect+needs+to+ensure+that+data+can+be+recovered+in+case+of+accidental+deletion&source=bl&ots=0NoIP5igY5&sig=ACfU3U3opL9Jha6jM2E18x7Eck4rigQHQ&hl=en&sa=X&ved=2ahUKEwiS9e3yy7vpAhVx73MBHZNoDnQQ6AEwAHoECBQQAQ#v=onepage&q=hosts%20a%20static%20website%20within%20an%20Amazon%20S3%20bucket.%20A%20solutions%20architect%20needs%20to%20ensure%20that%20data%20can%20be%20recovered%20in%20case%20of%20accidental%20deletion&f=false>

<https://aws.amazon.com/blogs/security/securing-access-to-aws-using-mfa-part-3/>

<https://docs.aws.amazon.com/AmazonS3/latest/dev/ObjectVersioning.html>

QUESTION 46

A company's production application runs online transaction processing (OLTP) transactions on an Amazon RDS MySQL DB instance. The company is launching a new reporting tool that will access the same data. The reporting tool must be highly available and not impact the performance of the production application.

How can this be achieved?

- A. Create hourly snapshots of the production RDS DB instance.
- B. Create a Multi-AZ RDS Read Replica of the production RDS DB instance.
- C. Create multiple RDS Read Replicas of the production RDS DB instance. Place the Read Replicas in an Auto Scaling group.
- D. Create a Single-AZ RDS Read Replica of the production RDS DB instance. Create a second Single-AZ RDS Read Replica from the replica.



Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:

Explanation:
Amazon RDS Read Replicas Now Support Multi-AZ Deployments
Amazon RDS Read Replicas enable you to create one or more read-only copies of your database instance within the same AWS Region or in a different AWS Region. Updates made to the source database are then asynchronously copied to your Read Replicas. In addition to providing scalability for read-heavy workloads, Read Replicas can be promoted to become a standalone database instance when needed.
Amazon RDS Multi-AZ deployments provide enhanced availability for database instances within a single AWS Region. With Multi-AZ, your data is synchronously replicated to a standby in a different Availability Zone (AZ). In the event of an infrastructure failure, Amazon RDS performs an automatic failover to the standby, minimizing disruption to your applications.
You can now use Read Replicas with Multi-AZ as part of a disaster recovery (DR) strategy for your production databases. A well-designed and tested DR plan is critical for maintaining business continuity after a disaster. A Read Replica in a different region than the source database can be used as a standby database and promoted to become the new production database in case of a regional disruption.

Reference:

<https://aws.amazon.com/about-aws/whats-new/2018/01/amazon-rds-read-replicas-now-support-multi-az-deployments/#:~>

QUESTION 47

A company runs an application in a branch office within a small data closet with no virtualized compute resources. The application data is stored on an NFS volume. Compliance standards require a daily offsite backup of the NFS volume. Which solution meets these requirements?

- A. Install an AWS Storage Gateway file gateway on premises to replicate the data to Amazon S3.
- B. Install an AWS Storage Gateway file gateway hardware appliance on premises to replicate the data to Amazon S3.
- C. Install an AWS Storage Gateway volume gateway with stored volumes on premises to replicate the data to Amazon S3.
- D. Install an AWS Storage Gateway volume gateway with cached volumes on premises to replicate the data to Amazon S3.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:

Explanation:
AWS Storage Gateway Hardware Appliance
Hardware Appliance
Storage Gateway is available as a hardware appliance, adding to the existing support for VMware ESXi, Microsoft Hyper-V, and Amazon EC2. This means that you can now make use of Storage Gateway in situations where you do not have a virtualized environment, server-class hardware or IT staff with the specialized skills that are needed to manage them. You can order appliances from Amazon.com for delivery to branch offices, warehouses, and "outpost" offices that lack dedicated IT resources. Setup (as you will see in a minute) is quick and easy, and gives you access to three storage solutions:
File Gateway – A file interface to Amazon S3, accessible via NFS or SMB. The files are stored as S3 objects, allowing you to make use of specialized S3 features such as lifecycle management and cross-region replication. You can trigger AWS Lambda functions, run Amazon Athena queries, and use Amazon Macie to discover and classify sensitive data.

Reference:

<https://aws.amazon.com/blogs/aws/new-aws-storage-gateway-hardware-appliance/>

<https://aws.amazon.com/storagegateway/file/>



QUESTION 48

A company's web application is using multiple Linux Amazon EC2 instances and storing data on Amazon Elastic Block Store (Amazon EBS) volumes. The company is looking for a solution to increase the resiliency of the application in case of a failure and to provide storage that complies with atomicity, consistency, isolation, and durability (ACID).

What should a solutions architect do to meet these requirements?

- A. Launch the application on EC2 instances in each Availability Zone. Attach EBS volumes to each EC2 instance.
- B. Create an Application Load Balancer with Auto Scaling groups across multiple Availability Zones. Mount an instance store on each EC2 instance.
- C. Create an Application Load Balancer with Auto Scaling groups across multiple Availability Zones. Store data on Amazon Elastic File System (Amazon EFS) and mount a target on each instance.
- D. Create an Application Load Balancer with Auto Scaling groups across multiple Availability Zones. Store data using Amazon S3 One Zone-Infrequent Access (S3 One Zone-IA).

Correct Answer: C
Section: (none)
Explanation

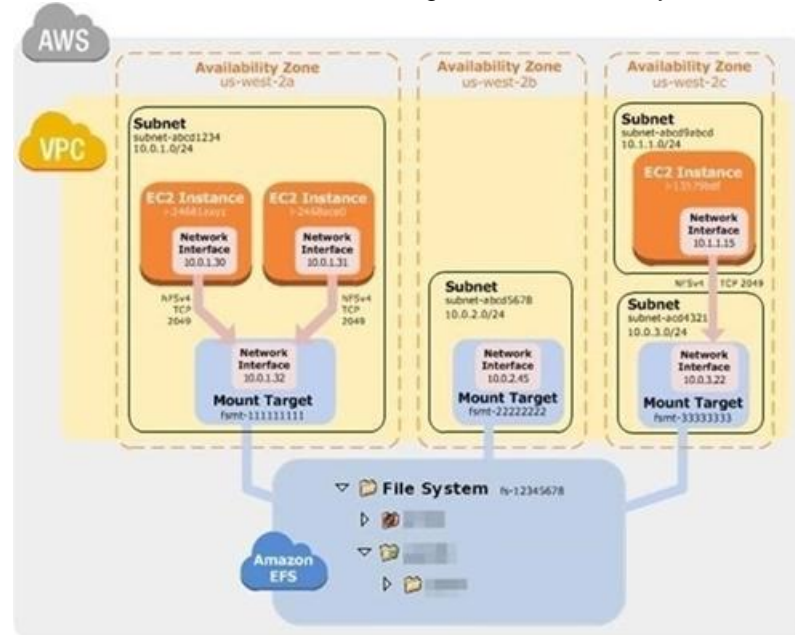
Explanation/Reference:

Explanation:

How Amazon EFS Works with Amazon EC2

The following illustration shows an example VPC accessing an Amazon EFS file system. Here, EC2 instances in the VPC have file systems mounted.

In this illustration, the VPC has three Availability Zones, and each has one mount target created in it. We recommend that you access the file system from a mount target within the same Availability Zone. One of the Availability Zones has two subnets. However, a mount target is created in only one of the subnets.



Benefits of Auto Scaling

Better fault tolerance. Amazon EC2 Auto Scaling can detect when an instance is unhealthy, terminate it, and launch an instance to replace it. You can also configure Amazon EC2 Auto Scaling to use multiple Availability Zones. If one Availability Zone becomes unavailable, Amazon EC2 Auto Scaling can launch instances in another one to compensate.

Better availability. Amazon EC2 Auto Scaling helps ensure that your application always has the right amount of capacity to handle the current traffic demand.

Better cost management. Amazon EC2 Auto Scaling can dynamically increase and decrease capacity as needed. Because you pay for the EC2 instances you use, you save money by launching instances when they are needed and terminating them when they aren't.

Reference:

<https://docs.aws.amazon.com/efs/latest/ug/how-it-works.html#how-it-works-ec2>

<https://docs.aws.amazon.com/autoscaling/ec2/userguide/auto-scaling-benefits.html>

QUESTION 49

A security team wants to limit access to specific services or actions in all of the team's AWS accounts. All accounts belong to a large organization in AWS Organizations. The solution must be scalable and there must be a single point where permissions can be maintained.

What should a solutions architect do to accomplish this?

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

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Service Control Policy concepts

SCPs offer central access controls for all IAM entities in your accounts. You can use them to enforce the permissions you want everyone in your business to follow. Using SCPs, you can give your developers more freedom to manage their own permissions because you know they can only operate within the boundaries you define.

You create and apply SCPs through AWS Organizations. When you create an organization, AWS Organizations automatically creates a root, which forms the parent container for all the accounts in your organization. Inside the root, you can group accounts in your organization into organizational units (OUs) to simplify management of these accounts. You can create multiple OUs within a single organization, and you can create OUs within other OUs to form a hierarchical structure. You can attach SCPs to the organization root, OUs, and individual accounts. SCPs attached to the root and OUs apply to all OUs and accounts inside of them.

SCPs use the AWS Identity and Access Management (IAM) policy language; however, they do not grant permissions. SCPs enable you set permission guardrails by defining the maximum available permissions for IAM entities in an account. If a SCP denies an action for an account, none of the entities in the account can take that action, even if their IAM permissions allow them to do so. The guardrails set in SCPs apply to all IAM entities in the account, which include all users, roles, and the account root user.

Reference:

[https://aws.amazon.com/blogs/security/how-to-use-service-control-policies-to-set-permissionguardrails-across-accounts-in-your-awsorganization/#:~:text=Central%20security%20administrators%20use%20service,users%20and%20roles\)%20adhere%20to.&text=Now%2C%20using%20SCPs%2C%20you%20can,your%20organization%20or%20organizational%20unit](https://aws.amazon.com/blogs/security/how-to-use-service-control-policies-to-set-permissionguardrails-across-accounts-in-your-awsorganization/#:~:text=Central%20security%20administrators%20use%20service,users%20and%20roles)%20adhere%20to.&text=Now%2C%20using%20SCPs%2C%20you%20can,your%20organization%20or%20organizational%20unit)
https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_policies_scp.html

QUESTION 50

A data science team requires storage for nightly log processing. The size and number of logs is unknown and will persist for 24 hours only. What is the MOST cost-effective solution?

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

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Reference:

https://aws.amazon.com/s3/storage-classes/#Unknown_or_changing_access

QUESTION 51

A company is hosting a web application on AWS using a single Amazon EC2 instance that stores user-uploaded documents in an Amazon Elastic Block Store (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 Elastic File System (Amazon EFS). Modify the application to save new documents to Amazon Elastic File System (Amazon EFS).
- D. Configure the Application Load Balancer to send the request to both servers. Return each document from the correct server.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

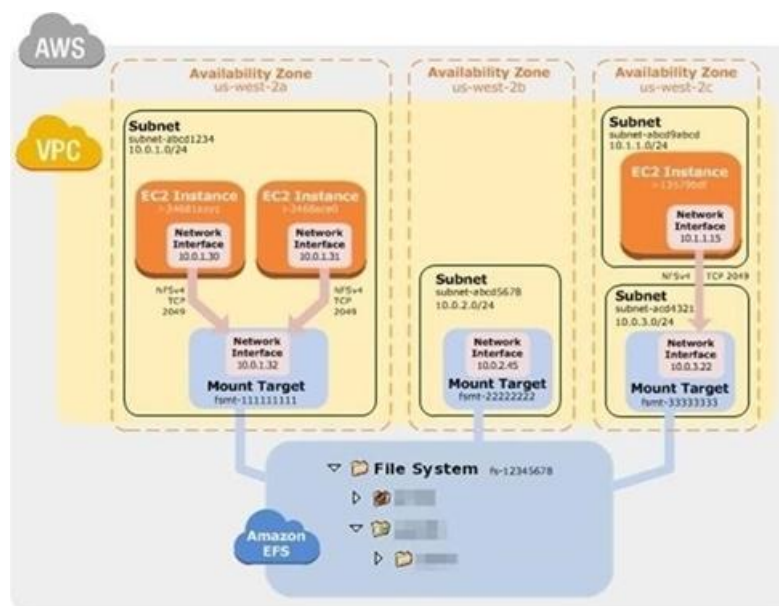
Explanation:

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-efs-utils Tools.

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

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

How Amazon EFS Works with Amazon EC2



Reference:
<https://docs.aws.amazon.com/efs/latest/ug/how-it-works.html#how-it-works-ec2>

QUESTION 52

A company is planning to use Amazon S3 to store images uploaded by its users. The images must be encrypted at rest in Amazon S3. The company does not want to spend time managing and rotating the keys, but it does want to control who can access those keys. What should a solutions architect use to accomplish this?

- A. Server-Side Encryption with keys stored in an S3 bucket
- B. Server-Side Encryption with Customer-Provided Keys (SSE-C)
- C. Server-Side Encryption with Amazon S3-Managed Keys (SSE-S3)
- D. Server-Side Encryption with AWS KMS-Managed Keys (SSE-KMS)



Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:

Explanation:
 "Server-Side Encryption with Customer Master Keys (CMKs) Stored in AWS Key Management Service (SSE-KMS) is similar to SSE-S3, but with some additional benefits and charges for using this service. There are separate permissions for the use of a CMK that provides added protection against unauthorized access of your objects in Amazon S3. SSE-KMS also provides you with an audit trail that shows when your CMK was used and by whom." Server-Side Encryption: Using SSE-KMS You can protect data at rest in Amazon S3 by using three different modes of server-side encryption: SSE-S3, SSE-C, or SSE-KMS. SSE-S3 requires that Amazon S3 manage the data and master encryption keys. For more information about SSE-S3, see Protecting Data Using Server-Side Encryption with Amazon S3- Managed Encryption Keys (SSE-S3). SSE-C requires that you manage the encryption key. For more information about SSE-C, see Protecting Data Using Server-Side Encryption with Customer-Provided Encryption Keys (SSE-C). SSE-KMS requires that AWS manage the data key but you manage the customer master key (CMK) in AWS KMS. The remainder of this topic discusses how to protect data by using server-side encryption with AWS KMS-managed keys (SSE-KMS). You can request encryption and select a CMK by using the Amazon S3 console or API. In the console, check the appropriate box to perform encryption and select your CMK from the list. For the Amazon S3 API, specify encryption and choose your CMK by setting the appropriate headers in a GET or PUT request.

Reference:
<https://aws.amazon.com/kms/faqs/>
<https://docs.aws.amazon.com/AmazonS3/latest/dev/serv-side-encryption.html>
<https://docs.aws.amazon.com/kms/latest/developerguide/services-s3.html#sse>

QUESTION 53

A company is running an ecommerce application on Amazon EC2. The application consists of a stateless web tier that requires a minimum of 10 instances, and a peak of 250 instances to support the application's usage. The application requires 50 instances 80% of the time. Which solution should be used to minimize costs?

- A. Purchase Reserved Instances to cover 250 instances.
- B. Purchase Reserved Instances to cover 80 instances. Use Spot Instances to cover the remaining instances.
- C. Purchase On-Demand Instances to cover 40 instances. Use Spot Instances to cover the remaining instances.

D. Purchase Reserved Instances to cover 50 instances. Use On-Demand and Spot Instances to cover the remaining instances.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Reserved Instances

Having 50 EC2 RIs provide a discounted hourly rate and an optional capacity reservation for EC2 instances. AWS Billing automatically applies your RI's discounted rate when attributes of EC2 instance usage match attributes of an active RI.

If an Availability Zone is specified, EC2 reserves capacity matching the attributes of the RI. The capacity reservation of an RI is automatically utilized by running instances matching these attributes.

You can also choose to forego the capacity reservation and purchase an RI that is scoped to a region. RIs that are scoped to a region automatically apply the RI's discount to instance usage across AZs and instance sizes in a region, making it easier for you to take advantage of the RI's discounted rate.

On-Demand Instance

On-Demand instances let you pay for compute capacity by the hour or second (minimum of 60 seconds) with no long-term commitments. This frees you from the costs and complexities of planning, purchasing, and maintaining hardware and transforms what are commonly large fixed costs into much smaller variable costs.

The pricing below includes the cost to run private and public AMIs on the specified operating system ("Windows Usage" prices apply to Windows Server 2003 R2, 2008, 2008 R2, 2012, 2012 R2, 2016, and 2019). Amazon also provides you with additional instances for Amazon EC2 running Microsoft Windows with SQL Server, Amazon EC2 running SUSE Linux Enterprise Server, Amazon EC2 running Red Hat Enterprise Linux and Amazon EC2 running IBM that are priced differently.

Spot Instances

A Spot Instance is an unused EC2 instance that is available for less than the On-Demand price.

Because Spot Instances enable you to request unused EC2 instances at steep discounts, you can lower your Amazon EC2 costs significantly. The hourly price for a Spot Instance is called a Spot price. The Spot price of each instance type in each Availability Zone is set by Amazon EC2, and adjusted gradually based on the long-term supply of and demand for Spot Instances. Your Spot Instance runs whenever capacity is available and the maximum price per hour for your request exceeds the Spot price.

Reference:

<https://aws.amazon.com/ec2/pricing/reserved-instances/>

<https://aws.amazon.com/ec2/pricing/on-demand/>

<https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/using-spot-instances.html>

QUESTION 54

A company has deployed an API in a VPC behind an internet-facing Application Load Balancer (ALB). An application that consumes the API as a client is deployed in a second account in private subnets behind a NAT gateway. When requests to the client application increase, the NAT gateway costs are higher than expected. A solutions architect has configured the ALB to be internal.

Which combination of architectural changes will reduce the NAT gateway costs? (Choose two.)

- A. Configure a VPC peering connection between the two VPCs. Access the API using the private address.
- B. Configure an AWS Direct Connect connection between the two VPCs. Access the API using the private address.
- C. Configure a ClassicLink connection for the API into the client VPC. Access the API using the ClassicLink address.
- D. Configure a PrivateLink connection for the API into the client VPC. Access the API using the PrivateLink address.
- E. Configure an AWS Resource Access Manager connection between the two accounts. Access the API using the private address.

Correct Answer: DE

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 55

A solutions architect is tasked with transferring 750 TB of data from an on-premises network-attached file system located at a branch office to Amazon S3 Glacier. The migration must not saturate the on-premises 1 Mbps internet connection. Which solution will meet these requirements?

- A. Create an AWS site-to-site VPN tunnel to an Amazon S3 bucket and transfer the files directly.
Transfer the files directly by using the AWS CLI.
- B. Order 10 AWS Snowball Edge Storage Optimized devices, and select an S3 Glacier vault as the destination.
- C. Mount the network-attached file system to an S3 bucket, and copy the files directly. Create a lifecycle policy to transition the S3 objects to Amazon S3 Glacier.
- D. Order 10 AWS Snowball Edge Storage Optimized devices, and select an Amazon S3 bucket as the destination. Create a lifecycle policy to transition the S3 objects to Amazon S3 Glacier.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 56

A company has a two-tier application architecture that runs in public and private subnets. Amazon EC2 instances running the web application are in the public subnet and a database runs on the private subnet. The web application instances and the database are running in a single Availability Zone (AZ).

Which combination of steps should a solutions architect take to provide high availability for this architecture? (Choose two.)

- A. Create new public and private subnets in the same AZ for high availability.
- B. Create an Amazon EC2 Auto Scaling group and Application Load Balancer spanning multiple AZs.
- C. Add the existing web application instances to an Auto Scaling group behind an Application Load Balancer.
- D. Create new public and private subnets in a new AZ. Create a database using Amazon EC2 in one AZ.
- E. Create new public and private subnets in the same VPC, each in a new AZ. Migrate the database to an Amazon RDS multi-AZ deployment.

Correct Answer: BE

Section: (none)

Explanation

Explanation/Reference:

Explanation:

You can take advantage of the safety and reliability of geographic redundancy by spanning your Auto Scaling group across multiple Availability Zones within a Region and then attaching a load balancer to distribute incoming traffic across those zones. Incoming traffic is distributed equally across all Availability Zones enabled for your load balancer.

Note

An Auto Scaling group can contain Amazon EC2 instances from multiple Availability Zones within the same Region. However, an Auto Scaling group can't contain instances from multiple Regions.

When one Availability Zone becomes unhealthy or unavailable, Amazon EC2 Auto Scaling launches new instances in an unaffected zone. When the unhealthy Availability Zone returns to a healthy state, Amazon EC2 Auto Scaling automatically redistributes the application instances evenly across all of the zones for your Auto Scaling group. Amazon EC2 Auto Scaling does this by attempting to launch new instances in the Availability Zone with the fewest instances. If the attempt fails, however, Amazon EC2 Auto Scaling attempts to launch in other Availability Zones until it succeeds.

You can expand the availability of your scaled and load-balanced application by adding an Availability Zone to your Auto Scaling group and then enabling that zone for your load balancer.

After you've enabled the new Availability Zone, the load balancer begins to route traffic equally among all the enabled zones.

High Availability (Multi-AZ) for Amazon RDS

Amazon RDS provides high availability and failover support for DB instances using Multi-AZ deployments. Amazon RDS uses several different technologies to provide failover support. Multi-AZ deployments for MariaDB, MySQL, Oracle, and PostgreSQL DB instances use Amazon's failover technology. SQL Server DB instances use SQL Server Database Mirroring (DBM) or Always On Availability Groups (AGs).

In a Multi-AZ deployment, Amazon RDS automatically provisions and maintains a synchronous standby replica in a different Availability Zone. The primary DB instance is synchronously replicated across Availability Zones to a standby replica to provide data redundancy, eliminate I/O freezes, and minimize latency spikes during system backups. Running a DB instance with high availability can enhance availability during planned system maintenance, and help protect your databases against DB instance failure and Availability Zone disruption. For more information on Availability Zones, see Regions, Availability Zones, and Local Zones.

Reference:

<https://docs.aws.amazon.com/autoscaling/ec2/userguide/as-add-availability-zone.html>

<https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/Concepts.MultiAZ.html>

QUESTION 57

A solutions architect is implementing a document review application using an Amazon S3 bucket for storage. The solution must prevent an accidental deletion of the documents and ensure that all versions of the documents are available.

Users must be able to download, modify, and upload documents.

Which combination of actions should be taken to meet these requirements? (Choose two.)

- A. Enable a read-only bucket ACL.
- B. Enable versioning on the bucket.
- C. Attach an IAM policy to the bucket.
- D. Enable MFA Delete on the bucket.
- E. Encrypt the bucket using AWS KMS.

Correct Answer: BD

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Object Versioning

Use Amazon S3 Versioning to keep multiple versions of an object in one bucket. For example, you could store my-image.jpg (version 111111) and my-image.jpg (version 222222) in a single bucket.

S3 Versioning protects you from the consequences of unintended overwrites and deletions. You can also use it to archive objects so that you have access to previous versions.

To customize your data retention approach and control storage costs, use object versioning with Object lifecycle management. For information about creating S3 Lifecycle policies using the AWS Management Console, see How Do I Create a Lifecycle Policy for an S3 Bucket? in the Amazon Simple Storage Service Console User Guide.

If you have an object expiration lifecycle policy in your non-versioned bucket and you want to maintain the same permanent delete behavior when you enable versioning, you must add a noncurrent expiration policy. The noncurrent expiration lifecycle policy will manage the deletes of the noncurrent object versions in the version-enabled bucket. (A version-enabled bucket maintains one current and zero or more noncurrent object versions.)

You must explicitly enable S3 Versioning on your bucket. By default, S3 Versioning is disabled.

Regardless of whether you have enabled Versioning, each object in your bucket has a version ID.

If you have not enabled Versioning, Amazon S3 sets the value of the version ID to null. If S3 Versioning is enabled, Amazon S3 assigns a version ID value for the object. This value distinguishes it from other versions of the same key.

Enabling and suspending versioning is done at the bucket level. When you enable versioning on an existing bucket, objects that are already stored in the bucket are unchanged. The version IDs

(null), contents, and permissions remain the same. After you enable S3 Versioning for a bucket, each object that is added to the bucket gets a version ID, which distinguishes it from other versions of the same key.

Only Amazon S3 generates version IDs, and they can't be edited. Version IDs are Unicode, UTF-8 encoded, URL-ready, opaque strings that are no more than 1,024 bytes long. The following is an example: 3/L4kqtJlcpXroDTDmJ+rmSpXd3dlbrHY+MTRCxf3vjVBH40Nr8X8gdRQBpUMLUo.

Using MFA delete

If a bucket's versioning configuration is MFA Delete-enabled, the bucket owner must include the x-amz-mfa request header in requests to permanently delete an object version or change the versioning state of the bucket. Requests that include x-amz-mfa must use HTTPS. The header's value is the concatenation of your authentication device's serial number, a space, and the authentication code displayed on it. If you do not include this request header, the request fails.

Reference:

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

<https://docs.aws.amazon.com/AmazonS3/latest/dev/ObjectVersioning.html>

<https://docs.aws.amazon.com/AmazonS3/latest/dev/UsingMFADelete.html>

QUESTION 58

An application hosted on AWS is experiencing performance problems, and the application vendor wants to perform an analysis of the log file to troubleshoot further. The log file is stored on Amazon S3 and is 10 GB in size. The application owner will make the log file available to the vendor for a limited time.

What is the MOST secure way to do this?

- A. Enable public read on the S3 object and provide the link to the vendor.
- B. Upload the file to Amazon WorkDocs and share the public link with the vendor.
- C. Generate a presigned URL and have the vendor download the log file before it expires.
- D. Create an IAM user for the vendor to provide access to the S3 bucket and the application. Enforce multi-factor authentication.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Share an object with others

All objects by default are private. Only the object owner has permission to access these objects.

However, the object owner can optionally share objects with others by creating a presigned URL, using their own security credentials, to grant time-limited permission to download the objects.

When you create a presigned URL for your object, you must provide your security credentials, specify a bucket name, an object key, specify the HTTP method (GET to download the object) and expiration date and time. The presigned URLs are valid only for the specified duration.

Anyone who receives the presigned URL can then access the object. For example, if you have a video in your bucket and both the bucket and the object are private, you can share the video with others by generating a presigned URL.

Reference:

<https://docs.aws.amazon.com/AmazonS3/latest/dev/ShareObjectPreSignedURL.html>

QUESTION 59

A solutions architect is designing a two-tier web application. The application consists of a publicfacing 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? (Choose 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.

Correct Answer: AC

Section: (none)

Explanation

Explanation/Reference:



Explanation:

QUESTION 60

A company allows its developers to attach existing IAM policies to existing IAM roles to enable faster experimentation and agility. However, the security operations team is concerned that the developers could attach the existing administrator policy, which would allow the developers to circumvent any other security policies. How should a solutions architect address this issue?

- A. Create an Amazon SNS topic to send an alert every time a developer creates a new policy.
- B. Use service control policies to disable IAM activity across all account in the organizational unit.
- C. Prevent the developers from attaching any policies and assign all IAM duties to the security operations team.
- D. Set an IAM permissions boundary on the developer IAM role that explicitly denies attaching the administrator policy.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Reference:

https://docs.aws.amazon.com/IAM/latest/UserGuide/access_policies_boundaries.html

QUESTION 61

A company has a multi-tier application that runs six front-end web servers in an Amazon EC2 Auto Scaling group in a single Availability Zone behind an Application Load Balancer (ALB). A solutions architect needs to modify the infrastructure to be highly available without modifying the application. Which architecture should the solutions architect choose that provides high availability?

- A. Create an Auto Scaling group that uses three instances across each of two Regions.
- B. Modify the Auto Scaling group to use three instances across each of two Availability Zones.
- C. Create an Auto Scaling template that can be used to quickly create more instances in another Region.
- D. Change the ALB in front of the Amazon EC2 instances in a round-robin configuration to balance traffic to the web tier.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Expanding Your Scaled and Load-Balanced Application to an Additional Availability Zone.

When one Availability Zone becomes unhealthy or unavailable, Amazon EC2 Auto Scaling launches new instances in an unaffected zone. When the unhealthy Availability Zone returns to a healthy state, Amazon EC2 Auto Scaling automatically redistributes the application instances evenly across all of the zones for your Auto Scaling group. Amazon EC2 Auto Scaling does this by attempting to launch new instances in the Availability Zone with the fewest instances. If the attempt fails, however, Amazon EC2 Auto Scaling attempts to launch in other Availability Zones until it succeeds.

You can expand the availability of your scaled and load-balanced application by adding an Availability Zone to your Auto Scaling group and then enabling that zone for your load balancer.

After you've enabled the new Availability Zone, the load balancer begins to route traffic equally among all the enabled zones.

Reference:

<https://docs.aws.amazon.com/autoscaling/ec2/userguide/as-add-availability-zone.html>

QUESTION 62

A company runs an application on a group of Amazon Linux EC2 instances. For compliance reasons, the company must retain all application log files for 7 years. The log files will be analyzed by a reporting tool that must access all files concurrently.

Which storage solution meets these requirements MOST cost-effectively?

- A. Amazon Elastic Block Store (Amazon EBS)
- B. Amazon Elastic File System (Amazon EFS)
- C. Amazon EC2 instance store
- D. Amazon S3

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Amazon S3

Requests to Amazon S3 can be authenticated or anonymous. Authenticated access requires credentials that AWS can use to authenticate your requests. When making REST API calls directly from your code, you create a signature using valid credentials and include the signature in your request. Amazon Simple Storage Service (Amazon S3) is an object storage service that offers industry-leading scalability, data availability, security, and performance. This means customers of all sizes and industries can use it to store and protect any amount of data for a range of use cases, such as websites, mobile applications, backup and restore, archive, enterprise applications, IoT devices, and big data analytics. Amazon S3 provides easy-to-use management features so you can organize your data and configure finely-tuned access controls to meet your specific business, organizational, and compliance requirements. Amazon S3 is designed for 99.99999999% (11 9's) of durability, and stores data for millions of applications for companies all around the world.

Reference:

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

QUESTION 63

A media streaming company collects real-time data and stores it in a disk-optimized database system. The company is not getting the expected throughput and wants an in-memory database storage solution that performs faster and provides high availability using data replication.

Which database should a solutions architect recommend?

- A. Amazon RDS for MySQL
- B. Amazon RDS for PostgreSQL
- C. Amazon ElastiCache for Redis
- D. Amazon ElastiCache for Memcached

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:


In-memory databases on AWS Amazon ElastiCache for Redis.

Amazon ElastiCache for Redis is a blazing fast in-memory data store that provides submillisecond latency to power internet-scale, real-time applications. Developers can use ElastiCache for Redis as an in-memory nonrelational database. The ElastiCache for Redis cluster configuration supports up to 15 shards and enables customers to run Redis workloads with up to 6.1 TB of in-memory capacity in a single cluster. ElastiCache for Redis also provides the ability to add and remove shards from a running cluster. You can dynamically scaleout and even scale in your Redis cluster workloads to adapt to changes in demand.

Reference:

<https://aws.amazon.com/elasticache/redis/faqs/>

<https://aws.amazon.com/nosql/in-memory/>



QUESTION 64

A company hosts its product information webpages on AWS. The existing solution uses multiple Amazon C2 instances behind an Application Load Balancer in an Auto Scaling group. The website also uses a custom DNS name and communicates with HTTPS only using a dedicated SSL certificate. The company is planning a new product launch and wants to be sure that users from around the world have the best possible experience on the new website.

What should a solutions architect do to meet these requirements?

- A. Redesign the application to use Amazon CloudFront.
- B. Redesign the application to use AWS Elastic Beanstalk.
- C. Redesign the application to use a Network Load Balancer.
- D. Redesign the application to use Amazon S3 static website hosting.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

What Is Amazon CloudFront?

Amazon CloudFront is a web service that speeds up distribution of your static and dynamic web content, such as .html, .css, .js, and image files, to your users. CloudFront delivers your content through a worldwide network of data centers called edge locations. When a user requests content that you're serving with CloudFront, the user is routed to the edge location that provides the lowest latency (time delay), so that content is delivered with the best possible performance.

If the content is already in the edge location with the lowest latency, CloudFront delivers it immediately.

If the content is not in that edge location, CloudFront retrieves it from an origin that you've defined – such as an Amazon S3 bucket, a MediaPackage channel, or an HTTP server (for example, a web server) that you have identified as the source for the definitive version of your content.

As an example, suppose that you're serving an image from a traditional web server, not from CloudFront. For example, you might serve an image, sunsetphoto.png, using the URL <http://example.com/sunsetphoto.png>.

Your users can easily navigate to this URL and see the image. But they probably don't know that their request was routed from one network to another – through the complex collection of interconnected networks that comprise the internet – until the image was found.

CloudFront speeds up the distribution of your content by routing each user request through the AWS backbone network to the edge location that can best serve your content. Typically, this is a CloudFront edge server that provides the fastest delivery to the viewer. Using the AWS network dramatically reduces the number of networks that your users' requests must pass through, which improves performance. Users get lower latency – the time it takes to load the first byte

of the file – and higher data transfer rates.

You also get increased reliability and availability because copies of your files (also known as objects) are now held (or cached) in multiple edge locations around the world.

Reference:

<https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/Introduction.html>

QUESTION 65

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.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Amazon Simple Queue Service

Amazon Simple Queue Service (SQS) is a fully managed message queuing service that enables you to decouple and scale microservices, distributed systems, and serverless applications. SQS eliminates the complexity and overhead associated with managing and operating message oriented middleware, and empowers developers to focus on differentiating work. Using SQS, you can send, store, and receive messages between software components at any volume, without losing messages or requiring other services to be available. Get started with SQS in minutes using the AWS console, Command Line Interface or SDK of your choice, and three simple commands.

SQS offers two types of message queues. Standard queues offer maximum throughput, best-effort ordering, and at-least-once delivery. SQS FIFO queues are designed to guarantee that messages are processed exactly once, in the exact order that they are sent.

Scaling Based on Amazon SQS

There are some scenarios where you might think about scaling in response to activity in an Amazon SQS queue. For example, suppose that you have a web app that lets users upload images and use them online. In this scenario, each image requires resizing and encoding before it can be published. The app runs on EC2 instances in an Auto Scaling group, and it's configured to handle your typical upload rates. Unhealthy instances are terminated and replaced to maintain current instance levels at all times. The app places the raw bitmap data of the images in an SQS queue for processing. It processes the images and then publishes the processed images where they can be viewed by users. The architecture for this scenario works well if the number of image uploads doesn't vary over time. But if the number of uploads changes over time, you might consider using dynamic scaling to scale the capacity of your Auto Scaling group.

Reference:

<https://aws.amazon.com/sqs/#:~:text=Amazon%20SQS%20leverages%20the%20AWS,queues%20provide%20nearly%20unlimited%20throughput>

<https://docs.aws.amazon.com/autoscaling/ec2/userguide/as-using-sqs-queue.html>

QUESTION 66

A marketing company is storing CSV files in an Amazon S3 bucket for statistical analysis. An application on an Amazon EC2 instance needs permission to efficiently process the CSV data stored in the S3 bucket.

Which action will MOST securely grant the EC2 instance access to the S3 bucket?

- A. Attach a resource-based policy to the S3 bucket.
- B. Create an IAM user for the application with specific permissions to the S3 bucket.
- C. Associate an IAM role with least privilege permissions to the EC2 instance profile.
- D. Store AWS credentials directly on the EC2 instance for applications on the instance to use for API calls.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 67

A company has on-premises servers running a relational database. The current database serves high read traffic for users in different locations. The company wants to migrate to AWS with the least amount of effort. The database solution should support disaster recovery and not affect the company's current traffic flow.

Which solution meets these requirements?

- A. Use a database in Amazon RDS with Multi-AZ and at least one read replica.
- B. Use a database in Amazon RDS with Multi-AZ and at least one standby replica.
- C. Use databases hosted on multiple Amazon EC2 instances in different AWS Regions.
- D. Use databases hosted on Amazon EC2 instances behind an Application Load Balancer in different Availability Zones.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Reference:

<https://aws.amazon.com/blogs/database/implementing-a-disaster-recovery-strategy-with-amazonrds/>

QUESTION 68

A company's application is running on Amazon EC2 instances within an Auto Scaling group behind an Elastic Load Balancer. Based on the application's history the company anticipates a spike in traffic during a holiday each year. A solutions architect must design a strategy to ensure that the Auto Scaling group proactively increases capacity to minimize any performance impact on application users. Which solution will meet these requirements?

- A. Create an Amazon CloudWatch alarm to scale up the EC2 instances when CPU utilization exceeds 90%.
- B. Create a recurring scheduled action to scale up the Auto Scaling group before the expected period of peak demand.
- C. Increase the minimum and maximum number of EC2 instances in the Auto Scaling group during the peak demand period.
- D. Configure an Amazon Simple Notification Service (Amazon SNS) notification to send alerts when there are autoscaling EC2_INSTANCE_LAUNCH events.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:



QUESTION 69

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

- A. Use the CreateQueue API call to create a new queue.
- B. Use the AddPermission API call to add appropriate permissions.
- C. Use the ReceiveMessage API call to set an appropriate wait time.
- D. Use the ChangeMessageVisibility API call to increase the visibility timeout.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 70

An Amazon EC2 administrator created the following policy associated with an IAM group containing several users:

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": "ec2:TerminateInstances",
      "Resource": "*",
      "Condition": {
        "IpAddress": {
          "aws:SourceIp": "10.100.100.0/24"
        }
      }
    },
    {
      "Effect": "Deny",
      "Action": "ec2:*",
      "Resource": "*",
      "Condition": {
        "StringNotEquals": {
          "ec2:Region": "us-east-1"
        }
      }
    }
  ]
}

```



What is the effect of this policy?

- A. Users can terminate an EC2 instance in any AWS Region except us-east-1.
- B. Users can terminate an EC2 instance with the IP address 10.100.100.1 in the us-east-1 Region.
- C. Users can terminate an EC2 instance in the us-east-1 Region when the user's source IP is 10.100.100.254.
- D. Users cannot terminate an EC2 instance in the us-east-1 Region when the user's source IP is 10.100.100.254.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
 Explanation:

QUESTION 71

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

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

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
 Explanation:
 Reference:

<https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/on-demand-streamingvideo.html>

QUESTION 72

A company has a three-tier image-sharing application. It uses an Amazon EC2 instance for the front-end layer, another for the backend tier, and a third for the MySQL database. A solutions architect has been tasked with designing a solution that is highly available, and requires the least amount of changes to the application. Which solution meets these requirements?

- A. Use Amazon S3 to host the front-end layer and AWS Lambda functions for the backend layer. Move the database to an Amazon DynamoDB table and use Amazon S3 to store and serve users' images.
- B. Use load-balanced Multi-AZ AWS Elastic Beanstalk environments for the front-end and backend layers. Move the database to an Amazon RDS instance with multiple read replicas to store and serve users' images.
- C. Use Amazon S3 to host the front-end layer and a fleet of Amazon EC2 instances in an Auto Scaling group for the backend layer. Move the database to a memory optimized instance type to store and serve users' images.
- D. Use load-balanced Multi-AZ AWS Elastic Beanstalk environments for the front-end and backend layers. Move the database to an Amazon RDS instance with a Multi-AZ deployment. Use Amazon S3 to store and serve users' images.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 73

A solutions architect is designing a system to analyze the performance of financial markets while the markets are closed. The system will run a series of compute-intensive jobs for 4 hours every night. The time to complete the compute jobs is expected to remain constant, and jobs cannot be interrupted once started. Once completed, the system is expected to run for a minimum of 1 year.

Which type of Amazon EC2 instances should be used to reduce the cost of the system?

- A. Spot Instances
- B. On-Demand Instances
- C. Standard Reserved Instances
- D. Scheduled Reserved Instances

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Reference:

<https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ec2-scheduled-instances.html>

QUESTION 74

A company built a food ordering application that captures user data and stores it for future analysis. The application's static front end is deployed on an Amazon EC2 instance. The front-end application sends the requests to the backend application running on separate EC2 instance. The backend application then stores the data in Amazon RDS.

What should a solutions architect do to decouple the architecture and make it scalable?

- A. Use Amazon S3 to serve the front-end application, which sends requests to Amazon EC2 to execute the backend application. The backend application will process and store the data in Amazon RDS.
- B. Use Amazon S3 to serve the front-end application and write requests to an Amazon Simple Notification Service (Amazon SNS) topic. Subscribe Amazon EC2 instances to the HTTP/HTTPS endpoint of the topic, and process and store the data in Amazon RDS.
- C. Use an EC2 instance to serve the front end and write requests to an Amazon SQS queue. Place the backend instance in an Auto Scaling group, and scale based on the queue depth to process and store the data in Amazon RDS.
- D. Use Amazon S3 to serve the static front-end application and send requests to Amazon API Gateway, which writes the requests to an Amazon SQS queue. Place the backend instances in an Auto Scaling group, and scale based on the queue depth to process and store the data in Amazon RDS.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 75

A solutions architect needs to design a managed storage solution for a company's application that includes high-performance machine learning functionality. This application runs on AWS Fargate and the connected storage needs to have concurrent access to files and deliver high performance.



Which storage option should the solutions architect recommend?

- A. Create an Amazon S3 bucket for the application and establish an IAM role for Fargate to communicate with Amazon S3.
- B. Create an Amazon FSx for Lustre file share and establish an IAM role that allows Fargate to communicate with FSx for Lustre.
- C. Create an Amazon Elastic File System (Amazon EFS) file share and establish an IAM role that allows Fargate to communicate with Amazon Elastic File System (Amazon EFS).
- D. Create an Amazon Elastic Block Store (Amazon EBS) volume for the application and establish an IAM role that allows Fargate to communicate with Amazon Elastic Block Store (Amazon EBS).

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 76

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

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

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

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Reference:

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



QUESTION 77

A solutions architect is designing a web application that will run on Amazon EC2 instances behind an Application Load Balancer (ALB). The company strictly requires that the application be resilient against malicious internet activity and attacks, and protect against new common vulnerabilities and exposures.

What should the solutions architect recommend?

- A. Leverage Amazon CloudFront with the ALB endpoint as the origin.
- B. Deploy an appropriate managed rule for AWS WAF and associate it with the ALB.
- C. Subscribe to AWS Shield Advanced and ensure common vulnerabilities and exposures are blocked.
- D. Configure network ACLs and security groups to allow only ports 80 and 443 to access the EC2 instances.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 78

A company has an application that calls AWS Lambda functions. A code review shows that database credentials are stored in a Lambda function's source code, which violates the company's security policy. The credentials must be securely stored and must be automatically rotated on an ongoing basis to meet security policy requirements.

What should a solutions architect recommend to meet these requirements in the MOST secure manner?

- A. Store the password in AWS CloudHSM. Associate the Lambda function with a role that can use the key ID to retrieve the password from CloudHSM. Use CloudHSM to automatically rotate the password.
- B. Store the password in AWS Secrets Manager. Associate the Lambda function with a role that can use the secret ID to retrieve the password from Secrets Manager. Use Secrets Manager to automatically rotate the password.
- C. Store the password in AWS Key Management Service (AWS KMS). Associate the Lambda function with a role that can use the key ID to retrieve the password from AWS KMS. Use AWS KMS to automatically rotate the uploaded password.
- D. Move the database password to an environment variable that is associated with the Lambda function. Retrieve the password from the environment variable by invoking the function. Create a deployment script to automatically rotate the

password.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:

Reference:

<https://aws.amazon.com/blogs/security/how-to-use-aws-secrets-manager-rotate-credentialsamazon-rds-database-types-oracle/>

QUESTION 79

A company is managing health records on-premises. The company must keep these records indefinitely, disable any modifications to the records once they are stored, and granularly audit access at all levels. The chief technology officer (CTO) is concerned because there are already millions of records not being used by any application, and the current infrastructure is running out of space. The CTO has requested a solutions architect design a solution to move existing data and support future records.

Which services can the solutions architect recommend to meet these requirements?

- A. Use AWS DataSync to move existing data to AWS. Use Amazon S3 to store existing and new data. Enable Amazon S3 object lock and enable AWS CloudTrail with data events.
- B. Use AWS Storage Gateway to move existing data to AWS. Use Amazon S3 to store existing and new data. Enable Amazon S3 object lock and enable AWS CloudTrail with management events.
- C. Use AWS DataSync to move existing data to AWS. Use Amazon S3 to store existing and new data. Enable Amazon S3 object lock and enable AWS CloudTrail with management events.
- D. Use AWS Storage Gateway to move existing data to AWS. Use Amazon Elastic Block Store (Amazon EBS) to store existing and new data. Enable Amazon S3 object lock and enable Amazon S3 server access logging.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:

Explanation:

QUESTION 80

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

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

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:

Explanation:

QUESTION 81

A company's operations team has an existing Amazon S3 bucket configured to notify an Amazon SQS queue when new objects are created within the bucket. The development team also wants to receive events when new objects are created. The existing operations team workflow must remain intact.

Which solution would satisfy these requirements?

- A. Create another SQS queue. Update the S3 events in the bucket to also update the new queue when a new object is created.
- B. Create a new SQS queue that only allows Amazon S3 to access the queue. Update Amazon S3 to update this queue when a new object is created.
- C. Create an Amazon SNS topic and SQS queue for the bucket updates. Update the bucket to send events to the new topic. Updates both queues to poll Amazon SNS.
- D. Create an Amazon SNS topic and SQS queue for the bucket updates. Update the bucket to send events to the new topic. Add subscriptions for both queues in the topic.

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:

Explanation:

QUESTION 82

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

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

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 83

A company built an application that lets users check in to places they visit, rank the places, and add reviews about their experiences. The application is successful with a rapid increase in the number of users every month. The chief technology officer fears the database supporting the current Infrastructure may not handle the new load the following month because the single Amazon RDS for MySQL instance has triggered alarms related to resource exhaustion due to read requests.

What can a solutions architect recommend to prevent service Interruptions at the database layer with minimal changes to code?

- A. Create RDS read replicas and redirect read-only traffic to the read replica endpoints. Enable a Multi-AZ deployment.
- B. Create an Amazon EMR cluster and migrate the data to a Hadoop Distributed File System (HDFS) with a replication factor of 3.
- C. Create an Amazon ElastiCache cluster and redirect all read-only traffic to the cluster. Set up the cluster to be deployed in three Availability Zones.
- D. Create an Amazon DynamoDB table to replace the RDS instance and redirect all read-only traffic to the DynamoDB table. Enable DynamoDB Accelerator to offload traffic from the main table.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 84

A company is looking for a solution that can store video archives in AWS from old news footage.

The company needs to minimize costs and will rarely need to restore these files. When the files are needed, they must be available in a maximum of five minutes.

What is the MOST cost-effective solution?

- A. Store the video archives in Amazon S3 Glacier and use Expedited retrievals.
- B. Store the video archives in Amazon S3 Glacier and use Standard retrievals.
- C. Store the video archives in Amazon S3 Standard-Infrequent Access (S3 Standard-IA).
- D. Store the video archives in Amazon S3 One Zone-Infrequent Access (S3 One Zone-IA).

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 85

A company has created a VPC with multiple private subnets in multiple Availability Zones (AZs) and one public subnet in one of the AZs. The public subnet is used to launch a NAT gateway.

There are instances in the private subnets that use a NAT gateway to connect to the internet. In case of an AZ failure, the company wants to ensure that the instances are not all experiencing internet connectivity issues and that there is a backup plan ready.

Which solution should a solutions architect recommend that is MOST highly available?

- A. Create a new public subnet with a NAT gateway in the same AZ. Distribute the traffic between the two NAT gateways.



- B. Create an Amazon EC2 NAT instance in a new public subnet. Distribute the traffic between the NAT gateway and the NAT instance.
- C. Create public subnets in each AZ and launch a NAT gateway in each subnet. Configure the traffic from the private subnets in each AZ to the respective NAT gateway.
- D. Create an Amazon EC2 NAT instance in the same public subnet. Replace the NAT gateway with the NAT instance and associate the instance with an Auto Scaling group with an appropriate scaling policy.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 86

A healthcare company stores highly sensitive patient records. Compliance requires that multiple copies be stored in different locations. Each record must be stored for 7 years. The company has a service level agreement (SLA) to provide records to government agencies immediately for the first 30 days and then within 4 hours of a request thereafter.

What should a solutions architect recommend?

- A. Use Amazon S3 with cross-Region replication enabled. After 30 days, transition the data to Amazon S3 Glacier using lifecycle policy.
- B. Use Amazon S3 with cross-origin resource sharing (CORS) enabled. After 30 days, transition the data to Amazon S3 Glacier using a lifecycle policy.
- C. Use Amazon S3 with cross-Region replication enabled. After 30 days, transition the data to Amazon S3 Glacier Deep Archive using a lifecycle policy.
- D. Use Amazon S3 with cross-origin resource sharing (CORS) enabled. After 30 days, transition the data to Amazon S3 Glacier Deep Archive using a lifecycle policy.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 87

A company recently deployed a new auditing system to centralize information about operating system versions, patching, and installed software for Amazon EC2 instances. A solutions architect must ensure all instances provisioned through EC2 Auto Scaling groups successfully send reports to the auditing system as soon as they are launched and terminated.

Which solution achieves these goals MOST efficiently?

- A. Use a scheduled AWS Lambda function and run a script remotely on all EC2 instances to send data to the audit system.
- B. Use EC2 Auto Scaling lifecycle hooks to run a custom script to send data to the audit system when instances are launched and terminated.
- C. Use an EC2 Auto Scaling launch configuration to run a custom script through user data to send data to the audit system when instances are launched and terminated.
- D. Run a custom script on the instance operating system to send data to the audit system. Configure the script to be executed by the EC2 Auto Scaling group when the instance starts and is terminated.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 88

A company recently implemented hybrid cloud connectivity using AWS Direct Connect and is migrating data to Amazon S3. The company is looking for a fully managed solution that will automate and accelerate the replication of data between the on-premises storage systems and AWS storage services.

Which solution should a solutions architect recommend to keep the data private?

- A. Deploy an AWS DataSync agent for the on-premises environment. Configure a sync job to replicate the data and connect it with an AWS service endpoint.
- B. Deploy an AWS DataSync agent for the on-premises environment. Schedule a batch job to replicate point-in-time snapshots to AWS.
- C. Deploy an AWS Storage Gateway volume gateway for the on-premises environment. Configure it to store data locally, and asynchronously back up point-in-time snapshots to AWS.
- D. Deploy an AWS Storage Gateway file gateway for the on-premises environment. Configure it to store data locally, and asynchronously back up point-in-time snapshots to AWS.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 89

A company has 150 TB of archived image data stored on-premises that needs to be moved to the AWS Cloud within the next month. The company's current network connection allows up to 100 Mbps uploads for this purpose during the night only.

What is the MOST cost-effective mechanism to move this data and meet the migration deadline?

- A. Use AWS Snowmobile to ship the data to AWS.
- B. Order multiple AWS Snowball devices to ship the data to AWS.
- C. Enable Amazon S3 Transfer Acceleration and securely upload the data.
- D. Create an Amazon S3 VPC endpoint and establish a VPN to upload the data.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 90

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? (Choose 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

Correct Answer: BE

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 91

A company is seeing access requests by some suspicious IP addresses. The security team discovers the requests are from different IP addresses under the same CIDR range.

What should a solutions architect recommend to the team?

- A. Add a rule in the inbound table of the security to deny the traffic from that CIDR range.
- B. Add a rule in the outbound table of the security group to deny the traffic from that CIDR range.
- C. Add a deny rule in the inbound table of the network ACL with a lower number than other rules.
- D. Add a deny rule in the outbound table of the network ACL with a lower rule number than other rules.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 92

A company recently expanded globally and wants to make its application accessible to users in those geographic locations. The application is deployed on Amazon EC2 instances behind an Application Load Balancer in an Auto Scaling group. The company needs the ability to shift traffic from resources in one region to another.

What should a solutions architect recommend?



- A. Configure an Amazon Route 53 latency routing policy.
- B. Configure an Amazon Route 53 geolocation routing policy.
- C. Configure an Amazon Route 53 geoproximity routing policy.
- D. Configure an Amazon Route 53 multivalue answer routing policy.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 93

A company wants to replicate its data to AWS to recover in the event of a disaster. Today, a system administrator has scripts that copy data to a NFS share. Individual backup files need to be accessed with low latency by application administrators to deal with errors in processing.

What should a solutions architect recommend to meet these requirements?

- A. Modify the script to copy data to an Amazon S3 bucket instead of the on-premises NFS share.
- B. Modify the script to copy data to an Amazon S3 Glacier Archive instead of the on-premises NFS share.
- C. Modify the script to copy data to an Amazon Elastic File System (Amazon EFS) volume instead of the on-premises NFS share.
- D. Modify the script to copy data to an AWS Storage Gateway for File Gateway virtual appliance instead of the on-premises NFS share.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 94

An application requires a development environment (DEV) and production environment (PROD) for several years. The DEV instances will run for 10 hours each day during normal business hours, while the PROD instances will run 24 hours each day. A solutions architect needs to determine a compute instance purchase strategy to minimize costs.

Which solution is the MOST cost-effective?

- A. DEV with Spot Instances and PROD with On-Demand Instances
- B. DEV with On-Demand Instances and PROD with Spot Instances
- C. DEV with Scheduled Reserved Instances and PROD with Reserved Instances
- D. DEV with On-Demand Instances and PROD with Scheduled Reserved Instances

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 95

A company runs multiple Amazon EC2 Linux instances in a VPC with applications that use a hierarchical directory structure. The applications need to rapidly and concurrently read and write to shared storage. How can this be achieved?

- A. Create an Amazon Elastic File System (Amazon EFS) file system and mount it from each EC2 instance.
- B. Create an Amazon S3 bucket and permit access from all the EC2 instances in the VPC.
- C. Create a file system on an Amazon Elastic Block Store (Amazon EBS) Provisioned IOPS SSD (io1) volume. Attach the volume to all the EC2 instances.
- D. Create file systems on Amazon Elastic Block Store (Amazon EBS) volumes attached to each EC2 instance. Synchronize the Amazon Elastic Block Store (Amazon EBS) volumes across the different EC2 instances.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 96

A solutions architect observes that a nightly batch processing job is automatically scaled up for 1 hour before the desired Amazon EC2 capacity is reached. The peak capacity is the same every night and the batch jobs always start at 1 AM. The solutions architect needs to find a cost-effective solution that will allow for the desired EC2 capacity to be reached quickly and allow the Auto Scaling group to scale down after the batch jobs are complete. What should the solutions architect do to meet these requirements?

- A. Increase the minimum capacity for the Auto Scaling group.
- B. Increase the maximum capacity for the Auto Scaling group.
- C. Configure scheduled scaling to scale up to the desired compute level.
- D. Change the scaling policy to add more EC2 instances during each scaling operation.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 97

A Solutions Architect must design a web application that will be hosted on AWS, allowing users to purchase access to premium, shared content that is stored in an S3 bucket. Upon payment, content will be available for download for 14 days before the user is denied access. Which of the following would be the LEAST complicated implementation?

- A. Use an Amazon CloudFront distribution with an origin access identity (OAI). Configure the distribution with an Amazon S3 origin to provide access to the file through signed URLs. Design a Lambda function to remove data that is older than 14 days.
- B. Use an S3 bucket and provide direct access to the file. Design the application to track purchases in a DynamoDB table. Configure a Lambda function to remove data that is older than 14 days based on a query to Amazon DynamoDB.
- C. Use an Amazon CloudFront distribution with an OAI. Configure the distribution with an Amazon S3 origin to provide access to the file through signed URLs. Design the application to set an expiration of 14 days for the URL.
- D. Use an Amazon CloudFront distribution with an OAI. Configure the distribution with an Amazon S3 origin to provide access to the file through signed URLs. Design the application to set an expiration of 60 minutes for the URL and recreate the URL as necessary.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 98

A solutions architect is designing a mission-critical web application. It will consist of Amazon EC2 instances behind an Application Load Balancer and a relational database. The database should be highly available and fault tolerant. Which database implementations will meet these requirements? (Choose two.)

- A. Amazon Redshift
- B. Amazon DynamoDB
- C. Amazon RDS for MySQL
- D. MySQL-compatible Amazon Aurora Multi-AZ
- E. Amazon RDS for SQL Server Standard Edition Multi-AZ

Correct Answer: DE

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 99

A company's web application is running on Amazon EC2 instances behind an Application Load Balancer. The company recently changed its policy, which now requires the application to be accessed from one specific country only. Which configuration will meet this requirement?

- A. Configure the security group for the EC2 instances.
- B. Configure the security group on the Application Load Balancer.
- C. Configure AWS WAF on the Application Load Balancer in a VPC.
- D. Configure the network ACL for the subnet that contains the EC2 instances.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Reference:

<https://aws.amazon.com/es/blogs/security/how-to-use-aws-waf-to-filter-incoming-traffic-fromembargoed-countries/>

QUESTION 100

A solutions architect has created two IAM policies: Policy1 and Policy2. Both policies are attached to an IAM group.

```
Policy1
{
  "Version": "2012-10-17", "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "iam:Get*",
        "iam:List*",
        "kms:List*",
        "ec2:*",
        "ds:*",
        "logs:Get*",
        "logs:Describe*"
      ],
      "Resource": "*"
    }
  ]
}

Policy2
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Deny",
      "Action": "ds:Delete*",
      "Resource": "*"
    }
  ]
}
```



A cloud engineer is added as an IAM user to the IAM group. Which action will the cloud engineer be able to perform?

- A. Deleting IAM users
- B. Deleting directories
- C. Deleting Amazon EC2 instances
- D. Deleting logs from Amazon CloudWatch Logs

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:
Explanation:

QUESTION 101

A company has an Amazon EC2 instance running on a private subnet that needs to access a public website to download patches and updates. The company does not want external websites to see the EC2 instance IP address or initiate connections to it.
How can a solutions architect achieve this objective?

- A. Create a site-to-site VPN connection between the private subnet and the network in which the public site is deployed.
- B. Create a NAT gateway in a public subnet. Route outbound traffic from the private subnet through the NAT gateway.
- C. Create a network ACL for the private subnet where the EC2 instance deployed only allows access from the IP address range of the public website.
- D. Create a security group that only allows connections from the IP address range of the public website. Attach the security group to the EC2 instance.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:

QUESTION 102

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

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

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:

QUESTION 103

A company has a website running on Amazon EC2 instances across two Availability Zones. The company is expecting spikes in traffic on specific holidays, and wants to provide a consistent user experience.
How can a solutions architect meet this requirement?

- A. Use step scaling.
- B. Use simple scaling.
- C. Use lifecycle hooks.
- D. Use scheduled scaling.

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Explanation:

QUESTION 104

An ecommerce company is running a multi-tier application on AWS. The front-end and backend tiers both run on Amazon EC2, and the database runs on Amazon RDS for MySQL. The backend tier communicates with the RDS instance. There are frequent calls to return identical datasets from the database that are causing performance slowdowns.
Which action should be taken to improve the performance of the backend?

- A. Implement Amazon SNS to store the database calls.



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

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 105

A company has an on-premises data center that is running out of storage capacity. The company wants to migrate its storage infrastructure to AWS while minimizing bandwidth costs. The solution must allow for immediate retrieval of data at no additional cost.

How can these requirements be met?

- A. Deploy Amazon S3 Glacier Vault and enable expedited retrieval. Enable provisioned retrieval capacity for the workload.
- B. Deploy AWS Storage Gateway using cached volumes. Use Storage Gateway to store data in Amazon S3 while retaining copies of frequently accessed data subsets locally.
- C. Deploy AWS Storage Gateway using stored volumes to store data locally. Use Storage Gateway to asynchronously back up point-in-time snapshots of the data to Amazon S3.
- D. Deploy AWS Direct Connect to connect with the on-premises data center. Configure AWS Storage Gateway to store data locally. Use Storage Gateway to asynchronously back up point-in-time snapshots of the data to Amazon S3.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 106

A company is processing data on a daily basis. The results of the operations are stored in an Amazon S3 bucket, analyzed daily for one week, and then must remain immediately accessible for occasional analysis. What is the MOST cost-effective storage solution alternative to the current configuration?

- A. Configure a lifecycle policy to delete the objects after 30 days.
- B. Configure a lifecycle policy to transition the objects to Amazon S3 Glacier after 30 days.
- C. Configure a lifecycle policy to transition the objects to Amazon S3 Standard-Infrequent Access (S3 Standard-IA) after 30 days.
- D. Configure a lifecycle policy to transition the objects to Amazon S3 One Zone-Infrequent Access (S3 One Zone-IA) after 30 days.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 107

A company delivers files in Amazon S3 to certain users who do not have AWS credentials. These users must be given access for a limited time.

What should a solutions architect do to securely meet these requirements?

- A. Enable public access on an Amazon S3 bucket.
- B. Generate a presigned URL to share with the users.
- C. Encrypt files using AWS KMS and provide keys to the users.
- D. Create and assign IAM roles that will grant GetObject permissions to the users.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 108

A company wants to run a hybrid workload for data processing. The data needs to be accessed by on-premises applications for local data processing using an NFS protocol, and must also be accessible from the AWS Cloud for further analytics and batch processing.

Which solution will meet these requirements?

- A. Use an AWS Storage Gateway file gateway to provide file storage to AWS, then perform analytics on this data in the AWS Cloud.
- B. Use an AWS Storage Gateway tape gateway to copy the backup of the local data to AWS, then perform analytics on this data in the AWS cloud.
- C. Use an AWS Storage Gateway volume gateway in a stored volume configuration to regularly take snapshots of the local data, then copy the data to AWS.
- D. Use an AWS Storage Gateway volume gateway in a cached volume configuration to back up all the local storage in the AWS cloud, then perform analytics on this data in the cloud.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Reference: <https://aws.amazon.com/storagegateway/file/>

QUESTION 109

A company plans to store sensitive user data on Amazon S3. Internal security compliance requirement mandate encryption of data before sending it to Amazon S3.

What should a solutions architect recommend to satisfy these requirements?

- A. Server-side encryption with customer-provided encryption keys
- B. Client-side encryption with Amazon S3 managed encryption keys
- C. Server-side encryption with keys stored in AWS Key Management Service (AWS KMS)
- D. Client-side encryption with a master key stored in AWS Key Management Service (AWS KMS)

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 110

A solutions architect is moving the static content from a public website hosted on Amazon EC2 instances to an Amazon S3 bucket. An Amazon CloudFront distribution will be used to deliver the static assets. The security group used by the EC2 instances restricts access to a limited set of IP ranges. Access to the static content should be similarly restricted.

Which combination of steps will meet these requirements? (Choose two.)

- A. Create an origin access identity (OAI) and associate it with the distribution. Change the permissions in the bucket policy so that only the OAI can read the objects.
- B. Create an AWS WAF web ACL that includes the same IP restrictions that exist in the EC2 security group. Associate this new web ACL with the CloudFront distribution.
- C. Create a new security group that includes the same IP restrictions that exist in the current EC2 security group. Associate this new security group with the CloudFront distribution.
- D. Create a new security group that includes the same IP restrictions that exist in the current EC2 security group. Associate this new security group with the S3 bucket hosting the static content.
- E. Create a new IAM role and associate the role with the distribution. Change the permissions either on the S3 bucket or on the files within the S3 bucket so that only the newly created IAM role has read and download permissions.

Correct Answer: AB

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 111

A company is investigating potential solutions that would collect, process, and store users' service usage data. The business objective is to create an analytics capability that will enable the company to gather operational insights quickly using standard SQL queries. The solution should be highly available and ensure Atomicity, Consistency, Isolation, and Durability (ACID) compliance in the data tier.

Which solution should a solutions architect recommend?

- A. Use an Amazon Timestream database.



- B. Use an Amazon Neptune database in a Multi-AZ design.
- C. Use a fully managed Amazon RDS for MySQL database in a Multi-AZ design.
- D. Deploy PostgreSQL on an Amazon EC2 instance that uses Amazon Elastic Block Store (Amazon EBS) Throughput Optimized HDD (st1) storage.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 112

A company recently launched its website to serve content to its global user base. The company wants to store and accelerate the delivery of static content to its users by leveraging Amazon CloudFront with an Amazon EC2 instance attached as its origin.

How should a solutions architect optimize high availability for the application?

- A. Use Lambda@Edge for CloudFront.
- B. Use Amazon S3 Transfer Acceleration for CloudFront.
- C. Configure another EC2 instance in a different Availability Zone as part of the origin group.
- D. Configure another EC2 instance as part of the origin server cluster in the same Availability Zone.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 113

An application running on an Amazon EC2 instance in VPC-A needs to access files in another EC2 instance in VPC-B. Both are in separate AWS accounts. The network administrator needs to design a solution to enable secure access to EC2 instance in VPC-B from VPC-A. The connectivity should not have a single point of failure or bandwidth concerns.

Which solution will meet these requirements?

- A. Set up a VPC peering connection between VPC-A and VPC-B.
- B. Set up VPC gateway endpoints for the EC2 instance running in VPC-B.
- C. Attach a virtual private gateway to VPC-B and enable routing from VPC-A.
- D. Create a private virtual interface (VIF) for the EC2 instance running in VPC-B and add appropriate routes from VPC-B.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 114

A company currently stores symmetric encryption keys in a hardware security module (HSM). A solutions architect must design a solution to migrate key management to AWS. The solution should allow for key rotation and support the use of customer provided keys.

Where should the key material be stored to meet these requirements?

- A. Amazon S3
- B. AWS Secrets Manager
- C. AWS Systems Manager Parameter store
- D. AWS Key Management Service (AWS KMS)

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 115

A recent analysis of a company's IT expenses highlights the need to reduce backup costs. The company's chief information officer wants to simplify the on-premises backup infrastructure and reduce costs by eliminating the use of physical backup tapes. The company must preserve the existing investment in the on-premises backup applications and workflows.

What should a solutions architect recommend?

- A. Set up AWS Storage Gateway to connect with the backup applications using the NFS interface.
- B. Set up an Amazon Elastic File System (Amazon EFS) file system that connects with the backup applications using the NFS interface.
- C. Set up an Amazon Elastic File System (Amazon EFS) file system that connects with the backup applications using the iSCSI interface.
- D. Set up AWS Storage Gateway to connect with the backup applications using the iSCSI-virtual tape library (VTL) interface.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 116

A company hosts an application on an Amazon EC2 instance that requires a maximum of 200 GB storage space. The application is used infrequently, with peaks during mornings and evenings.

Disk I/O varies, but peaks at 3,000 IOPS. The chief financial officer of the company is concerned about costs and has asked a solutions architect to recommend the most cost-effective storage option that does not sacrifice performance.

Which solution should the solutions architect recommend?

- A. Amazon Elastic Block Store (Amazon EBS) Cold HDD (sc1)
- B. Amazon Elastic Block Store (Amazon EBS) General Purpose SSD (gp2)
- C. Amazon Elastic Block Store (Amazon EBS) Provisioned IOPS SSD (io1)
- D. Amazon Elastic Block Store (Amazon EBS) Throughput Optimized HDD (st1)

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 117

A company's application hosted on Amazon EC2 instances needs to access an Amazon S3 bucket. Due to data sensitivity, traffic cannot traverse the internet.

How should a solutions architect configure access?

- A. Create a private hosted zone using Amazon Route 53.
- B. Configure a VPC gateway endpoint for Amazon S3 in the VPC.
- C. Configure AWS PrivateLink between the EC2 instance and the S3 bucket.
- D. Set up a site-to-site VPN connection between the VPC and the S3 bucket.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 118

A company has two applications it wants to migrate to AWS. Both applications process a large set of files by accessing the same files at the same time. Both applications need to read the files with low latency.

Which architecture should a solutions architect recommend for this situation?

- A. Configure two AWS Lambda functions to run the applications. Create an Amazon EC2 instance with an instance store volume to store the data.
- B. Configure two AWS Lambda functions to run the applications. Create an Amazon EC2 instance with an Amazon Elastic Block Store (Amazon EBS) volume to store the data.



- C. Configure one memory optimized Amazon EC2 instance to run both applications simultaneously.
Create an Amazon Elastic Block Store (Amazon EBS) volume with Provisioned IOPS to store the data.
- D. Configure two Amazon EC2 instances to run both applications. Configure Amazon Elastic File System (Amazon EFS) with General Purpose performance mode and Bursting Throughput mode to store the data.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 119

An ecommerce company has noticed performance degradation of its Amazon RDS based web application. The performance degradation is attributed to an increase in the number of read-only SQL queries triggered by business analysts. A solutions architect needs to solve the problem with minimal changes to the existing web application. What should the solutions architect recommend?

- A. Export the data to Amazon DynamoDB and have the business analysts run their queries.
- B. Load the data into Amazon ElastiCache and have the business analysts run their queries.
- C. Create a read replica of the primary database and have the business analysts run their queries.
- D. Copy the data into an Amazon Redshift cluster and have the business analysts run their queries.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 120

A company is running a highly sensitive application on Amazon EC2 backed by an Amazon RDS database. Compliance regulations mandate that all personally identifiable information (PII) be encrypted at rest. Which solution should a solutions architect recommend to meet this requirement with the LEAST amount of changes to the infrastructure?

- A. Deploy AWS Certificate Manager to generate certificates. Use the certificates to encrypt the database volume.
- B. Deploy AWS CloudHSM, generate encryption keys, and use the customer master key (CMK) to encrypt database volumes.
- C. Configure SSL encryption using AWS Key Management Service customer master keys (AWS KMS CMKs) to encrypt database volumes.
- D. Configure Amazon Elastic Block Store (Amazon EBS) encryption and Amazon RDS encryption with AWS Key Management Service (AWS KMS) keys to encrypt instance and database volumes.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 121

A company running an on-premises application is migrating the application to AWS to increase its elasticity and availability. The current architecture uses a Microsoft SQL Server database with heavy read activity. The company wants to explore alternate database options and migrate database engines, if needed. Every 4 hours, the development team does a full copy of the production database to populate a test database. During this period, users experience latency. What should a solutions architect recommend as replacement database?

- A. Use Amazon Aurora with Multi-AZ Aurora Replicas and restore from mysqldump for the test database.
- B. Use Amazon Aurora with Multi-AZ Aurora Replicas and restore snapshots from Amazon RDS for the test database.
- C. Use Amazon RDS for MySQL with a Multi-AZ deployment and read replicas, and use the standby instance for the test database.
- D. Use Amazon RDS for SQL Server with a Multi-AZ deployment and read replicas, and restore snapshots from RDS for the test database.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 122

A company has enabled AWS CloudTrail logs to deliver log files to an Amazon S3 bucket for each of its developer accounts. The company has created a central AWS account for streamlining management and audit reviews. An internal auditor needs to access the CloudTrail logs, yet access needs to be restricted for all developer account users. The solution must be secure and optimized. How should a solutions architect meet these requirements?

- A. Configure an AWS Lambda function in each developer account to copy the log files to the central account. Create an IAM role in the central account for the auditor. Attach an IAM policy providing read-only permissions to the bucket.
- B. Configure CloudTrail from each developer account to deliver the log files to an S3 bucket in the central account. Create an IAM user in the central account for the auditor. Attach an IAM policy providing full permissions to the bucket.
- C. Configure CloudTrail from each developer account to deliver the log files to an S3 bucket in the central account. Create an IAM role in the central account for the auditor. Attach an IAM policy providing read-only permissions to the bucket.
- D. Configure an AWS Lambda function in the central account to copy the log files from the S3 bucket in each developer account. Create an IAM user in the central account for the auditor. Attach an IAM policy providing full permissions to the bucket.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 123

A company has several business systems that require access to data stored in a file share. The business systems will access the file share using the Server Message Block (SMB) protocol. The file share solution should be accessible from both of the company's legacy on-premises environments and with AWS. Which services meet the business requirements? (Choose two.)

- A. Amazon Elastic Block Store (Amazon EBS)
- B. Amazon Elastic File System (Amazon EFS)
- C. Amazon FSx for Windows
- D. Amazon S3
- E. AWS Storage Gateway file gateway



Correct Answer: CE

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 124

A company is using Amazon EC2 to run its big data analytics workloads. These variable workloads run each night, and it is critical they finish by the start of business the following day. A solutions architect has been tasked with designing the MOST cost-effective solution. Which solution will accomplish this?

- A. Spot Fleet
- B. Spot Instances
- C. Reserved Instances
- D. On-Demand Instances

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 125

A company has a Microsoft Windows-based application that must be migrated to AWS. This application requires the use of a shared Windows file system attached to multiple Amazon EC2 Windows instances. What should a solutions architect do to accomplish this?

- A. Configure a volume using Amazon Elastic File System (Amazon EFS). Mount the EFS volume to each Windows instance.
- B. Configure AWS Storage Gateway in Volume Gateway mode. Mount the volume to each Windows instance.
- C. Configure Amazon FSx for Windows File Server. Mount the Amazon FSx volume to each Windows instance.
- D. Configure an Amazon Elastic Block Store (Amazon EBS) volume with the required size. Attach each EC2 instance to the volume. Mount the file system within the volume to each Windows instance.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 126

A company has created an isolated backup of its environment in another Region. The application is running in warm standby mode and is fronted by an Application Load Balancer (ALB). The current failover process is manual and requires updating a DNS alias record to point to the secondary ALB in another Region.

What should a solutions architect do to automate the failover process?

- A. Enable an ALB health check
- B. Enable an Amazon Route 53 health check.
- C. Create a CNAME record on Amazon Route 53 pointing to the ALB endpoint.
- D. Create conditional forwarding rules on Amazon Route 53 pointing to an internal BIND DNS server.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 127

A company has a mobile chat application with a data store based in Amazon DynamoDB. Users would like new messages to be read with as little latency as possible. A solutions architect needs to design an optimal solution that requires minimal application changes.

Which method should the solutions architect select?

- A. Configure Amazon DynamoDB Accelerator (DAX) for the new messages table. Update the code to use the DAX endpoint.
- B. Add DynamoDB read replicas to handle the increased read load. Update the application to point to the read endpoint for the read replicas.
- C. Double the number of read capacity units for the new messages table in DynamoDB. Continue to use the existing DynamoDB endpoint.
- D. Add an Amazon ElastiCache for Redis cache to the application stack. Update the application to point to the Redis cache endpoint instead of DynamoDB.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Reference: <https://aws.amazon.com/blogs/aws/amazon-dynamodb-accelerator-dax-in-memorycaching-for-read-intensive-workloads/>

QUESTION 128

A company is creating an architecture for a mobile app that requires minimal latency for its users.

The company's architecture consists of Amazon EC2 instances behind an Application Load Balancer running in an Auto Scaling group. The EC2 instances connect to Amazon RDS.

Application beta testing showed there was a slowdown when reading the data. However, the metrics indicate that the EC2 instances do not cross any CPU utilization thresholds.

How can this issue be addressed?

- A. Reduce the threshold for CPU utilization in the Auto Scaling group.
- B. Replace the Application Load Balancer with a Network Load Balancer.
- C. Add read replicas for the RDS instances and direct read traffic to the replica.
- D. Add Multi-AZ support to the RDS instances and direct read traffic to the new EC2 instance.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 129

A company has implemented one of its microservices on AWS Lambda that accesses an Amazon DynamoDB table named Books. A solutions architect is designing an IAM policy to be attached to the Lambda function's IAM role, giving it access to put, update, and delete items in the Books table. The IAM policy must prevent function from performing any other actions on the Books table or any other.

Which IAM policy would fulfill these needs and provide the LEAST privileged access?

- A.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "PutUpdateDeleteOnBooks",
      "Effect": "Allow",
      "Action": [
        "dynamodb: PutItem",
        "dynamodb: UpdateItem",
        "dynamodb: DeleteItem"
      ],
      "Resource": "arn:aws:dynamodb:us-west-2:123456789012:table/Books"
    }
  ]
}
```
- B.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "PutUpdateDeleteOnBooks",
      "Effect": "Allow",
      "Action": [
        "dynamodb: PutItem",
        "dynamodb: UpdateItem",
        "dynamodb: DeleteItem"
      ],
      "Resource": "arn:aws:dynamodb:us-west-2:123456789012:table/*"
    }
  ]
}
```
- C.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "PutUpdateDeleteOnBooks",
      "Effect": "Allow",
      "Action": "dynamodb:*",
      "Resource": "arn:aws:dynamodb:us-west-2:123456789012:table/Books"
    }
  ]
}
```
- D.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "PutUpdateDeleteOnBooks",
      "Effect": "Allow",
      "Action": "dynamodb:*",
      "Resource": "arn:aws:dynamodb:us-west-2:123456789012:table/Books"
    },
    {
      "Sid": "PutUpdateDeleteOnBooks",
      "Effect": "Deny",
      "Action": "dynamodb:*",
      "Resource": "arn:aws:dynamodb:us-west-2:123456789012:table/Books"
    }
  ]
}
```



Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 130

A company hosts its website on Amazon S3. The website serves petabytes of outbound traffic monthly, which accounts for most of the company's AWS costs. What should a solutions architect do to reduce costs?

- A. Configure Amazon CloudFront with the existing website as the origin.
- B. Move the website to Amazon EC2 with Amazon Elastic Block Store (Amazon EBS) volumes for storage.
- C. Use AWS Global Accelerator and specify the existing website as the endpoint.
- D. Rearchitect the website to run on a combination of Amazon API Gateway and AWS Lambda.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 131

A company runs a website on Amazon EC2 instances behind an ELB Application Load Balancer. Amazon Route 53 is used for the DNS. The company wants to set up a backup website with a message including a phone number and email address that users can reach if the primary website is down. How should the company deploy this solution?

- A. Use Amazon S3 website hosting for the backup website and Route 53 failover routing policy.
- B. Use Amazon S3 website hosting for the backup website and Route 53 latency routing policy.
- C. Deploy the application in another AWS Region and use ELB health checks for failover routing.
- D. Deploy the application in another AWS Region and use server-side redirection on the primary website.

U&dumps

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 132

A media company is evaluating the possibility of moving its systems to the AWS Cloud. The company needs at least 10 TB of storage with the maximum possible I/O performance for video processing, 300 TB of very durable storage for storing media content, and 900 TB of storage to meet requirements for archival media that is not in use anymore. Which set of services should a solutions architect recommend to meet these requirements?

- A. Amazon Elastic Block Store (Amazon EBS) for maximum performance, Amazon S3 for durable data storage, and Amazon S3 Glacier for archival storage
- B. Amazon Elastic Block Store (Amazon EBS) for maximum performance, Amazon Elastic File System (Amazon EFS) for durable data storage, and Amazon S3 Glacier for archival storage
- C. Amazon EC2 instance store for maximum performance, Amazon Elastic File System (Amazon EFS) for durable data storage, and Amazon S3 for archival storage
- D. Amazon EC2 instance store for maximum performance, Amazon S3 for durable data storage, and Amazon S3 Glacier for archival storage

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 133

A company uses Amazon S3 as its object storage solution. The company has thousands of S3 buckets it uses to store data. Some of the S3 buckets have data that is accessed less frequently than others. A solutions architect found that lifecycle policies are not consistently implemented or are implemented partially, resulting in data being stored in high-cost storage. Which solution will lower costs without compromising the availability of objects?

- A. Use S3 ACLs.
- B. Use Amazon Elastic Block Store (Amazon EBS) automated snapshots.
- C. Use S3 Intelligent-Tiering storage.
- D. Use S3 One Zone-Infrequent Access (S3 One Zone-IA).

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 134

An application is running on Amazon EC2 instances. Sensitive information required for the application is stored in an Amazon S3 bucket. The bucket needs to be protected from internet access while only allowing services within the VPC access to the bucket.

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

- A. Create a VPC endpoint for Amazon S3.
- B. Enable server access logging on the bucket.
- C. Apply a bucket policy to restrict access to the S3 endpoint.
- D. Add an S3 ACL to the bucket that has sensitive information.
- E. Restrict users using the IAM policy to use the specific bucket.

Correct Answer: AC

Section: (none)

Explanation

Explanation/Reference:

Explanation:



QUESTION 135

A web application runs on Amazon EC2 instances behind an Application Load Balancer. The application allows users to create custom reports of historical weather data. Generating a report can take up to 5 minutes. These long-running requests use many of the available incoming connections, making the system unresponsive to other users.

How can a solutions architect make the system more responsive?

- A. Use Amazon SQS with AWS Lambda to generate reports.
- B. Increase the idle timeout on the Application Load Balancer to 5 minutes.
- C. Update the client-side application code to increase its request timeout to 5 minutes.
- D. Publish the reports to Amazon S3 and use Amazon CloudFront for downloading to the user.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 136

A solutions architect must create a highly available bastion host architecture. The solution needs to be resilient within a single AWS Region and should require only minimal effort to maintain.

What should the solutions architect do to meet these requirements?

- A. Create a Network Load Balancer backed by an Auto Scaling group with a UDP listener.
- B. Create a Network Load Balancer backed by a Spot Fleet with instances in a partition placement group.
- C. Create a Network Load Balancer backed by the existing servers in different Availability Zones as the target.
- D. Create a Network Load Balancer backed by an Auto Scaling group with instances in multiple Availability Zones as the target.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 137

A three-tier web application processes orders from customers. The web tier consists of Amazon EC2 instances behind an Application Load Balancer, a middle tier of three EC2 instances decoupled from the web tier using Amazon SQS, and an Amazon DynamoDB backend. At peak times, customers who submit orders using the site have to wait much longer than normal to receive confirmations due to lengthy processing times. A solutions architect needs to reduce these processing times.

Which action will be MOST effective in accomplishing this?

- A. Replace the SQS queue with Amazon Kinesis Data Firehose.
- B. Use Amazon ElastiCache for Redis in front of the DynamoDB backend tier.
- C. Add an Amazon CloudFront distribution to cache the responses for the web tier.
- D. Use Amazon EC2 Auto Scaling to scale out the middle tier instances based on the SQS queue depth.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 138

A company relies on an application that needs at least 4 Amazon EC2 instances during regular traffic and must scale up to 12 EC2 instances during peak loads. The application is critical to the business and must be highly available. Which solution will meet these requirements?

- A. Deploy the EC2 instances in an Auto Scaling group. Set the minimum to 4 and the maximum to 12, with 2 in Availability Zone A and 2 in Availability Zone B.
- B. Deploy the EC2 instances in an Auto Scaling group. Set the minimum to 4 and the maximum to 12, with all 4 in Availability Zone A.
- C. Deploy the EC2 instances in an Auto Scaling group. Set the minimum to 8 and the maximum to 12, with 4 in Availability Zone A and 4 in Availability Zone B.
- D. Deploy the EC2 instances in an Auto Scaling group. Set the minimum to 8 and the maximum to 12, with all 8 in Availability Zone A.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 139

A solutions architect must design a solution for a persistent database that is being migrated from on-premises to AWS. The database requires 64,000 IOPS according to the database administrator. If possible, the database administrator wants to use a single Amazon Elastic Block Store (Amazon EBS) volume to host the database instance.

Which solution effectively meets the database administrator's criteria?

- A. Use an instance from the I3 I/O optimized family and leverage local ephemeral storage to achieve the IOPS requirement.
- B. Create a Nitro-based Amazon EC2 instance with an Amazon Elastic Block Store (Amazon EBS) Provisioned IOPS SSD (io1) volume attached. Configure the volume to have 64,000 IOPS.
- C. Create and map an Amazon Elastic File System (Amazon EFS) volume to the database instance and use the volume to achieve the required IOPS for the database.
- D. Provision two volumes and assign 32,000 IOPS to each. Create a logical volume at the operating system level that aggregates both volumes to achieve the IOPS requirements.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 140

A solutions architect is designing an architecture for a new application that requires low network latency and high network throughput between Amazon EC2 instances. Which component should be included in the architectural design?

- A. An Auto Scaling group with Spot Instance types.

- B. A placement group using a cluster placement strategy.
- C. A placement group using a partition placement strategy.
- D. An Auto Scaling group with On-Demand instance types.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 141

A company has global users accessing an application deployed in different AWS Regions, exposing public static IP addresses. The users are experiencing poor performance when accessing the application over the internet. What should a solutions architect recommend to reduce internet latency?

- A. Set up AWS Global Accelerator and add endpoints.
- B. Set up AWS Direct Connect locations in multiple Regions.
- C. Set up an Amazon CloudFront distribution to access an application.
- D. Set up an Amazon Route 53 geoproximity routing policy to route traffic.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 142

A company wants to migrate a workload to AWS. The chief information security officer requires that all data be encrypted at rest when stored in the cloud. The company wants complete control of encryption key lifecycle management. The company must be able to immediately remove the key material and audit key usage independently of AWS CloudTrail. The chosen services should integrate with other storage services that will be used on AWS. Which services satisfies these security requirements?

- A. AWS CloudHSM with the CloudHSM client
- B. AWS Key Management Service (AWS KMS) with AWS CloudHSM
- C. AWS Key Management Service (AWS KMS) with an external key material origin
- D. AWS Key Management Service (AWS KMS) with AWS managed customer master keys (CMKs)

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 143

A company recently deployed a two-tier application in two Availability Zones in the us-east-1 Region. The databases are deployed in a private subnet while the web servers are deployed in a public subnet. An internet gateway is attached to the VPC. The application and database run on Amazon EC2 instances. The database servers are unable to access patches on the internet. A solutions architect needs to design a solution that maintains database security with the least operational overhead.

Which solution meets these requirements?

- A. Deploy a NAT gateway inside the public subnet for each Availability Zone and associate it with an Elastic IP address. Update the routing table of the private subnet to use it as the default route.
- B. Deploy a NAT gateway inside the private subnet for each Availability Zone and associate it with an Elastic IP address. Update the routing table of the private subnet to use it as the default route.
- C. Deploy two NAT instances inside the public subnet for each Availability Zone and associate them with Elastic IP addresses. Update the routing table of the private subnet to use it as the default route.
- D. Deploy two NAT instances inside the private subnet for each Availability Zone and associate them with Elastic IP addresses. Update the routing table of the private subnet to use it as the default route.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 144

A company has an application with a REST-based interface that allows data to be received in near-real time from a third-party vendor. Once received, the application processes and stores the data for further analysis. The application is running on Amazon EC2 instances.

The third-party vendor has received many 503 Service Unavailable Errors when sending data to the application. When the data volume spikes, the compute capacity reaches its maximum limit and the application is unable to process all requests.

Which design should a solutions architect recommend to provide a more scalable solution?

- A. Use Amazon Kinesis Data Streams to ingest the data. Process the data using AWS Lambda functions.
- B. Use Amazon API Gateway on top of the existing application. Create a usage plan with a quota limit for the third-party vendor.
- C. Use Amazon Simple Notification Service (Amazon SNS) to ingest the data. Put the EC2 instances in an Auto Scaling group behind an Application Load Balancer.
- D. Repackage the application as a container. Deploy the application using Amazon Elastic Container Service (Amazon ECS) using the EC2 launch type with an Auto Scaling group.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 145

A solutions architect needs to design a low-latency solution for a static single-page application accessed by users utilizing a custom domain name. The solution must be serverless, encrypted in transit, and cost-effective.

Which combination of AWS services and features should the solutions architect use? (Choose two.)

- A. Amazon S3
- B. Amazon EC2
- C. AWS Fargate
- D. Amazon CloudFront
- E. Elastic Load Balancer

Correct Answer: AD

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 146

A company is migrating to the AWS Cloud. A file server is the first workload to migrate. Users must be able to access the file share using the Server Message Block (SMB) protocol.

Which AWS managed service meets these requirements?

- A. Amazon Elastic Block Store (Amazon EBS)
- B. Amazon EC2
- C. Amazon FSx
- D. Amazon S3

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 147

A solutions architect is designing a customer-facing application. The application is expected to have a variable amount of reads and writes depending on the time of year and clearly defined access patterns throughout the year. Management requires that database auditing and scaling be managed in the AWS Cloud. The Recovery Point Objective (RPO) must be less than 5 hours.

Which solutions can accomplish this? (Choose two.)



- A. Use Amazon DynamoDB with auto scaling. Use on-demand backups and AWS CloudTrail.
- B. Use Amazon DynamoDB with auto scaling. Use on-demand backups and Amazon DynamoDB Streams.
- C. Use Amazon Redshift Configure concurrency scaling. Enable audit logging. Perform database snapshots every 4 hours.
- D. Use Amazon RDS with Provisioned IOPS. Enable the database auditing parameter. Perform database snapshots every 5 hours.
- E. Use Amazon RDS with auto scaling. Enable the database auditing parameter. Configure the backup retention period to at least 1 day.

Correct Answer: AB

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 148

A company has migrated an on-premises Oracle database to an Amazon RDS for Oracle Multi-AZ DB instance in the us-east-1 Region. A solutions architect is designing a disaster recovery strategy to have the database provisioned in the us-west-2 Region in case the database becomes unavailable in the us-east-1 Region. The design must ensure the database is provisioned in the us-west-2 Region in a maximum of 2 hours, with a data loss window of no more than 3 hours. How can these requirements be met?

- A. Edit the DB instance and create a read replica in us-west-2. Promote the read replica to master in us-west-2 in case the disaster recovery environment needs to be activated.
- B. Select the multi-Region option to provision a standby instance in us-west-2. The standby instance will be automatically promoted to master in us-west-2 in case the disaster recovery environment needs to be created.
- C. Take automated snapshots of the database instance and copy them to us-west-2 every 3 hours.
Restore the latest snapshot to provision another database instance in us-west-2 in case the disaster recovery environment needs to be activated.
- D. Create a multimaster read/write instances across multiple AWS Regions. Select VPCs in us-east-1 and us-west-2 to make that deployment. Keep the master read/write instance in us-west-2 available to avoid having to activate a disaster recovery environment.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:



QUESTION 149

A monolithic application was recently migrated to AWS and is now running on a single Amazon EC2 instance. Due to application limitations, it is not possible to use automatic scaling to scale out the application. The chief technology officer (CTO) wants an automated solution to restore the EC2 instance in the unlikely event the underlying hardware fails. What would allow for automatic recovery of the EC2 instance as quickly as possible?

- A. Configure an Amazon CloudWatch alarm that triggers the recovery of the EC2 instance if it becomes impaired.
- B. Configure an Amazon CloudWatch alarm to trigger an SNS message that alerts the CTO when the EC2 instance is impaired.
- C. Configure AWS CloudTrail to monitor the health of the EC2 instance, and if it becomes impaired, trigger instance recovery.
- D. Configure an Amazon EventBridge event to trigger an AWS Lambda function once an hour that checks the health of the EC2 instance and triggers instance recovery if the EC2 instance is unhealthy.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 150

A solutions architect is working on optimizing a legacy document management application running on Microsoft Windows Server in an on-premises data center. The application stores a large number of files on a network file share. The chief information officer wants to reduce the on-premises data center footprint and minimize storage costs by moving on-premises storage to AWS. What should the solutions architect do to meet these requirements?

- A. Set up an AWS Storage Gateway file gateway.
- B. Set up Amazon Elastic File System (Amazon EFS)
- C. Set up AWS Storage Gateway as a volume gateway
- D. Set up an Amazon Elastic Block Store (Amazon EBS) volume.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 151

A solutions architect is designing a hybrid application using the AWS cloud. The network between the on-premises data center and AWS will use an AWS Direct Connect (DX) connection. The application connectivity between AWS and the on-premises data center must be highly resilient.

Which DX configuration should be implemented to meet these requirements?

- A. Configure a DX connection with a VPN on top of it.
- B. Configure DX connections at multiple DX locations.
- C. Configure a DX connection using the most reliable DX partner.
- D. Configure multiple virtual interfaces on top of a DX connection.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 152

A company runs an application on Amazon EC2 instances. The application is deployed in private subnets in three Availability Zones of the us-east-1 Region. The instances must be able to connect to the internet to download files. The company wants a design that is highly available across the Region.

Which solution should be implemented to ensure that there are no disruptions to internet connectivity?

- A. Deploy a NAT instance in a private subnet of each Availability Zone.
- B. Deploy a NAT gateway in a public subnet of each Availability Zone.
- C. Deploy a transit gateway in a private subnet of each Availability Zone.
- D. Deploy an internet gateway in a public subnet of each Availability Zone.



Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 153

Application developers have noticed that a production application is very slow when business reporting users run large production reports against the Amazon RDS instance backing the application. The CPU and memory utilization metrics for the RDS instance do not exceed 60% while the reporting queries are running. The business reporting users must be able to generate reports without affecting the application's performance.

Which action will accomplish this?

- A. Increase the size of the RDS instance.
- B. Create a read replica and connect the application to it.
- C. Enable multiple Availability Zones on the RDS instance.
- D. Create a read replica and connect the business reports to it.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 154

A company is running a two-tier ecommerce website using AWS services. The current architect uses a publish-facing Elastic Load Balancer that sends traffic to Amazon EC2 instances in a private subnet. The static content is hosted on EC2

instances, and the dynamic content is retrieved from a MySQL database. The application is running in the United States. The company recently started selling to users in Europe and Australia. A solutions architect needs to design solution so their international users have an improved browsing experience. Which solution is MOST cost-effective?

- A. Host the entire website on Amazon S3.
- B. Use Amazon CloudFront and Amazon S3 to host static images.
- C. Increase the number of public load balancers and EC2 instances.
- D. Deploy the two-tier website in AWS Regions in Europe and Australia.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 155

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

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

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:



QUESTION 156

A company wants to deploy a shared file system for its .NET application servers and Microsoft SQL Server databases running on Amazon EC2 instances with Windows Server 2016. The solution must be able to be integrated into the corporate Active Directory domain, be highly durable, be managed by AWS, and provide high levels of throughput and IOPS. Which solution meets these requirements?

- A. Use Amazon FSx for Windows File Server.
- B. Use Amazon Elastic File System (Amazon EFS).
- C. Use AWS Storage Gateway in file gateway mode.
- D. Deploy a Windows file server on two On Demand instances across two Availability Zones.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 157

A company that develops web applications has launched hundreds of Application Load Balancers (ALBs) in multiple Regions. The company wants to create an allow list for the IPs of all the load balancers on its firewall device. A solutions architect is looking for a one-time, highly available solution to address this request, which will also help reduce the number of IPs that need to be allowed by the firewall. What should the solutions architect recommend to meet these requirements?

- A. Create a AWS Lambda function to keep track of the IPs for all the ALBs in different Regions. Keep refreshing this list.
- B. Set up a Network Load Balancer (NLB) with Elastic IPs. Register the private IPs of all the ALBs as targets to this NLB.
- C. Launch AWS Global Accelerator and create endpoints for all the Regions. Register all the ALBs in different Regions to the corresponding endpoints.

D. Set up an Amazon EC2 instance, assign an Elastic IP to this EC2 instance, and configure the instance as a proxy to forward traffic to all the ALBs.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 158

A company runs an application using Amazon ECS. The application creates resized versions of an original image and then makes Amazon S3 API calls to store the resized images in Amazon S3. How can a solutions architect ensure that the application has permission to access Amazon S3?

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

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 159

A company is planning to migrate its virtual server-based workloads to AWS. The company has internet-facing load balancers backed by application servers. The application servers rely on patches from an internet-hosted repository. Which services should a solutions architect recommend be hosted on the public subnet? (Choose two.)

- A. NAT gateway
- B. Amazon RDS DB instances
- C. Application Load Balancers
- D. Amazon EC2 application servers
- E. Amazon Elastic File System (Amazon EFS) volumes

The logo for 'Udumps' features a stylized orange 'U' followed by the word 'dumps' in a grey, lowercase, sans-serif font.

Correct Answer: AC

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 160

A company has established a new AWS account. The account is newly provisioned and no changes have been made to the default settings. The company is concerned about the security of the AWS account root user. What should be done to secure the root user?

- A. Create IAM users for daily administrative tasks. Disable the root user.
- B. Create IAM users for daily administrative tasks. Enable multi-factor authentication on the root user.
- C. Generate an access key for the root user. Use the access key for daily administration tasks instead of the AWS Management Console.
- D. Provide the root user credentials to the most senior solutions architect. Have the solutions architect use the root user for daily administration tasks.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 161

A company is using a tape backup solution to store its key application data offsite. The daily data volume is around 50 TB. The company needs to retain the backups for 7 years for regulatory purposes. The backups are rarely accessed, and a week's notice is typically given if a backup needs to be restored. The company is now considering a cloud-based option to reduce the storage costs and operational burden of managing tapes. The company also wants to make sure that the transition from tape backups to the cloud minimizes disruptions. Which storage solution is MOST cost-effective?

- A. Use Amazon Storage Gateway to back up to Amazon Glacier Deep Archive.
- B. Use AWS Snowball Edge to directly integrate the backups with Amazon S3 Glacier.
- C. Copy the backup data to Amazon S3 and create a lifecycle policy to move the data to Amazon S3 Glacier.
- D. Use Amazon Storage Gateway to back up to Amazon S3 and create a lifecycle policy to move the backup to Amazon S3 Glacier.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:

QUESTION 162

A company requires a durable backup storage solution for its on-premises database servers while ensuring on-premises applications maintain access to these backups for quick recovery. The company will use AWS storage services as the destination for these backups. A solutions architect is designing a solution with minimal operational overhead. Which solution should the solutions architect implement?

- A. Deploy an AWS Storage Gateway file gateway on-premises and associate it with an Amazon S3 bucket.
- B. Back up the databases to an AWS Storage Gateway volume gateway and access it using the Amazon S3 API.
- C. Transfer the database backup files to an Amazon Elastic Block Store (Amazon EBS) volume attached to an Amazon EC2 instance.
- D. Back up the database directly to an AWS Snowball device and use lifecycle rules to move the data to Amazon S3 Glacier Deep Archive.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:



QUESTION 163

A company decides to migrate its three-tier web application from on-premises to the AWS Cloud. The new database must be capable of dynamically scaling storage capacity and performing table joins. Which AWS service meets these requirements?

- A. Amazon Aurora
- B. Amazon RDS for SqlServer
- C. Amazon DynamoDB Streams
- D. Amazon DynamoDB on-demand

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:

QUESTION 164

A company mandates that an Amazon S3 gateway endpoint must allow traffic to trusted buckets only. Which method should a solutions architect implement to meet this requirement?

- A. Create a bucket policy for each of the company's trusted S3 buckets that allows traffic only from the company's trusted VPCs.
- B. Create a bucket policy for each of the company's trusted S3 buckets that allows traffic only from the company's S3 gateway endpoint IDs.
- C. Create an S3 endpoint policy for each of the company's S3 gateway endpoints that blocks access from any VPC other than the company's trusted VPCs.
- D. Create an S3 endpoint policy for each of the company's S3 gateway endpoints that provides access to the Amazon Resource Name (ARN) of the trusted S3 buckets.

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Explanation:

QUESTION 165

A company is using a VPC peering strategy to connect its VPCs in a single Region to allow for cross-communication. A recent increase in account creations and VPCs has made it difficult to maintain the VPC peering strategy, and the company expects to grow to hundreds of VPCs. There are also new requests to create site-to-site VPNs with some of the VPCs. A solutions architect has been tasked with creating a centrally managed networking setup for multiple accounts, VPCs, and VPNs. Which networking solution meets these requirements?

- A. Configure shared VPCs and VPNs and share to each other.
- B. Configure a hub-and-spoke VPC and route all traffic through VPC peering.
- C. Configure an AWS Direct Connect connection between all VPCs and VPNs.
- D. Configure a transit gateway with AWS Transit Gateway and connect all VPCs and VPNs.

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Explanation:

QUESTION 166

A solutions architect is helping a developer design a new ecommerce shopping cart application using AWS services. The developer is unsure of the current database schema and expects to make changes as the ecommerce site grows. The solution needs to be highly resilient and capable of automatically scaling read and write capacity. Which database solution meets these requirements?

- A. Amazon Aurora PostgreSQL
- B. Amazon DynamoDB with on-demand enabled
- C. Amazon DynamoDB with DynamoDB Streams enabled
- D. Amazon SQS and Amazon Aurora PostgreSQL



Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:

QUESTION 167

A solutions architect must migrate a Windows internet information Services (IIS) web application to AWS. The application currently relies on a file share hosted in the user's on-premises networkattached storage (NAS). The solutions architect has proposed migrating the IIS web servers to Amazon EC2 instances in multiple Availability Zones that are connected to the storage solution, and configuring an Elastic Load Balancer attached to the instances. Which replacement to the on-premises file share is MOST resilient and durable?

- A. Migrate the file Share to Amazon RDS.
- B. Migrate the file Share to AWS Storage Gateway
- C. Migrate the file Share to Amazon FSx for Windows File Server.
- D. Migrate the file share to Amazon Elastic File System (Amazon EFS)

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:

QUESTION 168

A company needs to implement a relational database with a multi-Region disaster recovery Recovery Point Objective (RPO) of 1 second and a Recovery Time Objective (RTO) of 1 minute. Which AWS solution can achieve this?

- A. Amazon Aurora Global Database
- B. Amazon DynamoDB global tables
- C. Amazon RDS for MySQL with Multi-AZ enabled
- D. Amazon RDS for MySQL with a cross-Region snapshot copy

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 169

A company runs a web service on Amazon EC2 instances behind an Application Load Balancer.

The instances run in an Amazon EC2 Auto Scaling group across two Availability Zones. The company needs a minimum of four instances at all times to meet the required service level agreement (SLA) while keeping costs low. If an Availability Zone fails, how can the company remain compliant with the SLA?

- A. Add a target tracking scaling policy with a short cooldown period.
- B. Change the Auto Scaling group launch configuration to use a larger instance type.
- C. Change the Auto Scaling group to use six servers across three Availability Zones.
- D. Change the Auto Scaling group to use eight servers across two Availability Zones.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

**QUESTION 170**

A company is reviewing its AWS Cloud deployment to ensure its data is not accessed by anyone without appropriate authorization. A solutions architect is tasked with identifying all open Amazon S3 buckets and recording any S3 bucket configuration changes.

What should the solutions architect do to accomplish this?

- A. Enable AWS Config service with the appropriate rules
- B. Enable AWS Trusted Advisor with the appropriate checks.
- C. Write a script using an AWS SDK to generate a bucket report
- D. Enable Amazon S3 server access logging and configure Amazon CloudWatch Events.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 171

A company is planning to build a new web application on AWS. The company expects predictable traffic most of the year and very high traffic on occasion. The web application needs to be highly available and fault tolerant with minimal latency.

What should a solutions architect recommend to meet these requirements?

- A. Use an Amazon Route 53 routing policy to distribute requests to two AWS Regions, each with one Amazon EC2 instance.
- B. Use Amazon EC2 instances in an Auto Scaling group with an Application Load Balancer across multiple Availability Zones.
- C. Use Amazon EC2 instances in a cluster placement group with an Application Load Balancer across multiple Availability Zones.
- D. Use Amazon EC2 instances in a cluster placement group and include the cluster placement group within a new Auto Scaling group.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:

QUESTION 172

A company is designing a web application using AWS that processes insurance quotes. Users will request quotes from the application. Quotes must be separated by quote type must be responded to within 24 hours, and must not be lost. The solution should be simple to set up and maintain. Which solution meets these requirements?

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

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Explanation:

QUESTION 173

A solutions architect has configured the following IAM policy.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "lambda:*"
      ],
      "Resource": "*"
    },
    {
      "Effect": "Deny",
      "Action": [
        "lambda:CreateFunction",
        "lambda>DeleteFunction"
      ],
      "Resource": "*"
      "Condition": {
        "IpAddress": {
          "aws:SourceIp": "220.100.16.0/20"
        }
      }
    }
  ]
}
```

Which action will be allowed by the policy?

- A. An AWS Lambda function can be deleted from any network.
- B. An AWS Lambda function can be created from any network.



- C. An AWS Lambda function can be deleted from the 100.220.0.0/20 network.
- D. An AWS Lambda function can be deleted from the 220.100.16.0/20 network.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 174

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)

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 175

A company is running a three-tier web application to process credit card payments. The front-end user interface consists of static webpages. The application tier can have long-running processes. The database tier uses MySQL.

The application is currently running on a single, general purpose large Amazon EC2 instance. A solutions architect needs to decouple the services to make the web application highly available.

Which solution would provide the HIGHEST availability?

- A. Move static assets to Amazon CloudFront. Leave the application in EC2 in an Auto Scaling group. Move the database to Amazon RDS to deploy Multi-AZ.
- B. Move static assets and the application into a medium EC2 instance. Leave the database on the large instance. Place both instances in an Auto Scaling group.
- C. Move static assets to Amazon S3. Move the application to AWS Lambda with the concurrency limit set. Move the database to Amazon DynamoDB with on-demand enabled.
- D. Move static assets to Amazon S3. Move the application to Amazon Elastic Container Service (Amazon ECS) containers with Auto Scaling enabled. Move the database to Amazon RDS to deploy Multi-AZ.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 176

A media company stores video content in an Amazon Elastic Block Store (Amazon EBS) volume.

A certain video file has become popular and a large number of users across the world are accessing this content. This has resulted in a cost increase.

Which action will DECREASE cost without compromising user accessibility?

- A. Change the EBS volume to Provisioned IOPS (PIOPS).
- B. Store the video in an Amazon S3 bucket and create an Amazon CloudFront distribution.
- C. Split the video into multiple, smaller segments so users are routed to the requested video segments only.
- D. Clear an Amazon S3 bucket in each Region and upload the videos so users are routed to the nearest S3 bucket.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 177

A solutions architect is designing the cloud architecture for a new application being deployed to AWS. The application allows users to interactively download and upload files. Files older than 2 years will be accessed less frequently. The solutions architect needs to ensure that the application can scale to any number of files while maintaining high availability and durability.

Which scalable solutions should the solutions architect recommend? (Choose two.)

- A. Store the files on Amazon S3 with a lifecycle policy that moves objects older than 2 years to S3 Glacier.
- B. Store the files on Amazon S3 with a lifecycle policy that moves objects older than 2 years to S3 Standard-Infrequent Access (S3 Standard-IA)
- C. Store the files on Amazon Elastic File System (Amazon EFS) with a lifecycle policy that moves objects older than 2 years to EFS Infrequent Access (EFS IA).
- D. Store the files in Amazon Elastic Block Store (Amazon EBS) volumes. Schedule snapshots of the volumes. Use the snapshots to archive data older than 2 years.
- E. Store the files in RAID-striped Amazon Elastic Block Store (Amazon EBS) volumes. Schedule snapshots of the volumes. Use the snapshots to archive data older than 2 years.

Correct Answer: AC

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 178

A company has recently updated its internal security standards. The company must now ensure all Amazon S3 buckets and Amazon Elastic Block Store (Amazon EBS) volumes are encrypted with keys created and periodically rotated by internal security specialists. The company is looking for a native, software-based AWS service to accomplish this goal.

What should a solutions architect recommend as a solution?

- A. Use AWS Secrets Manager with customer master keys (CMKs) to store master key material and apply a routine to create a new CMK periodically and replace it in AWS Secrets Manager.
- B. Use AWS Key Management Service (AWS KMS) with customer master keys (CMKs) to store master key material and apply a routine to re-create a new key periodically and replace it in AWS KMS.
- C. Use an AWS CloudHSM cluster with customer master keys (CMKs) to store master key material and apply a routine to re-create a new key periodically and replace it in the CloudHSM cluster nodes.
- D. Use AWS Systems Manager Parameter Store with customer master keys (CMKs) to store master key material and apply a routine to re-create a new key periodically and replace it in the Parameter Store.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 179

A company's dynamic website is hosted using on-premises servers in the United States. The company is launching its product in Europe, and it wants to optimize site loading times for new European users. The site's backend must remain in the United States. The product is being launched in a few days, and an immediate solution is needed.

What should the solutions architect recommend?

- A. Launch an Amazon EC2 instance in us-east-1 and migrate the site to it.
- B. Move the website to Amazon S3. Use cross-Region replication between Regions.
- C. Use Amazon CloudFront with a custom origin pointing to the on-premises servers.
- D. Use an Amazon Route 53 geo-proximity routing policy pointing to on-premises servers.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 180

A development team needs to host a website that will be accessed by other teams. The website contents consist of HTML, CSS, client-side JavaScript, and images.

Which method is the MOST cost-effective for hosting the website?

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

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 181

A company is hosting multiple websites for several lines of business under its registered parent domain. Users accessing these websites will be routed to appropriate backend Amazon EC2 instances based on the subdomain. The websites host static webpages, images, and server-side scripts like PHP and JavaScript.

Some of the websites experience peak access during the first two hours of business with constant usage throughout the rest of the day. A solutions architect needs to design a solution that will automatically adjust capacity to these traffic patterns while keeping costs low.

Which combination of AWS services or features will meet these requirements? (Choose two.)

- A. AWS Batch
- B. Network Load Balancer
- C. Application Load Balancer
- D. Amazon EC2 Auto Scaling
- E. Amazon S3 website hosting

Correct Answer: DE

Section: (none)

Explanation

Explanation/Reference:

Explanation:



QUESTION 182

A company uses an Amazon S3 bucket to store static images for its website. The company configured permissions to allow access to Amazon S3 objects by privileged users only.

What should a solutions architect do to protect against data loss? (Choose two.)

- A. Enable versioning on the S3 bucket.
- B. Enable access logging on the S3 bucket.
- C. Enable server-side encryption on the S3 bucket.
- D. Configure an S3 lifecycle rule to transition objects to Amazon S3 Glacier.
- E. Use MFA Delete to require multi-factor authentication to delete an object.

Correct Answer: AE

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 183

An operations team has a standard that states IAM policies should not be applied directly to users.

Some new team members have not been following this standard. The operations manager needs a way to easily identify the users with attached policies.

What should a solutions architect do to accomplish this?

- A. Monitor using AWS CloudTrail.
- B. Create an AWS Config rule to run daily.
- C. Publish IAM user changes to Amazon SNS.
- D. Run AWS Lambda when a user is modified.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:

QUESTION 184

A company wants to use an AWS Region as a disaster recovery location for its on-premises infrastructure. The company has 10 TB of existing data, and the on-premises data center has a 1 Gbps internet connection. A solutions architect must find a solution so the company can have its existing data on AWS in 72 hours without transmitting it using an unencrypted channel. Which solution should the solutions architect select?

- A. Send the initial 10 TB of data to AWS using FTP.
- B. Send the initial 10 TB of data to AWS using AWS Snowball.
- C. Establish a VPN connection between Amazon VPC and the company's data center.
- D. Establish an AWS Direct Connect connection between Amazon VPC and the company's data center.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:

QUESTION 185

A company is building applications in containers. The company wants to migrate its on-premises development and operations services from its on-premises data center to AWS. Management states that production systems must be cloud agnostic and use the same configuration and administrator tools across production systems. A solutions architect needs to design a managed solution that will align open-source software. Which solution meets these requirements?

- A. Launch the containers on Amazon EC2 with EC2 instance worker nodes.
- B. Launch the containers on Amazon Elastic Kubernetes Service (Amazon EKS) and EKS worker nodes.
- C. Launch the containers on Amazon Elastic Containers service (Amazon ECS) with AWS Fargate instances.
- D. Launch the containers on Amazon Elastic Container Service (Amazon ECS) with Amazon EC2 instance worker nodes.



Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:

QUESTION 186

A company hosts its website on AWS. To address the highly variable demand, the company has implemented Amazon EC2 Auto Scaling. Management is concerned that the company is overprovisioning its infrastructure, especially at the front end of the three-tier application. A solutions architect needs to ensure costs are optimized without impacting performance. What should the solutions architect do to accomplish this?

- A. Use Auto Scaling with Reserved Instances.
- B. Use Auto Scaling with a scheduled scaling policy.
- C. Use Auto Scaling with the suspend-resume feature.
- D. Use Auto Scaling with a target tracking scaling policy.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:

QUESTION 187

A solutions architect is performing a security review of a recently migrated workload. The workload is a web application that consists of Amazon EC2 instances in an Auto Scaling group behind an Application Load Balancer. The solutions architect must improve the security posture and minimize the impact of a DDoS attack on resources. Which solution is MOST effective?

- A. Configure an AWS WAF ACL with rate-based rules. Create an Amazon CloudFront distribution that points to the Application Load Balancer. Enable the WAF ACL on the CloudFront distribution.
- B. Create a custom AWS Lambda function that adds identified attacks into a common vulnerability pool to capture a potential DDoS attack. Use the identified information to modify a network ACL to block access.
- C. Enable VPC Flow Logs and store them in Amazon S3. Create a custom AWS Lambda function that parses the logs looking for a DDoS attack. Modify a network ACL to block identified source IP addresses.
- D. Enable Amazon GuardDuty and configure findings written to Amazon CloudWatch. Create an event with CloudWatch Events for DDoS alerts that triggers Amazon Simple Notification Service (Amazon SNS). Have Amazon SNS invoke a custom AWS Lambda function that parses the logs, looking for a DDoS attack. Modify a network ACL to block identified source IP addresses.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 188

A company has multiple AWS accounts for various departments. One of the departments wants to share an Amazon S3 bucket with all other departments. Which solution will require the LEAST amount of effort?

- A. Enable cross-account S3 replication for the bucket.
- B. Create a pre-signed URL for the bucket and share it with other departments.
- C. Set the S3 bucket policy to allow cross-account access to other departments.
- D. Create IAM users for each of the departments and configure a read-only IAM policy.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Reference: <https://docs.aws.amazon.com/AmazonS3/latest/dev/example-walkthroughs-managing-access-example2.html>



QUESTION 189

A company needs to share an Amazon S3 bucket with an external vendor. The bucket owner must be able to access all objects. Which action should be taken to share the S3 bucket?

- A. Update the bucket to be a Requester Pays bucket.
- B. Update the bucket to enable cross-origin resource sharing (CORS).
- C. Create a bucket policy to require users to grant bucket-owner-full-control when uploading objects.
- D. Create an IAM policy to require users to grant bucket-owner-full-control when uploading objects.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

By default, an S3 object is owned by the AWS account that uploaded it. This is true even when the bucket is owned by another account. To get access to the object, the object owner must explicitly grant you (the bucket owner) access. The object owner can grant the bucket owner full control of the object by updating the access control list (ACL) of the object. The object owner can update the ACL either during a put or copy operation, or after the object is added to the bucket. Similar: <https://aws.amazon.com/it/premiumsupport/knowledge-center/s3-require-objectownership/Resolution> Add a bucket policy that grants users access to put objects in your bucket only when they grant you (the bucket owner) full control of the object.

Reference:

<https://aws.amazon.com/it/premiumsupport/knowledge-center/s3-bucket-owner-access/>

QUESTION 190

A company is developing a real-time multiplier game that uses UDP for communications between client and servers in an Auto Scaling group. Spikes in demand are anticipated during the day, so the game server platform must adapt accordingly. Developers want to store gamer scores and other non-relational data in a database solution that will scale without intervention. Which solution should a solutions architect recommend?

- A. Use Amazon Route 53 for traffic distribution and Amazon Aurora Serverless for data storage.
- B. Use a Network Load Balancer for traffic distribution and Amazon DynamoDB on-demand for data storage.
- C. Use a Network Load Balancer for traffic distribution and Amazon Aurora Global Database for data storage.
- D. Use an Application Load Balancer for traffic distribution and Amazon DynamoDB global tables for data storage.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 191

A company collects temperature, humidity, and atmospheric pressure data in cities across multiple continents. The average volume of data collected per site each day is 500 GB. Each site has a high-speed internet connection. The company's weather forecasting applications are based in a single Region and analyze the data daily.

What is the FASTEST way to aggregate data from all of these global sites?

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

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Step-1: To transfer to S3 from global sites: Amazon S3 Transfer Acceleration enables fast, easy, and secure transfers of files over long distances between your client and your Amazon S3 bucket.

S3 Transfer Acceleration leverages Amazon CloudFront's globally distributed AWS Edge Locations. Used to accelerate object uploads to S3 over long distances (latency). Transfer acceleration is as secure as a direct upload to S3.

Step-2: When the application analyze/aggregate the data from S3 and then again upload the results - Multipart upload

Reference: <http://lavnish.blogspot.com/2017/06/aws-s3-cross-region-replication.html>

<https://aws.amazon.com/s3/transfer-acceleration/>

QUESTION 192

A company has a custom application running on an Amazon EC instance that:

- Reads a large amount of data from Amazon S3
- Performs a multi-stage analysis
- Writes the results to Amazon DynamoDB

The application writes a significant number of large, temporary files during the multi-stage analysis. The process performance depends on the temporary storage performance.

What would be the fastest storage option for holding the temporary files?

- A. Multiple Amazon S3 buckets with Transfer Acceleration for storage.
- B. Multiple Amazon Elastic Block Store (Amazon EBS) drives with Provisioned IOPS and EBS optimization.
- C. Multiple Amazon Elastic File System (Amazon EFS) volumes using the Network File System version 4.1 (NFSv4.1) protocol.
- D. Multiple instance store volumes with software RAID 0.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 193

A leasing company generates and emails PDF statements every month for all its customers. Each statement is about 400 KB in size. Customers can download their statements from the website for up to 30 days from when the statements were generated. At the end of their 3-year lease, the customers are emailed a ZIP file that contains all the statements.

What is the MOST cost-effective storage solution for this situation?

- A. Store the statements using the Amazon S3 Standard storage class. Create a lifecycle policy to move the statements to Amazon S3 Glacier storage after 1 day.
- B. Store the statements using the Amazon S3 Glacier storage class. Create a lifecycle policy to move the statements to Amazon S3 Glacier Deep Archive storage after 30 days.
- C. Store the statements using the Amazon S3 Standard storage class. Create a lifecycle policy to move the statements to Amazon S3 One Zone-Infrequent Access (S3 One Zone-IA) storage after 30 days.
- D. Store the statements using the Amazon S3 Standard-Infrequent Access (S3 Standard-IA) storage class. Create a lifecycle policy to move the statements to Amazon S3 Glacier storage after 30 days.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 194

A company recently released a new type of internet-connected sensor. The company is expecting to sell thousands of sensors, which are designed to stream high volumes of data each second to a central location. A solutions architect must design a solution that ingests and stores data so that engineering teams can analyze it in near-real time with millisecond responsiveness.

Which solution should the solutions architect recommend?

- A. Use an Amazon SQS queue to ingest the data. Consume the data with an AWS Lambda function, which then stores the data in Amazon Redshift.
- B. Use an Amazon SQS queue to ingest the data. Consume the data with an AWS Lambda function, which then stores the data in Amazon DynamoDB.
- C. Use Amazon Kinesis Data Streams to ingest the data. Consume the data with an AWS Lambda function, which then stores the data in Amazon Redshift.
- D. Use Amazon Kinesis Data Streams to ingest the data. Consume the data with an AWS Lambda function, which then stores the data in Amazon DynamoDB.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 195

A website runs a web application that receives a burst of traffic each day at noon. The users upload new pictures and content daily, but have been complaining of timeouts. The architecture uses Amazon EC2 Auto Scaling groups, and the custom application consistently takes 1 minute to initiate upon boot up before responding to user requests.

How should a solutions architect redesign the architecture to better respond to changing traffic?

- A. Configure a Network Load Balancer with a slow start configuration.
- B. Configure AWS ElastiCache for Redis to offload direct requests to the servers.
- C. Configure an Auto Scaling step scaling policy with an instance warmup condition.
- D. Configure Amazon CloudFront to use an Application Load Balancer as the origin.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 196

A company is concerned that two NAT instances in use will no longer be able to support the traffic needed for the company's application. A solutions architect wants to implement a solution that is highly available fault tolerant, and automatically scalable.

What should the solutions architect recommend?

- A. Remove the two NAT instances and replace them with two NAT gateways in the same Availability Zone.
- B. Use Auto Scaling groups with Network Load Balancers for the NAT instances in different Availability Zones.
- C. Remove the two NAT instances and replace them with two NAT gateways in different Availability Zones.
- D. Replace the two NAT instances with Spot Instances in different Availability Zones and deploy a Network Load Balancer.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 197

A company operates a website on Amazon EC2 Linux instances. Some of the instances are failing. Troubleshooting points to insufficient swap space on the failed instances. The operations team lead needs a solution to monitor this. What should a solutions architect recommend?

- A. Configure an Amazon CloudWatch SwapUsage metric dimension. Monitor the SwapUsage dimension in the EC2 metrics in CloudWatch.
- B. Use EC2 metadata to collect information, then publish it to Amazon CloudWatch custom metrics. Monitor SwapUsage metrics in CloudWatch.
- C. Install an Amazon CloudWatch agent on the instances. Run an appropriate script on a set schedule. Monitor SwapUtilization metrics in CloudWatch.
- D. Enable detailed monitoring in the EC2 console. Create an Amazon CloudWatch SwapUtilization custom metric. Monitor SwapUtilization metrics in CloudWatch.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Reference:

<https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/mon-scripts.html>

QUESTION 198

A company has a web server running on an Amazon EC2 instance in a public subnet with an Elastic IP address. The default security group is assigned to the EC2 instance. The default network ACL has been modified to block all traffic. A solutions architect needs to make the web server accessible from everywhere on port 443.

Which combination of steps will accomplish this task? (Choose two.)

- A. Create a security group with a rule to allow TCP port 443 from source 0.0.0.0/0.
- B. Create a security group with a rule to allow TCP port 443 to destination 0.0.0.0/0.
- C. Update the network ACL to allow TCP port 443 from source 0.0.0.0/0.
- D. Update the network ACL to allow inbound/outbound TCP port 443 from source 0.0.0.0/0 and to destination 0.0.0.0/0.
- E. Update the network ACL to allow inbound TCP port 443 from source 0.0.0.0/0 and outbound TCP port 32768-65535 to destination 0.0.0.0/0.

Correct Answer: AE

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 199

A company must re-evaluate its need for the Amazon EC2 instances it currently has provisioned in an Auto Scaling group. At present, the Auto Scaling group is configured for a minimum of two instances and a maximum of four instances across two Availability Zones. A Solutions architect reviewed Amazon CloudWatch metrics and found that CPU utilization is consistently low for all the EC2 instances.

What should the solutions architect recommend to maximize utilization while ensuring the application remains fault tolerant?

- A. Remove some EC2 instances to increase the utilization of remaining instances.
- B. Increase the Amazon Elastic Block Store (Amazon EBS) capacity of instances with less CPU utilization.
- C. Modify the Auto Scaling group scaling policy to scale in and out based on a higher CPU utilization metric.
- D. Create a new launch configuration that uses smaller instance types. Update the existing Auto Scaling group.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

As the Launch Configuration can't be modified once created, only way to update the Launch Configuration for an Auto Scaling group is to create a new one and associate it with the Auto Scaling group

QUESTION 200

A company has an application that posts messages to Amazon SQS. Another application polls the queue and processes the messages in an I/O-intensive operation. The company has a service level agreement (SLA) that specifies the

maximum amount of time that can elapse between receiving the messages and responding to the users. Due to an increase in the number of messages, the company has difficulty meeting its SLA consistently. What should a solutions architect do to help improve the application's processing time and ensure it can handle the load at any level?

- A. Create an Amazon Machine Image (AMI) from the instance used for processing. Terminate the instance and replace it with a larger size.
- B. Create an Amazon Machine Image (AMI) from the instance used for processing. Terminate the instance and replace it with an Amazon EC2 Dedicated Instance.
- C. Create an Amazon Machine image (AMI) from the instance used for processing. Create an Auto Scaling group using this image in its launch configuration. Configure the group with a target tracking policy to keep its aggregate CPU utilization below 70%.
- D. Create an Amazon Machine Image (AMI) from the instance used for processing. Create an Auto Scaling group using this image in its launch configuration. Configure the group with a target tracking policy based on the age of the oldest message in the SQS queue.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 201

A company is designing a new web service that will run on Amazon EC2 instances behind an Elastic Load Balancer. However, many of the web service clients can only reach IP addresses whitelisted on their firewalls. What should a solutions architect recommend to meet the clients' needs?

- A. A Network Load Balancer with an associated Elastic IP address.
- B. An Application Load Balancer with an associated Elastic IP address
- C. An A record in an Amazon Route 53 hosted zone pointing to an Elastic IP address
- D. An EC2 instance with a public IP address running as a proxy in front of the load balancer

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:



QUESTION 202

A company wants to host a web application on AWS that will communicate to a database within a VPC. The application should be highly available. What should a solutions architect recommend?

- A. Create two Amazon EC2 instances to host the web servers behind a load balancer, and then deploy the database on a large instance.
- B. Deploy a load balancer in multiple Availability Zones with an Auto Scaling group for the web servers, and then deploy Amazon RDS in multiple Availability Zones.
- C. Deploy a load balancer in the public subnet with an Auto Scaling group for the web servers, and then deploy the database on an Amazon EC2 instance in the private subnet.
- D. Deploy two web servers with an Auto Scaling group, configure a domain that points to the two web servers, and then deploy a database architecture in multiple Availability Zones.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 203

A company's packaged application dynamically creates and returns single-use text files in response to user requests. The company is using Amazon CloudFront for distribution, but wants to further reduce data transfer costs. The company cannot modify the application's source code.

What should a solutions architect do to reduce costs?

- A. Use Lambda@Edge to compress the files as they are sent to users.
- B. Enable Amazon S3 Transfer Acceleration to reduce the response times.
- C. Enable caching on the CloudFront distribution to store generated files at the edge.
- D. Use Amazon S3 multipart uploads to move the files to Amazon S3 before returning them to users.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 204

A database is on an Amazon RDS MySQL 5.6 Multi-AZ DB instance that experiences highly dynamic reads. Application developers notice a significant slowdown when testing read performance from a secondary AWS Region. The developers want a solution that provides less than 1 second of read replication latency. What should the solutions architect recommend?

- A. Install MySQL on Amazon EC2 in the secondary Region.
- B. Migrate the database to Amazon Aurora with cross-Region replicas.
- C. Create another RDS for MySQL read replica in the secondary Region.
- D. Implement Amazon ElastiCache to improve database query performance.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Reference:

<https://aws.amazon.com/rds/aurora/global-database/>

QUESTION 205

A company is planning to deploy an Amazon RDS DB instance running Amazon Aurora. The company has a backup retention policy requirement of 90 days. Which solution should a solutions architect recommend?

- A. Set the backup retention period to 90 days when creating the RDS DB instance.
- B. Configure RDS to copy automated snapshots to a user-managed Amazon S3 bucket with a lifecycle policy set to delete after 90 days.
- C. Create an AWS Backup plan to perform a daily snapshot of the RDS database with the retention set to 90 days. Create an AWS Backup job to schedule the execution of the backup plan daily.
- D. Use a daily scheduled event with Amazon CloudWatch Events to execute a custom AWS Lambda function that makes a copy of the RDS automated snapshot. Purge snapshots older than 90 days.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Reference:

https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_WorkingWithAutomatedBackups.html

QUESTION 206

A company currently has 250 TB of backup files stored in Amazon S3 in a vendor's proprietary format. Using a Linux-based software application provided by the vendor, the company wants to retrieve files from Amazon S3, transform the files to an industry-standard format, and re-upload them to Amazon S3. The company wants to minimize the data transfer charges associated with this conversation. What should a solutions architect do to accomplish this?

- A. Install the conversion software as an Amazon S3 batch operation so the data is transformed without leaving Amazon S3.
- B. Install the conversion software onto an on-premises virtual machine. Perform the transformation and re-upload the files to Amazon S3 from the virtual machine.
- C. Use AWS Snowball Edge devices to export the data and install the conversion software onto the devices. Perform the data transformation and re-upload the files to Amazon S3 from the Snowball Edge devices.
- D. Launch an Amazon EC2 instance in the same Region as Amazon S3 and install the conversion software onto the instance. Perform the transformation and re-upload the files to Amazon S3 from the EC2 instance.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 207

A company is migrating a NoSQL database cluster to Amazon EC2. The database automatically replicates data to maintain at least three copies of the data. I/O throughput of the servers is the highest priority. Which instance type should a solutions architect recommend for the migration?

- A. Storage optimized instances with instance store
- B. Burstable general purpose instances with an Amazon Elastic Block Store (Amazon EBS) volume
- C. Memory optimized instances with Amazon Elastic Block Store (Amazon EBS) optimization enabled
- D. Compute optimized instances with Amazon Elastic Block Store (Amazon EBS) optimization enabled

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 208

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

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

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Reference:

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



QUESTION 209

A company has a web application with sporadic usage patterns. There is heavy usage at the beginning of each month, moderate usage at the start of each week, and unpredictable usage during the week. The application consists of a web server and a MySQL database server running inside the data center. The company would like to move the application to the AWS Cloud, and needs to select a cost-effective database platform that will not require database modifications. Which solution will meet these requirements?

- A. Amazon DynamoDB
- B. Amazon RDS for MySQL
- C. MySQL-compatible Amazon Aurora Serverless
- D. MySQL deployed on Amazon EC2 in an Auto Scaling group

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 210

A solutions architect is designing the storage architecture for a new web application used for storing and viewing engineering drawings. All application components will be deployed on the AWS infrastructure. The application design must support caching to minimize the amount of time that users wait for the engineering drawings to load. The application must be able to store petabytes of data. Which combination of storage and caching should the solutions architect use?

- A. Amazon S3 with Amazon CloudFront
- B. Amazon S3 Glacier with Amazon ElastiCache
- C. Amazon Elastic Block Store (Amazon EBS) volumes with Amazon CloudFront
- D. AWS Storage Gateway with Amazon ElastiCache

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:

QUESTION 211

A solutions architect is creating an application that will handle batch processing of large amounts of data. The input data will be held in Amazon S3 and the output data will be stored in a different S3 bucket. For processing, the application will transfer the data over the network between multiple Amazon EC2 instances. What should the solutions architect do to reduce the overall data transfer costs?

- A. Place all the EC2 instances in an Auto Scaling group.
- B. Place all the EC2 instances in the same AWS Region.
- C. Place all the EC2 instances in the same Availability Zone.
- D. Place all the EC2 instances in private subnets in multiple Availability Zones.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:

QUESTION 212

A company hosts its core network services, including directory services and DNS, in its on-premises data center. The data center is connected to the AWS Cloud using AWS Direct Connect (DX). Additional AWS accounts are planned that will require quick, cost-effective, and consistent access to these network services. What should a solutions architect implement to meet these requirements with the LEAST amount of operational overhead?

- A. Create a DX connection in each new account. Route the network traffic to the on-premises servers.
- B. Configure VPC endpoints in the DX VPC for all required services. Route the network traffic to the on-premises servers.
- C. Create a VPN connection between each new account and the DX VPC. Route the network traffic to the on-premises servers.
- D. Configure AWS Transit Gateway between the accounts. Assign DX to the transit gateway and route network traffic to the on-premises servers.

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Explanation:

QUESTION 213

A company operates an ecommerce website on Amazon EC2 instances behind an Application Load Balancer (ALB) in an Auto Scaling group. The site is experiencing performance issues related to a high request rate from illegitimate external systems with changing IP addresses. The security team is worried about potential DDoS attacks against the website. The company must block the illegitimate incoming requests in a way that has a minimal impact on legitimate users.

What should a solutions architect recommend?

- A. Deploy Amazon Inspector and associate it with the ALB.
- B. Deploy AWS WAF, associate it with the ALB, and configure a rate-limiting rule.
- C. Deploy rules to the network ACLs associated with the ALB to block the incoming traffic.
- D. Deploy Amazon GuardDuty and enable rate-limiting protection when configuring GuardDuty.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Reference:

<https://aws.amazon.com/blogs/aws/protect-web-sites-services-using-rate-based-rules-for-aws-waf/>

QUESTION 214

A company receives structured and semi-structured data from various sources once every day. A solutions architect needs to design a solution that leverages big data processing frameworks. The data should be accessible using SQL queries and business intelligence tools.

What should the solutions architect recommend to build the MOST high-performing solution?

- A. Use AWS Glue to process data and Amazon S3 to store data.
- B. Use Amazon EMR to process data and Amazon Redshift to store data.
- C. Use Amazon EC2 to process data and Amazon Elastic Block Store (Amazon EBS) to store data.
- D. Use Amazon Kinesis Data Analytics to process data and Amazon Elastic File System (Amazon EFS) to store data.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Reference:

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

QUESTION 215

A company is hosting an election reporting website on AWS for users around the world. The website uses Amazon EC2 instances for the web and application tiers in an Auto Scaling group with Application Load Balancers. The database tier uses an Amazon RDS for MySQL database.

The website is updated with election results once an hour and has historically observed hundreds of users accessing the reports.

The company is expecting a significant increase in demand because of upcoming elections in different countries. A solutions architect must improve the website's ability to handle additional demand while minimizing the need for additional EC2 instances.

Which solution will meet these requirements?

- A. Launch an Amazon ElastiCache cluster to cache common database queries.
- B. Launch an Amazon CloudFront web distribution to cache commonly requested website content.
- C. Enable disk-based caching on the EC2 instances to cache commonly requested website content.
- D. Deploy a reverse proxy into the design using an EC2 instance with caching enabled for commonly requested website content.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 216

A company is building a website that relies on reading and writing to an Amazon DynamoDB database. The traffic associated with the website predictably peaks during business hours on weekdays and declines overnight and during weekends. A solutions architect needs to design a cost-effective solution that can handle the load.

What should the solutions architect do to meet these requirements?

- A. Enable DynamoDB Accelerator (DAX) to cache the data.
- B. Enable Multi-AZ replication for the DynamoDB database.
- C. Enable DynamoDB auto scaling when creating the tables.
- D. Enable DynamoDB On-Demand capacity allocation when creating the tables.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 217

A company uses Amazon Redshift for its data warehouse. The company wants to ensure high durability for its data in case of any component failure.

What should a solutions architect recommend?

- A. Enable concurrency scaling.
- B. Enable cross-Region snapshots.
- C. Increase the data retention period.
- D. Deploy Amazon Redshift in Multi-AZ.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 218

A company has data stored in an on-premises data center that is used by several on-premises applications. The company wants to maintain its existing application environment and be able to use AWS services for data analytics and future visualizations.

Which storage service should a solutions architect recommend?

- A. Amazon Redshift
- B. AWS Storage Gateway for files
- C. Amazon Elastic Block Store (Amazon EBS)
- D. Amazon Elastic File System (Amazon EFS)

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 219

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

How should the solutions architect comply with these requirements?

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

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 220

A company has a 143 TB MySQL database that it wants to migrate to AWS. The plan is to use Amazon Aurora MySQL as the platform going forward. The company has a 100 Mbps AWS Direct Connect connection to Amazon VPC. Which solution meets the company's needs and takes the LEAST amount of time?

- A. Use a gateway endpoint for Amazon S3. Migrate the data to Amazon S3. Import the data into Aurora.
- B. Upgrade the Direct Connect link to 500 Mbps. Copy the data to Amazon S3. Import the data into Aurora.
- C. Order an AWS Snowmobile and copy the database backup to it. Have AWS import the data into Amazon S3. Import the backup into Aurora.
- D. Order four 50-TB AWS Snowball devices and copy the database backup onto them. Have AWS import the data into Amazon S3. Import the data into Aurora.

Correct Answer: D

Section: (none)

Explanation



Explanation/Reference:

Explanation:

QUESTION 221

A company hosts an online shopping application that stores all orders in an Amazon RDS for PostgreSQL Single-AZ DB instance. Management wants to eliminate single points of failure and has asked a solutions architect to recommend an approach to minimize database downtime without requiring any changes to the application code.

Which solution meets these requirements?

- A. Convert the existing database instance to a Multi-AZ deployment by modifying the database instance and specifying the Multi-AZ option.
- B. Create a new RDS Multi-AZ deployment. Take a snapshot of the current RDS instance and restore the new Multi-AZ deployment with the snapshot.
- C. Create a read-only replica of the PostgreSQL database in another Availability Zone. Use Amazon Route 53 weighted record sets to distribute requests across the databases.
- D. Place the RDS for PostgreSQL database in an Amazon EC2 Auto Scaling group with a minimum group size of two. Use Amazon Route 53 weighted record sets to distribute requests across instances.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 222

A company has a 10 Gbps AWS Direct Connect connection from its on-premises servers to AWS.

The workloads using the connection are critical. The company requires a disaster recovery strategy with maximum resiliency that maintains the current connection bandwidth at a minimum.

What should a solutions architect recommend?

- A. Set up a new Direct Connect connection in another AWS Region.
- B. Set up a new AWS managed VPN connection in another AWS Region.
- C. Set up two new Direct Connect connections: one in the current AWS Region and one in another Region.
- D. Set up two new AWS managed VPN connections: one in the current AWS Region and one in another Region.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 223

A solutions architect is designing a VPC with public and private subnets. The VPC and subnets use IPv4 CIDR blocks. There is one public subnet and one private subnet in each of three Availability Zones (AZs) for high availability. An internet gateway is used to provide internet access for the public subnets. The private subnets require access to the internet to allow Amazon EC2 instances to download software updates.

What should the solutions architect do to enable internet access for the private subnets?

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

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 224

As part of budget planning, management wants a report of AWS billed items listed by user. The data will be used to create department budgets. A solutions architect needs to determine the most efficient way to obtain this report information. Which solution meets these requirements?

- A. Run a query with Amazon Athena to generate the report.

- B. Create a report in Cost Explorer and download the report.
- C. Access the bill details from the billing dashboard and download the bill.
- D. Modify a cost budget in AWS Budgets to alert with Amazon Simple Email Service (Amazon SES).

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 225

A company with facilities in North America, Europe, and Asia is designing new distributed application to optimize its global supply chain and manufacturing process. The orders booked on one continent should be visible to all Regions in a second or less. The database should be able to support failover with a short Recovery Time Objective (RTO). The uptime of the application is important to ensure that manufacturing is not impacted. What should a solutions architect recommend?

- A. Use Amazon DynamoDB global tables.
- B. Use Amazon Aurora Global Database.
- C. Use Amazon RDS for MySQL with a cross-Region read replica.
- D. Use Amazon RDS for PostgreSQL with a cross-Region read replica.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 226

A company's near-real-time streaming application is running on AWS. As the data is ingested, a job runs on the data and takes 30 minutes to complete. The workload frequently experiences high latency due to large amounts of incoming data. A solutions architect needs to design a scalable and serverless solution to enhance performance. Which combination of steps should the solutions architect take? (Choose two.)

- A. Use Amazon Kinesis Data Firehose to ingest the data.
- B. Use AWS Lambda with AWS Step Functions to process the data.
- C. Use AWS Database Migration Service (AWS DMS) to ingest the data.
- D. Use Amazon EC2 instances in an Auto Scaling group to process the data.
- E. Use AWS Fargate with Amazon Elastic Container Service (Amazon ECS) to process the data.

Correct Answer: AE

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 227

An application running on an Amazon EC2 instance needs to access an Amazon DynamoDB table. Both the EC2 instance and the DynamoDB table are in the same AWS account. A solutions architect must configure the necessary permissions. Which solution will allow least privilege access to the DynamoDB table from the EC2 instance?

- A. Create an IAM role with the appropriate policy to allow access to the DynamoDB table. Create an instance profile to assign this IAM role to the EC2 instance.
- B. Create an IAM role with the appropriate policy to allow access to the DynamoDB table. Add the EC2 instance to the trust relationship policy document to allow it to assume the role.
- C. Create an IAM user with the appropriate policy to allow access to the DynamoDB table. Store the credentials in an Amazon S3 bucket and read them from within the application code directly.
- D. Create an IAM user with the appropriate policy to allow access to the DynamoDB table. Ensure that the application stores the IAM credentials securely on local storage and uses them to make the DynamoDB calls.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 228

A solutions architect is designing a solution that involves orchestrating a series of Amazon Elastic Container Service (Amazon ECS) task types running on Amazon EC2 instances that are part of an ECS cluster. The output and state data for all tasks needs to be stored. The amount of data output by each task is approximately 10 MB, and there could be hundreds of tasks running at a time. The system should be optimized for high-frequency reading and writing. As old outputs are archived and deleted, the storage size is not expected to exceed 1 TB.

Which storage solution should the solutions architect recommend?

- A. An Amazon DynamoDB table accessible by all ECS cluster instances.
- B. An Amazon Elastic File System (Amazon EFS) with Provisioned Throughput mode.
- C. An Amazon Elastic File System (Amazon EFS) file system with Bursting Throughput mode.
- D. An Amazon Elastic Block Store (Amazon EBS) volume mounted to the ECS cluster instances.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 229

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. Set the visibility timeout of the paid photos to zero. Configure Amazon EC2 instances to prioritize visibility settings so paid photos are processed first.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 230

A company wants to migrate its MySQL database from on premises to AWS. The company recently experienced a database outage that significantly impacted the business. To ensure this does not happen again, the company wants a reliable database solution on AWS that minimizes data loss and stores every transaction on at least two nodes.

Which solution meets these requirements?

- A. Create an Amazon RDS DB instance with synchronous replication to three nodes in three Availability Zones.
- B. Create an Amazon RDS MySQL DB instance with Multi-AZ functionality enabled to synchronously replicate the data.
- C. Create an Amazon RDS MySQL DB instance and then create a read replica in a separate AWS Region that synchronously replicates the data.
- D. Create an Amazon EC2 instance with a MySQL engine installed that triggers an AWS Lambda function to synchronously replicate the data to an Amazon RDS MySQL DB instance.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 231

A company stores user data in AWS. The data is used continuously with peak usage during business hours. Access patterns vary, with some data not being used for months at a time. A solutions architect must choose a cost-effective solution that maintains the highest level of durability while maintaining high availability.

Which storage solution meets these requirements?

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

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 232

A company receives inconsistent service from its data center provider because the company is headquartered in an area affected by natural disasters. The company is not ready to fully migrate to the AWS Cloud, but it wants a failure environment on AWS in case the on-premises data center fails.

The company runs web servers that connect to external vendors. The data available on AWS and on premises must be uniform.

Which solution should a solutions architect recommend that has the LEAST amount of downtime?

- A. Configure an Amazon Route 53 failover record. Run application servers on Amazon EC2 instances behind an Application Load Balancer in an Auto Scaling group. Set up AWS Storage Gateway with stored volumes to back up data to Amazon S3.
- B. Configure an Amazon Route 53 failover record. Execute an AWS CloudFormation template from a script to create Amazon EC2 instances behind an Application Load Balancer. Set up AWS Storage Gateway with stored volumes to back up data to Amazon S3.
- C. Configure an Amazon Route 53 failover record. Set up an AWS Direct Connect connection between a VPC and the data center. Run application servers on Amazon EC2 in an Auto Scaling group. Run an AWS Lambda function to execute an AWS CloudFormation template to create an Application Load Balancer.
- D. Configure an Amazon Route 53 failover record. Run an AWS Lambda function to execute an AWS CloudFormation template to launch two Amazon EC2 instances. Set up AWS Storage Gateway with stored volumes to back up data to Amazon S3. Set up an AWS Direct Connect connection between a VPC and the data center.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 233

A company has three VPCs named Development, Testing, and Production in the us-east-1 Region. The three VPCs need to be connected to an on-premises data center and are designed to be separate to maintain security and prevent any resource sharing. A solutions architect needs to find a scalable and secure solution.

What should the solutions architect recommend?

- A. Create an AWS Direct Connect connection and a VPN connection for each VPC to connect back to the data center.
- B. Create VPC peers from all the VPCs to the Production VPC. Use an AWS Direct Connect connection from the Production VPC back to the data center.
- C. Connect VPN connections from all the VPCs to a VPN in the Production VPC. Use a VPN connection from the Production VPC back to the data center.
- D. Create a new VPC called Network. Within the Network VPC, create an AWS Transit Gateway with an AWS Direct Connect connection back to the data center. Attach all the other VPCs to the Network VPC.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 234

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

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



Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Explanation:

QUESTION 235

A company needs a secure connection between its on-premises environment and AWS. This connection does not need high bandwidth and will handle a small amount of traffic. The connection should be set up quickly. What is the MOST cost-effective method to establish this type of connection?

- A. Implement a client VPN.
- B. Implement AWS Direct Connect.
- C. Implement a bastion host on Amazon EC2.
- D. Implement an AWS Site-to-Site VPN connection.

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Explanation:

QUESTION 236

A company uses Application Load Balancers (ALBs) in different AWS Regions. The ALBs receive inconsistent traffic that can spike and drop throughout the year. The company's networking team needs to allow the IP addresses of the ALBs in the on-premises firewall to enable connectivity.

Which solution is the MOST scalable with minimal configuration changes?

- A. Write an AWS Lambda script to get the IP addresses of the ALBs in different Regions. Update the on-premises firewall's rule to allow the IP addresses of the ALBs.
- B. Migrate all ALBs in different Regions to the Network Load Balancer (NLBs). Update the on-premises firewall's rule to allow the Elastic IP addresses of all the NLBs.
- C. Launch AWS Global Accelerator. Register the ALBs in different Regions to the accelerator.
Update the on-premises firewall's rule to allow static IP addresses associated with the accelerator.
- D. Launch a Network Load Balancer (NLB) in one Region. Register the private IP addresses of the ALBs in different Regions with the NLB. Update the on-premises firewall's rule to allow the Elastic IP address attached to the NLB.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:

QUESTION 237

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.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:

QUESTION 238

A company uses a legacy on-premises analytics application that operates on gigabytes of .csv files and represents months of data. The legacy application cannot handle the growing size of .csv files. New .csv files are added daily from various data sources to a central on-premises storage location. The company wants to continue to support the legacy application while users learn AWS analytics services. To achieve this, a solutions architect wants to maintain two synchronized copies of all the .csv files on-premises and in Amazon S3. Which solution should the solutions architect recommend?

- A. Deploy AWS DataSync on-premises. Configure DataSync to continuously replicate the .csv files between the company's on-premises storage and the company's S3 bucket.
- B. Deploy an on-premises file gateway. Configure data sources to write the .csv files to the file gateway. Point the legacy analytics application to the file gateway. The file gateway should replicate the .csv files to Amazon S3.
- C. Deploy an on-premises volume gateway. Configure data sources to write the .csv files to the volume gateway. Point the legacy analytics application to the volume gateway. The volume gateway should replicate data to Amazon S3.
- D. Deploy AWS DataSync on-premises. Configure DataSync to continuously replicate the .csv files between on-premises and Amazon Elastic File System (Amazon EFS). Enable replication from Amazon Elastic File System (Amazon EFS) to the company's S3 bucket.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:

QUESTION 239

A company has media and application files that need to be shared internally. Users currently are authenticated using Active Directory and access files from a Microsoft Windows platform. The chief executive officer wants to keep the same user permissions, but wants the company to improve the process as the company is reaching its storage capacity limit. What should a solutions architect recommend?

- A. Set up a corporate Amazon S3 bucket and move all media and application files.
- B. Configure Amazon FSx for Windows File Server and move all the media and application files.
- C. Configure Amazon Elastic File System (Amazon EFS) and move all media and application files.
- D. Set up Amazon EC2 on Windows, attach multiple Amazon Elastic Block Store (Amazon EBS) volumes, and move all media and application files.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Reference: <https://aws.amazon.com/fsx/windows/>



QUESTION 240

A company is deploying a web portal. The company wants to ensure that only the web portion of the application is publicly accessible. To accomplish this, the VPC was designed with two public subnets and two private subnets. The application will run on several Amazon EC2 instances in an Auto Scaling group. SSL termination must be offloaded from the EC2 instances. What should a solutions architect do to ensure these requirements are met?

- A. Configure the Network Load Balancer in the public subnets. Configure the Auto Scaling group in the private subnets and associate it with the Application Load Balancer.
- B. Configure the Network Load Balancer in the public subnets. Configure the Auto Scaling group in the public subnets and associate it with the Application Load Balancer.
- C. Configure the Application Load Balancer in the public subnets. Configure the Auto Scaling group in the private subnets and associate it with the Application Load Balancer.
- D. Configure the Application Load Balancer in the private subnets. Configure the Auto Scaling group in the private subnets and associate it with the Application Load Balancer.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:

QUESTION 241

A company is experiencing growth as demand for its product has increased. The company's existing purchasing application is slow when traffic spikes. The application is a monolithic threetier application that uses synchronous transactions and sometimes sees bottlenecks in the application tier. A solutions architect needs to design a solution that can meet required application response times while accounting for traffic volume spikes. Which solution will meet these requirements?

- A. Vertically scale the application instance using a larger Amazon EC2 instance size.
- B. Scale the application's persistence layer horizontally by introducing Oracle RAC on AWS.

- C. Scale the web and application tiers horizontally using Auto Scaling groups and an Application Load Balancer.
- D. Decouple the application and data tiers using Amazon Simple Queue Service (Amazon SQS) with asynchronous AWS Lambda calls.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 242

A company hosts an application used to upload files to an Amazon S3 bucket. Once uploaded, the files are processed to extract metadata, which takes less than 5 seconds. The volume and frequency of the uploads varies from a few files each hour to hundreds of concurrent uploads. The company has asked a solutions architect to design a cost-effective architecture that will meet these requirements. What should the solutions architect recommend?

- A. Configure AWS CloudTrail trails to log S3 API calls. Use AWS AppSync to process the files.
- B. Configure an object-created event notification within the S3 bucket to invoke an AWS Lambda function to process the files.
- C. Configure Amazon Kinesis Data Streams to process and send data to Amazon S3. Invoke an AWS Lambda function to process the files.
- D. Configure an Amazon Simple Notification Service (Amazon SNS) topic to process the files uploaded to Amazon S3. Invoke an AWS Lambda function to process the files.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 243

A company has copied 1 PB of data from a colocation facility to an Amazon S3 bucket in the useast-1 Region using an AWS Direct Connect link. The company now wants to copy the data to another S3 bucket in the us-west-2 Region. The colocation facility does not allow the use of AWS Snowball. What should a solutions architect recommend to accomplish this?

- A. Order a Snowball Edge device to copy the data from one Region to another Region.
- B. Transfer contents from the source S3 bucket to a target S3 bucket using the S3 console.
- C. Use the aws S3 sync command to copy data from the source bucket to the destination bucket.
- D. Add a cross-Region replication configuration to copy objects across S3 buckets in different Regions.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 244

A company is using a fleet of Amazon EC2 instances to ingest data from on-premises data sources. The data is in JSON format and ingestion rates can be as high as 1 MB/s. When an EC2 instance is rebooted, the data in-flight is lost. The company's data science team wants to query ingested data in near-real time. Which solution provides near-real-time data querying that is scalable with minimal data loss?

- A. Publish data to Amazon Kinesis Data Streams. Use Kinesis Data Analytics to query the data.
- B. Publish data to Amazon Kinesis Data Firehose with Amazon Redshift as the destination. Use Amazon Redshift to query the data.
- C. Store ingested data in an EC2 instance store. Publish data to Amazon Kinesis Data Firehose with Amazon S3 as the destination. Use Amazon Athena to query the data.
- D. Store ingested data in an Amazon Elastic Block Store (Amazon EBS) volume. Publish data to Amazon ElastiCache for Redis. Subscribe to the Redis channel to query the data.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 245

A company is deploying a multi-instance application within AWS that requires minimal latency between the instances. What should a solutions architect recommend?

- A. Use an Auto Scaling group with a cluster placement group.
- B. Use an Auto Scaling group with single Availability Zone in the same AWS Region.
- C. Use an Auto Scaling group with multiple Availability Zones in the same AWS Region.
- D. Use a Network Load Balancer with multiple Amazon EC2 Dedicated Hosts as the targets.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 246

A company is developing a mobile game that streams score updates to a backend processor and then posts results on a leaderboard. A solutions architect needs to design a solution that can handle large traffic spikes, process the mobile game updates in order of receipt, and store the processed updates in a highly available database. The company also wants to minimize the management overhead required to maintain the solution. What should the solutions architect do to meet these requirements?

- A. Push score updates to Amazon Kinesis Data Streams. Process the updates in Kinesis Data Streams with AWS Lambda. Store the processed updates in Amazon DynamoDB.
- B. Push score updates to Amazon Kinesis Data Streams. Process the updates with a fleet of Amazon EC2 instances set up for Auto Scaling. Store the processed updates in Amazon Redshift.
- C. Push score updates to an Amazon Simple Notification Service (Amazon SNS) topic. Subscribe an AWS Lambda function to the SNS topic to process the updates. Store the processed updates in a SQL database running on Amazon EC2.
- D. Push score updates to an Amazon Simple Queue Service (Amazon SQS) queue. Use a fleet of Amazon EC2 instances with Auto Scaling to process the updates in the SQS queue. Store the processed updates in an Amazon RDS Multi-AZ DB instance.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 247

A company is building a document storage application on AWS. The application runs on Amazon EC2 instances in multiple Availability Zones. The company requires the document store to be highly available. The documents need to be returned immediately when requested. The lead engineer has configured the application to use Amazon Elastic Block Store (Amazon EBS) to store the documents, but is willing to consider other options to meet the availability requirement. What should a solutions architect recommend?

- A. Snapshot the EBS volumes regularly and build new volumes using those snapshots in additional Availability Zones.
- B. Use Amazon Elastic Block Store (Amazon EBS) for the EC2 instance root volumes. Configure the application to build the document store on Amazon S3.
- C. Use Amazon Elastic Block Store (Amazon EBS) for the EC2 instance root volumes. Configure the application to build the document store on Amazon S3 Glacier.
- D. Use at least three Provisioned IOPS EBS volumes for EC2 instances. Mount the volumes to the EC2 instances in a RAID 5 configuration.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 248

A group requires permissions to list an Amazon S3 bucket and delete objects from that bucket. An administrator has created the following IAM policy to provide access to the bucket and applied that policy to the group. The group is not able to delete objects in the bucket. The company follows least-privilege access rules.



```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Action": [
        "s3:ListBucket",
        "s3:DeleteObject"
      ],
      "Resource": [
        "arn:aws:s3:::bucket-name"
      ],
      "Effect": "Allow"
    }
  ]
}

```

Which statement should a solutions architect add to the policy to correct bucket access?

- A. "Action": ["s3:*Object"], "Resource": ["arn:aws:s3:::bucket-name/*"], "Effect": "Allow"
- B. "Action": ["s3:*"], "Resource": ["arn:aws:s3:::bucket-name/*"], "Effect": "Allow"
- C. "Action": ["s3:DeleteObject"], "Resource": ["arn:aws:s3:::bucket-name*"], "Effect": "Allow"
- D. "Action": ["s3:DeleteObject"], "Resource": ["arn:aws:s3:::bucket-name/*"], "Effect": "Allow"



Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
 Explanation:

QUESTION 249

A solutions architect is designing a security solution for a company that wants to provide developers with individual AWS accounts through AWS Organizations, while also maintaining standard security controls. Because the individual developers will have AWS account root userlevel access to their own accounts, the solutions architect wants to ensure that the mandatory AWS CloudTrail configuration that is applied to new developer accounts is not modified. Which action meets these requirements?

- A. Create an IAM policy that prohibits changes to CloudTrail, and attach it to the root user.
- B. Create a new trail in CloudTrail from within the developer accounts with the organization trails option enabled.
- C. Create a service control policy (SCP) that prohibits changes to CloudTrail, and attach it to the developer accounts.
- D. Create a service-linked role for CloudTrail with a policy condition that allows changes only from an Amazon Resource Name (ARN) in the master account.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Reference:

https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_policies_scps_examples.html

QUESTION 250

A company wants to share forensic accounting data that is stored in an Amazon RDS DB instance with an external auditor. The auditor has its own AWS account and requires its own copy of the database. How should the company securely share the database with the auditor?

- A. Create a read replica of the database and configure IAM standard database authentication to grant the auditor access.
- B. Copy a snapshot of the database to Amazon S3 and assign an IAM role to the auditor to grant access to the object in that bucket.
- C. Export the database contents to text files, store the files in Amazon S3, and create a new IAM user for the auditor with access to that bucket.
- D. Make an encrypted snapshot of the database, share the snapshot, and allow access to the AWS Key Management Service (AWS KMS) encryption key.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:



QUESTION 251

A company has an automobile sales website that stores its listings in a database on Amazon RDS.

When an automobile is sold, the listing needs to be removed from the website and the data must be sent to multiple target systems.

Which design should a solutions architect recommend?

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

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 252

A company is building a media sharing application and decides to use Amazon S3 for storage.

When a media file is uploaded, the company starts a multi-step process to create thumbnails, identify objects in the images, transcode videos into standard formats and resolutions, and extract and store the metadata to an Amazon DynamoDB table. The metadata is used for searching and navigation.

The amount of traffic is variable. The solution must be able to scale to handle spikes in load without unnecessary expenses.

What should a solutions architect recommend to support this workload?

- A. Build the processing into the website or mobile app used to upload the content to Amazon S3.

- Save the required data to the DynamoDB table when the objects are uploaded.
- B. Trigger AWS Step Functions when an object is stored in the S3 bucket. Have the Step Functions perform the steps needed to process the object and then write the metadata to the DynamoDB table.
 - C. Trigger an AWS Lambda function when an object is stored in the S3 bucket. Have the Lambda function start AWS Batch to perform the steps to process the object. Place the object data in the DynamoDB table when complete.
 - D. Trigger an AWS Lambda function to store an initial entry in the DynamoDB table when an object is uploaded to Amazon S3. Use a program running on an Amazon EC2 instance in an Auto Scaling group to poll the index for unprocessed items, and use the program to perform the processing.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 253

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

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

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 254

An application is running on an Amazon EC2 instance and must have millisecond latency when running the workload. The application makes many small reads and writes to the file system, but the file system itself is small. Which Amazon Elastic Block Store (Amazon EBS) volume type should a solutions architect attach to their EC2 instance?

- A. Cold HDD (sc1)
- B. General Purpose SSD (gp2)
- C. Provisioned IOPS SSD (io1)
- D. Throughput Optimized HDD (st1)

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Reference: <https://aws.amazon.com/blogs/database/best-storage-practices-for-running-production-workloads-on-hosted-databases-with-amazon-rds-or-amazon-ec2/>

QUESTION 255

A solutions architect is designing a multi-Region disaster recovery solution for an application that will provide public API access. The application will use Amazon EC2 instances with a userdata script to load application code and an Amazon RDS for MySQL database. The Recovery Time Objective (RTO) is 3 hours and the Recovery Point Objective (RPO) is 24 hours. Which architecture would meet these requirements at the LOWEST cost?

- A. Use an Application Load Balancer for Region failover. Deploy new EC2 instances with the userdata script. Deploy separate RDS instances in each Region.
- B. Use Amazon Route 53 for Region failover. Deploy new EC2 instances with the userdata script. Create a read replica of the RDS instance in a backup Region.
- C. Use Amazon API Gateway for the public APIs and Region failover. Deploy new EC2 instances with the userdata script. Create a MySQL read replica of the RDS instance in a backup Region.
- D. Use Amazon Route 53 for Region failover. Deploy new EC2 instances with the userdata script for APIs, and create a snapshot of the RDS instance daily for a backup. Replicate the snapshot to a backup Region.

Correct Answer: D

Section: (none)



Explanation

Explanation/Reference:

Explanation:

QUESTION 256

A solutions architect needs to ensure that all Amazon Elastic Block Store (Amazon EBS) volumes restored from unencrypted EBC snapshots are encrypted. What should the solutions architect do to accomplish this?

- A. Enable EBS encryption by default for the AWS Region.
- B. Enable EBS encryption by default for the specific volumes.
- C. Create a new volume and specify the symmetric customer master key (CMK) to use for encryption.
- D. Create a new volume and specify the asymmetric customer master key (CMK) to use for encryption.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Reference:

<https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/EBSEncryption.html#volume-accountoff>

QUESTION 257

A company runs a static website through its on-premises data center. The company has multiple servers that handle all of its traffic, but on busy days, services are interrupted and the website becomes unavailable. The company wants to expand its presence globally and plans to triple its website traffic.

What should a solutions architect recommend to meet these requirements?

- A. Migrate the website content to Amazon S3 and host the website on Amazon CloudFront.
- B. Migrate the website content to Amazon EC2 instances with public Elastic IP addresses in multiple AWS Regions.
- C. Migrate the website content to Amazon EC2 instances and vertically scale as the load increases.
- D. Use Amazon Route 53 to distribute the loads across multiple Amazon CloudFront distributions for each AWS Region that exists globally.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 258

A company has a highly dynamic batch processing job that uses many Amazon EC2 instances to complete it. The job is stateless in nature, can be started and stopped at any given time with no negative impact, and typically takes upwards of 60 minutes total to complete. The company has asked a solutions architect to design a scalable and cost-effective solution that meets the requirements of the job.

What should the solutions architect recommend?

- A. Implement EC2 Spot Instances.
- B. Purchase EC2 Reserved Instances.
- C. Implement EC2 On-Demand Instances.
- D. Implement the processing on AWS Lambda.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 259

A company is hosting its static website in an Amazon S3 bucket, which is the origin for Amazon CloudFront. The company has users in the United States, Canada, and Europe and wants to reduce costs.

What should a solutions architect recommend?

- A. Adjust the CloudFront caching time to live (TTL) from the default to a longer timeframe.
- B. Implement CloudFront events with Lambda@Edge to run the website's data processing.
- C. Modify the CloudFront price class to include only the locations of the countries that are served.
- D. Implement a CloudFront Secure Sockets Layer (SSL) certificate to push security closer to the locations of the countries that are served.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 260

A company is planning to migrate a commercial off-the-shelf application from its on-premises data center to AWS. The software has a software licensing model using sockets and cores with predictable capacity and uptime requirements. The company wants to use its existing licenses, which were purchased earlier this year. Which Amazon EC2 pricing option is the MOST cost-effective?

- A. Dedicated Reserved Hosts
- B. Dedicated On-Demand Hosts
- C. Dedicated Reserved Instances
- D. Dedicated On-Demand Instances

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 261

A company is designing a website that uses an Amazon S3 bucket to store static images. The company wants all future requests to have faster response times while reducing both latency and cost. Which service configuration should a solutions architect recommend?

- A. Deploy a NAT server in front of Amazon S3.
- B. Deploy Amazon CloudFront in front of Amazon S3.
- C. Deploy a Network Load Balancer in front of Amazon S3.
- D. Configure Auto Scaling to automatically adjust the capacity of the website.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Reference: <https://aws.amazon.com/getting-started/hands-on/deliver-content-faster/>

QUESTION 262

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

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

Correct Answer: A

Section: (none)

Explanation



Explanation/Reference:

Reference: <https://aws.amazon.com/premiumsupport/knowledge-center/migrate-mysql-rds-dms/>

QUESTION 263

A company needs to comply with a regulatory requirement that states all emails must be stored and archived externally for 7 years. An administrator has created compressed email files on premises and wants a managed service to transfer the files to AWS storage.

Which managed service should a solutions architect recommend?

- A. Amazon Elastic File System (Amazon EFS)
- B. Amazon S3 Glacier
- C. AWS Backup
- D. AWS Storage Gateway

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Reference: <https://aws.amazon.com/storagegateway/faqs/>

QUESTION 264

A company has hired a new cloud engineer who should not have access to an Amazon S3 bucket named CompanyConfidential. The cloud engineer must be able to read from and write to an S3 bucket called AdminTools.

Which IAM policy will meet these requirements?

A. A.

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



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

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

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

Udumps

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:

QUESTION 265

A company that hosts its web application on AWS wants to ensure all Amazon EC2 instances, Amazon RDS DB instances, and Amazon Redshift clusters are configured with tags. The company wants to minimize the effort of configuring and operating this check. What should a solutions architect do to accomplish this?

- A. Use AWS Config rules to define and detect resources that are not properly tagged.
- B. Use Cost Explorer to display resources that are not properly tagged. Tag those resources manually.
- C. Write API calls to check all resources for proper tag allocation. Periodically run the code on an EC2 instance.
- D. Write API calls to check all resources for proper tag allocation. Schedule an AWS Lambda function through Amazon CloudWatch to periodically run the code.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Reference: <https://d1.awsstatic.com/whitepapers/aws-tagging-best-practices.pdf>

QUESTION 266

A company has a live chat application running on its on-premises servers that use WebSockets. The company wants to migrate the application to AWS. Application traffic is inconsistent, and the company expects there to be more traffic with sharp spikes in the future. The company wants a highly scalable solution with no server maintenance nor advanced capacity planning. Which solution meets these requirements?

- A. Use Amazon API Gateway and AWS Lambda with an Amazon DynamoDB table as the data store. Configure the DynamoDB table for provisioned capacity.
- B. Use Amazon API Gateway and AWS Lambda with an Amazon DynamoDB table as the data store. Configure the DynamoDB table for on-demand capacity.
- C. Run Amazon EC2 instances behind an Application Load Balancer in an Auto Scaling group with an Amazon DynamoDB table as the data store. Configure the DynamoDB table for on-demand capacity.
- D. Run Amazon EC2 instances behind a Network Load Balancer in an Auto Scaling group with an Amazon DynamoDB table as the data store. Configure the DynamoDB table for provisioned capacity.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:

QUESTION 267

A company hosts its static website content from an Amazon S3 bucket in the us-east-1 Region. Content is made available through an Amazon CloudFront origin pointing to that bucket. Cross-Region replication is set to create a second copy of the bucket in the ap-southeast-1 Region. Management wants a solution that provides greater availability for the website. Which combination of actions should a solutions architect take to increase availability? (Choose two.)

- A. Add both buckets to the CloudFront origin.
- B. Configure failover routing in Amazon Route 53.
- C. Create a record in Amazon Route 53 pointing to the replica bucket.
- D. Create an additional CloudFront origin pointing to the ap-southeast-1 bucket.
- E. Set up a CloudFront origin group with the us-east-1 bucket as the primary and the ap-southeast-1 bucket as the secondary.

Correct Answer: BE
Section: (none)
Explanation

Explanation/Reference:

Explanation:

QUESTION 268

A company hosts a training site on a fleet of Amazon EC2 instances. The company anticipates that its new course, which consists of dozens of training videos on the site, will be extremely popular when it is released in 1 week. What should a solutions architect do to minimize the anticipated server load?

- A. Store the videos in Amazon ElastiCache for Redis. Update the web servers to serve the videos using the ElastiCache API.
- B. Store the videos in Amazon Elastic File System (Amazon EFS). Create a user data script for the web servers to mount the EFS volume.
- C. Store the videos in an Amazon S3 bucket. Create an Amazon CloudFront distribution with an origin access identity (OAI) of that S3 bucket. Restrict Amazon S3 access to the OAI.
- D. Store the videos in an Amazon S3 bucket. Create an AWS Storage Gateway file gateway to access the S3 bucket. Create a user data script for the web servers to mount the file gateway.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 269

A company runs a production application on a fleet of Amazon EC2 instances. The application reads the data from an Amazon SQS queue and processes the messages in parallel. The message volume is unpredictable and often has intermittent traffic. This application should continually process messages without any downtime.

Which solution meets these requirements MOST cost-effectively?

- A. Use Spot Instances exclusively to handle the maximum capacity required.
- B. Use Reserved Instances exclusively to handle the maximum capacity required.
- C. Use Reserved Instances for the baseline capacity and use Spot Instances to handle additional capacity.
- D. Use Reserved Instances for the baseline capacity and use On-Demand Instances to handle additional capacity.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 270

A company has a hybrid application hosted on multiple on-premises servers with static IP addresses. There is already a VPN that provides connectivity between the VPC and the on-premises network. The company wants to distribute TCP traffic across the on-premises servers for internet users.

What should a solutions architect recommend to provide a highly available and scalable solution?

- A. Launch an internet-facing Network Load Balancer (NLB) and register on-premises IP addresses with the NLB.
- B. Launch an internet-facing Application Load Balancer (ALB) and register on-premises IP addresses with the ALB.
- C. Launch an Amazon EC2 instance, attach an Elastic IP address, and distribute traffic to the on-premises servers.
- D. Launch an Amazon EC2 instance with public IP addresses in an Auto Scaling group and distribute traffic to the on-premises servers.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 271

Management has decided to deploy all AWS VPCs with IPv6 enabled. After some time, a solutions architect tries to launch a new instance and receives an error stating that there is not enough IP address space available in the subnet. What should the solutions architect do to fix this?

- A. Check to make sure that only IPv6 was used during the VPC creation.

- B. Create a new IPv4 subnet with a larger range, and then launch the instance.
- C. Create a new IPv6-only subnet with a large range, and then launch the instance.
- D. Disable the IPv4 subnet and migrate all instances to IPv6 only. Once that is complete, launch the instance.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 272

A company has a build server that is in an Auto Scaling group and often has multiple Linux instances running. The build server requires consistent and mountable shared NFS storage for jobs and configurations. Which storage option should a solutions architect recommend?

- A. Amazon S3
- B. Amazon FSx
- C. Amazon Elastic Block Store (Amazon EBS)
- D. Amazon Elastic File System (Amazon EFS)

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Reference: <https://aws.amazon.com/efs/>

QUESTION 273

A company has an image processing workload running on Amazon Elastic Container Service (Amazon ECS) in two private subnets. Each private subnet uses a NAT instance for internet access. All images are stored in Amazon S3 buckets. The company is concerned about the data transfer costs between Amazon ECS and Amazon S3.

What should a solutions architect do to reduce costs?

- A. Configure a NAT gateway to replace the NAT instances.
- B. Configure a gateway endpoint for traffic destined to Amazon S3.
- C. Configure an interface endpoint for traffic destined to Amazon S3.
- D. Configure Amazon CloudFront for the S3 bucket storing the images.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 274

The financial application at a company stores monthly reports in an Amazon S3 bucket. The vice president of finance has mandated that all access to these reports be logged and that any modifications to the log files be detected. Which actions can a solutions architect take to meet these requirements?

- A. Use S3 server access logging on the bucket that houses the reports with the read and write data events and log file validation options enabled.
- B. Use S3 server access logging on the bucket that houses the reports with the read and write management events and log file validation options enabled.
- C. Use AWS CloudTrail to create a new trail. Configure the trail to log read and write data events on the S3 bucket that houses the reports. Log these events to a new bucket, and enable log file validation.
- D. Use AWS CloudTrail to create a new trail. Configure the trail to log read and write management events on the S3 bucket that houses the reports. Log these events to a new bucket, and enable log file validation.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Reference: <https://docs.aws.amazon.com/AmazonS3/latest/user-guide/enable-cloudtrailevents.html>

QUESTION 275

A company has an on-premises volume backup solution that has reached its end of life. The company wants to use AWS as part of a new backup solution and wants to maintain local access to all the data while it is backed up on AWS. The company wants to ensure that the data backed up on AWS is automatically and securely transferred.

Which solution meets these requirements?

- A. Use AWS Snowball to migrate data out of the on-premises solution to Amazon S3. Configure on-premises systems to mount the Snowball S3 endpoint to provide local access to the data.
- B. Use AWS Snowball Edge to migrate data out of the on-premises solution to Amazon S3. Use the Snowball Edge file interface to provide on-premises systems with local access to the data.
- C. Use AWS Storage Gateway and configure a cached volume gateway. Run the Storage Gateway software appliance on premises and configure a percentage of data to cache locally. Mount the gateway storage volumes to provide local access to the data.
- D. Use AWS Storage Gateway and configure a stored volume gateway. Run the Storage Gateway software appliance on premises and map the gateway storage volumes to on-premises storage. Mount the gateway storage volumes to provide local access to the data.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 276

A company is using a third-party vendor to manage its marketplace analytics. The vendor needs limited programmatic access to resources in the company's account. All the needed policies have been created to grant appropriate access. Which additional component will provide the vendor with the MOST secure access to the account?

- A. Create an IAM user.
- B. Implement a service control policy (SCP)
- C. Use a cross-account role with an external ID.
- D. Configure a single sign-on (SSO) identity provider.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Reference:

https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_policies_scps.html#data-from-iam

QUESTION 277

A company is developing an ecommerce application that will consist of a load-balanced front end, a container-based application, and a relational database. A solutions architect needs to create a highly available solution that operates with as little manual intervention as possible.

Which solutions meet these requirements? (Choose two.)

- A. Create an Amazon RDS DB instance in Multi-AZ mode.
- B. Create an Amazon RDS DB instance and one or more replicas in another Availability Zone.
- C. Create an Amazon EC2 instance-based Docker cluster to handle the dynamic application load.
- D. Create an Amazon Elastic Container Service (Amazon ECS) cluster with a Fargate launch type to handle the dynamic application load.
- E. Create an Amazon Elastic Container Service (Amazon ECS) cluster with an Amazon EC2 launch type to handle the dynamic application load.

Correct Answer: AD

Section: (none)

Explanation

Explanation/Reference:

Reference: https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_ReadRepl.html

QUESTION 278

A company has an ecommerce application that stores data in an on-premises SQL database. The company has decided to migrate this database to AWS. However, as part of the migration, the company wants to find a way to attain sub-millisecond responses to common read requests.



A solutions architect knows that the increase in speed is paramount and that a small percentage of stale data returned in the database reads is acceptable. What should the solutions architect recommend?

- A. Build Amazon RDS read replicas.
- B. Build the database as a larger instance type.
- C. Build a database cache using Amazon ElastiCache.
- D. Build a database cache using Amazon Elasticsearch Service (Amazon ES).

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Reference: <https://aws.amazon.com/redis/>

QUESTION 279

A company has an application that ingests incoming messages. These messages are then quickly consumed by dozens of other applications and microservices. The number of messages varies drastically and sometimes spikes as high as 100,000 each second. The company wants to decouple the solution and increase scalability. Which solution meets these requirements?

- A. Persist the messages to Amazon Kinesis Data Analytics. All the applications will read and process the messages.
- B. Deploy the application on Amazon EC2 instances in an Auto Scaling group, which scales the number of EC2 instances based on CPU metrics.
- C. Write the messages to Amazon Kinesis Data Streams with a single shard. All applications will read from the stream and process the messages.
- D. Publish the messages to an Amazon Simple Notification Service (Amazon SNS) topic with one or more Amazon Simple Queue Service (Amazon SQS) subscriptions. All applications then process the messages from the queues.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

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



QUESTION 280

A solutions architect is designing the cloud architecture for a company that needs to host hundreds of machine learning models for its users. During startup, the models need to load up to 10 GB of data from Amazon S3 into memory, but they do not need disk access. Most of the models are used sporadically, but the users expect all of them to be highly available and accessible with low latency. Which solution meets the requirements and is MOST cost-effective?

- A. Deploy models as AWS Lambda functions behind an Amazon API Gateway for each model.
- B. Deploy models as Amazon Elastic Container Service (Amazon ECS) services behind an Application Load Balancer for each model.
- C. Deploy models as AWS Lambda functions behind a single Amazon API Gateway with path-based routing where one path corresponds to each model.
- D. Deploy models as Amazon Elastic Container Service (Amazon ECS) services behind a single Application Load Balancer with path-based routing where one path corresponds to each model.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 281

A company has created a multi-tier application for its ecommerce website. The website uses an Application Load Balancer that resides in the public subnets, a web tier in the public subnets, and a MySQL cluster hosted on Amazon EC2 instances in the private subnets. The MySQL database needs to retrieve product catalog and pricing information that is hosted on the internet by a thirdparty provider. A solutions architect must devise a strategy that maximizes security without increasing operational overhead. What should the solutions architect do to meet these requirements?

- A. Deploy a NAT instance in the VPC. Route all the internet-based traffic through the NAT instance.
- B. Deploy a NAT gateway in the public subnets. Modify the private subnet route table to direct all internet-bound traffic to the NAT gateway.
- C. Configure an internet gateway and attach it to the VPC. Modify the private subnet route table to direct internet-bound traffic to the internet gateway.
- D. Configure a virtual private gateway and attach it to the VPC. Modify the private subnet route table to direct internet-bound traffic to the virtual private gateway.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:

QUESTION 282

A company is backing up on-premises databases to local file server shares using the SMB protocol. The company requires immediate access to 1 week of backup files to meet recovery objectives. Recovery after a week is less likely to occur, and the company can tolerate a delay in accessing those older backup files. What should a solutions architect do to meet these requirements with the LEAST operational effort?

- A. Deploy Amazon FSx for Windows File Server to create a file system with exposed file shares with sufficient storage to hold all the desired backups.
- B. Deploy an AWS Storage Gateway file gateway with sufficient storage to hold 1 week of backups. Point the backups to SMB shares from the file gateway.
- C. Deploy Amazon Elastic File System (Amazon EFS) to create a file system with exposed NFS shares with sufficient storage to hold all the desired backups.
- D. Continue to back up to the existing file shares. Deploy AWS Database Migration Service (AWS DMS) and define a copy task to copy backup files older than 1 week to Amazon S3, and delete the backup files from the local file store.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:

QUESTION 283

A company has developed a microservices application. It uses a client-facing API with Amazon API Gateway and multiple internal services hosted on Amazon EC2 instances to process user requests. The API is designed to support unpredictable surges in traffic, but internal services may become overwhelmed and unresponsive for a period of time during surges. A solutions architect needs to design a more reliable solution that reduces errors when internal services become unresponsive or unavailable. Which solution meets these requirements?

- A. Use AWS Auto Scaling to scale up internal services when there is a surge in traffic.
- B. Use different Availability Zones to host internal services. Send a notification to a system administrator when an internal service becomes unresponsive.
- C. Use an Elastic Load Balancer to distribute the traffic between internal services. Configure Amazon CloudWatch metrics to monitor traffic to internal services.
- D. Use Amazon Simple Queue Service (Amazon SQS) to store user requests as they arrive. Change the internal services to retrieve the requests from the queue for processing.

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Explanation:

QUESTION 284

A company is hosting 60 TB of production-level data in an Amazon S3 bucket. A solutions architect needs to bring that data on premises for quarterly audit requirements. This export of data must be encrypted while in transit. The company has low network bandwidth in place between AWS and its on-premises data center. What should the solutions architect do to meet these requirements?

- A. Deploy AWS Migration Hub with 90-day replication windows for data transfer.
- B. Deploy an AWS Storage Gateway volume gateway on AWS. Enable a 90-day replication window to transfer the data.
- C. Deploy Amazon Elastic File System (Amazon EFS), with lifecycle policies enabled, on AWS. Use it to transfer the data.
- D. Deploy an AWS Snowball device in the on-premises data center after completing an export job request in the AWS Snowball console.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:

QUESTION 285

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

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

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Reference: <https://docs.aws.amazon.com/AmazonS3/latest/dev/security-best-practices.html>

QUESTION 286

A solutions architect is designing a new API using Amazon API Gateway that will receive requests from users. The volume of requests is highly variable; several hours can pass without receiving a single request. The data processing will take place asynchronously, but should be completed within a few seconds after a request is made. Which compute service should the solutions architect have the API invoke to deliver the requirements at the lowest cost?

- A. An AWS Glue job
- B. An AWS Lambda function
- C. A containerized service hosted in Amazon Elastic Kubernetes Service (Amazon EKS)
- D. A containerized service hosted in Amazon ECS with Amazon EC2

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 287

A company hosts its application in the AWS Cloud. The application runs on Amazon EC2 instances behind an Elastic Load Balancer in an Auto Scaling group and with an Amazon DynamoDB table. The company wants to ensure the application can be made available in another AWS Region with minimal downtime. What should a solutions architect do to meet these requirements with the LEAST amount of downtime?

- A. Create an Auto Scaling group and a load balancer in the disaster recovery Region. Configure the DynamoDB table as a global table. Configure DNS failover to point to the new disaster recovery Region's load balancer.
- B. Create an AWS CloudFormation template to create EC2 instances, load balancers, and DynamoDB tables to be executed when needed. Configure DNS failover to point to the new disaster recovery Region's load balancer.
- C. Create an AWS CloudFormation template to create EC2 instances and a load balancer to be executed when needed. Configure the DynamoDB table as a global table. Configure DNS failover to point to the new disaster recovery Region's load balancer.
- D. Create an Auto Scaling group and load balancer in the disaster recovery Region. Configure the DynamoDB table as a global table. Create an Amazon CloudWatch alarm to trigger and AWS Lambda function that updates Amazon Route 53 pointing to the disaster recovery load balancer.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 288

A business application is hosted on Amazon EC2 and uses Amazon S3 for encrypted object storage. The chief information security officer has directed that no application traffic between the two services should traverse the public internet. Which capability should the solutions architect use to meet the compliance requirements?

- A. AWS Key Management Service (AWS KMS)



- B. VPC endpoint
- C. Private subnet
- D. Virtual private gateway

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:

QUESTION 289

A solutions architect is designing a solution that requires frequent updates to a website that is hosted on Amazon S3 with versioning enabled. For compliance reasons, the older versions of the objects will not be accessed frequently and will need to be deleted after 2 years.

What should the solutions architect recommend to meet these requirements at the LOWEST cost?

- A. Use S3 batch operations to replace object tags. Expire the objects based on the modified tags.
- B. Configure an S3 Lifecycle policy to transition older versions of objects to S3 Glacier. Expire the objects after 2 years.
- C. Enable S3 Event Notifications on the bucket that sends older objects to the Amazon Simple Queue Service (Amazon SQS) queue for further processing.
- D. Replicate older object versions to a new bucket. Use an S3 Lifecycle policy to expire the objects in the new bucket after 2 years.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:

QUESTION 290

A company runs an application on an Amazon EC2 instance backed by Amazon Elastic Block Store (Amazon EBS). The instance needs to be available for 12 hours daily. The company wants to save costs by making the instance unavailable outside the window required for the application.

However, the contents of the instance's memory must be preserved whenever the instance is unavailable.

What should a solutions architect do to meet this requirement?

- A. Stop the instance outside the application's availability window. Start up the instance again when required.
- B. Hibernate the instance outside the application's availability window. Start up the instance again when required.
- C. Use Auto Scaling to scale down the instance outside the application's availability window. Scale up the instance when required.
- D. Terminate the instance outside the application's availability window. Launch the instance by using a preconfigured Amazon Machine Image (AMI) when required.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:

QUESTION 291

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

Which additional configuration strategy should the solutions architect use to meet these requirements?

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

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:

Explanation:

QUESTION 292

A company hosts historical weather records in Amazon S3. The records are downloaded from the company's website by a way of a URL that resolves to a domain name. Users all over the world access this content through subscriptions. A third-party provider hosts the company's root domain name, but the company recently migrated some of its services to Amazon Route 53. The company wants to consolidate contracts, reduce latency for users, and reduce costs related to serving the application to subscribers.

Which solution meets these requirements?

- A. Create a web distribution on Amazon CloudFront to serve the S3 content for the application.
Create a CNAME record in a Route 53 hosted zone that points to the CloudFront distribution, resolving to the application's URL domain name.
- B. Create a web distribution on Amazon CloudFront to serve the S3 content for the application.
Create an ALIAS record in the Amazon Route 53 hosted zone that points to the CloudFront distribution, resolving to the application's URL domain name.
- C. Create an A record in a Route 53 hosted zone for the application. Create a Route 53 traffic policy for the web application, and configure a geolocation rule. Configure health checks to check the health of the endpoint and route DNS queries to other endpoints if an endpoint is unhealthy.
- D. Create an A record in a Route 53 hosted zone for the application. Create a Route 53 traffic policy for the web application, and configure a geoproximity rule. Configure health checks to check the health of the endpoint and route DNS queries to other endpoints if an endpoint is unhealthy.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 293

A company owns an asynchronous API that is used to ingest user requests and, based on the request type, dispatch requests to the appropriate microservice for processing. The company is using Amazon API Gateway to deploy the API front end, and an AWS Lambda function that invokes Amazon DynamoDB to store user requests before dispatching them to the processing microservices.

The company provisioned as much DynamoDB throughput as its budget allows, but the company is still experiencing availability issues and is losing user requests.

What should a solutions architect do to address this issue without impacting existing users?

- A. Add throttling on the API Gateway with server-side throttling limits.
- B. Use DynamoDB Accelerator (DAX) and Lambda to buffer writes to DynamoDB.
- C. Create a secondary index in DynamoDB for the table with the user requests.
- D. Use the Amazon Simple Queue Service (Amazon SQS) queue and Lambda to buffer writes to DynamoDB.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 294

A company is moving its on-premises applications to Amazon EC2 instances. However, as a result of fluctuating compute requirements, the EC2 instances must always be ready to use between 8 AM and 5 PM in specific Availability Zones.

Which EC2 instances should the company choose to run the applications?

- A. Scheduled Reserved Instances
- B. On-Demand Instances
- C. Spot Instances as part of a Spot Fleet
- D. EC2 instances in an Auto Scaling group

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 295

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

What should a solutions architect recommend?

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

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 296

A company is building an application on Amazon EC2 instances that generates temporary transactional data. The application requires access to data storage that can provide configurable and consistent IOPS.

What should a solutions architect recommend?

- A. Provision an EC2 instance with a Throughput Optimized HDD (st1) root volume and a Cold HDD (sc1) data volume.
- B. Provision an EC2 instance with a Throughput Optimized HDD (st1) volume that will serve as the root and data volume.
- C. Provision an EC2 instance with a General Purpose SSD (gp2) root volume and Provisioned IOPS SSD (io1) data volume.
- D. Provision an EC2 instance with a General Purpose SSD (gp2) root volume. Configure the application to store its data in an Amazon S3 bucket.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 297

A solutions architect needs to design a resilient solution for Windows users' home directories. The solution must provide fault tolerance, file-level backup and recovery, and access control, based upon the company's Active Directory.

Which storage solution meets these requirements?

- A. Configure Amazon S3 to store the users' home directories. Join Amazon S3 to Active Directory.
- B. Configure a Multi-AZ file system with Amazon FSx for Windows File Server. Join Amazon FSx to Active Directory.
- C. Configure Amazon Elastic File System (Amazon EFS) for the users' home directories. Configure AWS Single Sign-On with Active Directory.
- D. Configure Amazon Elastic Block Store (Amazon EBS) to store the users' home directories. Configure AWS Single Sign-On with Active Directory.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 298

A company wants to move a multi-tiered application from on premises to the AWS Cloud to improve the application's performance. The application consists of application tiers that communicate with each other by way of RESTful services.

Transactions are dropped when one tier becomes overloaded. A solutions architect must design a solution that resolves these issues and modernizes the application.

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

- A. Use Amazon API Gateway and direct transactions to the AWS Lambda functions as the application layer. Use Amazon Simple Queue Service (Amazon SQS) as the communication layer between application services.
- B. Use Amazon CloudWatch metrics to analyze the application performance history to determine the server's peak utilization during the performance failures. Increase the size of the application server's Amazon EC2 instances to meet the



peak requirements.

- C. Use Amazon Simple Notification Service (Amazon SNS) to handle the messaging between application servers running on Amazon EC2 in an Auto Scaling group. Use Amazon CloudWatch to monitor the SNS queue length and scale up and down as required.
- D. Use Amazon Simple Queue Service (Amazon SQS) to handle the messaging between application servers running on Amazon EC2 in an Auto Scaling group. Use Amazon CloudWatch to monitor the SQS queue length and scale up when communication failures are detected.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 299

A company serves a multilingual website from a fleet of Amazon EC2 instances behind an Application Load Balancer (ALB). This architecture is currently running in the us-west-1 Region but is exhibiting high request latency for users located in other parts of the world.

The website needs to serve requests quickly and efficiently regardless of a user's location.

However, the company does not want to recreate the existing architecture across multiple Regions.

How should a solutions architect accomplish this?

- A. Replace the existing architecture with a website served from an Amazon S3 bucket. Configure an Amazon CloudFront distribution with the S3 bucket as the origin.
- B. Configure an Amazon CloudFront distribution with the ALB as the origin. Set the cache behavior settings to only cache based on the Accept-Language request header.
- C. Set up Amazon API Gateway with the ALB as an integration. Configure API Gateway to use an HTTP integration type. Set up an API Gateway stage to enable the API cache.
- D. Launch an EC2 instance in each additional Region and configure NGINX to act as a cache server for that Region. Put all the instances plus the ALB behind an Amazon Route 53 record set with a geolocation routing policy.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:



QUESTION 300

A software vendor is deploying a new software-as-a-service (SaaS) solution that will be utilized by many AWS users. The service is hosted in a VPC behind a Network Load Balancer. The software vendor wants to provide access to this service to users with the least amount of administrative overhead and without exposing the service to the public internet.

What should a solutions architect do to accomplish this goal?

- A. Create a peering VPC connection from each user's VPC to the software vendor's VPC.
- B. Deploy a transit VPC in the software vendor's AWS account. Create a VPN connection with each user account.
- C. Connect the service in the VPC with an AWS Private Link endpoint. Have users subscribe to the endpoint.
- D. Deploy a transit VPC in the software vendor's AWS account. Create an AWS Direct Connect connection with each user account.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 301

A user wants to list the IAM role that is attached to their Amazon EC2 instance. The user has login access to the EC2 instance but does not have IAM permissions.

What should a solutions architect do to retrieve this information?

- A. Run the following EC2 command: `curl http://169.254.169.254/latest/meta-data/iam/info`
- B. Run the following EC2 command: `curl http://169.254.169.254/latest/user-data/iam/info`
- C. Run the following EC2 command: `http://169.254.169.254/latest/dynamic/instance-identity/`
- D. Run the following AWS CLI command: `aws iam get-instance-profile --instance-profile-name ExampleInstanceProfile`

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Reference: <https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/iam-roles-for-amazonec2.html>

QUESTION 302

A company has an application that is hosted on Amazon EC2 instances in two private subnets. A solutions architect must make the application available on the public internet with the least amount of administrative effort. What should the solutions architect recommend?

- A. Create a load balancer and associate two public subnets from the same Availability Zones as the private instances. Add the private instances to the load balancer.
- B. Create a load balancer and associate two private subnets from the same Availability Zones as the private instances. Add the private instances to the load balancer.
- C. Create an Amazon Machine Image (AMI) of the instances in the private subnet and restore in the public subnet. Create a load balancer and associate two public subnets from the same Availability Zones as the public instances.
- D. Create an Amazon Machine Image (AMI) of the instances in the private subnet and restore in the public subnet. Create a load balancer and associate two private subnets from the same Availability Zones as the public instances.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 303

A company has two applications: a sender application that sends messages with payloads to be processed and a processing application intended to receive 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. Configure both applications to use the instance. Store, process, and delete the messages, respectively.
- B. Use an Amazon Kinesis data stream to receive the messages from the sender application. Integrate the processing application with the Kinesis Client Library (KCL).
- C. Integrate the sender and processor applications with an Amazon Simple Queue Service (Amazon SQS) queue. Configure a dead-letter queue to collect the messages that failed to process.
- D. Subscribe the processing application to an Amazon Simple Notification Service (Amazon SNS) topic to receive notifications to process. Integrate the sender application to write to the SNS topic.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 304

A company's website hosted on Amazon EC2 instances processes classified data stored in Amazon S3. Due to security concerns, the company requires a private and secure connection between its EC2 resources and Amazon S3.

Which solution meets these requirements?

- A. Set up S3 bucket policies to allow access from a VPC endpoint.
- B. Set up an IAM policy to grant read-write access to the S3 bucket.
- C. Set up a NAT gateway to access resources outside the private subnet.
- D. Set up an access key ID and a secret access key to access the S3 bucket.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Reference: <https://docs.aws.amazon.com/AmazonS3/latest/dev/access-control-overview.html>

QUESTION 305

A company hosts its multi-tier, public web application in the AWS Cloud. The web application runs on Amazon EC2 instances and its database runs on Amazon RDS. The company is anticipating a large increase in sales during an upcoming holiday weekend. A solutions architect needs to build a solution to analyze the performance of the web application with a granularity of no more than 2 minutes. What should the solutions architect do to meet this requirement?

- A. Send Amazon CloudWatch logs to Amazon Redshift. Use Amazon QuickSight to perform further analysis.
- B. Enable detailed monitoring on all EC2 instances. Use Amazon CloudWatch metrics to perform further analysis.
- C. Create an AWS Lambda function to fetch EC2 logs from Amazon CloudWatch Logs. Use Amazon CloudWatch metrics to perform further analysis.
- D. Send EC2 logs to Amazon S3. Use Amazon Redshift to fetch logs from the S3 bucket to process raw data for further analysis with Amazon QuickSight.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 306

A company has developed a new video game as a web application. The application is in a three-tier architecture in a VPC with Amazon RDS for MySQL. In the database layer several players will compete concurrently online. The game's developers want to display a top-10 scoreboard in near-real time and offer the ability to stop and restore the game while preserving the current scores. What should a solutions architect do to meet these requirements?

- A. Set up an Amazon ElastiCache for Memcached cluster to cache the scores for the web application to display.
- B. Set up an Amazon ElastiCache for Redis cluster to compute and cache the scores for the web application to display.
- C. Place an Amazon CloudFront distribution in front of the web application to cache the scoreboard in a section of the application.
- D. Create a read replica on Amazon RDS for MySQL to run queries to compute the scoreboard and serve the read traffic to the web application.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:



QUESTION 307

A company is moving its on-premises Oracle database to Amazon Aurora PostgreSQL. The database has several applications that write to the same tables. The applications need to be migrated one by one with a month in between each migration. Management has expressed concerns that the database has a high number of reads and writes. The data must be kept in sync across both databases throughout the migration. What should a solutions architect recommend?

- A. Use AWS DataSync for the initial migration. Use AWS Database Migration Service (AWS DMS) to create a change data capture (CDC) replication task and a table mapping to select all tables.
- B. Use AWS DataSync for the initial migration. Use AWS Database Migration Service (AWS DMS) to create a full load plus change data capture (CDC) replication task and a table mapping to select all tables.
- C. Use the AWS Schema Conversion Tool with AWS Database Migration Service (AWS DMS) using a memory optimized replication instance. Create a full load plus change data capture (CDC) replication task and a table mapping to select all tables.
- D. Use the AWS Schema Conversion Tool with AWS Database Migration Service (AWS DMS) using a compute optimized replication instance. Create a full load plus change data capture (CDC) replication task and a table mapping to select the largest tables.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 308

A company recently migrated a message processing system to AWS. The system receives messages into an ActiveMQ queue running on an Amazon EC2 instance. Messages are processed by a consumer application running on Amazon EC2. The consumer application processes the messages and writes results to a MySQL database running on Amazon EC2. The company wants this application to be highly available with low operational complexity. Which architecture offers the HIGHEST availability?

- A. Add a second ActiveMQ server to another Availability Zone. Add an additional consumer EC2 instance in another Availability Zone. Replicate the MySQL database to another Availability Zone.
- B. Use Amazon MQ with active/standby brokers configured across two Availability Zones. Add an additional consumer EC2 instance in another Availability Zone. Replicate the MySQL database to another Availability Zone.

- C. Use Amazon MQ with active/standby brokers configured across two Availability Zones. Add an additional consumer EC2 instance in another Availability Zone. Use Amazon RDS for MySQL with Multi-AZ enabled.
- D. Use Amazon MQ with active/standby brokers configured across two Availability Zones. Add an Auto Scaling group for the consumer EC2 instances across two Availability Zones. Use Amazon RDS for MySQL with Multi-AZ enabled.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 309

A company is planning on deploying a newly built application on AWS in a default VPC. The application will consist of a web layer and database layer. The web server was created in public subnets, and the MySQL database was created in private subnets. All subnets are created with the default network ACL settings, and the default security group in the VPC will be replaced with new custom security groups.

The following are the key requirements:

The web servers must be accessible only to users on an SSL connection.

The database should be accessible to the web layer, which is created in a public subnet only.

All traffic to and from the IP range 182.20.0.0/16 subnet should be blocked.

Which combination of steps meets these requirements? (Choose two.)

- A. Create a database server security group with inbound and outbound rules for MySQL port 3306 traffic to and from anywhere (0.0.0.0/0).
- B. Create a database server security group with an inbound rule for MySQL port 3306 and specify the source as a web server security group.
- C. Create a web server security group with an inbound allow rule for HTTPS port 443 traffic from anywhere (0.0.0.0/0) and an inbound deny rule for IP range 182.20.0.0/16.
- D. Create a web server security group with an inbound rule for HTTPS port 443 traffic from anywhere (0.0.0.0/0). Create network ACL inbound and outbound deny rules for IP range 182.20.0.0/16.
- E. Create a web server security group with inbound and outbound rules for HTTPS port 443 traffic to and from anywhere (0.0.0.0/0). Create a network ACL inbound deny rule for IP range 182.20.0.0/16.

Correct Answer: BD

Section: (none)

Explanation

Explanation/Reference:

Explanation:



QUESTION 310

A company has an on-premises application that collects data and stores it to an on-premises NFS server. The company recently set up a 10 Gbps AWS Direct Connect connection. The company is running out of storage capacity on premises. The company needs to migrate the application data from on premises to the AWS Cloud while maintaining low-latency access to the data from the on-premises application.

What should a solutions architect do to meet these requirements?

- A. Deploy AWS Storage Gateway for the application data, and use the file gateway to store the data in Amazon S3. Connect the on-premises application servers to the file gateway using NFS.
- B. Attach an Amazon Elastic File System (Amazon EFS) file system to the NFS server, and copy the application data to the EFS file system. Then connect the on-premises application to Amazon EFS.
- C. Configure AWS Storage Gateway as a volume gateway. Make the application data available to the on-premises application from the NFS server and with Amazon Elastic Block Store (Amazon EBS) snapshots.
- D. Create an AWS DataSync agent with the NFS server as the source location and an Amazon Elastic File System (Amazon EFS) file system as the destination for application data transfer. Connect the on-premises application to the EFS file system.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 311

A solutions architect needs to design a network that will allow multiple Amazon EC2 instances to access a common data source used for mission-critical data that can be accessed by all the EC2 instances simultaneously. The solution must be highly scalable, easy to implement and support the NFS protocol.

Which solution meets these requirements?

- A. Create an Amazon Elastic File System (Amazon EFS) file system. Configure a mount target in each Availability Zone. Attach each instance to the appropriate mount target.
- B. Create an additional EC2 instance and configure it as a file server. Create a security group that allows communication between the Instances and apply that to the additional instance.
- C. Create an Amazon S3 bucket with the appropriate permissions. Create a role in AWS IAM that grants the correct permissions to the S3 bucket. Attach the role to the EC2 Instances that need access to the data.

- D. Create an Amazon Elastic Block Store (Amazon EBS) volume with the appropriate permissions.
Create a role in AWS IAM that grants the correct permissions to the EBS volume. Attach the role to the EC2 instances that need access to the data.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:

QUESTION 312

A company hosts its application using Amazon Elastic Container Service (Amazon ECS) and wants to ensure high availability. The company wants to be able to deploy updates to its application even if nodes in one Availability Zone are not accessible.

The expected request volume for the application is 100 requests per second, and each container task is able to serve at least 60 requests per second. The company set up Amazon ECS with a rolling update deployment type with the minimum healthy percent parameter set to 50% and the maximum percent set to 100%.

Which configuration of tasks and Availability Zones meets these requirements?

- A. Deploy the application across two Availability Zones, with one task in each Availability Zone.
- B. Deploy the application across two Availability Zones, with two tasks in each Availability Zone.
- C. Deploy the application across three Availability Zones, with one task in each Availability Zone.
- D. Deploy the application across three Availability Zones, with two tasks in each Availability Zone.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:

QUESTION 313

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

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

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:

QUESTION 314

A company wants to improve the availability and performance of its hybrid application. The application consists of a stateful TCP-based workload hosted on Amazon EC2 instances in different AWS Regions and a stateless UOP-based workload hosted on premises.

Which combination of actions should a solutions architect take to improve availability and performance? (Choose two.)

- A. Create an accelerator using AWS Global Accelerator. Add the load balancers as endpoints.
- B. Create an Amazon CloudFront distribution with an origin that uses Amazon Route 53 latencybased routing to route requests to the load balancers.
- C. Configure two Application Load Balancers in each Region. The first will route to the EC2 endpoints and the second will route to the on-premises endpoints.
- D. Configure a Network Load Balancer in each Region to address the EC2 endpoints. Configure a Network Load Balancer in each Region that routes to the on-premises endpoints.
- E. Configure a Network Load Balancer in each Region to address the EC2 endpoints. Configure an Application Load Balancer in each Region that routes to the on-premises endpoints

Correct Answer: AD
Section: (none)
Explanation

Explanation/Reference:

Explanation:

QUESTION 315

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.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 316

A company has an ecommerce application running in a single VPC. The application stack has a single web server and an Amazon RDS Multi-AZ DB instance.

The company launches new products twice a month. This increases website traffic by approximately 400% for a minimum of 72 hours. During product launches, users experience slow response times and frequent timeout errors in their browsers.

What should a solutions architect do to mitigate the slow response times and timeout errors while minimizing operational overhead?

- A. Increase the instance size of the web server.
- B. Add an Application Load Balancer and an additional web server.
- C. Add Amazon EC2 Auto Scaling and an Application Load Balancer.
- D. Deploy an Amazon ElastiCache cluster to store frequently accessed data.



Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 317

A solutions architect is designing an architecture to run a third-party database server. The database software is memory intensive and has a CPU-based licensing model where the cost increases with the number of vCPU cores within the operating system. The solutions architect must select an Amazon EC2 instance with sufficient memory to run the database software, but the selected instance has a large number of vCPUs. The solutions architect must ensure that the vCPUs will not be underutilized and must minimize costs.

Which solution meets these requirements?

- A. Select and launch a smaller EC2 instance with an appropriate number of vCPUs.
- B. Configure the CPU cores and threads on the selected EC2 instance during instance launch.
- C. Create a new EC2 instance and ensure multithreading is enabled when configuring the instance details.
- D. Create a new Capacity Reservation and select the appropriate instance type. Launch the instance into this new Capacity Reservation.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 318

A company receives 10 TB of instrumentation data each day from several machines located at a single factory. The data consists of JSON files stored on a storage area network (SAN) in an onpremises data center located within the factory. The company wants to send this data to Amazon S3 where it can be accessed by several additional systems that provide critical near-real-time analytics. A secure transfer is important because the data is considered sensitive.

Which solution offers the MOST reliable data transfer?

- A. AWS DataSync over public internet
- B. AWS DataSync over AWS Direct Connect
- C. AWS Database Migration Service (AWS DMS) over public internet
- D. AWS Database Migration Service (AWS DMS) over AWS Direct Connect

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 319

A company is creating a web application that will store a large number of images in Amazon S3.

The images will be accessed by users over variable periods of time. The company wants to:

Retain all the images

Incur no cost for retrieval.

Have minimal management overhead.

Have the images available with no impact on retrieval time.

Which solution meets these requirements?

- A. Implement S3 Intelligent-Tiering
- B. Implement S3 storage class analysis
- C. Implement an S3 Lifecycle policy to move data to S3 Standard-Infrequent Access (S3 Standard- IA).
- D. Implement an S3 Lifecycle policy to move data to S3 One Zone-Infrequent Access (S3 One Zone- IA).

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 320

A company hosts more than 300 global websites and applications. The company requires a platform to analyze more than 30 TB of clickstream data each day. What should a solutions architect do to transmit and process the clickstream data?

- A. Design an AWS Data Pipeline to archive the data to an Amazon S3 bucket and run an Amazon EMR cluster with the data to generate analytics.
- B. Create an Auto Scaling group of Amazon EC2 instances to process the data and send it to an Amazon S3 data lake for Amazon Redshift to use for analysis.
- C. Cache the data to Amazon CloudFront. Store the data in an Amazon S3 bucket. When an object is added to the S3 bucket, run an AWS Lambda function to process the data for analysis.
- D. Collect the data from Amazon Kinesis Data Streams. Use Amazon Kinesis Data Firehose to transmit the data to an Amazon S3 data lake. Load the data in Amazon Redshift for analysis.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 321

A company wants to build an online marketplace application on AWS as a set of loosely coupled microservices. For this application, when a customer submits a new order, two microservices should handle the event simultaneously. The Email microservice will send a confirmation email, and the OrderProcessing microservice will start the order delivery process. If a customer cancels an order, the OrderCancellation and Email microservices should handle the event simultaneously.

A solutions architect wants to use Amazon Simple Queue Service (Amazon SQS) and Amazon Simple Notification Service (Amazon SNS) to design the messaging between the microservices.

How should the solutions architect design the solution?

- A. Create a single SQS queue and publish order events to it. The Email OrderProcessing and OrderCancellation microservices can then consume messages of the queue.



- B. Create three SNS topics for each microservice. Publish order events to the three topics. Subscribe each of the Email OrderProcessing and OrderCancellation microservices to its own topic.
- C. Create an SNS topic and publish order events to it. Create three SQS queues for the Email OrderProcessing and OrderCancellation microservices. Subscribe all SQS queues to the SNS topic with message filtering.
- D. Create two SQS queues and publish order events to both queues simultaneously. One queue is for the Email and OrderProcessing microservices. The second queue is for the Email and OrderCancellation microservices.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 322

A company is running a multi-tier ecommerce web application in the AWS Cloud. The application runs on Amazon EC2 Instances with an Amazon RDS MySQL Multi-AZ DB instance. Amazon RDS is configured with the latest generation instance with 2,000 GB of storage in an Amazon Elastic Block Store (Amazon EBS) General Purpose SSD (gp2) volume. The database performance impacts the application during periods of high demand. After analyzing the logs in Amazon CloudWatch Logs, a database administrator finds that the application performance always degrades when the number of read and write IOPS is higher than 6,000. What should a solutions architect do to improve the application performance?

- A. Replace the volume with a Magnetic volume.
- B. Increase the number of IOPS on the gp2 volume.
- C. Replace the volume with a Provisioned IOPS (PIOPS) volume.
- D. Replace the 2,000 GB gp2 volume with two 1,000 GBgp2 volumes.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:



QUESTION 323

A company has an application that uses Amazon Elastic File System (Amazon EFS) to store data. The files are 1 GB in size or larger and are accessed often only for the first few days after creation. The application data is shared across a cluster of Linux servers. The company wants to reduce storage costs for the application. What should a solutions architect do to meet these requirements?

- A. Implement Amazon FSx and mount the network drive on each server.
- B. Move the files from Amazon Elastic File System (Amazon EFS) and store them locally on each Amazon EC2 instance.
- C. Configure a Lifecycle policy to move the files to the EFS Infrequent Access (IA) storage class after 7 days.
- D. Move the files to Amazon S3 with S3 lifecycle policies enabled. Rewrite the application to support mounting the S3 bucket.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 324

A company has a service that produces event data. The company wants to use AWS to process the event data as it is received. The data is written in a specific order that must be maintained throughout processing. The company wants to implement a solution that minimizes operational overhead. How should a solution architect accomplish this?

- A. Create an Amazon Simple Queue Service (Amazon SQS) FIFO queue to hold messages. Set up an AWS Lambda function to process messages from the queue.
- B. Create an Amazon Simple Notification Service (Amazon SNS) topic to deliver notifications containing payloads to process. Configure an AWS Lambda function as a subscriber.
- C. Create an Amazon Simple Queue Service (Amazon SQS) standard queue to hold messages. Set up an AWS Lambda function to process messages from the queue independently.
- D. Create an Amazon Simple Notification Service (Amazon SNS) topic to deliver notifications containing payloads to process. Configure an Amazon Simple Queue Service (Amazon SQS) queue as a subscriber.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 325

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.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 326

A company wants to migrate its web application to AWS. The legacy web application consists of a web tier, an application tier, and a MySQL database. The re-architected application must consist of technologies that do not require the administration team to manage instances or clusters.

Which combination of services should a solutions architect include in the overall architecture?

(Choose two.)

- A. Amazon Aurora Serverless
- B. Amazon EC2 Spot Instances
- C. Amazon Elasticsearch Service (Amazon ES)
- D. Amazon RDS for MySQL
- E. AWS Fargate



Correct Answer: DE

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 327

An ecommerce company is experiencing an increase in user traffic. The company's store is deployed on Amazon EC2 instances as a two-tier two application consisting of a web tier and a separate database tier. As traffic increases, the company notices that the architecture is causing significant delays in sending timely marketing and order confirmation email to users. The company wants to reduce the time it spends resolving complex email delivery issues and minimize operational overhead.

What should a solutions architect do to meet these requirements?

- A. Create a separate application tier using EC2 instances dedicated to email processing.
- B. Configure the web instance to send email through Amazon Simple Email Service (Amazon SES).
- C. Configure the web instance to send email through Amazon Simple Notification Service (Amazon SNS).
- D. Create a separate application tier using EC2 instances dedicated to email processing. Place the instances in an Auto Scaling group.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 328

A company recently started using Amazon Aurora as the data store for its global ecommerce application. When large reports are run, developers report that the ecommerce application is performing poorly. After reviewing metrics in

Amazon CloudWatch, a solutions architect finds that the ReadIOPS and CPUUtilization metrics are spiking when monthly reports run. What is the MOST cost-effective solution?

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

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 329

A company uses on-premises servers to host its applications. The company is running out of storage capacity. The applications use both block storage and NFS storage. The company needs a high-performing solution that supports local caching without re-architecting its existing applications.

Which combination of actions should a solutions architect take to meet these requirements?

(Choose two.)

- A. Mount Amazon S3 as a file system to the on-premises servers.
- B. Deploy an AWS Storage Gateway file gateway to replace NFS storage.
- C. Deploy AWS Snowball Edge to provision NFS mounts to on-premises servers.
- D. Deploy an AWS Storage Gateway volume gateway to replace the block storage.
- E. Deploy Amazon Elastic File System (Amazon EFS) volumes and mount them to on-premises servers.

Correct Answer: DE

Section: (none)

Explanation

Explanation/Reference:

Explanation:



QUESTION 330

A solution architect needs to design a highly available application consisting of web, application, and database tiers. HTTPS content delivery should be as close to the edge as possible, with the least delivery time.

Which solution meets these requirements and is MOST secure?

- A. Configure a public Application Load Balancer (ALB) with multiple redundant Amazon EC2 instances in public subnets. Configure Amazon CloudFront to deliver HTTPS content using the public ALB as the origin.
- B. Amazon EC2 instances in private subnets Configure. Configure a public Application Load Balancer with multiple redundant Amazon CloudFront to deliver HTTPS content using the EC2 instances as the origin.
- C. Configure a public Application Load Balancer (ALB) with multiple redundant Amazon EC2 instances in private subnets. Configure Amazon CloudFront to deliver HTTPS content using the public ALB as the origin.
- D. Configure a public Application Load Balancer with multiple redundant Amazon EC2 instances in public subnets. Configure Amazon CloudFront to deliver HTTPS content using the EC2 instances as the origin.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 331

A company has a popular gaming platform running on AWS. The application is sensitive to latency because latency can impact the user experience and introduce unfair advantages to some players.

The application is deployed in every AWS Region. It runs on Amazon EC2 instances that are part of Auto Scaling groups configured behind Application Load Balancers (ALBs). A solutions architect needs to implement a mechanism to monitor the health of the application and redirect traffic to healthy endpoints.

Which solution meets these requirements?

- A. Configure an accelerator in AWS Global Accelerator. Add a listener for the port that the application listens on and attach it to a Regional endpoint in each Region. Add the ALB as the endpoint.
- B. Create an Amazon CloudFront distribution and specify the ALB as the origin server. Configure the cache behavior to use origin cache headers. Use AWS Lambda functions to optimize the traffic.
- C. Create an Amazon CloudFront distribution and specify Amazon S3 as the origin server. Configure the cache behavior to use origin cache headers. Use AWS Lambda functions to optimize the traffic.

D. Configure an Amazon DynamoDB database to serve as the data store for the application. Create a DynamoDB Accelerator (DAX) cluster to act as the in-memory cache for DynamoDB hosting the application data.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 332

A company is designing an internet-facing web application. The application runs on Amazon EC2 for Linux-based instances that store sensitive user data in Amazon RDS MySQL Multi-AZ DB instances. The EC2 instances are in public subnets, and the RDS DB instances are in private subnets. The security team has mandated that the DB instances be secured against web-based attacks.

What should a solutions architect recommend?

- A. Ensure the EC2 instances are part of an Auto Scaling group and are behind an Application Load Balancer. Configure the EC2 instance iptables rules to drop suspicious web traffic. Create a security group for the DB instances. Configure the RDS security group to only allow port 3306 inbound from the individual EC2 instances.
- B. Ensure the EC2 instances are part of an Auto Scaling group and are behind an Application Load Balancer. Move DB instances to the same subnets that EC2 instances are located in. Create a security group for the DB instances. Configure the RDS security group to only allow port 3306 inbound from the individual EC2 instances.
- C. Ensure the EC2 instances are part of an Auto Scaling group and are behind an Application Load Balancer. Use AWS WAF to monitor inbound web traffic for threats. Create a security group for the web application servers and a security group for the DB instances. Configure the RDS security group to only allow port 3306 inbound from the web application server security group.
- D. Ensure the EC2 instances are part of an Auto Scaling group and are behind an Application Load Balancer. Use AWS WAF to monitor inbound web traffic for threats. Configure the Auto Scaling group to automatically create new DB instances under heavy traffic. Create a security group for the RDS DB instances. Configure the RDS security group to only allow port 3306 inbound.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 333

A development team stores its Amazon RDS MySQL DB instance user name and password credentials in a configuration file. The configuration file is stored as plaintext on the root device volume of the team's Amazon EC2 instance. When the team's application needs to reach the database, it reads the file and loads the credentials into the code. The team has modified the permissions of the configuration file so that only the application can read its content. A solution architect must design a more secure solution.

What should the solutions architect do to meet this requirement?

- A. Store the configuration file in Amazon S3. Grant the application access to read the configuration file.
- B. Create an IAM role with permission to access the database. Attach this IAM role to the EC2 instance.
- C. Enable SSL connections on the database instance. Alter the database user to require SSL when logging in.
- D. Move the configuration file to an EC2 instance store, and create an Amazon Machine Image (AMI) of the instance. Launch new instances from this AMI.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 334

A company wants a storage option that enables its data science team to analyze its data on premises and in the AWS Cloud. The team needs to be able to run statistical analyses by using the data on premises and by using a fleet of Amazon EC2 instances across multiple Availability Zones.

What should a solutions architect do to meet these requirements?

- A. Use an AWS Storage Gateway tape gateway to copy the on-premises files into Amazon S3.
- B. Use an AWS Storage Gateway volume gateway to copy the on-premises files into Amazon S3.
- C. Use an AWS Storage Gateway file gateway to copy the on-premises files to Amazon Elastic Block Store (Amazon EBS).
- D. Attach an Amazon Elastic File System (Amazon EFS) file system to the on-premises servers. Copy the files to Amazon EFS.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 335

A company wants to improve the availability and performance of its stateless UDP-based workload. The workload is deployed on Amazon EC2 instances in multiple AWS Regions. What should a solutions architect recommend to accomplish this?

- A. Place the EC2 instances behind Network Load Balancers (NLBs) in each Region. Create an accelerator using AWS Global Accelerator. Use the NLBs as endpoints for the accelerator.
- B. Place the EC2 instances behind Application Load Balancers (ALBs) in each Region. Create an accelerator using AWS Global Accelerator. Use the ALBs as endpoints for the accelerator.
- C. Place the EC2 instances behind Network Load Balancers (NLBs) in each Region. Create an Amazon CloudFront distribution with an origin that uses Amazon Route 53 latency-based routing to route requests to the NLBs.
- D. Place the EC2 instances behind Application Load Balancers (ALBs) in each Region. Create an Amazon CloudFront distribution with an origin that uses Amazon Route 53 latency-based routing to route requests to the ALBs.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 336

A company wants to use high performance computing (HPC) infrastructure on AWS for financial risk modeling. The company's HPC workloads run on Linux. Each HPC workflow runs on hundreds of AmazonEC2 Spot Instances, is short-lived, and generates thousands of output files that are ultimately stored in persistent storage for analytics and long-term future use.

The company seeks a cloud storage solution that permits the copying of on premises data to longterm persistent storage to make data available for processing by all EC2 instances. The solution should also be a high performance file system that is integrated with persistent storage to read and write datasets and output files.

Which combination of AWS services meets these requirements?

- A. Amazon FSx for Lustre integrated with Amazon S3
- B. Amazon FSx for Windows File Server integrated with Amazon S3
- C. Amazon S3 Glacier integrated with Amazon Elastic Block Store (Amazon EBS)
- D. Amazon S3 bucket with a VPC endpoint integrated with an Amazon Elastic Block Store (Amazon EBS) General Purpose SSD (gp2) volume



Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 337

A solutions architect must design a database solution for a high-traffic ecommerce web application. The database stores customer profiles and shopping cart information. The database must support a peak load of several million requests each second and deliver responses in milliseconds. The operational overhead for managing and scaling the database must be minimized.

Which database solution should the solutions architect recommend?

- A. Amazon Aurora
- B. Amazon DynamoDB
- C. Amazon RDS
- D. Amazon Redshift

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 338

A company is working with an external vendor that requires write access to the company's Amazon Simple Queue Service (Amazon SQS) queue. The vendor has its own AWS account.

What should a solutions architect do to implement least privilege access?

- A. Update the permission policy on the SQS queue to give write access to the vendor's AWS account.
- B. Create an IAM user with write access to the SQS queue and share the credentials for the IAM user.
- C. Update AWS Resource Access Manager to provide write access to the SQS queue from the vendor's AWS account.
- D. Create a cross-account role with access to all SQS queues and use the vendor's AWS account in the trust document for the role.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 339

A company is creating a three-tier web application consisting of a web server, an application server, and a database server. The application will track GPS coordinates of packages as they are being delivered. The application will update the database every 0-5 seconds.

The tracking will need to read as fast as possible for users to check the status of their packages.

Only a few packages might be tracked on some days, whereas millions of package might be tracked on other days. Tracking will need to be searchable by tracking ID customer ID and order ID. Order than 1 month no longer read to be tracked.

What should a solution architect recommend to accomplish this with minimal cost of ownership?

- A. Use Amazon DynamoDB Enable Auto Scaling on the DynamoDB table. Schedule an automatic deletion script for items older than 1 month.
- B. Use Amazon DynamoDB with global secondary indexes. Enable Auto Scaling on the DynamoDB table and the global secondary indexes. Enable TTL on the DynamoDB table.
- C. Use an Amazon RDS On-Demand instance with Provisioned IOPS (PIOPS). Enable Amazon CloudWatch alarms to send notifications when PIOPS are exceeded. Increase and decrease PIOPS as needed.
- D. Use an Amazon RDS Reserved Instance with Provisioned IOPS (PIOPS). Enable Amazon CloudWatch alarms to send notification when PIOPS are exceeded. Increase and decrease PIOPS as needed.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:



QUESTION 340

A solutions architect is creating a data processing job that runs once daily and can take up to 2 hours to complete. If the job is interrupted, it has to restart from the beginning.

How should the solutions architect address this issue in the MOST cost-effective manner?

- A. Create a script that runs locally on an Amazon EC2 Reserved Instance that is triggered by a cron job.
- B. Create an AWS Lambda function triggered by an Amazon EventBridge (Amazon CloudWatch Events) scheduled event.
- C. Use an Amazon Elastic Container Service (Amazon ECS) Fargate task triggered by an Amazon EventBridge (Amazon CloudWatch Events) scheduled event.
- D. Use an Amazon Elastic Container Service (Amazon ECS) task running on Amazon EC2 triggered by an Amazon EventBridge (Amazon CloudWatch Events) scheduled event.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 341

A company needs to store data in Amazon S3. A compliance requirement states that when any changes are made to objects the previous state of the object with any changes must be preserved.

Additionally, files older than 5 years should not be accessed but need to be archived for auditing.

What should a solutions architect recommend that is MOST cost-effective?

- A. Enable object-level versioning and S3 Object Lock in governance mode
- B. Enable object-level versioning and S3 Object Lock in compliance mode
- C. Enable object-level versioning. Enable a lifecycle policy to move data older than 5 years to S3 Glacier Deep Archive
- D. Enable object-level versioning. Enable a lifecycle policy to move data older than 5 years to S3 Standard-Infrequent Access (S3 Standard-IA)

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:

QUESTION 342

A new employee has joined a company as a deployment engineer. The deployment engineer will be using AWS CloudFormation templates to create multiple AWS resources. A solutions architect wants the deployment engineer to perform job activities while following the principle of least privilege.

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

- A. Have the deployment engineer use AWS account root user credentials for performing AWS CloudFormation stack operations.
- B. Create a new IAM user for the deployment engineer and add the IAM user to a group that has the PowerUsers IAM policy attached.
- C. Create a new IAM user for the deployment engineer and add the IAM user to a group that has the Administrate/Access IAM policy attached.
- D. Create a new IAM User for the deployment engineer and add the IAM user to a group that has an IAM policy that allows AWS CloudFormation actions only.
- E. Create an IAM role for the deployment engineer to explicitly define the permissions specific to the AWS CloudFormation stack and launch stacks using Dial IAM role.

Correct Answer: DE
Section: (none)
Explanation

Explanation/Reference:
Explanation:

QUESTION 343

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

What should a solutions architect recommend?

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



Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:

QUESTION 344

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

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

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

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:

QUESTION 345

A company is preparing to deploy a new serverless workload. A solutions architect needs to configure permissions for invoking an AWS Lambda function. The function will be triggered by an Amazon EventBridge (Amazon CloudWatch Events) rule. Permissions should be configured using the principle of least privilege. Which solution will meet these requirements?

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

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 346

A company is building its web application using containers on AWS. The company requires three instances of the web application to run at all times. The application must be able to scale to meet increases in demand. Management is extremely sensitive to cost but agrees that the application should be highly available. What should a solutions architect recommend?

- A. Create an Amazon Elastic Container Service (Amazon ECS) cluster using the Fargate launch type. Create a task definition for the web application. Create an ECS service with a desired count of three tasks.
- B. Create an Amazon Elastic Container Service (Amazon ECS) cluster using the Amazon EC2 launch type with three container instances in one Availability Zone. Create a task definition for the web application. Place one task for each container instance.
- C. Create an Amazon Elastic Container Service (Amazon ECS) cluster using the Fargate launch type with one container instance in three different Availability Zones. Create a task definition for the web application. Create an ECS service with a desired count of three tasks.
- D. Create an Amazon Elastic Container Service (Amazon ECS) cluster using the Amazon EC2 launch type with one container instance in two different Availability Zones. Create a task definition for the web application. Place two tasks on one container instance and one task on the remaining container instance.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 347

A company is Re-architecting a strongly coupled application to be loosely coupled. Previously the application used a request/response pattern to communicate between tiers. The company plans to use Amazon Simple Queue Service (Amazon SQS) to achieve decoupling requirements. The initial design contains one queue for requests and one for responses. However, this approach is not processing all the messages as the application scales. What should a solutions architect do to resolve this issue?

- A. Configure a dead-letter queue on the `ReceiveMessage` API action of the SQS queue.
- B. Configure a FIFO queue, and use the message deduplication ID and message group ID.
- C. Create a temporary queue, with the Temporary Queue Client to receive each response message.
- D. Create a queue for each request and response on startup for each producer, and use a correlation ID message attribute.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 348

A company is launching an ecommerce website on AWS. This website is built with a three-tier architecture that includes a MySQL database in a Multi-AZ deployment of Amazon Aurora MySQL. The website application must be highly available and will initially be launched in an AWS Region with three Availability Zones. The application produces a metric that describes the load the application experiences. Which solution meets these requirements?

- A. Configure an Application Load Balancer (ALB) with Amazon EC2 Auto Scaling behind the ALB with scheduled scaling
- B. Configure an Application Load Balancer (ALB) and Amazon EC2 Auto Scaling behind the ALB with a simple scaling policy.
- C. Configure a Network Load Balancer (NLB) and launch a Spot Fleet with Amazon EC2 Auto Scaling behind the NLB.
- D. Configure an Application Load Balancer (ALB) and Amazon EC2 Auto Scaling behind the ALB with a target tracking scaling policy.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 349

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 a CloudFront and set the Origin Protocol Policy setting to HTTPS. Only for the ViewerProtocol Pokey.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 350

A solutions architect is redesigning a monolithic application to be a loosely coupled application composed of two microservices: Microservice A and Microservice B.

Microservice A places messages in a main Amazon Simple Queue Service (Amazon SQS) queue for Microservice B to consume. When Microservice B fails to process a message after four retries, the message needs to be removed from the queue and stored for further investigation.

What should the solutions architect do to meet these requirements?

- A. Create an SQS dead-letter queue. Microservice B adds failed messages to that queue after it receives and fails to process the message four times.
- B. Create an SQS dead-letter queue. Configure the main SQS queue to deliver messages to the dead-letter queue after the message has been received four times.
- C. Create an SQS queue for failed messages. Microservice A adds failed messages to that queue after Microservice B receives and fails to process the message four times.
- D. Create an SQS queue for failed messages. Configure the SQS queue for failed messages to pull messages from the main SQS queue after the original message has been received four times.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 351

A company has NFS servers in an on-premises data center that need to periodically back up small amounts of data to Amazon S3. Which solution meets these requirements and is MOST costeffective?

- A. Set up AWS Glue to copy the data from the on-premises servers to Amazon S3.
- B. Set up an AWS DataSync agent on the on-premises servers, and sync the data to Amazon S3.
- C. Set up an SFTP sync using AWS Transfer for SFTP to sync data from on-premises to Amazon S3.
- D. Set up an AWS Direct Connect connection between the on-premises data center and a VPC, and copy the data to Amazon S3.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 352

A company runs its production workload on an Amazon Aurora MySQL DB cluster that includes six Aurora Replicas. The company wants near-real-time reporting queries from one of its departments to be automatically distributed across three of the Aurora Replicas. Those three replicas have a different compute and memory specification from the rest of the DB cluster.

Which solution meets these requirements?

- A. Create and use a custom endpoint for the workload.
- B. Create a three-node cluster clone and use the reader endpoint.
- C. Use any of the instance endpoints for the selected three nodes.
- D. Use the reader endpoint to automatically distribute the read-only workload.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 353

A company has multiple applications that use Amazon RDS for MySQL as its database. The company recently discovered that a new custom reporting application has increased the number of Queries on the database. This is slowing down performance.

How should a solutions architect resolve this issue with the LEAST amount of application changes?

- A. Add a secondary DB instance using Multi-AZ.
- B. Set up a read replica and Multi-AZ on Amazon RDS.
- C. Set up a standby replica and Multi-AZ on Amazon RDS.
- D. Use caching on Amazon RDS to improve the overall performance.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 354

A company wants to automate the security assessment of its Amazon EC2 instances. The company needs to validate and demonstrate that security and compliance standards are being followed throughout the development process.

What should a solutions architect do to meet these requirements?

- A. Use Amazon Macie to automatically discover, classify and protect the EC2 instances.
- B. Use Amazon GuardDuty to publish Amazon Simple Notification Service (Amazon SNS) notifications.
- C. Use Amazon Inspector with Amazon CloudWatch to publish Amazon Simple Notification Service (Amazon SNS) notifications
- D. Use Amazon EventBridge (Amazon CloudWatch Events) to detect and react to changes in the status of AWS Trusted Advisor checks.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 355

A company stores 200 GB of data each month in Amazon S3. The company needs to perform analytics on this data at the end of each month to determine the number of items sold in each sales region for the previous month.

Which analytics strategy is MOST cost-effective for the company to use?

- A. Create an Amazon Elasticsearch Service (Amazon ES) cluster. Query the data in Amazon ES. Visualize the data by using Kibana.



- B. Create a table in the AWS Glue Data Catalog. Query the data in Amazon S3 by using Amazon Athena. Visualize the data in Amazon QuickSight.
- C. Create an Amazon EMR cluster. Query the data by using Amazon EMR, and store the results in Amazon S3. Visualize the data in Amazon QuickSight.
- D. Create an Amazon Redshift cluster. Query the data in Amazon Redshift, and upload the results to Amazon S3. Visualize the data in Amazon QuickSight.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 356

A company wants to move its on-premises network attached storage (NAS) to AWS. The company wants to make the data available to any Linux instances within its VPC and ensure changes are automatically synchronized across all instances accessing the data store. The majority of the data is accessed very rarely, and some files are accessed by multiple users at the same time.

Which solution meets these requirements and is MOST cost-effective?

- A. Create an Amazon Elastic Block Store (Amazon EBS) snapshot containing the data. Share it with users within the VPC.
- B. Create an Amazon S3 bucket that has a lifecycle policy set to transition the data to S3 Standard- Infrequent Access (S3 Standard-IA) after the appropriate number of days.
- C. Create an Amazon Elastic File System (Amazon EFS) file system within the VPC. Set the throughput mode to Provisioned and to the required amount of IOPS to support concurrent usage.
- D. Create an Amazon Elastic File System (Amazon EFS) file system within the VPC. Set the lifecycle policy to transition the data to EFS Infrequent Access (EFS IA) after the appropriate number of days.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 357

A company plans to host a survey website on AWS. The company anticipates an unpredictable amount of traffic. This traffic results in asynchronous updates to the database. The company wants to ensure that writes to the database hosted on AWS do not get dropped.

How should the company write its application to handle these database requests?

- A. Configure the application to publish to an Amazon Simple Notification Service (Amazon SNS) topic. Subscribe the database to the SNS topic.
- B. Configure the application to subscribe to an Amazon Simple Notification Service (Amazon SNS) topic. Publish the database updates to the SNS topic.
- C. Use Amazon Simple Queue Service (Amazon SQS) FIFO queues to queue the database connection until the database has resources to write the data.
- D. Use Amazon Simple Queue Service (Amazon SQS) FIFO queues for capturing the writes and draining the queue as each write is made to the database.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 358

A company that recently started using AWS establishes a Site-to-Site VPN between its onpremises datacenter and AWS. The company's security mandate states that traffic originating from on premises should stay within the company's private IP space when communicating with an Amazon Elastic Container Service (Amazon ECS) cluster that is hosting a sample web application.

Which solution meets this requirement?

- A. Configure a gateway endpoint for Amazon ECS. Modify the route table to include an entry pointing to the ECS cluster.
- B. Create a Network Load Balancer and AWS PrivateLink endpoint for Amazon ECS in the same VPC that is hosting the ECS cluster.
- C. Create a Network Load Balancer in one VPC and an AWS PrivateLink endpoint for Amazon ECS in another VPC. Connect the two VPCs by using VPC peering.
- D. Configure an Amazon Route 53 record with Amazon ECS as the target. Apply a server certificate to Route 53 from AWS Certificate Manager (ACM) for SSL offloading.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 359

A solutions architect must analyze and update a company's existing IAM policies prior to deploying a new workload. The solutions architect created the following policy:

```
{
  "Version": "2012-10-17",
  "Statement": [{
    "Effect": "Deny",
    "NotAction": "s3:PutObject",
    "Resource": "*",
    "Condition": {"BoolIfExists": {"aws:MultiFactorAuthPresent": "false"}}
  }]
}
```

What is the net effect of this policy?

- A. Users will be allowed all actions except s3:PutObject if multi-factor authentication (MFA) is enabled.
- B. Users will be allowed all actions except s3:PutObject if multi-factor authentication (MFA) is not enabled.
- C. Users will be denied all actions except s3:PutObject if multi-factor authentication (MFA) is enabled.
- D. Users will be denied all actions except s3:PutObject if multi-factor authentication (MFA) is not enabled.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 360

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

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

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

Correct Answer: CD

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 361

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

What should the solutions architect recommend?

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

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 362

A company is using Amazon DynamoDB with provisioned throughput for the database tier of its ecommerce website. During flash sales, customers experience periods of time when the database cannot handle the high number of transactions taking place. This causes the company to lose transactions. During normal periods, the database performs appropriately. Which solution solves the performance problem the company faces?

- A. Switch DynamoDB to on-demand mode during flash sales.
- B. Implement DynamoDB Accelerator for fast in memory performance.
- C. Use Amazon Kinesis to queue transactions for processing to DynamoDB.
- D. Use Amazon Simple Queue Service (Amazon SQS) to queue transactions to DynamoDB.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 363

A company is reviewing a recent migration of a three-tier application to a VPC. The security team discovers that the principle of least privilege is not being applied to Amazon EC2 security group ingress and egress rules between the application tiers.

What should a solutions architect do to correct this issue?

- A. Create security group rules using the instance ID as the source or destination.
- B. Create security group rules using the security group ID as the source or destination.
- C. Create security group rules using the VPC CIDR blocks as the source or destination.
- D. Create security group rules using the subnet CIDR blocks as the source or destination.



Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 364

A company requires that all versions of objects in its Amazon S3 bucket be retained. Current object versions will be frequently accessed during the first 30 days, after which they will be rarely accessed and must be retrievable within 5 minutes. Previous object versions need to be kept forever, will be rarely accessed, and can be retrieved within 1 week. All storage solutions must be highly available and highly durable.

What should a solutions architect recommend to meet these requirements in the MOST costeffective manner?

- A. Create an S3 lifecycle policy for the bucket that moves current object versions from S3 Standard storage to S3 Glacier after 30 days and moves previous object versions to S3 Glacier after 1 day.
- B. Create an S3 lifecycle policy for the bucket that moves current object versions from S3 Standard storage to S3 Glacier after 30 days and moves previous object versions to S3 Glacier Deep Archive after 1 day.
- C. Create an S3 lifecycle policy for the bucket that moves current object versions from S3 Standard storage to S3 Standard-infrequent Access (S3 Standard-IA) after 30 days and moves previous object versions to S3 Glacier Deep Archive after 1 day.
- D. Create an S3 lifecycle policy for the bucket that moves current object versions from S3 Standard storage to S3 One Zone-Infrequent Access (S3 One Zone-IA) after 30 days and moves previous object versions to S3 Glacier Deep Archive after 1 day.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 365

A development team is collaborating with another company to create an integrated product. The other company needs to access an Amazon Simple Queue Service (Amazon SQS) queue that is contained in the development team's account. The other company wants to poll the queue without giving up its own account permissions to do so. How should a solutions architect provide access to the SQS queue?

- A. Create an instance profile that provides the other company access to the SQS queue.
- B. Create an IAM policy that provides the other company access to the SQS queue.
- C. Create an SQS access policy that provides the other company access to the SQS queue.
- D. Create an Amazon Simple Notification Service (Amazon SNS) access policy that provides the other company access to the SQS queue.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 366

A company is developing a video conversion application hosted on AWS. The application will be available in two tiers: a free tier and a paid tier. Users in the paid tier will have their videos converted first, and then the free tier users will have their videos converted.

Which solution meets these requirements and is MOST cost-effective?

- A. One FIFO queue for the paid tier and one standard queue for the free tier.
- B. A single FIFO Amazon Simple Queue Service (Amazon SQS) queue for all file types.
- C. A single standard Amazon Simple Queue Service (Amazon SQS) queue for all file types.
- D. Two standard Amazon Simple Queue Service (Amazon SQS) queues with one for the paid tier and one for the free tier.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 367

An administrator of a large company wants to monitor for and prevent any cryptocurrency-related attacks on the company's AWS accounts.

Which AWS service can the administrator use to protect the company against attacks?

- A. Amazon Cognito
- B. Amazon GuardDuty
- C. Amazon Inspector
- D. Amazon Macie

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 368

A company has applications hosted on Amazon EC2 instances with IPv6 addresses. The applications must initiate communications with other external applications using the internet.

However, the company's security policy states that any external service cannot initiate a connection to the EC2 instances. What should a solutions architect recommend to resolve this issue?

- A. Create a NAT gateway and make it the destination of the subnet's route table.
- B. Create an internet gateway and make it the destination of the subnet's route table.
- C. Create a virtual private gateway and make it the destination of the subnet's route table.
- D. Create an egress-only internet gateway and make it the destination of the subnet's route table.



Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 369

A company provides an online service for posting video content and transcoding it for use by any mobile platform. The application architecture uses Amazon Elastic File System (Amazon EFS) Standard to collect and store the videos so that multiple Amazon EC2 Linux instances can access the video content for processing. As the popularity of the service has grown over time, the storage costs have become too expensive. Which storage solution is MOST cost-effective?

- A. Use AWS Storage Gateway for files to store and process the video content.
- B. Use AWS Storage Gateway for volumes to store and process the video content.
- C. Use Amazon Elastic File System (Amazon EFS) for storing the video content. Once processing is complete, transfer the files to Amazon Elastic Block Store (Amazon EBS).
- D. Use Amazon S3 for storing the video content. Move the files temporarily over to an Amazon ElasticBlock Store (Amazon EBS) volume attached to the server for processing.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 370

A company wants to host its web application on AWS using multiple Amazon EC2 instances across different AWS Regions. Since the application content will be specific to each geographic region, the client requests need to be routed to the server that hosts the content for that clients Region. What should a solutions architect do to accomplish this?

- A. Configure Amazon Route 53 with a latency routing policy.
- B. Configure Amazon Route 53 with a weighted routing policy.
- C. Configure Amazon Route 53 with a geolocation routing policy.
- D. Configure Amazon Route 53 with a multivalue answer routing policy



Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 371

A solutions architect is planning the deployment of a new static website. The solution must minimize costs and provide at least 99% availability. Which solution meets these requirements?

- A. Deploy the application to an Amazon S3 bucket in one AWS Region that has versioning disabled.
- B. Deploy the application to Amazon EC2 instances that run in two AWS Regions and two Availability Zones.
- C. Deploy the application to an Amazon S3 bucket that has versioning and cross-Region replication enabled.
- D. Deploy the application to an Amazon EC2 instance that runs in one AWS Region and one Availability Zone.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 372

A recently created startup built a three-tier web application. The front end has static content. The application layer is based on microservices. User data is stored as JSON documents that need to be accessed with low latency. The company expects regular traffic to be low during the first year, with peaks in traffic when it publicizes new features every month. The startup team needs to minimize operational overhead costs. What should a solutions architect recommend to accomplish this?

- A. Use Amazon S3 static website hosting to store and serve the front end. Use AWS Elastic Beanstalk for the application layer. Use Amazon DynamoDB to store user data.
- B. Use Amazon S3 static website hosting to store and serve the front end. Use Amazon Elastic KubernetesService (Amazon EKS) for the application layer. Use Amazon DynamoDB to store user data.
- C. Use Amazon S3 static website hosting to store and serve the front end. Use Amazon API Gateway and AWS Lambda functions for the application layer. Use Amazon DynamoDB to store user data.
- D. Use Amazon S3 static website hosting to store and serve the front end. Use Amazon API Gateway and AWS Lambda functions for the application layer. Use Amazon RDS with read replicas to store user data.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 373

A company is building a payment application that must be highly available even during regional service disruptions. A solutions architect must design a data storage solution that can be easily replicated and used in other AWS Regions. The application also requires low-latency atomicity, consistency, isolation, and durability (ACID) transactions that need to be immediately available to generate reports. The development team also needs to use SQL. Which data storage solution meets these requirements?

- A. Amazon Aurora Global Database
- B. Amazon DynamoDB global tables
- C. Amazon S3 with cross-Region replication and Amazon Athena
- D. MySQL on Amazon EC2 instances with Amazon Elastic Block Store (Amazon EBS) snapshot replication

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 374

A company stores call recordings on a monthly basis. Statistically, the recorded data may be referenced randomly within a year but accessed rarely after 1 year. Files that are newer than 1 year old must be queried and retrieved as quickly as possible. A delay in retrieving older files is acceptable. A solutions architect needs to store the recorded data at a minimal cost. Which solution is MOST cost-effective?

- A. Store individual files in Amazon S3 Glacier and store search metadata in object tags created in S3 Glacier Query. S3 Glacier tags and retrieve the files from S3 Glacier.
- B. Store individual files in Amazon S3. Use lifecycle policies to move the files to Amazon S3 Glacier after 1 year. Query and retrieve the files from Amazon S3 or S3 Glacier.
- C. Archive individual files and store search metadata for each archive in Amazon S3. Use lifecycle policies to move the files to Amazon S3 Glacier after 1 year. Query and retrieve the files by searching for metadata from Amazon S3.
- D. Archive individual files in Amazon S3. Use lifecycle policies to move the files to Amazon S3 Glacier after 1 year. Store search metadata in Amazon DynamoDB. Query the files from DynamoDB and retrieve them from Amazon S3 or S3 Glacier.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 375

A company is developing a new machine learning model solution in AWS. The models are developed as independent microservices that fetch about 1 GB 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 solution meets these requirements?

- A. The requests from the API are sent to an Application Load Balancer (ALB). Models are deployed as AWS Lambda functions invoked by the ALB.
- B. The requests from the API are sent to the models Amazon Simple Queue Service (Amazon SQS) queue. Models are deployed as AWS Lambda functions triggered by SQS events. AWS Auto Scaling is enabled on Lambda to increase the number of vCPUs based on the SQS queue size.
- C. The requests from the API are sent to the model's Amazon Simple Queue Service (Amazon SQS) queue. Models are deployed as Amazon Elastic Container Service (Amazon ECS) services reading from the queue. AWS App Mesh

scales the instances of the ECS cluster based on the SQS queue size.

D. The requests from the API are sent to the models Amazon Simple Queue Service (Amazon SQS) queue. Models are deployed as Amazon Elastic Container Service (Amazon ECS) services reading from the queue AWS Auto Scaling is enabled on Amazon ECS for both the cluster and copies of the service based on the queue size.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 376

A company has no existing file share services. A new project requires access to file storage that is mountable as a drive for on-premises desktops. The file server must authenticate users to an Active Directory domain before they are able to access the storage.

Which service will allow Active Directory users to mount storage as a drive on their desktops?

- A. Amazon S3 Glacier
- B. AWS DataSync
- C. AWS Snowball Edge
- D. AWS Storage Gateway

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 377

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 largescale 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.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 378

A company has a custom application with embedded credentials that retrieves information from an Amazon RDS MySQL DB instance. Management says the application must be made more secure with the least amount of programming effort.

What should a solutions architect do to meet these requirements?

- A. Use AWS Key Management Service (AWS KMS) customer master keys (CMKs) to create keys. Configure the application to load the database credentials from AWS KMS. Enable automatic key rotation.
- B. Create credentials on the RDS for MySQL database for the application user and store the credentials in AWS Secrets Manager. Configure the application to load the database credentials from Secrets Manager. Create an AWS Lambda function that rotates the credentials in Secret Manager.
- C. Create credentials on the RDS for MySQL database for the application user and store the credentials in AWS Secrets Manager. Configure the application to load the database credentials from Secrets Manager. Set up a credentials rotation schedule for the application user in the RDS for MySQL database using Secrets Manager.
- D. Create credentials on the RDS for MySQL database for the application user and store the credentials in AWS Systems Manager Parameter Store. Configure the application to load the database credentials from Parameter Store. Set up a credentials rotation schedule for the application user in the RDS for MySQL database using Parameter Store.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 379

A company is running a multi-tier web application on AWS. The application runs its database tier on Amazon Aurora MySQL. The application and database tiers are in the us-east-1 Region. A database administrator who regularly monitors the Aurora DB cluster finds that an intermittent increase in read traffic is creating high CPU utilization on the read replica and causing increased read latency of the application. What should a solutions architect do to improve read scalability?

- A. Reboot the Aurora DB cluster.
- B. Create a cross-Region read replica
- C. Increase the instance class of the read replica.
- D. Configure Aurora Auto Scaling for the read replica.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 380

A company's order fulfillment service uses a MySQL database. The database needs to support a large number of concurrent queries and transactions. Developers are spending time patching and tuning the database. This is causing delays in releasing new product features.

The company wants to use cloud-based services to help address this new challenge. The solution must allow the developers to migrate the database with little or no code changes and must optimize performance. Which service should a solutions architect use to meet these requirements?

- A. Amazon Aurora
- B. Amazon DynamoDB
- C. Amazon ElastiCache
- D. MySQL on Amazon EC2



Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 381

A company is planning to transfer multiple terabytes of data to AWS. The data is collected offline from ships. The company wants to run complex transformation before transferring the data. Which AWS service should a solutions architect recommend for this migration?

- A. AWS Snowball
- B. AWS Snowmobile
- C. AWS Snowball Edge Storage Optimize
- D. AWS Snowball Edge Compute Optimize

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 382

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

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

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

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 383

A company is selling up an application to use an Amazon RDS MySQL DB instance. The database must be architected for high availability across Availability Zones and AWS Regions with minimal downtime. How should a solutions architect meet this requirement?

- A. Set up an RDS MySQL Multi-AZ DB instance. Configure an appropriate backup window.
- B. Set up an RDS MySQL Multi-AZ DB instance. Configure a read replica in a different Region.
- C. Set up an RDS MySQL Single-AZ DB instance. Configure a read replica in a different Region.
- D. Set up an RDS MySQL Single-AZ DB instance. Copy automated snapshots to at least one other Region.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:



QUESTION 384

A company hosts its web application on AWS using seven Amazon EC2 instances. The company requires that the IP addresses of all healthy EC2 instances be returned in response to DNS queries. Which policy should be used to meet this requirement?

- A. Simple routing policy
- B. Latency routing policy
- C. Multi-value routing policy
- D. Geolocation routing policy

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 385

A company has 700 TB of backup data stored in network attached storage (NAS) in its data center. This backup data needs to be accessible for infrequent regulatory requests and must be retained 7 years. The company has decided to migrate this backup data from its data center to AWS. The migration must be complete within 1 month. The company has 500 Mbps of dedicated bandwidth on its public internet connection available for data transfer. What should a solutions architect do to migrate and store the data at the LOWEST cost?

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

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 386

A company is preparing to deploy a data lake on AWS. A solutions architect must define the encryption strategy for data at rest in Amazon S3. The company's security policy states:

Keys must be rotated every 90 days.

Strict separation of duties between key users and key administrators must be implemented.

Auditing key usage must be possible.

What should the solutions architect recommend?

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

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 387

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

Which storage solution is MOST cost-effective?

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

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 388

A company previously migrated its data warehouse solution to AWS. The company also has an AWS Direct Connect connection. Corporate office users query the data warehouse using a visualization tool. The average size of a query returned by the data warehouse is 50 MB and each webpage sent by the visualization tool is approximately 500 KB. Result sets returned by the data warehouse are not cached.

Which solution provides the LOWEST data transfer egress cost for the company?

- A. Host the visualization tool on premises and query the data warehouse directly over the internet.
- B. Host the visualization tool in the same AWS Region as the data warehouse. Access it over the internet.
- C. Host the visualization tool on premises and query the data warehouse directly over a Direct Connect connection at a location in the same AWS Region.
- D. Host the visualization tool in the same AWS Region as the data warehouse and access it over a DirectConnect connection at a location in the same Region.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 389

A mobile gaming company runs application servers on Amazon EC2 instances. The servers receive updates from players every 15 minutes. The mobile game creates a JSON object of the progress made in the game since the last update,

and sends the JSON object to an Application Load Balancer. As the mobile game is played, game updates are being lost. The company wants to create a durable way to get the updates in order. What should a solutions architect recommend to decouple the system?

- A. Use Amazon Kinesis Data Streams to capture the data and store the JSON object in Amazon S3.
- B. Use Amazon Kinesis Data Firehose to capture the data and store the JSON object in Amazon S3.
- C. Use Amazon Simple Queue Service (Amazon SQS) FIFO queues to capture the data and EC2 instances to process the messages in the queue.
- D. Use Amazon Simple Notification Service (Amazon SNS) to capture the data and EC2 instances to process the messages sent to the Application Load Balancer.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 390

A company has an application that runs on Amazon EC2 instances within a private subnet in a VPC. The instances access data in an Amazon S3 bucket in the same AWS Region. The VPC contains a NAT gateway in a public subnet to access the S3 bucket. The company wants to reduce costs by replacing the NAT gateway without compromising security or redundancy.

Which solution meets these requirements?

- A. Replace the NAT gateway with a NAT instance.
- B. Replace the NAT gateway with an internet gateway.
- C. Replace the NAT gateway with a gateway VPC endpoint.
- D. Replace the NAT gateway with an AWS Direct Connect connection.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:



QUESTION 391

A company hosts a website on premises and wants to migrate it to the AWS Cloud. The website exposes a single hostname to the internet but it routes its functions to different on-premises server groups based on the path of the URL. The server groups are scaled independently depending on the needs of the functions they support. The company has an AWS Direct Connect connection configured to its on-premises network.

What should a solutions architect do to provide path-based routing to send the traffic to the correct group of servers?

- A. Route all traffic to an internet gateway. Configure pattern matching rules at the internet gateway to route traffic to the group of servers supporting that path.
- B. Route all traffic to a Network Load Balancer (NLB) with target groups for each group of servers. Use pattern matching rules at the NLB to route traffic to the correct target group.
- C. Route all traffic to an Application Load Balancer (ALB). Configure path-based routing at the ALB to route traffic to the correct target group for the servers supporting that path.
- D. Use Amazon Route 53 as the DNS server. Configure Route 53 path-based alias records to route traffic to the correct Elastic Load Balancer for the group of servers supporting that path.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 392

An application uses an Amazon RDS MySQL DB instance. The RDS database is becoming low on disk space. A solutions architect wants to increase the disk space without downtime. Which solution meets these requirements with the LEAST amount of effort?

- A. Enable storage auto scaling in RDS.
- B. Increase the RDS database instance size.
- C. Change the RDS database instance storage type to Provisioned IOPS.
- D. Back up the RDS database, increase the storage capacity, restore the database and stop the previous instance.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 393

An ecommerce website is deploying its web application as Amazon Elastic Container Service

(Amazon ECS) container instances behind an Application Load Balancer (ALB). During periods of high activity, the website slows down and availability is reduced. A solutions architect uses Amazon CloudWatch alarms to receive notifications whenever there is an availability issue so they can scale out resources. Company management wants a solution that automatically responds to such events.

Which solution meets these requirements?

- A. Set up AWS Auto Scaling to scale out the ECS service when there are timeouts on the ALB. Set up AWS Auto Scaling to scale out the ECS cluster when the CPU or memory reservation is too high.
- B. Set up AWS Auto Scaling to scale out the ECS service when the ALB CPU utilization is too high.
Setup AWS Auto Scaling to scale out the ECS cluster when the CPU or memory reservation is too high.
- C. Set up AWS Auto Scaling to scale out the ECS service when the service's CPU utilization is too high. Set up AWS Auto Scaling to scale out the ECS cluster when the CPU or memory reservation is too high.
- D. Set up AWS Auto Scaling to scale out the ECS service when the ALB target group CPU utilization is too high. Set up AWS Auto Scaling to scale out the ECS cluster when the CPU or memory reservation is too high.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 394

A company has a website deployed on AWS. The database backend is hosted on Amazon RDS for MySQL with a primary instance and five read replicas to support scaling needs. The read replicas should lag no more than 1 second behind the primary instance to support the user experience.

As traffic on the website continues to increase, the replicas are falling further behind during periods of peak load, resulting in complaints from users when searches yield inconsistent results.

A solutions architect needs to reduce the replication lag as much as possible, with minimal changes to the application code or operational requirements.

Which solution meets these requirements?

- A. Migrate the database to Amazon Aurora MySQL. Replace the MySQL read replicas with Aurora Replicas and enable Aurora Auto Scaling
- B. Deploy an Amazon ElastiCache for Redis cluster in front of the database. Modify the website to check the cache before querying the database read endpoints.
- C. Migrate the database from Amazon RDS to MySQL running on Amazon EC2 compute instances.
Choose very large compute optimized instances for all replica nodes.
- D. Migrate the database to Amazon DynamoDB. Initially provision a large number of read capacity units (RCUs) to support the required throughput with on-demand capacity scaling enabled.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 395

A company has an API-based inventory reporting application running on Amazon EC2 instances.

The application stores information in an Amazon DynamoDB table. The company's distribution centers have an on-premises shipping application that calls an API to update the inventory before printing shipping labels. The company has been experiencing application interruptions several times each day, resulting in lost transactions.

What should a solutions architect recommend to improve application resiliency?

- A. Modify the shipping application to write to a local database.
- B. Modify the application APIs to run serverless using AWS Lambda
- C. Configure Amazon API Gateway to call the EC2 inventory application APIs.
- D. Modify the application to send inventory updates using Amazon Simple Queue Service (Amazon SQS).

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 396

A company has a three-tier environment on AWS that ingests sensor data from its users' devices.

The traffic flows through a Network Load Balancer (NLB) then to Amazon EC2 instances for the web tier, and finally to EC2 instances for the application tier that makes database calls.

What should a solutions architect do to improve the security of data in transit to the web tier?

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

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 397

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. Other applications can consume the data stored in Amazon S3.
- C. 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.
- D. Store the batched transactions data in Amazon S3 as files. 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.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 398

A company has a dynamic web application hosted on two Amazon EC2 instances. The company has its own SSL certificate, which is on each instance to perform SSL termination.

There has been an increase in traffic recently, and the operations team determined that SSL encryption and decryption is causing the compute capacity of the web servers to reach their maximum limit.

What should a solutions architect do to increase the application's performance?

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

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 399

A web application must persist order data to Amazon S3 to support near-real time processing. A solutions architect needs create an architecture that is both scalable and fault tolerant. Which solutions meet these requirements? (Choose two.)

- A. Write the order event to an Amazon DynamoDB table. Use DynamoDB Streams to trigger an AWS Lambda function that parses the payload and writes the data to Amazon S3.
- B. Write the order event to an Amazon Simple Queue Service (Amazon SQS) queue. Use the queue to trigger an AWS Lambda function that parses the payload and writes the data to Amazon S3.
- C. Write the order event to an Amazon Simple Notification Service (Amazon SNS) topic. Use the SNS topic to trigger an AWS Lambda function that parses the payload and writes the data to Amazon S3.
- D. Write the order event to an Amazon Simple Queue Service (Amazon SQS) queue. Use an Amazon EventBridge (Amazon CloudWatch Events) rule to trigger an AWS Lambda function that parses the payload and writes the data to Amazon S3.
- E. Write the order event to an Amazon Simple Notification Service (Amazon SNS) topic. Use an Amazon EventBridge (Amazon CloudWatch Events) rule to trigger an AWS Lambda function that parses the payload and writes the data to Amazon S3.

Correct Answer: BE

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 400

A company has an application hosted on Amazon EC2 instances in two VPCs across different AWS Regions. To communicate with each other, the instances use the internet for connectivity. The security team wants to ensure that no communication between the instances happens over the internet. What should a solutions architect do to accomplish this?

- A. Create a NAT gateway and update the route table of the EC2 instances' subnet.
- B. Create a VPC endpoint and update the route table of the EC2 instances' subnet.
- C. Create a VPN connection and update the route table of the EC2 instances' subnet.
- D. Create a VPC peering connection and update the route table of the EC2 instances' subnet.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 401

An online shopping application accesses an Amazon RDS Multi-AZ DB instance. Database performance is slowing down the application. After upgrading to the next-generation instance type, there was no significant performance improvement.

Analysis shows approximately 700 IOPS are sustained, common queries run for long durations and memory utilization is high.

Which application change should a solutions architect recommend to resolve these issues?

- A. Migrate the RDS instance to an Amazon Redshift cluster and enable weekly garbage collection.
- B. Separate the long-running queries into a new Multi-AZ RDS database and modify the application to query whichever database is needed.
- C. Deploy a two-node Amazon ElastiCache cluster and modify the application to query the cluster first and query the database only if needed.
- D. Create an Amazon Simple Queue Service (Amazon SQS) FIFO queue for common queries and query it first and query the database only if needed.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 402

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

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 403

A company is preparing to migrate its on-premises application to AWS. The application consists of application servers and a Microsoft SQL Server database. The database cannot be migrated to a different engine because SQL Server features are used in the application's .NET code. The company wants to attain the greatest availability possible while minimizing operational and management overhead.

What should a solutions architect do to accomplish this?

- A. Install SQL Server on Amazon EC2 in a Multi-AZ deployment.
- B. Migrate the data to Amazon RDS for SQL Server in a Multi-AZ deployment.
- C. Deploy the database on Amazon RDS for SQL Server with Multi-AZ Replicas.
- D. Migrate the data to Amazon RDS for SQL Server in a cross-Region Multi-AZ deployment.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 404

A company has an application running on Amazon EC2 instances in a private subnet. The application needs to store and retrieve data in Amazon S3. To reduce costs, the company wants to configure its AWS resources in a cost-effective manner.

How should the company accomplish this?

- A. Deploy a NAT gateway to access the S3 buckets.
- B. Deploy AWS Storage Gateway to access the S3 buckets.
- C. Deploy an S3 gateway endpoint to access the S3 buckets.
- D. Deploy an S3 interface endpoint to access the S3 buckets.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 405

A media company has an application that tracks user clicks on its websites and performs analytics to provide near-real time recommendations. The application has a fleet of Amazon EC2 instances that receive data from the websites and send the data to an Amazon RDS DB instance. Another fleet of EC2 instances hosts the portion of the application that is continuously checking changes in the database and executing SQL queries to provide recommendations. Management has requested a redesign to decouple the infrastructure. The solution must ensure that data analysts are writing SQL to analyze the data only. No data can be lost during the deployment.

What should a solutions architect recommend?

- A. Use Amazon Kinesis Data Streams to capture the data from the websites, Kinesis Data Firehose to persist the data on Amazon S3, and Amazon Athena to query the data.
- B. Use Amazon Kinesis Data Streams to capture the data from the websites, Kinesis Data Analytics to query the data, and Kinesis Data Firehose to persist the data on Amazon S3.
- C. Use Amazon Simple Queue Service (Amazon SQS) to capture the data from the websites, keep the fleet of EC2 instances, and change to a bigger instance type in the Auto Scaling group configuration.
- D. Use Amazon Simple Notification Service (Amazon SNS) to receive data from the websites and proxy the messages to AWS Lambda functions that execute the queries and persist the data.
Change Amazon RDS to Amazon Aurora Serverless to persist the data.

Correct Answer: B

Section: (none)

Explanation



Explanation/Reference:

Explanation:

QUESTION 406

A company runs an application that uses multiple Amazon EC2 instances to gather data from its users. The data is then processed and transferred to Amazon S3 for long-term storage. A review of the application shows that there were long periods of time when the EC2 instances were not being used. A solutions architect needs to design a solution that optimizes utilization and reduces costs.

Which solution meets these requirements?

- A. Use Amazon EC2 in an Auto Scaling group with On-Demand instances.
- B. Build the application to use Amazon Lightsail with On-Demand Instances.
- C. Create an Amazon CloudWatch cron job to automatically stop the EC2 instances when there is no activity.
- D. Redesign the application to use an event-driven design with Amazon Simple Queue Service (Amazon SQS) and AWS Lambda.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 407

A company is using Site-to-Site VPN connections for secure connectivity to its AWS Cloud resources from on premises. Due to an increase in traffic across the VPN connections to the Amazon EC2 instances, users are experiencing slower VPN connectivity.

Which solution will improve the VPN throughput?

- A. Implement multiple customer gateways for the same network to scale the throughput.
- B. Use a transit gateway with equal cost multipath routing and add additional VPN tunnels.
- C. Configure a virtual private gateway with equal cost multipath routing and multiple channels.
- D. Increase the number of tunnels in the VPN configuration to scale the throughput beyond the default limit.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 408

A company has a mobile game that reads most of its metadata from an Amazon RDS DB instance. As the game increased in popularity developers noticed slowdowns related to the game's metadata load times. Performance metrics indicate that simply scaling the database will not help. A solutions architect must explore all options that include capabilities for snapshots replication and sub-millisecond response times.

What should the solutions architect recommend to solve these issues?

- A. Migrate the database to Amazon Aurora with Aurora Replicas.
- B. Migrate the database to Amazon DynamoDB with global tables.
- C. Add an Amazon ElastiCache for Redis layer in front of the database.
- D. Add an Amazon ElastiCache for Memcached layer in front of the database.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 409

A company has several Amazon EC2 instances set up in a private subnet for security reasons.

These instances host applications that read and write large amounts of data to and from Amazon S3 regularly. Currently, subnet routing directs all the traffic destined for the internet through a NAT gateway. The company wants to optimize the overall cost without impacting the ability of the application to communicate with Amazon S3 or the outside internet.

What should a solutions architect do to optimize costs?

- A. Create an additional NAT gateway. Update the route table to route to the NAT gateway. Update the network ACL to allow S3 traffic.
- B. Create an internet gateway. Update the route table to route traffic to the internet gateway. Update the network ACL to allow S3 traffic.
- C. Create a VPC endpoint for Amazon S3. Attach an endpoint policy to the endpoint. Update the route table to direct traffic to the VPC endpoint.
- D. Create an AWS Lambda function outside of the VPC to handle S3 requests. Attach an IAM policy to the EC2 instances, allowing them to invoke the Lambda function.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 410

A company is deploying an application in three AWS Regions using an Application Load Balancer. Amazon Route 53 will be used to distribute traffic between these Regions. Which Route 53 configuration should a solutions architect use to provide the MOST highperforming experience?

- A. Create an A record with a latency policy.
- B. Create an A record with a geolocation policy.
- C. Create a CNAME record with a failover policy.
- D. Create a CNAME record with a geoproximity policy.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:



QUESTION 411

A company has an application workflow that uses an AWS Lambda function to download and decrypt files from Amazon S3. These files are encrypted using AWS Key Management Service Customer Master Keys (AWS KMS CMKs). A solutions architect needs to design a solution that will ensure the required permissions are set correctly. Which combination of actions accomplish this? (Choose two.)

- A. Attach the kms:decrypt permission to the Lambda function's resource policy.
- B. Grant the decrypt permission for the Lambda IAM role in the KMS key's policy.
- C. Grant the decrypt permission for the Lambda resource policy in the KMS key's policy.
- D. Create a new IAM policy with the kms:decrypt permission and attach the policy to the Lambda function.
- E. Create a new IAM role with the kms:decrypt permission and attach the execution role to the Lambda function.

Correct Answer: BE

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 412

A company is migrating a Linux-based web server group to AWS. The web servers must access files in a shared file store for some content. To meet the migration date, minimal changes can be made. What should a solutions architect do to meet these requirements?

- A. Create an Amazon S3 Standard bucket with access to the web server.
- B. Configure an Amazon CloudFront distribution with an Amazon S3 bucket as the origin.
- C. Create an Amazon Elastic File System (Amazon EFS) volume and mount it on all web servers.
- D. Configure Amazon Elastic Block Store (Amazon EBS) Provisioned IOPS SSD (io1) volumes and mount them on all web servers.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 413

A company that operates a web application on premises is preparing to launch a newer version of the application on AWS. The company needs to route requests to either the AWS-hosted or the on-premises-hosted application based on the URL query string. The on-premises application is not available from the internet, and a VPN connection is established between Amazon VPC and the company's data center. The company wants to use an Application Load Balancer (ALB) for this launch.

Which solution meets these requirements?

- A. Use two ALBs: one for on-premises and one for the AWS resource. Add hosts to each target group of each ALB. Route with Amazon Route 53 based on the URL query string.
- B. Use two ALBs: one for on-premises and one for the AWS resource. Add hosts to the target group of each ALB. Create a software router on an EC2 instance based on the URL query string.
- C. Use one ALB with two target groups: one for the AWS resource and one for on premises. Add hosts to each target group of the ALB. Configure listener rules based on the URL query string.
- D. Use one ALB with two AWS Auto Scaling groups: one for the AWS resource and one for on premises. Add hosts to each Auto Scaling group. Route with Amazon Route 53 based on the URL query string.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 414

A solutions architect is developing a multiple-subnet VPC architecture. The solution will consist of six subnets in two Availability Zones. The subnets are defined as public, private and dedicated for databases. Only the Amazon EC2 instances running in the private subnets should be able to access a database.

Which solution meets these requirements?

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

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 415

A disaster response team is using drones to collect images of recent storm damage. The response team's laptops lack the storage and compute capacity to transfer the images and process the data. While the team has Amazon EC2 instances for processing and Amazon S3 buckets for storage, network connectivity is intermittent and unreliable. The images need to be processed to evaluate the damage.

What should a solutions architect recommend?

- A. Use AWS Snowball Edge devices to process and store the images.
- B. Upload the images to Amazon Simple Queue Service (Amazon SQS) during intermittent connectivity to EC2 instances.
- C. Configure Amazon Kinesis Data Firehose to create multiple delivery streams aimed separately at the S3 buckets for storage and the EC2 instances for processing the images.
- D. Use AWS Storage Gateway pre-installed on a hardware appliance to cache the images locally for Amazon S3 to process the images when connectivity becomes available.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 416

A company has a multi-tier application deployed on several Amazon EC2 instances in an Auto Scaling group. An Amazon RDS for Oracle instance is the application's data layer that uses Oracle-specific PL/SQL functions. Traffic to the application has been steadily increasing. This is causing the EC2 instances to become overloaded and the RDS instance to run out of storage. The Auto Scaling group does not have any scaling metrics and defines the minimum healthy

instance count only. The company predicts that traffic will continue to increase at a steady but unpredictable rate before leveling off. What should a solutions architect do to ensure the system can automatically scale for the increased traffic? (Choose two.)

- A. Configure storage Auto Scaling on the RDS for Oracle instance.
- B. Migrate the database to Amazon Aurora to use Auto Scaling storage.
- C. Configure an alarm on the RDS for Oracle instance for low free storage space.
- D. Configure the Auto Scaling group to use the average CPU as the scaling metric.
- E. Configure the Auto Scaling group to use the average free memory as the scaling metric.

Correct Answer: AC

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 417

An engineering team is developing and deploying AWS Lambda functions. The team needs to create roles and manage policies in AWS IAM to configure the permissions of the Lambda functions. How should the permissions for the team be configured so they also adhere to the concept of least privilege?

- A. Create an IAM role with a managed policy attached. Allow the engineering team and the Lambda functions to assume this role.
- B. Create an IAM group for the engineering team with an IAMFullAccess policy attached. Add all the users from the team to this IAM group.
- C. Create an execution role for the Lambda functions. Attach a managed policy that has permission boundaries specific to these Lambda functions.
- D. Create an IAM role with a managed policy attached that has permission boundaries specific to the Lambda functions. Allow the engineering team to assume this role.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:



QUESTION 418

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 over 10 million rows. The database has 2 TB of General Purpose SSD (gp2) storage. There are millions of updates against this data every day through the company's website. The company has noticed some operations are taking 10 seconds or longer and has determined that the database storage performance is the bottleneck.

Which solution addresses the performance issue?

- A. Change the storage type to Provisioned IOPS SSD (io1).
- B. Change the instance to a memory-optimized instance class.
- C. Change the instance to a burstable performance DB instance class.
- D. Enable Multi-AZ RDS read replicas with MySQL native asynchronous replication.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 419

A company has an Amazon S3 bucket that contains mission-critical data. The company wants to ensure this data is protected from accidental deletion. The data should still be accessible, and a user should be able to delete the data intentionally.

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

Correct Answer: AB

Section: (none)

Explanation

Explanation/Reference:

Reference: <https://acloud.guru/forums/aws-certified-solutions-architect-associate/discussion/-LNMknp7BP01IYVSlee8/Final%20practice%20exam>

QUESTION 420

A company has an on-premises business application that generates hundreds of files each day.

These files are stored on an SMB file share and require a low-latency connection to the application servers. A new company policy states all application-generated files must be copied to AWS.

There is already a VPN connection to AWS.

The application development team does not have time to make the necessary code modifications to move the application to AWS.

Which service should a solutions architect recommend to allow the application to copy files to AWS?

- A. Amazon Elastic File System (Amazon EFS)
- B. Amazon FSx for Windows File Server
- C. AWS Snowball
- D. AWS Storage Gateway

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Reference: <https://aws.amazon.com/blogs/storage/accessing-smb-file-shares-remotely-withamazon-fsx-for-windows-file-server/>

QUESTION 421

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? (Choose two.)

- A. Configure a VPC gateway endpoint for 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.

Correct Answer: AC

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 422

A solutions architect plans to convert a company's monolithic web application into a multi-tier application. The company wants to avoid managing its own infrastructure. The minimum requirements for the web application are high availability, scalability, and regional low latency during peak hours. The solution should also store and retrieve data with millisecond latency using the application's API.

Which solution meets these requirements?

- A. Use AWS Fargate to host the web application with backend Amazon RDS Multi-AZ DB instances.
- B. Use Amazon API Gateway with an edge-optimized API endpoint, AWS Lambda for compute, and Amazon DynamoDB as the data store.
- C. Use an Amazon Route 53 routing policy with geolocation that points to an Amazon S3 bucket with static website hosting and Amazon DynamoDB as the data store.
- D. Use an Amazon CloudFront distribution that points to an Elastic Load Balancer with an Amazon EC2 Auto Scaling group, along with Amazon RDS Multi-AZ DB instances.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 423

A team has an application that detects new objects being uploaded into an Amazon S3 bucket. The uploads trigger AWS Lambda function to write object metadata into an Amazon DynamoDB table and an Amazon RDS for PostgreSQL database. Which action should the team take to ensure high availability?

- A. Enable Cross-Region Replication in the S3 bucket.
- B. Create a Lambda function for each Availability Zone the application is deployed in.
- C. Enable Multi-AZ on the RDS for PostgreSQL database.
- D. Create a DynamoDB stream for the DynamoDB table.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 424

A company is planning to migrate a legacy application to AWS. The application currently uses NFS to communicate to an on-premises storage solution to store application data. The application cannot be modified to use any other communication protocols other than NFS for this purpose. Which storage solution should a solutions architect recommend for use after the migration?

- A. AWS DataSync
- B. Amazon Elastic Block Store (Amazon EBS)
- C. Amazon Elastic File System (Amazon EFS)
- D. Amazon EMR File System (Amazon EMRFS)

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 425

An application calls a service run by a vendor. The vendor charges based on the number of calls. The finance department needs to know the number of calls that are made to the service to validate the billing statements. How can a solutions architect design a system to durably store the number of calls without requiring changes to the application?

- A. Call the service through an internet gateway.
- B. Decouple the application from the service with an Amazon Simple Queue Service (Amazon SQS) queue.
- C. Publish a custom Amazon CloudWatch metric that counts calls to the service.
- D. Call the service through a VPC peering connection.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

There are 2 main types of monitoring you can do on AWS EC2 Instances as follows:

Basic Monitoring for Amazon EC2 instances: Seven pre-selected metrics at five-minute frequency and three status check metrics at one-minute frequency, for no additional charge.

Detailed Monitoring for Amazon EC2 instances: All metrics available to Basic Monitoring at oneminute frequency, for an additional charge. Instances with Detailed Monitoring enabled allows data aggregation by Amazon EC2 AMI ID and instance type.

Reference: <https://datanextsolutions.com/blog/how-to-collect-custom-metrics-from-aws-ec2-instances/>

QUESTION 426

A company wants to reduce its Amazon S3 storage costs in its production environment without impacting durability or performance of the stored objects. What is the FIRST step the company should take to meet these objectives?

- A. Enable Amazon Macie on the business-critical S3 buckets to classify the sensitivity of the objects.
- B. Enable S3 analytics to identify S3 buckets that are candidates for transitioning to S3 Standard- Infrequent Access (S3 Standard-IA).
- C. Enable versioning on all business-critical S3 buckets.
- D. Migrate the objects in all S3 buckets to S3 Intelligent-Tiering.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 427

A company is building a web-based application running on Amazon EC2 instances in multiple Availability Zones. The web application will provide access to a repository of text documents totaling about 900 TB in size. The company anticipates that the web application will experience periods of high demand. A solutions architect must ensure that the storage component for the text documents can scale to meet the demand of the application at all times. The company is concerned about the overall cost of the solution.

Which storage solution meets these requirements MOST cost-effectively?

- A. Amazon Elastic Block Store (Amazon EBS)
- B. Amazon Elastic File System (Amazon EFS)
- C. Amazon Elasticsearch Service (Amazon ES)
- D. Amazon S3

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Reference: <https://www.missioncloud.com/blog/resource-amazon-ebs-vs-efs-vs-s3-picking-thebest-aws-storage-option-for-your-business>

QUESTION 428

A company hosts multiple production applications. One of the applications consists of resources from Amazon EC2, AWS Lambda, Amazon RDS, Amazon Simple Notification Service (Amazon SNS), and Amazon Simple Queue Service (Amazon SQS) across multiple AWS Regions. All company resources are tagged with a tag name of "application" and a value that corresponds to each application. A solutions architect must provide the quickest solution for identifying all of the tagged components.

Which solution meets these requirements?

- A. Use AWS CloudTrail to generate a list of resources with the application tag.
- B. Use the AWS CLI to query each service across all Regions to report the tagged components.
- C. Run a query in Amazon CloudWatch Logs Insights to report on the components with the application tag.
- D. Run a query with the AWS Resource Groups Tag Editor to report on the resources globally with the application tag.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 429

A development team is deploying a new product on AWS and is using AWS Lambda as part of the deployment. The team allocates 512 MB of memory for one of the Lambda functions. With this memory allocation, the function is completed in 2 minutes. The function runs millions of times monthly, and the development team is concerned about cost. The team conducts tests to see how different Lambda memory allocations affect the cost of the function.

Which steps will reduce the Lambda costs for the product? (Choose two.)

- A. Increase the memory allocation for this Lambda function to 1,024 MB if this change causes the execution time of each function to be less than 1 minute.
- B. Increase the memory allocation for this Lambda function to 1,024 MB if this change causes the execution time of each function to be less than 90 seconds.
- C. Reduce the memory allocation for this Lambda function to 256 MB if this change causes the execution time of each function to be less than 4 minutes.



- D. Increase the memory allocation for this Lambda function to 2,048 MB if this change causes the execution time of each function to be less than 1 minute.
- E. Reduce the memory allocation for this Lambda function to 256 MB if this change causes the execution time of each function to be less than 5 minutes.

Correct Answer: AE

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 430

A company recently launched Linux-based application instances on Amazon EC2 in a private subnet and launched a Linux-based bastion host on an Amazon EC2 instance in a public subnet of a VPC. A solutions architect needs to connect from the on-premises network, through the company's internet connection, to the bastion host, and to the application servers. The solutions architect must make sure that the security groups of all the EC2 instances will allow that access. Which combination of steps should the solutions architect take to meet these requirements? (Choose two.)

- A. Replace the current security group of the bastion host with one that only allows inbound access from the application instances.
- B. Replace the current security group of the bastion host with one that only allows inbound access from the internal IP range for the company.
- C. Replace the current security group of the bastion host with one that only allows inbound access from the external IP range for the company.
- D. Replace the current security group of the application instances with one that allows inbound SSH access from only the private IP address of the bastion host.
- E. Replace the current security group of the application instances with one that allows inbound SSH access from only the public IP address of the bastion host.

Correct Answer: AC

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 431

A user owns a MySQL database that is accessed by various clients who expect, at most, 100 ms latency on requests. Once a record is stored in the database, it is rarely changed. Clients only access one record at a time. Database access has been increasing exponentially due to increased client demand. The resultant load will soon exceed the capacity of the most expensive hardware available for purchase. The user wants to migrate to AWS, and is willing to change database systems. Which service would alleviate the database load issue and offer virtually unlimited scalability for the future?

- A. Amazon RDS
- B. Amazon DynamoDB
- C. Amazon Redshift
- D. AWS Data Pipeline

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Reference: <https://aws.amazon.com/blogs/big-data/near-zero-downtime-migration-from-mysql-to-dynamodb/>

QUESTION 432

A company designs a mobile app for its customers to upload photos to a website. The app needs a secure login with multi-factor authentication (MFA). The company wants to limit the initial build time and the maintenance of the solution. Which solution should a solutions architect recommend to meet these requirements?

- A. Use Amazon Cognito Identity with SMS-based MFA.
- B. Edit IAM policies to require MFA for all users.
- C. Federate IAM against the corporate Active Directory that requires MFA.
- D. Use Amazon API Gateway and require server-side encryption (SSE) for photos.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Reference: <https://aws.amazon.com/cognito/>

QUESTION 433

A company has an application that uses overnight digital images of products on store shelves to analyze inventory data. The application runs on Amazon EC2 instances behind an Application Load Balancer (ALB) and obtains the images from an Amazon S3 bucket for its metadata to be processed by worker nodes for analysis. A solutions architect needs to ensure that every image is processed by the worker nodes.

What should the solutions architect do to meet this requirement in the MOST cost-efficient way?

- A. Send the image metadata from the application directly to a second ALB for the worker nodes that use an Auto Scaling group of EC2 Spot Instances as the target group.
- B. Process the image metadata by sending it directly to EC2 Reserved Instances in an Auto Scaling group. With a dynamic scaling policy, use an Amazon CloudWatch metric for average CPU utilization of the Auto Scaling group as soon as the front-end application obtains the images.
- C. Write messages to Amazon Simple Queue Service (Amazon SQS) when the front-end application obtains an image. Process the images with EC2 On-Demand instances in an Auto Scaling group with instance scale-in protection and a fixed number of instances with periodic health checks.
- D. Write messages to Amazon Simple Queue Service (Amazon SQS) when the application obtains an image. Process the images with EC2 Spot Instances in an Auto Scaling group with instance scale-in protection and a dynamic scaling policy using a custom Amazon CloudWatch metric for the current number of messages in the queue.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 434

A solutions architect needs to host a high performance computing (HPC) workload in the AWS Cloud. The workload will run on hundreds of Amazon EC2 instances and will require parallel access to a shared file system to enable distributed processing of large datasets. Datasets will be accessed across multiple instances simultaneously. The workload requires access latency within 1 ms. After processing has completed, engineers will need access to the dataset for manual postprocessing.

Which solution will meet these requirements?

- A. Use Amazon Elastic File System (Amazon EFS) as a shared file system. Access the dataset from Amazon EFS.
- B. Mount an Amazon S3 bucket to serve as the shared file system. Perform postprocessing directly from the S3 bucket.
- C. Use Amazon FSx for Lustre as a shared file system. Link the file system to an Amazon S3 bucket for postprocessing.
- D. Configure AWS Resource Access Manager to share an Amazon S3 bucket so that it can be mounted to all instances for processing and postprocessing.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Reference: <https://jayendrapatil.com/aws-fsx-for-lustre/>

QUESTION 435

A company is using Amazon Route 53 latency-based routing to route requests to its UDP-based application for users around the world. The application is hosted on redundant servers in the company's on-premises data centers in the United States, Asia, and Europe. The company's compliance requirements state that the application must be hosted on premises. The company wants to improve the performance and availability of the application.

What should a solutions architect do to meet these requirements?

- A. Configure three Network Load Balancers (NLBs) in the three AWS Regions to address the onpremises endpoints. Create an accelerator by using AWS Global Accelerator, and register the NLBs as its endpoints. Provide access to the application by using a CNAME that points to the accelerator DNS.
- B. Configure three Application Load Balancers (ALBs) in the three AWS Regions to address the onpremises endpoints. Create an accelerator by using AWS Global Accelerator, and register the ALBs as its endpoints. Provide access to the application by using a CNAME that points to the accelerator DNS.
- C. Configure three Network Load Balancers (NLBs) in the three AWS Regions to address the onpremises endpoints. In Route 53, create a latency-based record that points to the three NLBs, and use it as an origin for an Amazon CloudFront distribution. Provide access to the application by using a CNAME that points to the CloudFront DNS.
- D. Configure three Application Load Balancers (ALBs) in the three AWS Regions to address the onpremises endpoints. In Route 53, create a latency-based record that points to the three ALBs, and use it as an origin for an Amazon CloudFront distribution. Provide access to the application by using a CNAME that points to the CloudFront DNS.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 436

A company manages its own Amazon EC2 instances that run MySQL databases. The company is manually managing replication and scaling as demand increases or decreases. The company needs a new solution that simplifies the process of adding or removing compute capacity to or from its database tier as needed. The solution also must offer improved performance, scaling, and durability with minimal effort from operations. Which solution meets these requirements?

- A. Migrate the databases to Amazon Aurora Serverless for Aurora MySQL.
- B. Migrate the databases to Amazon Aurora Serverless for Aurora PostgreSQL.
- C. Combine the databases into one larger MySQL database. Run the larger database on larger EC2 instances.
- D. Create an EC2 Auto Scaling group for the database tier. Migrate the existing databases to the new environment.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Reference: <https://aws-quickstart.s3.amazonaws.com/quickstart-drupal/doc/drupal-on-the-awscloud.pdf> (p.6)

QUESTION 437

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.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 438

A company has two AWS accounts: Production and Development. There are code changes ready in the Development account to push to the Production account. In the alpha phase, only two senior developers on the development team need access to the Production account. In the beta phase, more developers might need access to perform testing as well.

What should a solutions architect recommend?

- A. Create two policy documents using the AWS Management Console in each account. Assign the policy to developers who need access.
- B. Create an IAM role in the Development account. Give one IAM role access to the Production account. Allow developers to assume the role.
- C. Create an IAM role in the Production account with the trust policy that specifies the Development account. Allow developers to assume the role.
- D. Create an IAM group in the Production account and add it as a principal in the trust policy that specifies the Production account. Add developers to the group.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 439

A company is using an Amazon S3 bucket to store data uploaded by different departments from multiple locations. During an AWS Well-Architected review, the financial manager notices that 10 TB of S3 Standard storage data has been charged each month. However, in the AWS Management Console for Amazon S3, using the command to select all files and folders shows a total size of 5 TB. What are the possible causes for this difference? (Choose two.)

- A. Some files are stored with deduplication.
- B. The S3 bucket has versioning enabled.
- C. There are incomplete S3 multipart uploads.
- D. The S3 bucket has AWS Key Management Service (AWS KMS) enabled.
- E. The S3 bucket has Intelligent-Tiering enabled.

Correct Answer: AB

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 440

A company is using a centralized AWS account to store log data in various Amazon S3 buckets. A solutions architect needs to ensure that the data is encrypted at rest before the data is uploaded to the S3 buckets. The data also must be encrypted in transit. Which solution meets these requirements?

- A. Use client-side encryption to encrypt the data that is being uploaded to the S3 buckets.
- B. Use server-side encryption to encrypt the data that is being uploaded to the S3 buckets.
- C. Create bucket policies that require the use of server-side encryption with S3 managed encryption keys (SSE-S3) for S3 uploads.
- D. Enable the security option to encrypt the S3 buckets through the use of a default AWS Key Management Service (AWS KMS) key.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

**QUESTION 441**

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

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

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 442

A company has two VPCs that are located in the us-west-2 Region within the same AWS account. The company needs to allow network traffic between these VPCs. Approximately 500 GB of data transfer will occur between the VPCs each month. What is the MOST cost-effective solution to connect these VPCs?

- A. Implement AWS Transit Gateway to connect the VPCs. Update the route tables of each VPC to use the transit gateway for inter-VPC communication.
- B. Implement an AWS Site-to-Site VPN tunnel between the VPCs. Update the route tables of each VPC to use the VPN tunnel for inter-VPC communication.
- C. Set up a VPC peering connection between the VPCs. Update the route tables of each VPC to use the VPC peering connection for inter-VPC communication.
- D. Set up a 1 GB AWS Direct Connect connection between the VPCs. Update the route tables of each VPC to use the Direct Connect connection for inter-VPC communication.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 443

A company is deploying an application that processes streaming data in near-real time. The company plans to use Amazon EC2 instances for the workload. The network architecture must be configurable to provide the lowest possible latency between nodes.

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

- A. Enable and configure enhanced networking on each EC2 instance.
- B. Group the EC2 instances in separate accounts.
- C. Run the EC2 instances in a cluster placement group.
- D. Attach multiple elastic network interfaces to each EC2 instance.
- E. Use Amazon Elastic Block Store (Amazon EBS) optimized instance types.

Correct Answer: CD

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 444

A company is running a global application. The application's users submit multiple videos that are then merged into a single video file. The application uses a single Amazon S3 bucket in the useast-1 Region to receive uploads from users. The same S3 bucket provides the download location of the single video file that is produced. The final video file output has an average size of 250 GB.

The company needs to develop a solution that delivers faster uploads and downloads of the video files that are stored in Amazon S2. The company will offer the solution as a subscription to users who want to pay for the increased speed. What should a solutions architect do to meet these requirements?

- A. Enable AWS Global Accelerator for the S3 endpoint. Adjust the application's upload and download links to use the Global Accelerator S3 endpoint for users who have a subscription.
- B. Enable S3 Cross-Region Replication to S3 buckets in all other AWS Regions. Use an Amazon Route 53 geolocation routing policy to route S3 requests based on the location of users who have a subscription.
- C. Create an Amazon CloudFront distribution and use the S3 bucket in us-east-1 as an origin. Adjust the application to use the CloudFront URL as the upload and download links for users who have a subscription.
- D. Enable S3 Transfer Acceleration for the S3 bucket in us-east-1. Configure the application to use the bucket's S3-accelerate endpoint domain name for the upload and download links for users who have a subscription.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 445

The following IAM policy is attached to an IAM group. This is the only policy applied to the group.



```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "1",
      "Effect": "Allow",
      "Action": "ec2:*",
      "Resource": "*",
      "Condition": {
        "StringEquals": {
          "ec2:Region": "us-east-1"
        }
      }
    },
    {
      "Sid": "2",
      "Effect": "Deny",
      "Action": [
        "ec2:StopInstances",
        "ec2:TerminateInstances"
      ],
      "Resource": "*",
      "Condition": {
        "BoolIfExists": {"aws:MultiFactorAuthPresent": false}
      }
    }
  ]
}

```

What are the effective IAM permissions of this policy for group members?

- A. Group members are permitted any Amazon EC2 action within the us-east-1 Region. Statements after the Allow permission are not applied.
- B. Group members are denied any Amazon EC2 permissions in the us-east-1 Region unless they are logged in with multi-factor authentication (MFA).
- C. Group members are allowed the ec2:StopInstances and ec2:TerminateInstances permissions for all Regions when logged in with multi-factor authentication (MFA). Group members are permitted any other Amazon EC2 action.
- D. Group members are allowed the ec2:StopInstances and ec2:TerminateInstances permissions for the us-east-1 Region only when logged in with multi-factor authentication (MFA). Group members are permitted any other Amazon EC2 action within the us-east-1 Region.

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Explanation:

QUESTION 446

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.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:

QUESTION 447

A company is using AWS Organizations with two AWS accounts: Logistics and Sales. The Logistics account operates an Amazon Redshift cluster. The Sales account includes Amazon EC2 instances. The Sales account needs to access the Logistics account's Amazon Redshift cluster.

What should a solutions architect recommend to meet this requirement MOST cost-effectively?

- A. Set up VPC sharing with the Logistics account as the owner and the Sales account as the participant to transfer the data.
- B. Create an AWS Lambda function in the Logistics account to transfer data to the Amazon EC2 instances in the Sales account.
- C. Create a snapshot of the Amazon Redshift cluster, and share the snapshot with the Sales account. In the Sales account, restore the cluster by using the snapshot ID that is shared by the Logistics account.
- D. Run COPY commands to load data from Amazon Redshift into Amazon S3 buckets in the Logistics account. Grant permissions to the Sales account to access the S3 buckets of the Logistics account.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Reference: <https://www.sqlshack.com/share-aws-redshift-data-across-accounts/>

QUESTION 448

A company is using Amazon Redshift for analytics and to generate customer reports. The company recently acquired 50 TB of additional customer demographic data. The data is stored in .csv files in Amazon S3. The company needs a solution that joins the data and visualizes the results with the least possible cost and effort.

What should a solutions architect recommend to meet these requirements?

- A. Use Amazon Redshift Spectrum to query the data in Amazon S3 directly and join that data with the existing data in Amazon Redshift. Use Amazon QuickSight to build the visualizations.
- B. Use Amazon Athena to query the data in Amazon S3. Use Amazon QuickSight to join the data from Athena with the existing data in Amazon Redshift and to build the visualizations.
- C. Increase the size of the Amazon Redshift cluster, and load the data from Amazon S3. Use Amazon EMR Notebooks to query the data and build the visualizations in Amazon Redshift.
- D. Export the data from the Amazon Redshift cluster into Apache Parquet files in Amazon S3. Use Amazon Elasticsearch Service (Amazon ES) to query the data. Use Kibana to visualize the results.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

**QUESTION 449**

A solutions architect must provide a fully managed replacement for an on-premises solution that allows employees and partners to exchange files. The solution must be easily accessible to employees connecting from on-premises systems, remote employees, and external partners.

Which solution meets these requirements?

- A. Use AWS Transfer for SFTP to transfer files into and out of Amazon S3.
- B. Use AWS Snowball Edge for local storage and large-scale data transfers.
- C. Use Amazon FSx to store and transfer files to make them available remotely.
- D. Use AWS Storage Gateway to create a volume gateway to store and transfer files to Amazon S3.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Reference: <https://aws.amazon.com/aws-transfer-family/?whats-new-cards.sortby=item.additionalFields.postDateTime&whats-new-cards.sort-order=desc>

QUESTION 450

A company's database is hosted on an Amazon Aurora MySQL DB cluster in the us-east-1 Region. The database is 4 TB in size. The company needs to expand its disaster recovery strategy to the us-west-2 Region. The company must have the ability to fail over to us-west-2 with a recovery time objective (RTO) of 15 minutes.

What should a solutions architect recommend to meet these requirements?

- A. Create a Multi-Region Aurora MySQL DB cluster in us-east-1 and us-west-2. Use an Amazon Route 53 health check to monitor us-east-1 and fail over to us-west-2 upon failure.
- B. Take a snapshot of the DB cluster in us-east-1. Configure an Amazon EventBridge (Amazon CloudWatch Events) rule that invokes an AWS Lambda function upon receipt of resource events. Configure the Lambda function to copy the snapshot to us-west-2 and restore the snapshot in us-west-2 when failure is detected.

- C. Create an AWS CloudFormation script to create another Aurora MySQL DB cluster in us-west-2 in case of failure. Configure an Amazon EventBridge (Amazon CloudWatch Events) rule that invokes an AWS Lambda function upon receipt of resource events. Configure the Lambda function to deploy the AWS CloudFormation stack in us-west-2 when failure is detected.
- D. Recreate the database as an Aurora global database with the primary DB cluster in us-east-1 and a secondary DB cluster in us-west-2. Configure an Amazon EventBridge (Amazon CloudWatch Events) rule that invokes an AWS Lambda function upon receipt of resource events. Configure the Lambda function to promote the DB cluster in us-west-2 when failure is detected.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Reference: <https://docs.aws.amazon.com/aws-backup/latest/devguide/eventbridge.html>

QUESTION 451

A company is migrating its applications to AWS. Currently, applications that run on premises generate hundreds of terabytes of data that is stored on a shared file system. The company is running an analytics application in the cloud that runs hourly to generate insights from this data.

The company needs a solution to handle the ongoing data transfer between the on-premises shared file system and Amazon S3. The solution also must be able to handle occasional interruptions in internet connectivity.

Which solutions should the company use for the data transfer to meet these requirements?

- A. AWS DataSync
- B. AWS Migration Hub
- C. AWS Snowball Edge Storage Optimized
- D. AWS Transfer for SFTP

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Reference: <https://aws.amazon.com/cloud-data-migration/>

QUESTION 452

A solutions architect is designing the architecture for a new web application. The application will run on AWS Fargate containers with an Application Load Balancer (ALB) and an Amazon Aurora PostgreSQL database. The web application will perform primarily read queries against the database.

What should the solutions architect do to ensure that the website can scale with increasing traffic?

(Choose two.)

- A. Enable auto scaling on the ALB to scale the load balancer horizontally.
- B. Configure Aurora Auto Scaling to adjust the number of Aurora Replicas in the Aurora cluster dynamically.
- C. Enable cross-zone load balancing on the ALB to distribute the load evenly across containers in all Availability Zones.
- D. Configure an Amazon Elastic Container Service (Amazon ECS) cluster in each Availability Zone to distribute the load across multiple Availability Zones.
- E. Configure Amazon Elastic Container Service (Amazon ECS) Service Auto Scaling with a target tracking scaling policy that is based on CPU utilization.

Correct Answer: BE

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 453

A company captures ordered clickstream data from multiple websites and uses batch processing to analyze the data. The company receives 100 million event records, all approximately 1 KB in size, each day. The company loads the data into Amazon Redshift each night, and business analysts consume the data.

The company wants to move toward near-real-time data processing for timely insights. The solution should process the streaming data while requiring the least possible operational overhead.

Which combination of AWS services will meet these requirements MOST cost-effectively?

(Choose two.)

- A. Amazon EC2
- B. AWS Batch
- C. Amazon Simple Queue Service (Amazon SQS)

- D. Amazon Kinesis Data Firehose
- E. Amazon Kinesis Data Analytics

Correct Answer: CE

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 454

A company has a customer relationship management (CRM) application that stores data in an Amazon RDS DB instance that runs Microsoft SQL Server. The company's IT staff has administrative access to the database. The database contains sensitive data. The company wants to ensure that the data is not accessible to the IT staff and that only authorized personnel can view the data.

What should a solutions architect do to secure the data?

- A. Use client-side encryption with an Amazon RDS managed key.
- B. Use client-side encryption with an AWS Key Management Service (AWS KMS) customer managed key.
- C. Use Amazon RDS encryption with an AWS Key Management Service (AWS KMS) default encryption key.
- D. Use Amazon RDS encryption with an AWS Key Management Service (AWS KMS) customer managed key.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 455

A company with a single AWS account runs its internet-facing containerized web application on an Amazon Elastic Kubernetes Service (Amazon EKS) cluster. The EKS cluster is placed in a private subnet of a VPC. System administrators access the EKS cluster through a bastion host on a public subnet.

A new corporate security policy requires the company to avoid the use of bastion hosts. The company also must not allow internet connectivity to the EKS cluster.

Which solution meets these requirements MOST cost-effectively?

- A. Set up an AWS Direct Connect connection.
- B. Create a transit gateway.
- C. Establish a VPN connection.
- D. Use AWS Storage Gateway.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 456

A company has deployed a multiplayer game for mobile devices. The game requires live location tracking of players based on latitude and longitude. The data store for the game must support rapid updates and retrieval of locations.

The game uses an Amazon RDS for PostgreSQL DB instance with read replicas to store the location data. During peak usage periods, the database is unable to maintain the performance that is needed for reading and writing updates. The game's user base is increasing rapidly.

What should a solutions architect do to improve the performance of the data tier?

- A. Take a snapshot of the existing DB instance. Restore the snapshot with Multi-AZ enabled.
- B. Migrate from Amazon RDS to Amazon Elasticsearch Service (Amazon ES) with Kibana.
- C. Deploy Amazon DynamoDB Accelerator (DAX) in front of the existing DB instance. Modify the game to use DAX.
- D. Deploy an Amazon ElastiCache for Redis cluster in front of the existing DB instance. Modify the game to use Redis.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 457

A company is migrating a large, mission-critical database to AWS. A solutions architect has decided to use an Amazon RDS for MySQL Multi-AZ DB instance that is deployed with 80,000 Provisioned IOPS for storage. The solutions architect is using AWS Database Migration Service (AWS DMS) to perform the data migration. The migration is taking longer than expected, and the company wants to speed up the process. The company's network team has ruled out bandwidth as a limiting factor. Which actions should the solutions architect take to speed up the migration? (Choose two.)

- A. Disable Multi-AZ on the target DB instance.
- B. Create a new DMS instance that has a larger instance size.
- C. Turn off logging on the target DB instance until the initial load is complete.
- D. Restart the DMS task on a new DMS instance with transfer acceleration enabled.
- E. Change the storage type on the target DB instance to Amazon Elastic Block Store (Amazon EBS) General Purpose SSD (gp2).

Correct Answer: CD

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 458

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 these 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).

udumps

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Reference: <https://aws.amazon.com/fargate/?whats-new-cards.sortby=item.additionalFields.postDateTime&whats-new-cards.sort-order=desc&fargate-blogs.sortby=item.additionalFields.createdDate&fargate-blogs.sort-order=desc>

QUESTION 459

A company is designing a new application that runs in a VPC on Amazon EC2 instances. The application stores data in Amazon S3 and uses Amazon DynamoDB as its database. For compliance reasons, the company prohibits all traffic between the EC2 instances and other AWS services from passing over the public internet. What can a solutions architect do to meet this requirement?

- A. Configure gateway VPC endpoints to Amazon S3 and DynamoDB.
- B. Configure interface VPC endpoints to Amazon S3 and DynamoDB.
- C. Configure a gateway VPC endpoint to Amazon S3. Configure an interface VPC endpoint to DynamoDB.
- D. Configure a gateway VPC endpoint to DynamoDB. Configure an interface VPC endpoint to Amazon S3.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 460

A company's security team requests that network traffic be captured in VPC Flow Logs. The logs will be frequently accessed for 90 days and then accessed intermittently. What should a solutions architect do to meet these requirements when configuring the logs?

- A. Use Amazon CloudWatch as the target. Set the CloudWatch log group with an expiration of 90 days.
- B. Use Amazon Kinesis as the target. Configure the Kinesis stream to always retain the logs for 90 days.
- C. Use AWS CloudTrail as the target. Configure CloudTrail to save to an Amazon S3 bucket, and enable S3 Intelligent-Tiering.
- D. Use Amazon S3 as the target. Enable an S3 Lifecycle policy to transition the logs to S3 Standard- Infrequent Access (S3 Standard-IA) after 90 days.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 461

A company needs to provide its employees with secure access to confidential and sensitive files.

The company wants to ensure that the files can be accessed only by authorized users. The files must be downloaded securely to the employees' devices.

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

Which solution will meet these requirements?

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

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:



QUESTION 462

A company hosts a multi-tier web application that uses an Amazon Aurora MySQL DB cluster for storage. The application tier is hosted on Amazon EC2 instances. The company's IT security guidelines mandate that the database credentials be encrypted and rotated every 14 days.

What should a solutions architect do to meet this requirement with the LEAST operational effort?

- A. Create a new AWS Key Management Service (AWS KMS) encryption key. Use AWS Secrets Manager to create a new secret that uses the KMS key with the appropriate credentials. Associate the secret with the Aurora DB cluster. Configure a custom rotation period of 14 days.
- B. Create two parameters in AWS Systems Manager Parameter Store: one for the user name as a string parameter and one that uses the SecureString type for the password. Select AWS Key Management Service (AWS KMS) encryption for the password parameter, and load these parameters in the application tier. Implement an AWS Lambda function that rotates the password every 14 days.
- C. Store a file that contains the credentials in an AWS Key Management Service (AWS KMS) encrypted Amazon Elastic File System (Amazon EFS) file system. Mount the EFS file system in all EC2 instances of the application tier. Restrict the access to the file on the file system so that the application can read the file and that only super users can modify the file. Implement an AWS Lambda function that rotates the key in Aurora every 14 days and writes new credentials into the file.
- D. Store a file that contains the credentials in an AWS Key Management Service (AWS KMS) encrypted Amazon S3 bucket that the application uses to load the credentials. Download the file to the application regularly to ensure that the correct credentials are used. Implement an AWS Lambda function that rotates the Aurora credentials every 14 days and uploads these credentials to the file in the S3 bucket.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Reference: <https://docs.aws.amazon.com/systems-manager/latest/userguide/systems-managerparameter-store.html>

QUESTION 463

A company is building an application that consists of several microservices. The company has decided to use container technologies to deploy its software on AWS. The company needs a solution that minimizes the amount of ongoing effort for maintenance and scaling. The company cannot manage additional infrastructure.

Which combination of actions should a solutions architect take to meet these requirements?

(Choose two.)

- A. Deploy an Amazon Elastic Container Service (Amazon ECS) cluster.

- B. Deploy the Kubernetes control plane on Amazon EC2 instances that span multiple Availability Zones.
- C. Deploy an Amazon Elastic Container Service (Amazon ECS) service with an Amazon EC2 launch type. Specify a desired task number level of greater than or equal to 2.
- D. Deploy an Amazon Elastic Container Service (Amazon ECS) service with a Fargate launch type. Specify a desired task number level of greater than or equal to 2.
- E. Deploy Kubernetes worker nodes on Amazon EC2 instances that span multiple Availability Zones. Create a deployment that specifies two or more replicas for each microservice.

Correct Answer: AB

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 464

A company recently launched a new service that involves medical images. The company scans the images and sends them from its on-premises data center through an AWS Direct Connect connection to Amazon EC2 instances. After processing is complete, the images are stored in an Amazon S3 bucket.

A company requirement states that the EC2 instances cannot be accessible through the internet.

The EC2 instances run in a private subnet, which has a default route back to the on-premises data center for outbound internet access.

Usage of the new service is increasing rapidly. A solutions architect must recommend a solution that meets the company's requirements and reduces the Direct Connect charges.

Which solution accomplishes these goals MOST cost-effectively?

- A. Configure a VPC endpoint for Amazon S3. Add an entry to the private subnet's route table for the S3 endpoint.
- B. Configure a NAT gateway in a public subnet. Configure the private subnet's route table to use the NAT gateway.
- C. Configure Amazon S3 as a file system mount point on the EC2 instances. Access Amazon S3 through the mount.
- D. Move the EC2 instances into a public subnet. Configure the public subnet route table to point to an internet gateway.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:



QUESTION 465

A company is building an online multiplayer game. The game communicates by using UDP, and low latency between the client and the backend is important. The backend is hosted on Amazon EC2 instances that can be deployed to multiple AWS Regions to meet demand. The company needs the game to be highly available so that users around the world can access the game at all times.

What should a solutions architect do to meet these requirements?

- A. Deploy Amazon CloudFront to support the global traffic. Configure CloudFront with an origin group to allow access to EC2 instances in multiple Regions.
- B. Deploy an Application Load Balancer in one Region to distribute traffic to EC2 instances in each Region that hosts the game's backend instances.
- C. Deploy Amazon CloudFront to support an origin access identity (OAI). Associate the OAI with EC2 instances in each Region to support global traffic.
- D. Deploy a Network Load Balancer in each Region to distribute the traffic. Use AWS Global Accelerator to route traffic to the correct Regional endpoint.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Reference: <https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/privatecontent-restricting-access-to-s3.html>

QUESTION 466

A company runs its two-tier ecommerce 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 DB 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 thirdparty web service. The application must be highly available.

Which combination of configuration options will meet these requirements? (Choose 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.

Correct Answer: AB

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 467

A security team needs to enforce the rotation of all IAM users' access keys every 90 days. If an access key is found to be older, the key must be made inactive and removed. A solutions architect must create a solution that will check for and remediate any keys older than 90 days.

Which solution meets these requirements with the LEAST operational effort?

- A. Create an AWS Config rule to check for the key age. Configure the AWS Config rule to run an AWS Batch job to remove the key.
- B. Create an Amazon EventBridge (Amazon CloudWatch Events) rule to check for the key age. Configure the rule to run an AWS Batch job to remove the key.
- C. Create an AWS Config rule to check for the key age. Define an Amazon EventBridge (Amazon CloudWatch Events) rule to schedule an AWS Lambda function to remove the key.
- D. Create an Amazon EventBridge (Amazon CloudWatch Events) rule to check for the key age. Define an EventBridge (CloudWatch Events) rule to run an AWS Batch job to remove the key.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Reference: <https://aws.amazon.com/blogs/mt/managing-aged-access-keys-through-aws-configremediations/>

QUESTION 468

A solutions architect must provide an automated solution for a company's compliance policy that states security groups cannot include a rule that allows SSH from 0.0.0.0/0. The company needs to be notified if there is any breach in the policy. A solution is needed as soon as possible.

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

- A. Write an AWS Lambda script that monitors security groups for SSH being open to 0.0.0.0/0 addresses and creates a notification every time it finds one.
- B. Enable the restricted-ssh AWS Config managed rule and generate an Amazon Simple Notification Service (Amazon SNS) notification when a noncompliant rule is created.
- C. Create an IAM role with permissions to globally open security groups and network ACLs. Create an Amazon Simple Notification Service (Amazon SNS) topic to generate a notification every time the role is assumed by a user.
- D. Configure a service control policy (SCP) that prevents non-administrative users from creating or editing security groups. Create a notification in the ticketing system when a user requests a rule that needs administrator permissions.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Reference: <https://www.stratoscale.com/blog/compute/aws-security-groups-5-best-practices/>

QUESTION 469

A media company is using two video conversion tools that run on Amazon EC2 instances. One tool runs on Windows instances, and the other tool runs on Linux instances. Each video file is large in size and must be processed by both tools. The company needs a storage solution that can provide a centralized file system that can be mounted on all the EC2 instances that are used in this process.

Which solution meets these requirements?

- A. Use Amazon FSx for Windows File Server for the Windows instances. Use Amazon Elastic File System (Amazon EFS) with Max I/O performance mode for the Linux instances.
- B. Use Amazon FSx for Windows File Server for the Windows instances. Use Amazon FSx for Lustre for the Linux instances. Link both Amazon FSx file systems to the same Amazon S3 bucket.
- C. Use Amazon Elastic File System (Amazon EFS) with General Purpose performance mode for the Windows instances and the Linux instances
- D. Use Amazon FSx for Windows File Server for the Windows instances and the Linux instances.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 470

A company operates a two-tier application for image processing. The application uses two Availability Zones, each with one public subnet and one private subnet. An Application Load Balancer (ALB) for the web tier uses the public subnets. Amazon EC2 instances for the application tier use the private subnets.

Users report that the application is running more slowly than expected. A security audit of the web server log files shows that the application is receiving millions of illegitimate requests from a small number of IP addresses. A solutions architect needs to resolve the immediate performance problem while the company investigates a more permanent solution.

What should the solutions architect recommend to meet this requirement?

- A. Modify the inbound security group for the web tier. Add a deny rule for the IP addresses that are consuming resources.
- B. Modify the network ACL for the web tier subnets. Add an inbound deny rule for the IP addresses that are consuming resources.
- C. Modify the inbound security group for the application tier. Add a deny rule for the IP addresses that are consuming resources.
- D. Modify the network ACL for the application tier subnets. Add an inbound deny rule for the IP addresses that are consuming resources.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 471

A company is planning to migrate a TCP-based application into the company's VPC. The application is publicly accessible on a nonstandard TCP port through a hardware appliance in the company's data center. This public endpoint can process up to 3 million requests per second with low latency. The company requires the same level of performance for the new public endpoint in AWS.

What should a solutions architect recommend to meet this requirement?

- A. Deploy a Network Load Balancer (NLB). Configure the NLB to be publicly accessible over the TCP port that the application requires.
- B. Deploy an Application Load Balancer (ALB). Configure the ALB to be publicly accessible over the TCP port that the application requires.
- C. Deploy an Amazon CloudFront distribution that listens on the TCP port that the application requires. Use an Application Load Balancer as the origin.
- D. Deploy an Amazon API Gateway API that is configured with the TCP port that the application requires. Configure AWS Lambda functions with provisioned concurrency to process the requests.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 472

An ecommerce company is creating an application that requires a connection to a third-party payment service to process payments. The payment service needs to explicitly allow the public IP address of the server that is making the payment request. However, the company's security policies do not allow any server to be exposed directly to the public internet.

Which solution will meet these requirements?

- A. Provision an Elastic IP address. Host the application servers on Amazon EC2 instances in a private subnet. Assign the public IP address to the application servers.
- B. Create a NAT gateway in a public subnet. Host the application servers on Amazon EC2 instances in a private subnet. Route payment requests through the NAT gateway.
- C. Deploy an Application Load Balancer (ALB). Host the application servers on Amazon EC2 instances in a private subnet. Route the payment requests through the ALB.
- D. Set up an AWS Client VPN connection to the payment service. Host the application servers on Amazon EC2 instances in a private subnet. Route the payment requests through the VPN.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 473

A company is running an ASP.NET MVC application on a single Amazon EC2 instance. A recent increase in application traffic is causing slow response times for users during lunch hours. The company needs to resolve this concern with the

least amount of configuration.

What should a solutions architect recommend to meet these requirements?

- A. Move the application to AWS Elastic Beanstalk. Configure load-based auto scaling and timebased scaling to handle scaling during lunch hours.
- B. Move the application to Amazon Elastic Container Service (Amazon ECS). Create an AWS Lambda function to handle scaling during lunch hours.
- C. Move the application to Amazon Elastic Container Service (Amazon ECS). Configure scheduled scaling for AWS Application Auto Scaling during lunch hours.
- D. Move the application to AWS Elastic Beanstalk. Configure load-based auto scaling, and create an AWS Lambda function to handle scaling during lunch hours.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Reference: <https://docs.aws.amazon.com/elasticbeanstalk/latest/dg/environments-cfg-autoscalingscheduledactions.html>

QUESTION 474

An online gaming company is designing a game that is expected to be popular all over the world.

A solutions architect needs to define an AWS Cloud architecture that supports near-real-time recording and displaying of current game statistics for each player, along with the names of the top 25 players in the world, at any given time.

Which AWS database solution and configuration should the solutions architect use to meet these requirements?

- A. Use Amazon RDS for MySQL as the data store for player activity. Configure the RDS DB instance for Multi-AZ support.
- B. Use Amazon DynamoDB as the data store for player activity. Configure DynamoDB Accelerator (DAX) for the player data.
- C. Use Amazon DynamoDB as the data store for player activity. Configure global tables in each required AWS Region for the player data.
- D. Use Amazon RDS for MySQL as the data store for player activity. Configure cross-Region read replicas in each required AWS Region based on player proximity.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:



QUESTION 475

A company uses Amazon RDS for PostgreSQL databases for its data tier. The company must implement password rotation for the databases.

Which solution meets this requirement with the LEAST operational overhead?

- A. Store the password in AWS Secrets Manager. Enable automatic rotation on the secret.
- B. Store the password in AWS Systems Manager Parameter Store. Enable automatic rotation on the parameter.
- C. Store the password in AWS Systems Manager Parameter Store. Write an AWS Lambda function that rotates the password.
- D. Store the password in AWS Key Management Service (AWS KMS). Enable automatic rotation on the customer master key (CMK).

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

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

QUESTION 476

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

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

Which system architecture should the solutions architect recommend?

- A. Launch an Amazon EC2 instance to serve as the HTTPS endpoint and to process the messages. Configure the EC2 instance to save the results to an Amazon S3 bucket.
- B. Create an HTTPS endpoint in Amazon API Gateway. Configure the API Gateway endpoint to invoke an AWS Lambda function to process the messages and save the results to an Amazon DynamoDB table.
- C. Use Amazon Route 53 to direct incoming sensor messages to an AWS Lambda function. Configure the Lambda function to process the messages and save the results to an Amazon DynamoDB table.

D. Create a gateway VPC endpoint for Amazon S3. Configure a Site-to-Site VPN connection from the facility network to the VPC so that sensor data can be written directly to an S3 bucket by way of the VPC endpoint.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:

QUESTION 477

An Amazon EC2 instance is located in a private subnet in a new VPC. This subnet does not have outbound internet access, but the EC2 instance needs the ability to download monthly security updates from an outside vendor. What should a solutions architect do to meet these requirements?

- A. Create an internet gateway, and attach it to the VPC. Configure the private subnet route table to use the internet gateway as the default route.
- B. Create a NAT gateway, and place it in a public subnet. Configure the private subnet route table to use the NAT gateway as the default route.
- C. Create a NAT instance, and place it in the same subnet where the EC2 instance is located. Configure the private subnet route table to use the NAT instance as the default route.
- D. Create an internet gateway, and attach it to the VPC. Create a NAT instance, and place it in the same subnet where the EC2 instance is located. Configure the private subnet route table to use the internet gateway as the default route.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:

QUESTION 478

A company has been running a web application with an Oracle relational database in an on-premises data center for the past 15 years. The company must migrate the database to AWS. The company needs to reduce operational overhead without having to modify the application's code. Which solution meets these requirements?

- A. Use AWS Database Migration Service (AWS DMS) to migrate the database servers to Amazon RDS.
- B. Use Amazon EC2 instances to migrate and operate the database servers.
- C. Use AWS Database Migration Service (AWS DMS) to migrate the database servers to Amazon DynamoDB.
- D. Use an AWS Snowball Edge Storage Optimized device to migrate the data from Oracle to Amazon Aurora.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:

Reference: <https://docs.aws.amazon.com/prescriptive-guidance/latest/patterns/migrate-an-onpremises-oracle-database-to-amazon-rds-for-oracle.html>

QUESTION 479

A company is running an application on Amazon EC2 instances. Traffic to the workload increases substantially during business hours and decreases afterward. The CPU utilization of an EC2 instance is a strong indicator of end-user demand on the application. The company has configured an Auto Scaling group to have a minimum group size of 2 EC2 instances and a maximum group size of 10 EC2 instances. The company is concerned that the current scaling policy that is associated with the Auto Scaling group might not be correct. The company must avoid over-provisioning EC2 instances and incurring unnecessary costs. What should a solutions architect recommend to meet these requirements?

- A. Configure Amazon EC2 Auto Scaling to use a scheduled scaling plan and launch an additional 8 EC2 instances during business hours.
- B. Configure AWS Auto Scaling to use a scaling plan that enables predictive scaling. Configure predictive scaling with a scaling mode of forecast and scale, and to enforce the maximum capacity setting during scaling.
- C. Configure a step scaling policy to add 4 EC2 instances at 50% CPU utilization and add another 4 EC2 instances at 90% CPU utilization. Configure scale-in policies to perform the reverse and remove EC2 instances based on the two values.
- D. Configure AWS Auto Scaling to have a desired capacity of 5 EC2 instances, and disable any existing scaling policies. Monitor the CPU utilization metric for 1 week. Then create dynamic scaling policies that are based on the observed values.

Correct Answer: D
Section: (none)
Explanation



Explanation/Reference:

Reference: <https://docs.aws.amazon.com/autoscaling/ec2/userguide/as-scaling-simple-step.html>

QUESTION 480

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

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

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Reference: https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_PIT.html

QUESTION 481

A company wants to use a custom distributed application that calculates various profit and loss scenarios. To achieve this goal, the company needs to provide a network connection between its Amazon EC2 instances. The connection must minimize latency and must maximize throughput. Which solution will meet these requirements?

- A. Provision the application to use EC2 Dedicated Hosts of the same instance type.
- B. Configure a placement group for EC2 instances that have the same instance type.
- C. Use multiple AWS elastic network interfaces and link aggregation.
- D. Configure AWS PrivateLink for the EC2 instances.

Correct Answer: B

Section: (none)

Explanation

**Explanation/Reference:**

Reference: <https://aws.amazon.com/premiumsupport/knowledge-center/network-throughputbenchmark-linux-ec2/>

QUESTION 482

A company designed a stateless two-tier application that uses Amazon EC2 in a single Availability Zone and an Amazon RDS Multi-AZ DB instance. New company management wants to ensure the application is highly available. What should a solutions architect do to meet this requirement?

- A. Configure the application to use Multi-AZ EC2 Auto Scaling and create an Application Load Balancer.
- B. Configure the application to take snapshots of the EC2 instances and send them to a different AWS Region.
- C. Configure the application to use Amazon Route 53 latency-based routing to feed requests to the application.
- D. Configure Amazon Route 53 rules to handle incoming requests and create a Multi-AZ Application Load Balancer.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 483

A company is relocating its data center and wants to securely transfer 50 TB of data to AWS within 2 weeks. The existing data center has a Site-to-Site VPN connection to AWS that is 90% utilized. Which AWS service should a solutions architect use to meet these requirements?

- A. AWS DataSync with a VPC endpoint
- B. AWS Direct Connect
- C. AWS Snowball Edge Storage Optimized

D. AWS Storage Gateway

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:

QUESTION 484

An entertainment company is using Amazon DynamoDB to store media metadata. The application is read intensive and experiencing delays. The company does not have staff to handle additional operational overhead and needs to improve the performance efficiency of DynamoDB without reconfiguring the application. What should a solutions architect recommend to meet this requirement?

- A. Use Amazon ElastiCache for Redis.
- B. Use Amazon DynamoDB Accelerator (DAX).
- C. Replicate data by using DynamoDB global tables.
- D. Use Amazon ElastiCache for Memcached with Auto Discovery enabled.

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Explanation:

QUESTION 485

A company wants to build a scalable key management infrastructure to support developers who need to encrypt data in their applications. What should a solutions architect do to reduce the operational burden?

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



Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Reference: <https://aws.amazon.com/kms/faqs/>

QUESTION 486

A company uses AWS Organizations to manage multiple AWS accounts for different departments. The management account has an Amazon S3 bucket that contains project reports. The company wants to limit access to this S3 bucket to only users of accounts within the organization in AWS Organizations. Which solution meets these requirements with the LEAST amount of operational overhead?

- A. Add the aws:PrincipalOrgID global condition key with a reference to the organization ID to the S3 bucket policy.
- B. Create an organizational unit (OU) for each department. Add the aws:PrincipalOrgPaths global condition key to the S3 bucket policy.
- C. Use AWS CloudTrail to monitor the CreateAccount, InviteAccountToOrganization, LeaveOrganization, and RemoveAccountFromOrganization events. Update the S3 bucket policy accordingly.
- D. Tag each user that needs access to the S3 bucket. Add the aws:PrincipalTag global condition key to the S3 bucket policy.

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Explanation:

QUESTION 487

A company runs an application in the AWS Cloud and uses Amazon DynamoDB as the database.

The company deploys Amazon EC2 instances to a private network to process data from the database. The company uses two NAT instances to provide connectivity to DynamoDB.

The company wants to retire the NAT instances. A solutions architect must implement a solution that provides connectivity to DynamoDB and that does not require ongoing management.

What is the MOST cost-effective solution that meets these requirements?

- A. Create a gateway VPC endpoint to provide connectivity to DynamoDB.
- B. Configure a managed NAT gateway to provide connectivity to DynamoDB.
- C. Establish an AWS Direct Connect connection between the private network and DynamoDB.
- D. Deploy an AWS PrivateLink endpoint service between the private network and DynamoDB.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 488

A solutions architect is designing a two-tiered architecture that has separate private subnets for compute resources and the database. An AWS Lambda function that is deployed in the compute subnets needs connectivity to the database.

Which solution will provide this connectivity in the MOST secure way?

- A. Configure the Lambda function to use Amazon RDS Proxy outside the VPC.
- B. Associate a security group with the Lambda function. Authorize this security group in the database's security group.
- C. Authorize the compute subnet's CIDR ranges in the database's security group.
- D. During the initialization phase, authorize all IP addresses in the database's security group temporarily. Remove the rule after the initialization is complete.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 489

A ride-sharing company stores historical service usage data as structured .csv data files in Amazon S3. A data analyst needs to perform SQL queries on this data. A solutions architect must recommend a solution that optimizes cost-effectiveness for the queries.

Which solution meets these requirements?

- A. Create an Amazon EMR cluster. Load the data. Perform the queries.
- B. Create an Amazon Redshift cluster. Import the data. Perform the queries.
- C. Create an Amazon Aurora PostgreSQL DB cluster. Import the data. Perform the queries.
- D. Create an Amazon Athena database. Associate the data in Amazon S3. Perform the queries.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Reference: <https://searchcloudcomputing.techtarget.com/answer/Compare-EMR-Redshift-and-Athena-for-data-analysis-on-AWS>

QUESTION 490

A company is running a publicly accessible serverless application that uses Amazon API Gateway and AWS Lambda. The application's traffic recently spiked due to fraudulent requests from botnets.

Which steps should a solutions architect take to block requests from unauthorized users? (Choose two.)

- A. Create a usage plan with an API key that is shared with genuine users only.
- B. Integrate logic within the Lambda function to ignore the requests from fraudulent IP addresses.
- C. Implement an AWS WAF rule to target malicious requests and trigger actions to filter them out.
- D. Convert the existing public API to a private API. Update the DNS records to redirect users to the new API endpoint.



E. Create an IAM role for each user attempting to access the API. A user will assume the role when making the API call.

Correct Answer: AE

Section: (none)

Explanation

Explanation/Reference:

Reference:
https://www.google.com/url?sa=t&rct=j&q=&esrc=s&source=web&cd=&ved=2ahUKEwiR_oHhu5DyAhVvA2MBHWEbDhgQFjAGegQIIhAD&url=https%3A%2F%2Faws.amazon.com%2Fblogs%2Fmedia%2Ftag%2Famazon-api-gateway%2Ffeed%2F&usq=AOvVaw2OaNncetRRtgvJ-d60ePyu

QUESTION 491

A company has hired a solutions architect to design a reliable architecture for its application. The application consists of one Amazon RDS DB instance and two manually provisioned Amazon EC2 instances that run web servers. The EC2 instances are located in a single Availability Zone.

An employee recently deleted the DB instance, and the application was unavailable for 24 hours as a result. The company is concerned with the overall reliability of its environment.

What should the solutions architect do to maximize reliability of the application's infrastructure?

- A. Delete one EC2 instance and enable termination protection on the other EC2 instance. Update the DB instance to be Multi-AZ, and enable deletion protection.
- B. Update the DB instance to be Multi-AZ, and enable deletion protection. Place the EC2 instances behind an Application Load Balancer, and run them in an EC2 Auto Scaling group across multiple Availability Zones.
- C. Create an additional DB instance along with an Amazon API Gateway and an AWS Lambda function. Configure the application to invoke the Lambda function through API Gateway. Have the Lambda function write the data to the two DB instances.
- D. Place the EC2 instances in an EC2 Auto Scaling group that has multiple subnets located in multiple Availability Zones. Use Spot Instances instead of On-Demand Instances. Set up Amazon CloudWatch alarms to monitor the health of the instances. Update the DB instance to be Multi-AZ, and enable deletion protection.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:



QUESTION 492

An online photo-sharing company stores its photos in an Amazon S3 bucket that exists in the uswest-1 Region. The company needs to store a copy of all existing and new photos in another geographical location.

Which solution will meet this requirement with the LEAST operational effort?

- A. Create a second S3 bucket in us-east-1. Enable S3 Cross-Region Replication from the existing S3 bucket to the second S3 bucket.
- B. Create a cross-origin resource sharing (CORS) configuration of the existing S3 bucket. Specify useast-1 in the CORS rule's AllowedOrigin element.
- C. Create a second S3 bucket in us-east-1 across multiple Availability Zones. Create an S3 Lifecycle management rule to save photos into the second S3 bucket.
- D. Create a second S3 bucket in us-east-1 to store the replicated photos. Configure S3 event notifications on object creation and update events that invoke an AWS Lambda function to copy photos from the existing S3 bucket to the second S3 bucket.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 493

A company wants to migrate its accounting system from an on-premises data center to the AWS Cloud in a single AWS Region. Data security and an immutable audit log are the top priorities. The company must monitor all AWS activities for compliance auditing. The company has enabled AWS CloudTrail but wants to make sure it meets these requirements.

Which actions should a solutions architect take to protect and secure CloudTrail? (Choose two.)

- A. Enable CloudTrail log file validation.
- B. Install the CloudTrail Processing Library.
- C. Enable logging of Insights events in CloudTrail.
- D. Enable custom logging from the on-premises resources.
- E. Create an AWS Config rule to monitor whether CloudTrail is configured to use server-side encryption with AWS KMS managed encryption keys (SSE-KMS).

Correct Answer: AC

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 494

A company needs to ingest and handle large amounts of streaming data that its application generates. The application runs on Amazon EC2 instances and sends data to Amazon Kinesis Data Streams, which is configured with default settings. Every other day, the application consumes the data and writes the data to an Amazon S3 bucket for business intelligence (BI) processing. The company observes that Amazon S3 is not receiving all the data that the application sends to Kinesis Data Streams.

What should a solutions architect do to resolve this issue?

- A. Update the Kinesis Data Streams default settings by modifying the data retention period.
- B. Update the application to use the Kinesis Producer Library (KPL) to send the data to Kinesis Data Streams.
- C. Update the number of Kinesis shards to handle the throughput of the data that is sent to Kinesis Data Streams.
- D. Turn on S3 Versioning within the S3 bucket to preserve every version of every object that is ingested in the S3 bucket.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Reference: <https://aws.amazon.com/kinesis/data-firehose/faqs/>

QUESTION 495

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. Configure new Lambda functions to scan the local storage to save the customer data to the database.
- D. 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.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 496

A company is developing a file-sharing application that will use an Amazon S3 bucket for storage.

The company wants to serve all the files through an Amazon CloudFront distribution. The company does not want the files to be accessible through direct navigation to the S3 URL.

What should a solutions architect do to meet these requirements?

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

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Reference: <https://aws.amazon.com/premiumsupport/knowledge-center/cloudfront-access-toamazon-s3/>

QUESTION 497

A company has primary and secondary data centers that are 500 miles (804.7 km) apart and interconnected with high-speed fiber-optic cable. The company needs a highly available and secure network connection between its data centers and a VPC on AWS for a mission-critical workload. A solutions architect must choose a connection solution that provides maximum resiliency. Which solution meets these requirements?

- A. Two AWS Direct Connect connections from the primary data center terminating at two Direct Connect locations on two separate devices
- B. A single AWS Direct Connect connection from each of the primary and secondary data centers terminating at one Direct Connect location on the same device
- C. Two AWS Direct Connect connections from each of the primary and secondary data centers terminating at two Direct Connect locations on two separate devices
- D. A single AWS Direct Connect connection from each of the primary and secondary data centers terminating at one Direct Connect location on two separate devices

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 498

A company runs a fleet of web servers using an Amazon RDS for PostgreSQL DB instance. After a routine compliance check, the company sets a standard that requires a recovery point objective (RPO) of less than 1 second for all its production databases. Which solution meets these requirements?

- A. Enable a Multi-AZ deployment for the DB instance.
- B. Enable auto scaling for the DB instance in one Availability Zone.
- C. Configure the DB instance in one Availability Zone, and create multiple read replicas in a separate Availability Zone.
- D. Configure the DB instance in one Availability Zone, and configure AWS Database Migration Service (AWS DMS) change data capture (CDC) tasks.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Reference: <https://aws.amazon.com/blogs/database/implementing-a-disaster-recovery-strategywith-amazon-rds/>

**QUESTION 499**

A company is hosting its website by using Amazon EC2 instances behind an Elastic Load Balancer across multiple Availability Zones. The instances run in an EC2 Auto Scaling group. The website uses Amazon Elastic Block Store (Amazon EBS) volumes to store product manuals for users to download. The company updates the product content often, so new instances launched by the Auto Scaling group often have old data. It can take up to 30 minutes for the new instances to receive all the updates. The updates also require the EBS volumes to be resized during business hours.

The company wants to ensure that the product manuals are always up to date on all instances and that the architecture adjusts quickly to increased user demand. A solutions architect needs to meet these requirements without causing the company to update its application code or adjust its website.

What should the solutions architect do to accomplish this goal?

- A. Store the product manuals in an EBS volume. Mount that volume to the EC2 instances.
- B. Store the product manuals in an Amazon S3 bucket. Redirect the downloads to this bucket.
- C. Store the product manuals in an Amazon Elastic File System (Amazon EFS) volume. Mount that volume to the EC2 instances.
- D. Store the product manuals in an Amazon S3 Standard-Infrequent Access (S3 Standard-IA) bucket. Redirect the downloads to this bucket.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 500

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.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Reference: <https://aws.amazon.com/getting-started/hands-on/deliver-content-faster/>

QUESTION 501

A company is building its web application by using containers on AWS. The company requires three instances of the web application to run at all times. The application must be highly available and must be able to scale to meet increases in demand.

Which solution meets these requirements?

- A. Use the AWS Fargate launch type to create an Amazon Elastic Container Service (Amazon ECS) cluster. Create a task definition for the web application. Create an ECS service that has a desired count of three tasks.
- B. Use the Amazon EC2 launch type to create an Amazon Elastic Container Service (Amazon ECS) cluster that has three container instances in one Availability Zone. Create a task definition for the web application. Place one task for each container instance.
- C. Use the AWS Fargate launch type to create an Amazon Elastic Container Service (Amazon ECS) cluster that has three container instances in three different Availability Zones. Create a task definition for the web application. Create an ECS service that has a desired count of three tasks.
- D. Use the Amazon EC2 launch type to create an Amazon Elastic Container Service (Amazon ECS) cluster that has one container instance in two different Availability Zones. Create a task definition for the web application. Place two tasks on one container instance. Place one task on the remaining container instance.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Reference:

<https://docs.aws.amazon.com/AmazonECS/latest/developerguide/task-placement.html>

<https://aws.amazon.com/blogs/containers/amazon-ecs-availability-best-practices/>



QUESTION 502

An online learning company is migrating to the AWS Cloud. The company maintains its student records in a PostgreSQL database. The company needs a solution in which its data is available and online across multiple AWS Regions at all times.

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

- A. Migrate the PostgreSQL database to a PostgreSQL cluster on Amazon EC2 instances.
- B. Migrate the PostgreSQL database to an Amazon RDS for PostgreSQL DB instance with the Multi-AZ feature turned on.
- C. Migrate the PostgreSQL database to an Amazon RDS for PostgreSQL DB instance. Create a read replica in another Region.
- D. Migrate the PostgreSQL database to an Amazon RDS for PostgreSQL DB instance. Set up DB snapshots to be copied to another Region.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 503

A solutions architect is designing a new hybrid architecture to extend a company's on-premises infrastructure to AWS. The company requires a highly available connection with consistent low latency to an AWS Region. The company needs to minimize costs and is willing to accept slower traffic if the primary connection fails.

What should the solutions architect do to meet these requirements?

- A. Provision an AWS Direct Connect connection to a Region. Provision a VPN connection as a backup if the primary Direct Connect connection fails.
- B. Provision a VPN tunnel connection to a Region for private connectivity. Provision a second VPN tunnel for private connectivity and as a backup if the primary VPN connection fails.

- C. Provision an AWS Direct Connect connection to a Region. Provision a second Direct Connect connection to the same Region as a backup if the primary Direct Connect connection fails.
- D. Provision an AWS Direct Connect connection to a Region. Use the Direct Connect failover attribute from the AWS CLI to automatically create a backup connection if the primary Direct Connect connection fails.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 504

A financial company hosts a web application on AWS. The application uses an Amazon API Gateway Regional API endpoint to give users the ability to retrieve current stock prices. The company's security team has noticed an increase in the number of API requests. The security team is concerned that HTTP flood attacks might take the application offline.

A solutions architect must design a solution to protect the application from this type of attack.

Which solution meets these requirements with the LEAST operational overhead?

- A. Create an Amazon CloudFront distribution in front of the API Gateway Regional API endpoint with a maximum TTL of 24 hours.
- B. Create a Regional AWS WAF web ACL with a rate-based rule. Associate the web ACL with the API Gateway stage.
- C. Use Amazon CloudWatch metrics to monitor the Count metric and alert the security team when the predefined rate is reached.
- D. Create an Amazon CloudFront distribution with Lambda@Edge in front of the API Gateway Regional API endpoint. Create an AWS Lambda function to block requests from IP addresses that exceed the predefined rate.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 505

A company is running an application on AWS to process weather sensor data that is stored in an Amazon S3 bucket. Three batch jobs run hourly to process the data in the S3 bucket for different purposes. The company wants to reduce the overall processing time by running the three applications in parallel using an event-based approach.

What should a solutions architect do to meet these requirements?

- A. Enable S3 Event Notifications for new objects to an Amazon Simple Queue Service (Amazon SQS) FIFO queue. Subscribe all applications to the queue for processing.
- B. Enable S3 Event Notifications for new objects to an Amazon Simple Queue Service (Amazon SQS) standard queue. Create an additional SQS queue for all applications, and subscribe all applications to the initial queue for processing.
- C. Enable S3 Event Notifications for new objects to separate Amazon Simple Queue Service (Amazon SQS) FIFO queues. Create an additional SQS queue for each application, and subscribe each queue to the initial topic for processing.
- D. Enable S3 Event Notifications for new objects to an Amazon Simple Notification Service (Amazon SNS) topic. Create an Amazon Simple Queue Service (Amazon SQS) queue for each application, and subscribe each queue to the topic for processing.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Reference: <https://docs.aws.amazon.com/AmazonS3/latest/userguide/ways-to-add-notificationconfig-to-bucket.html>

QUESTION 506

A company has an AWS Lambda function that needs read access to an Amazon S3 bucket that is located in the same AWS account.

Which solution will meet these requirements in the MOST secure manner?

- A. Apply an S3 bucket policy that grants read access to the S3 bucket.
- B. Apply an IAM role to the Lambda function. Apply an IAM policy to the role to grant read access to the S3 bucket.
- C. Embed an access key and a secret key in the Lambda function's code to grant the required IAM permissions for read access to the S3 bucket.
- D. Apply an IAM role to the Lambda function. Apply an IAM policy to the role to grant read access to all S3 buckets in the account.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Reference: <https://aws.amazon.com/premiumsupport/knowledge-center/access-denied-lambdas3-bucket/>

QUESTION 507

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.

Which solution will meet these requirements MOST cost-effectively?

- 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. Create a DB instance with the Multi-AZ option activated.
- C. Deploy EC2 instances across at least two Availability Zones within the same Region. Create a DB instance in a single Availability Zone.
- D. Deploy EC2 instances across at least two Availability Zones within the same Region. Create a DB instance with the Multi-AZ option activated.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 508

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 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. Increase the storage capacity to 24 TiB. Change the storage type to Provisioned IOPS.
- C. Create an Amazon S3 bucket. Update the application to store documents in the S3 bucket. Store the object metadata in the existing database.
- D. Create an Amazon DynamoDB table. Update the application to use DynamoDB. Use AWS Database Migration Service (AWS DMS) to migrate data from the Oracle database to DynamoDB.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:**QUESTION 509**

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

Which solution meets these requirements?

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

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 510

A company runs its infrastructure on AWS and has a registered base of 700,000 users for its document management application. The company intends to create a product that converts large .pdf files to .jpg image files. The .pdf files average 5 MB in size. The company needs to store the original files and the converted files. A solutions architect must design a scalable solution to accommodate demand that will grow rapidly over time.

Which solution meets these requirements MOST cost-effectively?

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

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 511

A company runs an application on a large fleet of Amazon EC2 instances. The application reads and write entries into an Amazon DynamoDB table. The size of the DynamoDB table continuously grows, but the application needs only data from the last 30 days. The company needs a solution that minimizes cost and development effort.

Which solution meets these requirements?

- A. Use an AWS CloudFormation template to deploy the complete solution. Redeploy the CloudFormation stack every 30 days, and delete the original stack.
- B. Use an EC2 instance that runs a monitoring application from AWS Marketplace. Configure the monitoring application to use Amazon DynamoDB Streams to store the timestamp when a new item is created in the table. Use a script that runs on the EC2 instance to delete items that have a timestamp that is older than 30 days.
- C. Configure Amazon DynamoDB Streams to invoke an AWS Lambda function when a new item is created in the table. Configure the Lambda function to delete items in the table that are older than 30 days.
- D. Extend the application to add an attribute that has a value of the current timestamp plus 30 days to each new item that is created in the table. Configure DynamoDB to use the attribute as the TTL attribute.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:



QUESTION 512

A company has many projects that run in multiple AWS Regions. The projects usually have a three-tier architecture with Amazon EC2 instances that run behind an Application Load Balancer. The instances run in an Auto Scaling group and share Amazon Elastic File System (Amazon EFS) storage and Amazon RDS databases. Some projects have resources in more than one Region. A solutions architect needs to identify each project's individual costs.

Which solution will provide this information with the LEAST amount of operational effort?

- A. Use Cost Explorer to perform one-time queries for each Region and create a report that filters by project.
- B. Use the AWS Billing and Cost Management details page to see the actual usage costs of the resources by project.
- C. Use AWS Systems Manager to group resources by project and monitor each project's resources and cost.
- D. Use AWS Billing and Cost Management to activate cost allocation tags and create reports that are based on the project tags.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Reference: <https://docs.aws.amazon.com/awsaccountbilling/latest/aboutv2/cost-alloc-tags.html>

QUESTION 513

A company wants to perform an online migration of active datasets from an on-premises NFS server to an Amazon S3 bucket that is named DOC-EXAMPLE-BUCKET. Data integrity verification is required during the transfer and at the end of the transfer. The data also must be encrypted.

A solutions architect is using an AWS solution to migrate the data.

Which solution meets these requirements?

- A. AWS Storage Gateway file gateway
- B. S3 Transfer Acceleration

- C. AWS DataSync
- D. AWS Snowball Edge Storage Optimized

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Reference: <https://aws.amazon.com/blogs/storage/transferring-data-between-aws-accounts-usingaws-datasync/>

QUESTION 514

A news company that has reporters all over the world is hosting its broadcast system on AWS.

The reporters send live broadcasts to the broadcast system. The reporters use software on their phones to send live streams through the Real Time Messaging Protocol (RTMP).

A solutions architect must design a solution that gives the reporters the ability to send the highest quality streams. The solution must provide accelerated TCP connections back to the broadcast system.

What should the solutions architect use to meet these requirements?

- A. Amazon CloudFront
- B. AWS Global Accelerator
- C. AWS Client VPN
- D. Amazon EC2 instances and AWS Elastic IP addresses

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Reference: <https://aws.amazon.com/solutions/implementations/live-streaming-on-aws/>

QUESTION 515

A company is using a VPC that is provisioned with a 10.10.1.0/24 CIDR block. Because of continued growth, IP address space in this block might be depleted soon. A solutions architect must add more IP address capacity to the VPC. Which solution will meet these requirements with the LEAST operational overhead?

- A. Create a new VPC. Associate a larger CIDR block.
- B. Add a secondary CIDR block of 10.10.2.0/24 to the VPC.
- C. Resize the existing VPC CIDR block from 10.10.1.0/24 to 10.10.1.0/16.
- D. Establish VPC peering with a new VPC that has a CIDR block of 10.10.1.0/16.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 516

A company is designing a shared storage solution for a gaming application that is hosted in the AWS Cloud. The company needs the ability to use SMB clients to access data. The solution must be fully managed. Which AWS solution meets these requirements?

- A. Create an AWS DataSync task that shares the data as a mountable file system. Mount the file system to the application server.
- B. Create an Amazon EC2 Windows instance. Install and configure a Windows file share role on the instance. Connect the application server to the file share.
- C. Create an Amazon FSx for Windows File Server file system. Attach the file system to the origin server. Connect the application server to the file system.
- D. Create an Amazon S3 bucket. Assign an IAM role to the application to grant access to the S3 bucket. Mount the S3 bucket to the application server.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Reference: <https://aws.amazon.com/fsx/windows/>

QUESTION 517

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. Create a task definition for the container image with an Amazon Elastic File System (Amazon EFS) volume. Create a service with that task definition.
- D. 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. Create a task definition for the container image. Create a service with that task definition.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Reference: <https://docs.aws.amazon.com/lambda/latest/dg/images-create.html>

QUESTION 518

A company is running an application on Amazon EC2 instances hosted in a private subnet of a VPC. The EC2 instances are configured in an Auto Scaling group behind an Elastic Load Balancer (ELB). The EC2 instances use a NAT gateway for outbound internet access. However, the EC2 instances are not able to connect to the public internet to download software updates.

What are the possible root causes of this issue? (Choose two.)

- A. The ELB is not configured with a proper health check.
- B. The route tables in the VPC are configured incorrectly.
- C. The EC2 instances are not associated with an Elastic IP address.
- D. The security group attached to the NAT gateway is configured incorrectly.
- E. The outbound rules on the security group attached to the EC2 instances are configured incorrectly.



Correct Answer: BD

Section: (none)

Explanation

Explanation/Reference:

Reference: <https://docs.aws.amazon.com/vpc/latest/peering/vpc-peering-routing.html>

<https://forums.aws.amazon.com/thread.jspa?threadID=226927>

QUESTION 519

A company wants to share data that is collected from self-driving cars with the automobile community. The data will be made available from within an Amazon S3 bucket. The company wants to minimize its cost of making this data available to other AWS accounts.

What should a solutions architect do to accomplish this goal?

- A. Create an S3 VPC endpoint for the bucket.
- B. Configure the S3 bucket to be a Requester Pays bucket.
- C. Create an Amazon CloudFront distribution in front of the S3 bucket.
- D. Require that the files be accessible only with the use of the BitTorrent protocol.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

You can configure an Amazon S3 bucket to be a Requester Pays bucket so that the requester pays the cost of the request and data download instead of the bucket owner.

Reference:

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

QUESTION 520

A company is running several business applications in three separate VPCs within the eu-east-1 Region. The applications must be able to communicate between VPCs. The applications also must be able to consistently send hundreds of 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 these 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. Create a transit gateway, and attach each VPC to the transit gateway. Establish connectivity between the Direct Connect connection and the transit gateway.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Reference: <https://docs.aws.amazon.com/directconnect/latest/UserGuide/direct-connectgateways-intro.html>

QUESTION 521

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

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

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Reference: <https://docs.aws.amazon.com/secretsmanager/latest/userguide/intro.html>



QUESTION 522

A company is hosting a three-tier ecommerce application in the AWS Cloud. The company hosts the website on Amazon S3 and integrates the website with an API that handles sales requests. The company hosts the API on three Amazon EC2 instances behind an Application Load Balancer (ALB). The API consists of static and dynamic front-end content along with backend workers that process sales requests asynchronously. The company is expecting a significant and sudden increase in the number of sales requests during events for the launch of new products. What should a solutions architect recommend to ensure that all the requests are processed successfully?

- A. Add an Amazon CloudFront distribution for the dynamic content. Increase the number of EC2 instances to handle the increase in traffic.
- B. Add an Amazon CloudFront distribution for the static content. Place the EC2 instances in an Auto Scaling group to launch new instances based on network traffic.
- C. Add an Amazon CloudFront distribution for the dynamic content. Add an Amazon ElastiCache instance in front of the ALB to reduce traffic for the API to handle.
- D. Add an Amazon CloudFront distribution for the static content. Add an Amazon Simple Queue Service (Amazon SQS) queue to receive requests from the website for later processing by the EC2 instances.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Reference: <https://aws.amazon.com/sqs/>

QUESTION 523

A company has an application that scans millions of connected devices for security threats and pushes the scan logs to an Amazon S3 bucket. A total of 70 GB of data is generated each week, and the company needs to store 3 years of data for historical reporting. The company must process, aggregate, and enrich the data from Amazon S3 by performing complex analytical queries and joins in the least amount of time. The aggregated dataset is visualized on an Amazon QuickSight dashboard. What should a solutions architect recommend to meet these requirements?

- A. Create and run an ETL job in AWS Glue to process the data from Amazon S3 and load it into Amazon Redshift. Perform the aggregation queries on Amazon Redshift.

- B. Use AWS Lambda functions based on S3 PutObject event triggers to copy the incremental changes to Amazon DynamoDB. Perform the aggregation queries on DynamoDB.
- C. Use AWS Lambda functions based on S3 PutObject event triggers to copy the incremental changes to Amazon Aurora MySQL. Perform the aggregation queries on Aurora MySQL.
- D. Use AWS Glue to catalog the data in Amazon S3. Perform the aggregation queries on the cataloged tables by using Amazon Athena. Query the data directly from Amazon S3.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Reference: <https://docs.aws.amazon.com/prescriptive-guidance/latest/patterns/build-an-etl-service-pipeline-to-load-data-incrementally-from-amazon-s3-to-amazon-redshift-using-aws-glue.html>

QUESTION 524

A company has an AWS account used for software engineering. The AWS account has access to the company's on-premises data center through a pair of AWS Direct Connect connections. All non-VPC traffic routes to the virtual private gateway.

A development team recently created an AWS Lambda function through the console. The development team needs to allow the function to access a database that runs in a private subnet in the company's data center.

Which solution will meet these requirements?

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

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Reference: <https://docs.aws.amazon.com/lambda/latest/dg/configuration-vpc.html>

QUESTION 525

A company has multiple AWS accounts with applications deployed in the us-west-2 Region.

Application logs are stored within Amazon S3 buckets in each account. The company wants to build a centralized log analysis solution that uses a single S3 bucket. Logs must not leave uswest-2, and the company wants to incur minimal operational overhead.

Which solution meets these requirements and is MOST cost-effective?

- A. Create an S3 Lifecycle policy that copies the objects from one of the application S3 buckets to the centralized S3 bucket.
- B. Use S3 Same-Region Replication to replicate logs from the S3 buckets to another S3 bucket in uswest-2. Use this S3 bucket for log analysis.
- C. Write a script that uses the PutObject API operation every day to copy the entire contents of the buckets to another S3 bucket in us-west-2. Use this S3 bucket for log analysis.
- D. Write AWS Lambda functions in these accounts that are triggered every time logs are delivered to the S3 buckets (s3:ObjectCreated:* event). Copy the logs to another S3 bucket in us-west-2. Use this S3 bucket for log analysis.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Reference: <https://www.varonis.com/blog/how-to-use-aws-s3/>

QUESTION 526

A company is building a web application that serves a content management system. The content management system runs on Amazon EC2 instances behind an Application Load Balancer (ALB).

The EC2 instances run in an Auto Scaling group across multiple Availability Zones. Users are constantly adding and updating files, blogs, and other website assets in the content management system.

A solutions architect must implement a solution in which all the EC2 instances share up-to-date website content with the least possible lag time.

Which solution meets these requirements?

- A. Update the EC2 user data in the Auto Scaling group lifecycle policy to copy the website assets from the EC2 instance that was launched most recently. Configure the ALB to make changes to the website assets only in the newest EC2 instance.
- B. Copy the website assets to an Amazon Elastic File System (Amazon EFS) file system. Configure each EC2 instance to mount the EFS file system locally. Configure the website hosting application to reference the website assets that are stored in the EFS file system.
- C. Copy the website assets to an Amazon S3 bucket. Ensure that each EC2 instance downloads the website assets from the S3 bucket to the attached Amazon Elastic Block Store (Amazon EBS) volume. Run the S3 sync command once each hour to keep files up to date.

D. Restore an Amazon Elastic Block Store (Amazon EBS) snapshot with the website assets. Attach the EBS snapshot as a secondary EBS volume when a new EC2 instance is launched. Configure the website hosting application to reference the website assets that are stored in the secondary EBS volume.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Reference: <https://docs.aws.amazon.com/efs/latest/ug/how-it-works.html>

QUESTION 527

A company is running a database on Amazon Aurora. The database is idle every evening. An application that performs extensive reads on the database experiences performance issues during morning hours when user traffic spikes. During these peak periods, the application receives timeout errors when reading from the database. The company does not have a dedicated operations team and needs an automated solution to address the performance issues.

Which actions should a solutions architect take to automatically adjust to the increased read load on the database? (Choose two.)

- A. Migrate the database to Aurora Serverless.
- B. Increase the instance size of the Aurora database.
- C. Configure Aurora Auto Scaling with Aurora Replicas.
- D. Migrate the database to an Aurora multi-master cluster.
- E. Migrate the database to an Amazon RDS for MySQL Multi-AZ deployment.

Correct Answer: CD

Section: (none)

Explanation

Explanation/Reference:

Reference:

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

<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/aurora-multi-master.html>

Udumps

QUESTION 528

A company is designing an application where users upload small files into Amazon S3. After a user uploads a file, the file requires one-time simple processing to transform the data and save the data in JSON format for later analysis. Each file must be processed as quickly as possible after it is uploaded. Demand will vary. On some days, users will upload a high number of files. On other days, users will upload a few files or no files.

Which solution meets these requirements with the LEAST operational overhead?

- A. Configure Amazon EMR to read text files from Amazon S3. Run processing scripts to transform the data. Store the resulting JSON file in an Amazon Aurora DB cluster.
- B. Configure Amazon S3 to send an event notification to an Amazon Simple Queue Service (Amazon SQS) queue. Use Amazon EC2 instances to read from the queue and process the data. Store the resulting JSON file in Amazon DynamoDB.
- C. Configure Amazon S3 to send an event notification to an Amazon Simple Queue Service (Amazon SQS) queue. Use an AWS Lambda function to read from the queue and process the data. Store the resulting JSON file in Amazon DynamoDB.
- D. Configure Amazon EventBridge (Amazon CloudWatch Events) to send an event to Amazon Kinesis Data Streams when a new file is uploaded. Use an AWS Lambda function to consume the event from the stream and process the data. Store the resulting JSON file in Amazon Aurora DB cluster.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Reference: <https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/aurora-ug.pdf>

QUESTION 529

A company has a Microsoft .NET application that runs on an on-premises Windows Server. The application stores data by using an Oracle Database Standard Edition server. The company is planning a migration to AWS and wants to minimize development changes while moving the application. The AWS application environment should be highly available.

Which combination of actions should the company take to meet these requirements? (Choose two.)

- A. Refactor the application as serverless with AWS Lambda functions running .NET Core.
- B. Rehost the application in AWS Elastic Beanstalk with the .NET platform in a Multi-AZ deployment.
- C. Replatform the application to run on Amazon EC2 with the Amazon Linux Amazon Machine Image (AMI).
- D. Use AWS Database Migration Service (AWS DMS) to migrate from the Oracle database to Amazon DynamoDB in a Multi-AZ deployment.

E. Use AWS Database Migration Service (AWS DMS) to migrate from the Oracle database to Oracle on Amazon RDS in a Multi-AZ deployment.

Correct Answer: BD

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The .NET web application is deployed to AWS Elastic Beanstalk, which runs in an Amazon EC2 Auto Scaling Group. You can set up a scaling policy based on Amazon CloudWatch metrics such as CPU utilization. For a database, you can use Amazon RDS in a Multi-AZ environment.

Reference: <https://docs.aws.amazon.com/prescriptive-guidance/latest/patterns/migrate-a-netapplication-from-microsoft-azure-app-service-to-aws-elastic-beanstalk.html>

<https://aws.amazon.com/dms/>

QUESTION 530

A company's security policy requires that all AWS API activity in its AWS accounts be recorded for periodic auditing. The company needs to ensure that AWS CloudTrail is enabled on all of its current and future AWS accounts using AWS Organizations.

Which solution is MOST secure?

- A. At the organization's root, define and attach a service control policy (SCP) that permits enabling CloudTrail only.
- B. Create IAM groups in the organization's management account as needed. Define and attach an IAM policy to the groups that prevents users from disabling CloudTrail.
- C. Organize accounts into organizational units (OUs). At the organization's root, define and attach a service control policy (SCP) that prevents users from disabling CloudTrail.
- D. Add all existing accounts under the organization's root. Define and attach a service control policy (SCP) to every account that prevents users from disabling CloudTrail.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:



QUESTION 531

A company is deploying an application that processes large quantities of data in batches as needed. The company plans to use Amazon EC2 instances for the workload. The network architecture must support a highly scalable solution and prevent groups of nodes from sharing the same underlying hardware.

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

- A. Create Capacity Reservations for the EC2 instances to run in a placement group.
- B. Run the EC2 instances in a spread placement group.
- C. Run the EC2 instances in a cluster placement group.
- D. Place the EC2 instances in an EC2 Auto Scaling group.
- E. Run the EC2 instances in a partition placement group.

Correct Answer: BC

Section: (none)

Explanation

Explanation/Reference:

Reference: <https://ipwithease.com/ec2-placement-groups-aws/>

QUESTION 532

A company is running a media store across multiple Amazon EC2 instances distributed across multiple Availability Zones in a single VPC. The company wants a high-performing solution to share data between all the EC2 instances, and prefers to keep the data within the VPC only.

What should a solutions architect recommend?

- A. Create an Amazon S3 bucket and call the service APIs from each instance's application.
- B. Create an Amazon S3 bucket and configure all instances to access it as a mounted volume.
- C. Configure an Amazon Elastic Block Store (Amazon EBS) volume and mount it across all instances.
- D. Configure an Amazon Elastic File System (Amazon EFS) file system and mount it across all instances.

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:

Reference: <https://docs.aws.amazon.com/efs/latest/ug/wt1-test.html>

QUESTION 533

An image-hosting company stores its objects in Amazon S3 buckets. The company wants to avoid accidental exposure of the objects in the S3 buckets to the public. All S3 objects in the entire AWS account need to remain private. Which solution will meet these requirements?

- A. Use Amazon GuardDuty to monitor S3 bucket policies. Create an automatic remediation action rule that uses an AWS Lambda function to remediate any change that makes the objects public.
- B. Use AWS Trusted Advisor to find publicly accessible S3 buckets. Configure email notifications in Trusted Advisor when a change is detected. Manually change the S3 bucket policy if it allows public access.
- C. Use AWS Resource Access Manager to find publicly accessible S3 buckets. Use Amazon Simple Notification Service (Amazon SNS) to invoke an AWS Lambda function when a change is detected. Deploy a Lambda function that programmatically remediates the change.
- D. Use the S3 Block Public Access feature on the account level. Use AWS Organizations to create a service control policy (SCP) that prevents IAM users from changing the setting. Apply the SCP to the account.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:

Reference: <https://docs.aws.amazon.com/guardduty/latest/ug/guardduty-ug.pdf>

QUESTION 534

A company has a website hosted on AWS. The website is behind an Application Load Balancer (ALB) that is configured to handle HTTP and HTTPS separately. The company wants to forward all requests to the website so that the requests will use HTTPS. What solution should a solutions architect do to meet this requirement?

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



Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:

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

QUESTION 535

A company used an AWS Direct Connect connection to copy 1 PB of data from a colocation facility to an Amazon S3 bucket in the us-east-1 Region. The company now wants to copy the data to another S3 bucket in the us-west-2 Region. Which solution will meet this requirement?

- A. Use an AWS Snowball Edge Storage Optimized device to copy the data from the colocation facility to us-west-2.
- B. Use the S3 console to copy the data from the source S3 bucket to the target S3 bucket.
- C. Use S3 Transfer Acceleration and the S3 copy-object command to copy the data from the source S3 bucket to the target S3 bucket.
- D. Add an S3 Cross-Region Replication configuration to copy the data from the source S3 bucket to the target S3 bucket.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:

Reference: <https://aws.amazon.com/premiumsupport/knowledge-center/move-objects-s3-bucket/>

QUESTION 536

A company recently announced the deployment of its retail website to a global audience. The website runs on multiple Amazon EC2 instances behind an Elastic Load Balancer. The instances run in an Auto Scaling group across multiple Availability Zones.

The company wants to provide its customers with different versions of content based on the devices that the customers use to access the website. Which combination of actions should a solutions architect take to meet these requirements? (Choose two.)

- A. Configure Amazon CloudFront to cache multiple versions of the content.
- B. Configure a host header in a Network Load Balancer to forward traffic to different instances.
- C. Configure a Lambda@Edge function to send specific objects to users based on the User-Agent header.
- D. Configure AWS Global Accelerator. Forward requests to a Network Load Balancer (NLB). Configure the NLB to set up host-based routing to different EC2 instances.
- E. Configure AWS Global Accelerator. Forward requests to a Network Load Balancer (NLB). Configure the NLB to set up path-based routing to different EC2 instances.

Correct Answer: CE

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 537

An ecommerce company needs to run a scheduled daily job to aggregate and filter sales records for analytics. The company stores the sales records in an Amazon S3 bucket. Each object can be up to 10 GB in size. Based on the number of sales events, the job can take up to an hour to complete. The CPU and memory usage of the job are constant and are known in advance.

A solutions architect needs to minimize the amount of operational effort that is needed for the job to run.

Which solution meets these requirements?

- A. Create an AWS Lambda function that has an Amazon EventBridge (Amazon CloudWatch Events) notification. Schedule the EventBridge (CloudWatch Events) event to run once a day.
- B. Create an AWS Lambda function. Create an Amazon API Gateway HTTP API. and integrate the API with the function. Create an Amazon EventBridge (Amazon CloudWatch Events) scheduled event that calls the API and invokes the function.
- C. Create an Amazon Elastic Container Service (Amazon ECS) cluster with an AWS Fargate launch type. Create an Amazon EventBridge (Amazon CloudWatch Events) scheduled event that launches an ECS task on the cluster to run the job.
- D. Create an Amazon Elastic Container Service (Amazon ECS) cluster with an Amazon EC2 launch type and an Auto Scaling group with at least one EC2 instance. Create an Amazon EventBridge (Amazon CloudWatch Events) scheduled event that launches an ECS task on the cluster to run the job.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Reference: <https://docs.aws.amazon.com/eventbridge/latest/userguide/eb-run-lambdaschedule.html>

QUESTION 538

A company wants to enforce strict security guidelines on accessing AWS Cloud resources as the company migrates production workloads from its data centers. Company management wants all users to receive permissions according to their job roles and functions.

Which solution meets these requirements with the LEAST operational overhead?

- A. Create an AWS Single Sign-On deployment. Connect to the on-premises Active Directory to centrally manage users and permissions across the company.
- B. Create an IAM role for each job function. Require each employee to call the sts:AssumeRole action in the AWS Management Console to perform their job role.
- C. Create individual IAM user accounts for each employee. Create an IAM policy for each job function, and attach the policy to all IAM users based on their job role.
- D. Create individual IAM user accounts for each employee. Create IAM policies for each job function. Create IAM groups, and attach associated policies to each group. Assign the IAM users to a group based on their job role.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Reference: <https://d1.awsstatic.com/whitepapers/adds-on-aws.pdf>

QUESTION 539

A law firm needs to share information with the public. The information includes hundreds of files that must be publicly readable. Modifications or deletions of the files by anyone before a designated future date are prohibited.

Which solution will meet these requirements in the MOST secure way?

- A. Upload all files to an Amazon S3 bucket that is configured for static website hosting. Grant readonly IAM permissions to any AWS principals that access the S3 bucket until the designated date.
- B. Create a new Amazon S3 bucket with S3 Versioning enabled. Use S3 Object Lock with a retention period in accordance with the designated date. Configure the S3 bucket for static website hosting. Set an S3 bucket policy to allow read-only access to the objects.
- C. Create a new Amazon S3 bucket with S3 Versioning enabled. Configure an event trigger to run an AWS Lambda function in case of object modification or deletion. Configure the Lambda function to replace the objects with the original versions from a private S3 bucket.
- D. Upload all files to an Amazon S3 bucket that is configured for static website hosting. Select the folder that contains the files. Use S3 Object Lock with a retention period in accordance with the designated date. Grant read-only IAM permissions to any AWS principals that access the S3 bucket.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Reference:

<https://docs.aws.amazon.com/AmazonS3/latest/userguide/HostingWebsiteOnS3Setup.html>

QUESTION 540

A solutions architect needs to design a solution that retrieves data every 2 minutes from a thirdparty web service that is accessible through the internet. A Python script runs the data retrieval in less than 100 milliseconds for each retrieval. The response is a JSON object that contains sensor data that is less than 1 KB in size. The solutions architect needs to store the JSON object along with the timestamp.

Which solution meets these requirements MOST cost-effectively?

- A. Deploy an Amazon EC2 instance with a Linux operating system. Configure a cron job to run the script every 2 minutes. Extend the script to store the JSON object along with the timestamp in a MySQL database that is hosted on an Amazon RDS DB instance.
- B. Deploy an Amazon EC2 instance with a Linux operating system to extend the script to run in an infinite loop every 2 minutes. Store the JSON object along with the timestamp in an Amazon DynamoDB table that uses the timestamp as the primary key. Run the script on the EC2 instance.
- C. Deploy an AWS Lambda function to extend the script to store the JSON object along with the timestamp in an Amazon DynamoDB table that uses the timestamp as the primary key. Use an Amazon EventBridge (Amazon CloudWatch Events) scheduled event that is initiated every 2 minutes to invoke the Lambda function.
- D. Deploy an AWS Lambda function to extend the script to run in an infinite loop every 2 minutes. Store the JSON object along with the timestamp in an Amazon DynamoDB table that uses the timestamp as the primary key. Ensure that the script is called by the handler function that is configured for the Lambda function.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Reference: <https://docs.aws.amazon.com/connect/latest/adminguide/connect-ag.pdf>

QUESTION 541

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 indepth 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.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Reference: <https://docs.aws.amazon.com/awsaccountbilling/latest/aboutv2/awsaccountbillingaboutv2.pdf>

QUESTION 542

A company serves its website by using an Auto Scaling group of Amazon EC2 instances in a single AWS Region. The website does not require a database.

The company is expanding, and the company's engineering team deploys the website to a second Region. The company wants to distribute traffic across both Regions to accommodate growth and for disaster recovery purposes. The solution should not serve traffic from a Region in which the website is unhealthy.

Which policy or resource should the company use to meet these requirements?

- A. An Amazon Route 53 simple routing policy
- B. An Amazon Route 53 multivalued answer routing policy
- C. An Application Load Balancer in one Region with a target group that specifies the EC2 instance IDs from both Regions
- D. An Application Load Balancer in one Region with a target group that specifies the IP addresses of the EC2 instances from both Regions

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Reference: <https://aws.amazon.com/premiumsupport/knowledge-center/multivalued-versus-simplepolicies/>

QUESTION 543

A company is running a photo hosting service in the us-east-1 Region. The service enables users across multiple countries to upload and view photos. Some photos are heavily viewed for months, and others are viewed for less than a week. The application allows uploads of up to 20 MB for each photo. The service uses the photo metadata to determine which photos to display to each user.

Which solution provides the appropriate user access MOST cost-effectively?

- A. Store the photos in Amazon DynamoDB. Turn on DynamoDB Accelerator (DAX) to cache frequently viewed items.
- B. Store the photos in the Amazon S3 Intelligent-Tiering storage class. Store the photo metadata and its S3 location in DynamoDB.
- C. Store the photos in the Amazon S3 Standard storage class. Set up an S3 Lifecycle policy to move photos older than 30 days to the S3 Standard-Infrequent Access (S3 Standard-IA) storage class. Use the object tags to keep track of metadata.
- D. Store the photos in the Amazon S3 Glacier storage class. Set up an S3 Lifecycle policy to move photos older than 30 days to the S3 Glacier Deep Archive storage class. Store the photo metadata and its S3 location in Amazon Elasticsearch Service (Amazon ES).

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

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



QUESTION 544

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

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

Which storage solution meets these requirements MOST cost-effectively?

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

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Reference: <https://aws.amazon.com/blogs/aws/s3-intelligent-tiering-adds-archive-access-tiers/>

QUESTION 545

A trucking company is deploying an application that will track the GPS coordinates of all the company's trucks. The company needs a solution that will generate real-time statistics based on metadata lookups with high read throughput and microsecond latency. The database must be fault tolerant and must minimize operational overhead and development effort.

Which combination of steps should a solutions architect take to meet these requirements?

(Choose two.)

- A. Use Amazon DynamoDB as the database.

- B. Use Amazon Aurora MySQL as the database.
- C. Use Amazon RDS for MySQL as the database
- D. Use Amazon ElastiCache as the caching layer.
- E. Use Amazon DynamoDB Accelerator (DAX) as the caching layer.

Correct Answer: AB

Section: (none)

Explanation

Explanation/Reference:

Reference: <https://aws.amazon.com/blogs/publicsector/creating-a-serverless-gps-monitoring-andalerting-solution/>
<https://aws.amazon.com/rds/aurora/customers/>

QUESTION 546

A company is deploying an application that processes large quantities of data in parallel. The company plans to use Amazon EC2 instances for the workload. The network architecture must be configurable to prevent groups of nodes from sharing the same underlying hardware.

Which networking solution meets these requirements?

- A. Run the EC2 instances in a spread placement group.
- B. Group the EC2 instances in separate accounts.
- C. Configure the EC2 instances with dedicated tenancy.
- D. Configure the EC2 instances with shared tenancy.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Reference: <https://aws.amazon.com/blogs/compute/using-partition-placement-groups-for-largedistributed-and-replicated-workloads-in-amazon-ec2/>

QUESTION 547

A solutions architect is designing the cloud architecture for a new application that is being deployed on AWS. The application's users will interactively download and upload files. Files that are more than 90 days old will be accessed less frequently than newer files, but all files need to be instantly available. The solutions architect must ensure that the application can scale to store petabytes of data with maximum durability.

Which solution meets these requirements?

- A. Store the files in Amazon S3 Standard. Create an S3 Lifecycle policy that moves objects that are more than 90 days old to S3 Glacier.
- B. Store the files in Amazon S3 Standard. Create an S3 Lifecycle policy that moves objects that are more than 90 days old to S3 Standard-Infrequent Access (S3 Standard-IA).
- C. Store the files in Amazon Elastic Block Store (Amazon EBS) volumes. Schedule snapshots of the volumes. Use the snapshots to archive data that is more than 90 days old.
- D. Store the files in RAID-striped Amazon Elastic Block Store (Amazon EBS) volumes. Schedule snapshots of the volumes. Use the snapshots to archive data that is more than 90 days old.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

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

QUESTION 548

A company is developing a serverless web application that gives users the ability to interact with real-time analytics from online games. The data from the games must be streamed in real life. The company needs a durable, low-latency database option for user data. The company does not know how many users will use the application. Any design considerations must provide response times of single-digit milliseconds as the application scales.

Which combination of AWS services will meet these requirements? (Choose two.)

- A. Amazon CloudFront
- B. Amazon DynamoDB
- C. Amazon Kinesis
- D. Amazon RDS
- E. AWS Global Accelerator

Correct Answer: AB

Section: (none)

Explanation

Explanation/Reference:

Reference: <https://docs.aws.amazon.com/AmazonS3/latest/userguide/optimizingperformance.html>
<https://aws.amazon.com/products/databases/real-time-apps-elasticache-for-redis/>

QUESTION 549

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

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

Which solution will meet these requirements?

- A. Refactor the Lambda function code to Apache Tomcat code that runs on Amazon EC2 instances. Connect the database by using native Java Database Connectivity (JDBC) drivers.
- B. Change the platform from Aurora to Amazon DynamoDB. Provision a DynamoDB Accelerator (DAX) cluster. Use the DAX client SDK to point the existing DynamoDB API calls at the DAX cluster.
- C. Set up two Lambda functions. Configure one function to receive the information. Configure the other function to load the information into the database. Integrate the Lambda functions by using Amazon Simple Notification Service (Amazon SNS).
- D. Set up two Lambda functions. Configure one function to receive the information. Configure the other function to load the information into the database. Integrate the Lambda functions by using an Amazon Simple Queue Service (Amazon SQS) queue.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Reference: <https://aws.amazon.com/dynamodb/dax/>

QUESTION 550

A solutions architect at a company is designing the architecture for a two-tiered web application.

The web application is composed of an internet-facing Application Load Balancer (ALB) that forwards traffic to an Auto Scaling group of Amazon EC2 instances. The EC2 instances must be able to access a database that runs on Amazon RDS.

The company has requested a defense-in-depth approach to the network layout. The company does not want to rely solely on security groups or network ACLs. Only the minimum resources that are necessary should be routable from the internet.

Which network design should the solutions architect recommend to meet these requirements?

- A. Place the ALB, EC2 instances, and RDS database in private subnets.
- B. Place the ALB in public subnets. Place the EC2 instances and RDS database in private subnets.
- C. Place the ALB and EC2 instances in public subnets. Place the RDS database in private subnets.
- D. Place the ALB outside the VPC. Place the EC2 instances and RDS database in private subnets.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Reference: <https://aws.amazon.com/premiumsupport/knowledge-center/public-load-balancerprivate-ec2/>

QUESTION 551

A company needs to build a reporting solution on AWS. The solution must support SQL queries that data analysts run on the data. The data analysts will run fewer than 10 total queries each day.

The company generates 3 GB of new data daily in an on-premises relational database. This data needs to be transferred to AWS to perform reporting tasks.

What should a solutions architect recommend to meet these requirements at the LOWEST cost?

- A. Use AWS Database Migration Service (AWS DMS) to replicate the data from the on-premises database into Amazon S3. Use Amazon Athena to query the data.
- B. Use an Amazon Kinesis Data Firehose delivery stream to deliver the data into an Amazon Elasticsearch Service (Amazon ES) cluster. Run the queries in Amazon ES.
- C. Export a daily copy of the data from the on-premises database. Use an AWS Storage Gateway file gateway to store and copy the export into Amazon S3. Use an Amazon EMR cluster to query the data.
- D. Use AWS Database Migration Service (AWS DMS) to replicate the data from the on-premises database and load it into an Amazon Redshift cluster. Use the Amazon Redshift cluster to query the data.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Reference: https://docs.aws.amazon.com/dms/latest/userguide/CHAP_Target.Redshift.html

QUESTION 552

A company is using an Application Load Balancer (ALB) to present its application to the internet.

The company finds abnormal traffic access patterns across the application. A solutions architect needs to improve visibility into the infrastructure to help the company understand these abnormalities better.

What is the MOST operationally efficient solution that meets these requirements?

- A. Create a table in Amazon Athena for AWS CloudTrail logs. Create a query for the relevant information.
- B. Enable ALB access logging to Amazon S3. Create a table in Amazon Athena, and query the logs.
- C. Enable ALB access logging to Amazon S3. Open each file in a text editor, and search each line for the relevant information.
- D. Use Amazon EMR on a dedicated Amazon EC2 instance to directly query the ALB to acquire traffic access log information.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Reference: <https://aws.amazon.com/blogs/big-data/catalog-and-analyze-application-load-balancer-logs-more-efficiently-with-aws-glue-custom-classifiers-and-amazon-athena/>

QUESTION 553

A company wants to monitor its AWS costs for financial review. The cloud operations team is designing an architecture in the AWS Organizations management account to query AWS Cost and Usage Reports for all member accounts. The team must run this query once a month and provide a detailed analysis of the bill.

Which solution is the MOST scalable and cost-effective way to meet these requirements?

- A. Enable Cost and Usage Reports in the management account. Deliver reports to Amazon Kinesis. Use Amazon EMR for analysis.
- B. Enable Cost and Usage Reports in the management account. Deliver the reports to Amazon S3. Use Amazon Athena for analysis.
- C. Enable Cost and Usage Reports for member accounts. Deliver the reports to Amazon S3. Use Amazon Redshift for analysis.
- D. Enable Cost and Usage Reports for member accounts. Deliver the reports to Amazon Kinesis. Use Amazon QuickSight for analysis.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Reference: <https://docs.aws.amazon.com/cur/latest/userguide/what-is-cur.html>

QUESTION 554

A company is planning to run a group of Amazon EC2 instances that connect to an Amazon Aurora database. The company has built an AWS Cloud Formation template to deploy the EC2 instances and the Aurora DB cluster. The company wants to allow the instances to authenticate to the database in a secure way. The company does not want to maintain static database credentials.

Which solution meets these requirements with the LEAST operational effort?

- A. Create a database user with a user name and password. Add parameters for the database user name and password to the CloudFormation template. Pass the parameters to the EC2 instances when the instances are launched.
- B. Create a database user with a user name and password. Store the user name and password in AWS Systems Manager Parameter Store. Configure the EC2 instances to retrieve the database credentials from Parameter Store.
- C. Configure the DB cluster to use IAM database authentication. Create a database user to use with IAM authentication. Associate a role with the EC2 instances to allow applications on the instances to access the database.
- D. Configure the DB cluster to use IAM database authentication with an IAM user. Create a database user that has a name that matches the IAM user. Associate the IAM user with the EC2 instances to allow applications on the instances to access the database.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Reference:



<https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/deploying.applications.html>

QUESTION 555

A company wants to move from many standalone AWS accounts to a consolidated, multi-account architecture. The company plans to create many new AWS accounts for different business units. The company needs to authenticate access to these AWS accounts by using a centralized corporate directory service. Which combination of actions should a solutions architect recommend to meet these requirements? (Choose two.)

- A. Create a new organization in AWS Organizations with all features turned on. Create the new AWS accounts in the organization.
- B. Set up an Amazon Cognito identity pool. Configure AWS Single Sign-On to accept Amazon Cognito authentication.
- C. Configure a service control policy (SCP) to manage the AWS accounts. Add AWS Single Sign-On to AWS Directory Service.
- D. Create a new organization in AWS Organizations. Configure the organization's authentication mechanism to use AWS Directory Service directly.
- E. Set up AWS Single Sign-On (AWS SSO) in the organization. Configure AWS SSO, and integrate it with the company's corporate directory service.

Correct Answer: BC

Section: (none)

Explanation

Explanation/Reference:

Reference: <https://aws.amazon.com/cognito/>
https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_policies_scps.html

QUESTION 556

A company sells datasets to customers who do research in artificial intelligence and machine learning (AI/ML). The datasets are large, formatted files that are stored in an Amazon S3 bucket in the us-east-1 Region. The company hosts a web application that the customers use to purchase access to a given dataset. The web application is deployed on multiple Amazon EC2 instances behind an Application Load Balancer. After a purchase is made, customers receive an S3 signed URL that allows access to the files. The customers are distributed across North America and Europe. The company wants to reduce the cost that is associated with data transfers and wants to maintain or improve performance. What should a solutions architect do to meet these requirements?

- A. Configure S3 Transfer Acceleration on the existing S3 bucket. Direct customer requests to the S3 Transfer Acceleration endpoint. Continue to use S3 signed URLs for access control.
- B. Deploy an Amazon CloudFront distribution with the existing S3 bucket as the origin. Direct customer requests to the CloudFront URL. Switch to CloudFront signed URLs for access control.
- C. Set up a second S3 bucket in the eu-central-1 Region with S3 Cross-Region Replication between the buckets. Direct customer requests to the closest Region. Continue to use S3 signed URLs for access control.
- D. Modify the web application to enable streaming of the datasets to end users. Configure the web application to read the data from the existing S3 bucket. Implement access control directly in the application.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

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

QUESTION 557

A solutions architect needs to design a system to store client case files. The files are core company assets and are important. The number of files will grow over time. The files must be simultaneously accessible from multiple application servers that run on Amazon EC2 instances. The solution must have built-in redundancy. Which solution meets these requirements?

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

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Reference: <https://aws.amazon.com/efs/faq/>

QUESTION 558

A company runs a three-tier web application in a VPC across multiple Availability Zones. Amazon EC2 instances run in an Auto Scaling group for the application tier. The company needs to make an automated scaling plan that will analyze each resource's daily and weekly historical workload trends. The configuration must scale resources appropriately according to both the forecast and live changes in

utilization.

Which scaling strategy should a solutions architect recommend to meet these requirements?

- A. Implement dynamic scaling with step scaling based on average CPU utilization from the EC2 instances.
- B. Enable predictive scaling to forecast and scale. Configure dynamic scaling with target tracking.
- C. Create an automated scheduled scaling action based on the traffic patterns of the web application.
- D. Set up a simple scaling policy. Increase the cooldown period based on the EC2 instance startup time.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Reference: <https://aws.amazon.com/blogs/aws/new-predictive-scaling-for-ec2-powered-by-machine-learning/>

QUESTION 559

A solutions architect is designing a shared storage solution for a web application that is deployed across multiple Availability Zones. The web application runs on Amazon EC2 instances that are in an Auto Scaling group. The company plans to make frequent changes to the content. The solution must have strong consistency in returning the new content as soon as the changes occur.

Which solutions meet these requirements? (Choose two.)

- A. Use AWS Storage Gateway Volume Gateway Internet Small Computer Systems Interface (iSCSI) block storage that is mounted to the individual EC2 instances.
- B. Create an Amazon Elastic File System (Amazon EFS) file system. Mount the EFS file system on the individual EC2 instances.
- C. Create a shared Amazon Elastic Block Store (Amazon EBS) volume. Mount the EBS volume on the individual EC2 instances.
- D. Use AWS DataSync to perform continuous synchronization of data between EC2 hosts in the Auto Scaling group.
- E. Create an Amazon S3 bucket to store the web content. Set the metadata for the Cache-Control header to no-cache. Use Amazon CloudFront to deliver the content.

Correct Answer: AB

Section: (none)

Explanation

Explanation/Reference:

Reference:

<https://docs.aws.amazon.com/storagegateway/latest/userguide/WhatIsStorageGateway.html>

<https://docs.aws.amazon.com/efs/latest/ug/how-it-works.html>



QUESTION 560

A company receives data from millions of users totaling about 1 each day. The company provides its users with usage reports going back 12 months. All usage data must be stored for at least 5 years to comply with regulatory and auditing requirements.

Which storage solution is MOST cost-effective?

- A. Store the data in Amazon S3 Standard. Set a lifecycle rule to transition the data to S3 Glacier Deep Archive after 1 year. Set a lifecycle rule to delete the data after 5 years.
- B. Store the data in Amazon S3 One Zone-Infrequent Access (S3 One Zone-IA). Set a lifecycle rule to transition the data to S3 Glacier after 1 year. Set the lifecycle rule to delete the data after 5 years.
- C. Store the data in Amazon S3 Standard. Set a lifecycle rule to transition the data to S3 Standard- Infrequent Access (S3 Standard-IA) after 1 year. Set a lifecycle rule to delete the data after 5 years.
- D. Store the data in Amazon S3 Standard. Set a lifecycle rule to transition the data to S3 One Zone- Infrequent Access (S3 One Zone-IA) after 1 year. Set a lifecycle rule to delete the data after 5 years.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

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

QUESTION 561

A company has a legacy data processing application that runs on Amazon EC2 instances. Data is processed sequentially, but the order of results does not matter. The application uses a monolithic architecture. The only way that the company can scale the application to meet increased demand is to increase the size of the instances.

The company's developers have decided to rewrite the application to use a microservices architecture on Amazon Elastic Container Service (Amazon ECS).

What should a solutions architect recommend for communication between the microservices?

- A. Create an Amazon Simple Queue Service (Amazon SQS) queue. Add code to the data producers, and send data to the queue. Add code to the data consumers to process data from the queue.

- B. Create an Amazon Simple Notification Service (Amazon SNS) topic. Add code to the data producers, and publish notifications to the topic. Add code to the data consumers to subscribe to the topic.
- C. Create an AWS Lambda function to pass messages. Add code to the data producers to call the Lambda function with a data object. Add code to the data consumers to receive a data object that is passed from the Lambda function.
- D. Create an Amazon DynamoDB table. Enable DynamoDB Streams. Add code to the data producers to insert data into the table. Add code to the data consumers to use the DynamoDB Streams API to detect new table entries and retrieve the data.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Reference: <https://aws.amazon.com/blogs/database/dynamodb-streams-use-cases-and-designpatterns/>

QUESTION 562

A solutions architect is creating an application. The application will run on Amazon EC2 instances in private subnets across multiple Availability Zones in a VPC. The EC2 instances will frequently access large files that contain confidential information. These files are stored in Amazon S3 buckets for processing. The solutions architect must optimize the network architecture to minimize data transfer costs.

What should the solutions architect do to meet these requirements?

- A. Create a gateway endpoint for Amazon S3 in the VPC. In the route tables for the private subnets, add an entry for the gateway endpoint.
- B. Create a single NAT gateway in a public subnet. In the route tables for the private subnets, add a default route that points to the NAT gateway.
- C. Create an AWS PrivateLink interface endpoint for Amazon S3 in the VPC. In the route tables for the private subnets, add an entry for the interface endpoint.
- D. Create one NAT gateway for each Availability Zone in public subnets. In each of the route tables for the private subnets, add a default route that points to the NAT gateway in the same Availability Zone.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

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

QUESTION 563

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

Which solution meets these requirements?

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

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

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

QUESTION 564

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

Which solution will meet these requirements?

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

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 565

A solutions architect is designing a workload that will store hourly energy consumption by business tenants in a building. The sensors will feed a database through HTTP requests that will add up usage for each tenant. The solutions architect must use managed services when possible. The workload will receive more features in the future as the solutions architect adds independent components.

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

- A. Use Amazon API Gateway with AWS Lambda functions to receive the data from the sensors, process the data, and store the data in an Amazon DynamoDB table.
- B. Use an Elastic Load Balancer that is supported by an Auto Scaling group of Amazon EC2 instances to receive and process the data from the sensors. Use an Amazon S3 bucket to store the processed data.
- C. Use Amazon API Gateway with AWS Lambda functions to receive the data from the sensors, process the data, and store the data in a Microsoft SQL Server Express database on an Amazon EC2 instance.
- D. Use an Elastic Load Balancer that is supported by an Auto Scaling group of Amazon EC2 instances to receive and process the data from the sensors. Use an Amazon Elastic File System (Amazon EFS) shared file system to store the processed data.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Reference: <https://aws.amazon.com/about-aws/whats-new/2021/>

QUESTION 566

A company has deployed a database in Amazon RDS for MySQL. Due to increased transactions, the database support team is reporting slow reads against the DB instance and recommends adding a read replica.

Which combination of actions should a solutions architect take before implementing this change?

(Choose two.)

- A. Enable binlog replication on the RDS primary node.
- B. Choose a failover priority for the source DB instance.
- C. Allow long-running transactions to complete on the source DB instance.
- D. Create a global table and specify the AWS Regions where the table will be available.
- E. Enable automatic backups on the source instance by setting the backup retention period to a value other than 0.



Correct Answer: CD

Section: (none)

Explanation

Explanation/Reference:

Reference: <https://hevo.com/learn/aws-rds-postgres-replication/>

<https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/GlobalTables.html>

QUESTION 567

A solutions architect is migrating a document management workload to AWS. The workload keeps 7 TiB of contract documents on a shared storage file system and tracks them on an external database. Most of the documents are stored and retrieved eventually for reference in the future.

The application cannot be modified during the migration, and the storage solution must be highly available.

Documents are retrieved and stored by web servers that run on Amazon EC2 instances in an Auto Scaling group. The Auto Scaling group can have up to 12 instances.

Which solution meets these requirements MOST cost-effectively?

- A. Provision an enhanced networking optimized EC2 instance to serve as a shared NFS storage system.
- B. Create an Amazon S3 bucket that uses the S3 Standard-Infrequent Access (S3 Standard-IA) storage class. Mount the S3 bucket to the EC2 instances in the Auto Scaling group.
- C. Create an SFTP server endpoint by using AWS Transfer for SFTP and an Amazon S3 bucket.
Configure the EC2 instances in the Auto Scaling group to connect to the SFTP server.
- D. Create an Amazon Elastic File System (Amazon EFS) file system that uses the EFS Standard- Infrequent Access (EFS Standard-IA) storage class. Mount the file system to the EC2 instances in the Auto Scaling group.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Reference: <https://docs.aws.amazon.com/prescriptive-guidance/latest/patterns/migrate-an-onpremises-sftp-server-to-aws-using-aws-transfer-for-sftp.html>

QUESTION 568

A company wants to relocate its on-premises MySQL database to AWS. The database accepts regular imports from a client-facing application, which causes a high volume of write operations. The company is concerned that the amount of traffic might be causing performance issues within the application. How should a solutions architect design the architecture on AWS?

- A. Provision an Amazon RDS for MySQL DB instance with Provisioned IOPS SSD storage. Monitor write operation metrics by using Amazon CloudWatch. Adjust the provisioned IOPS if necessary.
- B. Provision an Amazon RDS for MySQL DB instance with General Purpose SSD storage. Place an Amazon ElastiCache cluster in front of the DB instance. Configure the application to query ElastiCache instead.
- C. Provision an Amazon DocumentDB (with MongoDB compatibility) instance with a memory optimized instance type. Monitor Amazon CloudWatch for performance-related issues. Change the instance class if necessary.
- D. Provision an Amazon Elastic File System (Amazon EFS) file system in General Purpose performance mode. Monitor Amazon CloudWatch for IOPS bottlenecks. Change to Provisioned Throughput performance mode if necessary.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Reference: <https://docs.aws.amazon.com/documentdb/latest/developerguide/db-instanceclasses.html>

QUESTION 569

A company has a remote factory that has unreliable connectivity. The factory needs to gather and process machine data and sensor data so that it can sense products on its conveyor belts and initiate a robotic movement to direct the products to the right location. Predictable low-latency compute processing is essential for the on-premises control systems. Which solution should the factory use to process the data?

- A. Amazon CloudFront Lambda@Edge functions
- B. An Amazon EC2 instance that has enhanced networking enabled
- C. An Amazon EC2 instance that uses an AWS Global Accelerator
- D. An Amazon Elastic Block Store (Amazon EBS) volume on an AWS Snowball Edge cluster

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 570

A company is creating a prototype of an ecommerce website on AWS. The website consists of an Application Load Balancer, an Auto Scaling group of Amazon EC2 instances for web servers, and an Amazon RDS for MySQL DB instance that runs with the Single-AZ configuration.

The website is slow to respond during searches of the product catalog. The product catalog is a group of tables in the MySQL database that the company does not update frequently. A solutions architect has determined that the CPU utilization on the DB instance is high when product catalog searches occur.

What should the solutions architect recommend to improve the performance of the website during searches of the product catalog?

- A. Migrate the product catalog to an Amazon Redshift database. Use the COPY command to load the product catalog tables.
- B. Implement an Amazon ElastiCache for Redis cluster to cache the product catalog. Use lazy loading to populate the cache.
- C. Add an additional scaling policy to the Auto Scaling group to launch additional EC2 instances when database response is slow.
- D. Turn on the Multi-AZ configuration for the DB instance. Configure the EC2 instances to throttle the product catalog queries that are sent to the database.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:**QUESTION 571**

A company wants to use AWS Systems Manager to manage a fleet of Amazon EC2 instances.

According to the company's security requirements, no EC2 instances can have internet access. A solutions architect needs to design network connectivity from the EC2 instances to Systems Manager while fulfilling this security obligation.

Which solution will meet these requirements?



- A. Deploy the EC2 instances into a private subnet with no route to the internet.
- B. Configure an interface VPC endpoint for Systems Manager. Update routes to use the endpoint.
- C. Deploy a NAT gateway into a public subnet. Configure private subnets with a default route to the NAT gateway.
- D. Deploy an internet gateway. Configure a network ACL to deny traffic to all destinations except Systems Manager.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

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

QUESTION 572

A development team is creating an event-based application that uses AWS Lambda functions.

Events will be generated when files are added to an Amazon S3 bucket. The development team currently has Amazon Simple Notification Service (Amazon SNS) configured as the event target form Amazon S3.

What should a solutions architect do to process the events form Amazon S3 in a scalable way?

- A. Create an SNS subscription that processes the event in Amazon Elastic Container Service (Amazon ECS) before the event runs in Lambda.
- B. Create an SNS subscription that processes the event in Amazon Elastic Kubernetes Service (Amazon EKS) before the event runs in Lambda.
- C. Create an SNS subscription that sends the event to Amazon Simple Queue Service (Amazon SQS). Configure the SQS queue to trigger a Lambda function.
- D. Create an SNS subscription that sends the event to AWS Server Migration Service (AWS SMS). Configure the Lambda function to poll from the SMS event.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Reference:

<https://docs.aws.amazon.com/AWSSimpleQueueService/latest/SQSDeveloperGuide/sqsconfigure-subscribe-queue-sns-topic.html>



QUESTION 573

A company is running a multi-tier ecommerce web application in the AWS Cloud. The web application is running on Amazon EC2 instances. The database tier is on a provisioned Amazon Aurora MySQL DB cluster with a writer and a reader in a Multi-AZ environment. The new requirement for the database tier is to serve the application to achieve continuous write availability through an instance failover.

What should a solutions architect do to meet this new requirement?

- A. Add a new AWS Region to the DB cluster for multiple writes.
- B. Add a new reader in the same Availability Zone as the writer.
- C. Migrate the database tier to an Aurora multi-master cluster.
- D. Migrate the database tier to an Aurora DB cluster with parallel query enabled.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Reference: <https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/aurora-multimaster.html>

QUESTION 574

A company uses AWS to run all components of its three-tier application. The company wants to automatically detect any potential security breaches within the environment. The company wants to track any findings and notify administrators if a potential breach occurs.

Which solution meets these requirements?

- A. Set up AWS WAF to evaluate suspicious web traffic. Create AWS Lambda functions to log any findings in Amazon CloudWatch and send email notifications to administrators.
- B. Set up AWS Shield to evaluate suspicious web traffic. Create AWS Lambda functions to log any findings in Amazon CloudWatch and send email notifications to administrators.
- C. Deploy Amazon Inspector to monitor the environment and generate findings in Amazon CloudWatch. Configure an Amazon EventBridge (Amazon CloudWatch Events) rule to publish a message to an Amazon Simple Notification Service

(Amazon SNS) topic to notify administrators by email.

D. Deploy Amazon GuardDuty to monitor the environment and generate findings in Amazon CloudWatch. Configure an Amazon EventBridge (Amazon CloudWatch Events) rule to publish a message to an Amazon Simple Notification Service (Amazon SNS) topic to notify administrators by email.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Reference: <https://docs.aws.amazon.com/sns/latest/dg/sns-event-sources.html>

QUESTION 575

A company hosts its enterprise content management platform in one AWS Region but needs to operate the platform across multiple Regions. The company has an Amazon Elastic Kubernetes Service (Amazon EKS) cluster that runs its microservices. The EKS cluster stores and retrieves objects from Amazon S3. The EKS cluster also stores and retrieves metadata from Amazon DynamoDB.

Which combination of steps should a solutions architect take to deploy the platform across multiple Regions? (Choose two.)

- A. Replicate the EKS cluster with cross-Region replication.
- B. Use Amazon API Gateway to create a global endpoint to the EKS cluster.
- C. Use AWS Global Accelerator endpoints to distribute the traffic to multiple Regions.
- D. Use Amazon S3 access points to give access to the objects across multiple Regions. Configure DynamoDB Accelerator (DAX). Connect DAX to the relevant tables.
- E. Deploy an EKS cluster and an S3 bucket in another Region. Configure cross-Region replication on both S3 buckets. Turn on global tables for DynamoDB.

Correct Answer: BE

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 576

A company is running a multi-tier ecommerce web application in the AWS Cloud. The application runs on Amazon EC2 Instances with an Amazon RDS MySQL Multi-AZ DB instance. Amazon RDS is configured with the latest generation instance with 2,000 GB of storage in a General Purpose SSD (gp3) Amazon Elastic Block Store (Amazon EBS) volume. The database performance affects the application during periods of high demand.

A database administrator analyzes the logs in Amazon CloudWatch Logs and discovers that the application performance always degrades when the number of read and write IOPS is higher than 20,000.

What should a solutions architect do to improve the application performance?

- A. Replace the volume with a magnetic volume.
- B. Increase the number of IOPS on the gp3 volume.
- C. Replace the volume with a Provisioned IOPS SSD (io2) volume.
- D. Replace the 2,000 GB volume with two 1,000 GB gp3 volumes.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 577

A company has an application that provides marketing services to stores. The services 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. Use Amazon Inspector to scan the objects in the bucket. If objects contain PII, trigger an S3 Lifecycle policy to remove the objects that contain PII.
- B. Use an Amazon S3 bucket as a secure transfer point. Use Amazon Macie to scan the objects in the bucket. If objects contain PII, use Amazon Simple Notification Service (Amazon SNS) to trigger a notification to the administrators to remove the objects that contain PII.
- C. Implement custom scanning algorithms in an AWS Lambda function. Trigger the function when objects are loaded into the bucket. If objects contain PII, use Amazon Simple Notification Service (Amazon SNS) to trigger a notification to the administrators to remove the objects that contain PII.

D. Implement custom scanning algorithms in an AWS Lambda function. Trigger the function when objects are loaded into the bucket. If objects contain PII, use Amazon Simple Email Service (Amazon SES) to trigger a notification to the administrators and trigger an S3 Lifecycle policy to remove the objects that contain PII.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 578

A company has a three-tier image-sharing application. The application uses an Amazon EC2 instance for the front-end layer, another EC2 instance for the application layer, and a third EC2 instance for a MySQL database. A solutions architect must design a scalable and highly available solution that requires the least amount of change to the application.

Which solution meets these requirements?

A. Use Amazon S3 to host the front-end layer and AWS Lambda functions for the application layer.

Move the database to an Amazon DynamoDB table and use Amazon S3 to store and serve users' images.

B. Use load-balanced Multi-AZ AWS Elastic Beanstalk environments for the front-end and the application layer. Move the database to an Amazon RDS instance with multiple read replicas to serve users' images.

C. Use Amazon S3 to host the front-end layer and a fleet of Amazon EC2 instances in an Auto Scaling group for the application layer. Move the database to a memory optimized instance type to store and serve users' images.

D. Use load-balanced Multi-AZ AWS Elastic Beanstalk environments for the front-end layer and the application layer. Move the database to an Amazon RDS Multi-AZ DB instance. Use Amazon S3 to store and serve users' images.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 579

A company has a web application that runs on Amazon EC2 instances. The company wants end users to authenticate themselves before they use the web application. The web application accesses AWS resources, such as Amazon S3 buckets, on behalf of users who are logged on.

Which combination of actions must a solutions architect take to meet these requirements?

(Choose two.)

A. Configure AWS App Mesh to log on users.

B. Enable and configure AWS Single Sign-On in AWS Identity and Access Management (IAM).

C. Define a default IAM role for authenticated users.

D. Use AWS Identity and Access Management (IAM) for user authentication.

E. Use Amazon Cognito for user authentication.

Correct Answer: BD

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 580

A company is building an ecommerce 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. Use EBS encryption to encrypt the data. Use an IAM instance role to restrict access.

B. Store sensitive data in Amazon RDS for MySQL. Use AWS Key Management Service (AWS KMS) client-side encryption to encrypt the data.

C. Store sensitive data in Amazon S3. Use AWS Key Management Service (AWS KMS) server-side encryption to encrypt the data. Use S3 bucket policies to restrict access.

D. Store sensitive data in Amazon FSx for Windows Server. Mount the file share on application servers. Use Windows file permissions to restrict access.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 581

A company has a production web application in which users upload documents through a web interface or a mobile app. According to a new regulatory requirement, new documents cannot be modified or deleted after they are stored. What should a solutions architect do to meet this requirement?

- A. Store the uploaded documents in an Amazon S3 bucket with S3 Versioning and S3 Object Lock enabled.
- B. Store the uploaded documents in an Amazon S3 bucket. Configure an S3 Lifecycle policy to archive the documents periodically.
- C. Store the uploaded documents in an Amazon S3 bucket with S3 Versioning enabled. Configure an ACL to restrict all access to read-only.
- D. Store the uploaded documents on an Amazon Elastic File System (Amazon EFS) volume. Access the data by mounting the volume in read-only mode.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 582

A company needs to migrate a large amount of data from an on-premises storage area network (SAN) to Amazon S3. The SAN currently has 200 TB of data and is receiving an additional 20 TB of data each month. The company has a 500 Mbps connection to the internet. What should the company do to migrate the data to Amazon S3 in the LEAST amount of time?

- A. Use a file syncing application to sync the data to Amazon S3 over the internet through a public S3 endpoint. Sync any changed data the same way until the SAN is decommissioned.
- B. Use an AWS Snowball Edge Storage Optimized device to migrate the initial 200 TB of data to Amazon S3. Sync any changes data by using AWS DataSync until the SAN is decommissioned.
- C. Set up an AWS Site-to-Site VPN connection. Use a file syncing application to sync the data to Amazon S3 through a private S3 endpoint. Sync any changed data the same way until the SAN is decommissioned.
- D. Set up a 10 Gbps AWS Direct Connection connection. Migrate the initial 200 TB of data to Amazon S3 by using a file syncing application. Sync any changed data the same way until the SAN is decommissioned.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 583

A company is planning to store data on Amazon RDS DB instances. The company must encrypt the data at rest. What should a solutions architect do to meet this requirement?

- A. Create an encryption key, and store the key in AWS Secrets Manager. Use the key to encrypt the DB instances.
- B. Generate a certificate in AWS Certificate Manager (ACM). Enable SSL/TLS on the DB instances by using the certificate.
- C. Create a customer master key (CMK) in AWS Key Management Service (AWS KMS). Enable encryption for the DB instances.
- D. Generate a certificate in AWS Identity and Access Management (IAM). Enable SSL/TLS on the DB instances by using the certificate.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 584

A solutions architect is using an AWS CloudFormation template to deploy a three-tier web application. The web application consists of a web tier and an application tier that stores and retrieves user data in Amazon DynamoDB tables. The web and application tiers are hosted on Amazon EC2 instances, and the database tier is not publicly accessible. The application EC2 instances need to access the DynamoDB tables without exposing API credentials in the template. What should the solutions architect do to meet these requirements?

- A. Create an IAM role to read the DynamoDB tables. Associate the role with the application instances by referencing an instance profile.

- B. Create an IAM role that has the required permissions to read and write from the DynamoDB tables. Add the role to the EC2 instance profile, and associate the instance profile with the application instances.
- C. Use the parameter section in the AWS CloudFormation template to have the user input access and secret keys from an already-created IAM user that has the required permissions to read and write from the DynamoDB tables.
- D. Create an IAM user in the AWS CloudFormation template that has the required permissions to read and write from the DynamoDB tables. Use the GetAtt function to retrieve the access and secret keys, and pass them to the application instances through the user data.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 585

A solutions architect is designing the architecture for a company website that is composed of static component. The company's target customers are located in the United States and Europe. Which architecture should the solutions architect recommend to MINIMIZE cost?

- A. Store the website files in Amazon S3 in the us-east-2 Region. Use an Amazon CloudFront distribution with the price class configured to limit the edge locations in use.
- B. Store the website files on Amazon S3 in the us-east-2 Region. Use an Amazon CloudFront distribution with the price class configured to maximize the use of edge locations.
- C. Store the website files on Amazon S3 in the us-east-2 Region and the eu-west-1 Region. Use an Amazon CloudFront geolocation routing policy to route requests to the closest Region to the user.
- D. Store the website files on Amazon S3 in the us-east-2 Region and the eu-west-1 Region. Use an Amazon CloudFront distribution with an Amazon Route 53 latency routing policy to route requests to the closest Region to the user.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 586

A company has a business-critical application that runs on Amazon EC2 instances. The application stores data in an Amazon DynamoDB table. The company must be able to revert the table to any point within the last 24 hours. Which solution meets these requirements with the LEAST operational overhead?

- A. Configure point-in-time recovery for the table.
- B. Use AWS Backup for the table.
- C. Use an AWS Lambda function to make an on-demand backup of the table every hour.
- D. Turn on streams on the table to capture a log of all changes to the table in the last 24 hours. Store a copy of the stream in an Amazon S3 bucket.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 587

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

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

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 588

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

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

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 589

A company is deploying an application that processes streaming data in near-real time. The company plans to use Amazon EC2 instances for the workload. The network architecture must be configurable to provide the lowest possible latency between nodes.

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

- A. Place the EC2 instances in multiple VPCs, and configure VPC peering.
- B. Run the EC2 instances in a spread placement group.
- C. Attach an Elastic Fabric Adapter (EFA) to each EC2 instance.
- D. Use Amazon Elastic Block Store (Amazon EBS) optimized instance types.

Correct Answer: BD

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 590

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. Add Route 53 geolocation routing entries.
- B. Provision accelerators in AWS Global Accelerator. Associate the supplied IP addresses with the S3 bucket. Edit the Route 53 entries to point to the IP addresses of the accelerators.
- C. Add an Amazon CloudFront distribution in front of the S3 bucket. Edit the Route 53 entries to point to the CloudFront distribution.
- D. Enable S3 Transfer Acceleration on the bucket. Edit the Route 53 entries to point to the new endpoint.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 591

A company needs to retain application log 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.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 592

A company wants to run applications in containers in the AWS Cloud. These applications are stateless and can tolerate disruptions within the underlying infrastructure. The company needs a solution that minimizes cost and operational overhead.

What should a solutions architect do to meet these requirements?

- A. Use Spot Instances in an Amazon EC2 Auto Scaling group to run the application containers.
- B. Use Spot Instances in an Amazon Elastic Kubernetes Service (Amazon EKS) managed node group.
- C. Use On-Demand Instances in an Amazon EC2 Auto Scaling group to run the application containers.
- D. Use On-Demand Instances in an Amazon Elastic Kubernetes Service (Amazon EKS) managed node group.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 593

A company wants to reduce the cost of its existing three-tier web architecture. The web, application, and database servers are running on Amazon EC2 instances for the development, test, and production environments. The EC2 instances average 30% CPU utilization during peak hours and 10% CPU utilization during non-peak hours. The production EC2 instances run 24 hours a day. The development and test EC2 instances run for at least 8 hours each day. The company plans to implement automation to stop the development and test EC2 instances when they are not in use.

Which EC2 instance purchasing solution will meet the company's requirements MOST costeffectively?

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

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 594

A company must save all the email messages that its employees send to customers for a period of 12 months. The messages are stored in a binary format and vary in size from 1 KB to 20 KB. The company has selected Amazon S3 as the storage service for the messages.

Which combination of steps will meet these requirements MOST cost-effectively? (Choose two.)

- A. Create an S3 bucket policy that denies the s3:DeleteObject action.
- B. Create an S3 Lifecycle configuration that deletes the messages after 12 months.
- C. Upload the messages to Amazon S3. Use S3 Object Lock in governance mode.
- D. Upload the messages to Amazon S3. Use S3 Object Lock in compliance mode.
- E. Use S3 Inventory. Create an AWS Batch job that periodically scans the inventory and deletes the messages after 12 months.

Correct Answer: AC

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 595

A company is implementing a shared storage solution for gaming application that is hosted in an on-premises data center. The company needs to ability to use Lustre clients to access data. The solution must be fully managed. Which solution meets these requirements?

- A. Create an AWS Storage Gateway file gateway. Create a file share that uses the required client protocol. Connect the application server to the file share.
- B. Create an Amazon EC2 Windows instance. Install and configure a Windows file share role on the instance. Connect the application server to the file share.
- C. Create an Amazon Elastic File System (Amazon EFS) file system, and configure it to support Lustre. Attach the file system to the origin server. Connect the application server to the file system.
- D. Create an Amazon FSx for Lustre file system. Attach the file system to the origin server. Connect the application server to the file system.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 596

A company has an application that gives users the ability to upload images to an Amazon S3 bucket. Each night, the company launches an Amazon EC2 Spot Fleet that processes all the images that the company received that day. The code to process the images is small (less than 200 MB), and each image takes less than 5 minutes to process.

The company wants to change the application to process the images when the images are uploaded. However, the company is concerned about the cost of this new functionality.

What should a solutions architect do to implement this change MOST cost-effectively?

- A. Use S3 events to invoke an AWS Lambda function to process the images.
- B. Use S3 events to launch an EC2 Reserved instance to process the images.
- C. Use S3 events to launch a container in AWS Fargate to process the images.
- D. Use S3 events to deploy an AWS Elastic Beanstalk application to process the images.



Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 597

A company has a Windows-based application that must be migrated to AWS. The application requires the use of a shared Windows file system attached to multiple Amazon EC2 Windows instances that are deployed across multiple Availability Zones.

What should a solutions architect do to meet this requirement?

- A. Configure AWS Storage Gateway in volume gateway mode. Mount the volume to each Windows instances.
- B. Configure Amazon FSx for Windows File Server. Mount the Amazon FSx file system to each Windows instance.
- C. Configure a file system by using Amazon Elastic File System (Amazon EFS). Mount the EFS file system to each Windows instance.
- D. Configure an Amazon Elastic Block Store (Amazon EBS) volume with the required size. Attach each EC2 instance to the volume. Mount the file system within the volume to each Windows instance.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 598

A company has one million users that use its mobile app. The company must analyze the data usage in near-real time. The company also must encrypt the data in near-real time and must store the data in a centralized location in Apache Parquet format for further processing.

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

- A. Create an Amazon Kinesis data stream to store the data in Amazon S3. Create an Amazon Kinesis Data Analytics application to analyze the data. Invoke an AWS Lambda function to send the data to the Kinesis Data Analytics application.
- B. Create an Amazon Kinesis data stream to store the data in Amazon S3. Create an Amazon EMR cluster to analyze the data. Invoke an AWS Lambda function to send the data to the EMR cluster.
- C. Create an Amazon Kinesis Data Firehose delivery stream to store the data in Amazon S3. Create an Amazon EMR cluster to analyze the data.
- D. Create an Amazon Kinesis Data Firehose delivery stream to store the data in Amazon S3. Create an Amazon Kinesis Data Analytics application to analyze the data.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 599

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.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:



QUESTION 600

A solutions architect is implementing a complex Java application with a MySQL database. The Java application must be deployed on Apache Tomcat and must be highly available.

What should the solutions architect do to meet these requirements?

- A. Deploy the application in AWS Lambda. Configure an Amazon API Gateway API to connect with the Lambda functions.
- B. Deploy the application by using AWS Elastic Beanstalk. Configure a load-balanced environment and a rolling deployment policy.
- C. Migrate the database to Amazon ElastiCache. Configure the ElastiCache security group to allow access from the application.
- D. Launch an Amazon EC2 instance. Install a MySQL server on the EC2 instance. Configure the application on the server. Create an AMI. Use the AMI to create a launch template with an Auto Scaling group.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 601

A company hosts a web-based application on AWS in a departmental level VPC. The company has used AWS CloudFormation stacks to provision all the required resources in the AWS Cloud.

The company needs to migrate this application to the company level VPC. The application and network components need to be managed as separate entities. Before the company implements changes, the company wants to understand how the changes will affect resources that are already running.

Which combination of actions should a solutions architect take to meet these requirements?

(Choose two.)

- A. Perform drift detection on the CloudFormation stacks to track the changes.
- B. Enable AWS Config to track the changes to the CloudFormation stacks.
- C. Create change sets before updating the CloudFormation stacks.
- D. Create a CloudFormation root stack that has two nested stacks: one stack for the application layer and one stack for the network layer.
- E. Create a CloudFormation cross-stack reference. Export company level VPC stack outputs. Import the stack outputs into the application stack.

Correct Answer: CE
Section: (none)
Explanation

Explanation/Reference:
Explanation:

QUESTION 602

A company hosts a website analytics application on a single Amazon EC2 On-Demand Instance.

The analytics software is written in PHP and uses a MySQL database. The analytics software, the web server that provides PHP, and the database server are all hosted on the EC2 instance. The application is showing signs of performance degradation during busy times and is presenting 5xx errors. The company needs to make the application scale seamlessly.

Which solution will meet these requirements MOST cost-effectively?

- A. Migrate the database to an Amazon RDS for MySQL DB instance. Create an AMI of the web application. Use the AMI to launch a second EC2 On-Demand Instance. Use an Application Load Balancer to distribute the load to each EC2 instance.
- B. Migrate the database to an Amazon RDS for MySQL DB instance. Create an AMI of the web application. Use the AMI to launch a second EC2 On-Demand Instance. Use Amazon Route 53 weighted routing to distribute the load across the two EC2 instances.
- C. Migrate the database to an Amazon Aurora MySQL DB instance. Create an AWS Lambda function to stop the EC2 instance and change the instance type. Create an Amazon CloudWatch alarm to invoke the Lambda function when CPU utilization surpasses 75%.
- D. Migrate the database to an Amazon Aurora MySQL DB instance. Create an AMI of the web application. Apply the AMI to launch template. Create an Auto Scaling group with the launch template. Configure the launch template to use a Spot Fleet. Attach an Application Load Balancer to the Auto Scaling group.

Correct Answer: AD
Section: (none)
Explanation

Explanation/Reference:
Explanation:

QUESTION 603

A company hosts its web applications in the AWS Cloud. The company configures Elastic Load Balancers to use certificates that are imported into AWS Certificate Manager (ACM). The company's security team must be notified 30 days before the expiration of each certificate.

What should a solutions architect recommend to meet this requirement?

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

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Explanation:

QUESTION 604

A company collects data for temperature, humidity, and atmospheric pressure in cities across multiple continents. The average volume of data that the company collects from each site daily is 500 GB. Each site has a high-speed internet connection.

The company wants to aggregate the data from all these global sites as quickly as possible in a single Amazon S3 bucket. The solution must minimize operational complexity.

Which solution meets these requirements?

- A. Turn on S3 Transfer Acceleration on the destination S3 bucket. Use multipart uploads to directly upload site data to the destination S3 bucket.
- B. Upload the data from each site to an S3 bucket in the closest Region. Use S3 Cross-Region Replication to copy objects to the destination S3 bucket. Then remove the data from the origin S3 bucket.

- C. Schedule AWS Snowball Edge Storage Optimized device jobs daily to transfer data from each site to the closest Region. Use S3 Cross-Region Replication to copy objects to the destination S3 bucket.
- D. Upload the data from each site to an Amazon EC2 instance in the closest Region. Store the data in an Amazon Elastic Block Store (Amazon EBS) volume. At regular intervals, take an EBS snapshot and copy it to the Region that contains the destination S3 bucket. Restore the EBS volume in that Region.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 605

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 Logs. Export the logs to the S3 bucket.
- C. Create an instance profile on Amazon EC2 to allow S3 access.
- D. Create an Amazon API Gateway API with a private link to access the S3 endpoint.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 606

A company has a customer relationship management (CRM) application that stores data in an Amazon RDS DB instance that runs Microsoft SQL Server. The company's IT staff has administrative access to the database. The database contains sensitive data. The company wants to ensure that the data is not accessible to the IT staff and that only authorized personnel can view the data. What should a solutions architect do to secure the data?

- A. Use client-side encryption with an Amazon RDS managed key.
- B. Use client-side encryption with an AWS Key Management Service (AWS KMS) customer managed key.
- C. Use Amazon RDS encryption with an AWS Key Management Service (AWS KMS) default encryption key.
- D. Use AWS Secrets Manager to manage database users. Encrypt secrets with an AWS Key Management Service (AWS KMS) customer managed key. Enable RDS encryption.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 607

A company has multiple AWS accounts that are configured in an AWS Organizations organizational unit (OU). The company wants users to be able to create Amazon S3 buckets in the us-east-2 Region only. The FullAWSAccess service control policy (SCP) is attached to the OU. Which SCP should a solutions architect use to meet these requirements?

A.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": "s3:CreateBucket",
      "Resource": "arn:aws:s3:*:*:*",
      "Condition": {
        "StringLike": {
          "s3:LocationConstraint": "us-east-2"
        }
      }
    }
  ]
}
```

B.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Deny",
      "Action": "s3:CreateBucket",
      "Resource": "arn:aws:s3:*:*:*",
      "Condition": {
        "StringNotLike": {
          "s3:LocationConstraint": "us-east-2"
        }
      }
    }
  ]
}
```

udumps

C.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Deny",
      "Action": "s3:CreateBucket",
      "Resource": "arn:aws:s3:*:*:*",
      "Condition": {
        "StringLike": {
          "aws:RequestedRegion": "us-east-2"
        }
      }
    }
  ]
}
```

D.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "NotAction": "s3:CreateBucket",
      "Resource": "arn:aws:s3:*:*:*",
      "Condition": {
        "StringLike": {
          "aws:RequestedRegion": "us-east-2"
        }
      }
    }
  ]
}
```



Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:

QUESTION 608

A startup company is hosting a website for its customers on an Amazon EC2 instance. The website consists of a stateless Python application and a MySQL database. The website servers only a small amount of traffic. The company is concerned about the reliability of the instance and needs to migrate to a highly available architecture. The company cannot modify the application code. Which combination of actions should a solutions architect take to achieve high availability for the website?

- A. Provision an internet gateway in each Availability Zone in use.
- B. Migrate the database to an Amazon RDS for MySQL Multi-AZ DB instance.
- C. Migrate the database to Amazon DynamoDB, and enable DynamoDB auto scaling.
- D. Use AWS DataSync to synchronize the database data across multiple EC2 instances.
- E. Create an Application Load Balancer to distribute traffic to an Auto Scaling group of EC2 instances that are distributed across two Availability Zones.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 609

A company runs an application on Amazon EC2 instances that are part of an Auto Scaling group.

Traffic to the application increases substantially during business hours. A solutions architect needs to implement an Auto Scaling policy that addresses user latency concerns during periods of high traffic. The company does not want to provision more compute than is necessary.

What should the solutions architect do to meet these requirements?

- A. Configure a predictive scaling policy with the appropriate scaling metric.
- B. Configure a dynamic target tracking scaling policy with the appropriate scaling metric.
- C. Configure a scheduled scaling policy that launches additional EC2 instances during business hours.
- D. Configure dynamic step or simple policies with Amazon CloudWatch alarms to add and remove EC2 instances based on alarm status.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:



QUESTION 610

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. Configure EC2 Auto Scaling to use scheduled scaling.
- B. 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.
- C. Implement the primary server and the compute nodes with Amazon EC2 instances that are managed in an Auto Scaling group. Configure AWS CloudTrail as a destination for the jobs. Configure EC2 Auto Scaling based on the load on the primary server.
- D. 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.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 611

A company's web application consists of multiple Amazon EC2 instances that run behind an Application Load Balancer in a VPC. An Amazon RDS for MySQL DB instance contains the data.

The company needs the ability to automatically detect and respond to suspicious or unexpected behavior in its AWS environment. The company already has added AWS WAF to its architecture.

What should a solutions architect do next to protect against threats?

- A. Use Amazon GuardDuty to perform threat detection. Configure Amazon EventBridge (Amazon CloudWatch Events) to filter for GuardDuty findings and to invoke an AWS Lambda function to adjust the AWS WAF rules.
- B. Use AWS Firewall Manager to perform threat detection. Configure Amazon EventBridge (Amazon CloudWatch Events) to filter for Firewall Manager findings and to invoke an AWS Lambda function to adjust the AWS WAF web ACL.

- C. Use Amazon Inspector to perform threat detection and to update the AWS WAF rules. Create a VPC network ACL to limit access to the web application.
- D. Use Amazon Macie to perform threat detection and to update the AWS WAF rules. Create a VPC network ACL to limit access to the web application.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 612

A company runs a web-based portal that provides users with global breaking news, local alerts, and weather updates. The portal delivers each user a personalized view by using a mixture of static and dynamic content. Content is served over HTTPS through an API server running on an Amazon EC2 instance behind an Application Load Balancer (ALB). The company wants the portal to provide this content to its users across the world as quickly as possible. How should a solutions architect design the application to ensure the LEAST amount of latency for all users?

- A. Deploy the application stack in a single AWS Region. Use Amazon CloudFront to serve all static and dynamic content by specifying the ALB as an origin.
- B. Deploy the application stack in two AWS Regions. Use an Amazon Route 53 latency routing policy to serve all content from the ALB in the closest Region.
- C. Deploy the application stack in a single AWS Region. Use Amazon CloudFront to serve the static content. Serve the dynamic content directly from the ALB.
- D. Deploy the application stack in two AWS Regions. Use an Amazon Route 53 geolocation routing policy to serve all content from the ALB in the closest Region.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 613

A company has an application that loads documents into an Amazon S3 bucket and converts the documents into another format. The application stores the converted documents in another S3 bucket and saves the document name and URLs in an Amazon DynamoDB table. The DynamoDB entries are used during subsequent days to access the documents. The company uses a DynamoDB Accelerator (DAX) cluster in front of the table. Recently, traffic to the application has increased. Document processing tasks are timing out during the scheduled DAX maintenance window. A solutions architect must ensure that the documents continue to load during the maintenance window.

What should the solutions architect do to accomplish this goal?

- A. Modify the application to write to the DAX cluster. Configure the DAX cluster to write to the DynamoDB table when the maintenance window is complete.
- B. Enable Amazon DynamoDB Streams for the DynamoDB table. Modify the application to write to the stream. Configure the stream to load the data when the maintenance window is complete.
- C. Convert the application to an AWS Lambda function. Configure the Lambda function runtime to be longer than the maintenance window. Create an Amazon CloudWatch alarm to monitor Lambda timeouts.
- D. Modify the application to write the document name and URLs to an Amazon Simple Queue Service (Amazon SQS) queue. Create an AWS Lambda function to read the SQS queue and write to DynamoDB.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 614

A company is designing a distributed application to optimize its global supply chain and manufacturing process. The company has facilities near the us-east-1 Region, the eu-west-1 Region, and the ap-south-1 Region.

According to the application requirements, orders that are booked in one Region must be visible in the other two Regions in 1 second or less. The database must be able to support failover with a recovery time objective (RTO) of less than 5 minutes. The application must avoid downtime so that the manufacturing process is not negatively affected.

Which solution meets these requirements?

- A. Use Amazon DynamoDB to invoke an AWS Lambda function.
- B. Use an Amazon Aurora global database.
- C. Use Amazon RDS for MySQL with a cross-Region read replica.
- D. Use Amazon RDS for PostgreSQL with a cross-Region read replica.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 615

A company is developing an internal application that uses a PostgreSQL database. The company decides to host the database on Amazon Aurora. The application does not need to be highly available, but data must be stored in multiple Availability Zones.

Which database configuration meets these requirements MOST cost-effectively?

- A. An Aurora PostgreSQL global database cluster
- B. An Aurora PostgreSQL DB cluster that has a single DB instance
- C. An Aurora PostgreSQL DB cluster that has a primary DB instance and a read replica
- D. Two Aurora PostgreSQL DB clusters, with each DB cluster in a separate Availability Zone

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 616

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. Write orders to Amazon Simple Notification Service (Amazon SNS). Subscribe the database endpoint to the SNS topic.
- B. 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.
- C. Write orders to Amazon Simple Notification Service (Amazon SNS). Subscribe the database endpoint to the SNS topic. Use EC2 instances in an Auto Scaling group behind an Application Load Balancer to read from the SNS topic.
- D. Write orders to an Amazon Simple Queue Service (Amazon SQS) queue when the EC2 instance reaches CPU threshold limits. 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.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 617

A company runs a critical, customer-facing application on Amazon Elastic Kubernetes Service

(Amazon EKS). The application has a microservices architecture. The company needs to implement a solution that collects, aggregates, and summarizes metrics and logs from the application in a centralized location.

Which solution meets these requirements?

- A. Run the Amazon CloudWatch agent in the existing EKS cluster. View the metrics and logs in the CloudWatch console.
- B. Run AWS App Mesh in the existing EKS cluster. View the metrics and logs in the App Mesh console.
- C. Configure AWS CloudTrail to capture data events. Query CloudTrail by using Amazon OpenSearch Service (Amazon Elasticsearch Service).
- D. Configure Amazon CloudWatch Container Insights in the existing EKS cluster. View the metrics and logs in the CloudWatch console.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 618

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).

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 619

A company's infrastructure consists of hundreds of Amazon EC2 instances that use Amazon Elastic Block Store (Amazon EBS) storage. A solutions architect must ensure that every EC2 instance can be recovered after a disaster. What should the solutions architect do to meet this requirement with the LEAST amount of effort?

- A. Take a snapshot of the EBS storage that is attached to each EC2 instance. Create an AWS CloudFormation template to launch new EC2 instances from the EBS storage.
- B. Take a snapshot of the EBS storage that is attached to each EC2 instance. Create AWS Elastic Beanstalk to set the environment based on the EC2 template and attach the EBS storage.
- C. Use AWS Backup to set up a backup plan for the entire group of EC2 instances. Use the AWS Backup API or the AWS CLI to speed up the restore process for multiple EC2 instances.
- D. Create an AWS Lambda function to take a snapshot of the EBS storage that is attached to each EC2 instance and copy the Amazon Machine Images (AMIs). Create another Lambda function to perform the restores with the copied AMIs and attach the EBS storage.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:



QUESTION 620

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

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 621

A company runs a photo processing application that needs to frequently upload and download pictures from Amazon S3 buckets that are located in the same AWS Region. A solutions architect has noticed an increased cost in data transfer fees and needs to implement a solution to reduce these costs. How can the solutions architect meet this requirement?

- A. Deploy Amazon API Gateway into a public subnet and adjust the route table to route S3 calls through it.
- B. Deploy a NAT gateway into a public subnet and attach an endpoint policy that allows access to the S3 buckets.
- C. Deploy the application into a public subnet and allow it to route through an internet gateway to access the S3 buckets.

D. Deploy an S3 VPC gateway endpoint into the VPC and attach an endpoint policy that allows access to the S3 buckets.

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Explanation:

QUESTION 622

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

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

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:

QUESTION 623

To meet security requirements, a company needs to encrypt all of its application data in transit while communicating with an Amazon RDS MySQL DB instance. A recent security audit revealed that encryption at rest is enabled using AWS Key Management Service (AWS KMS), but data in transit is not enabled. What should a solutions architect do to satisfy the security requirements?



- A. Enable IAM database authentication on the database.
- B. Provide self-signed certificates. Use the certificates in all connections to the RDS instance.
- C. Take a snapshot of the RDS instance. Restore the snapshot to a new instance with encryption enabled.
- D. Download AWS-provided root certificates. Provide the certificates in all connections to the RDS instance.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:

QUESTION 624

A company is designing an application to run in a VPC on AWS. The application consists of Amazon EC2 instances that run in private subnets as part of an Auto Scaling group. The application also includes a Network Load Balancer that extends across public subnets. The application stores data in an Amazon RDS DB instance. The company has attached a security group that is named "web-servers" to the EC2 instances. The company has attached a security group that is named "database" to the DB instance. How should a solutions architect configure the communication between the EC2 instances and the DB instance?

- A. Configure the "web-servers" security group to allow access to the DB instance's current IP addresses. Configure the "database" security group to allow access from the current set of IP addresses in use by the EC2 instances.
- B. Configure the "web-servers" security group to allow access to the "database" security group. Configure the "database" security group to allow access from the "web-servers" security group.
- C. Configure the "web-servers" security group to allow access to the DB instance's current IP addresses. Configure the "database" security group to allow access from the Auto Scaling group.
- D. Configure the "web-servers" security group to allow access to the "database" security group. Configure the "database" security group to allow access from the Auto Scaling group.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:

Explanation:

QUESTION 625

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 for multiple EC2 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. Configure the ECS task definitions to mount the Amazon FSx instance volume at launch.
- B. Launch an Amazon Elastic File System (Amazon EFS) instance. Configure the ECS task definitions to mount the EFS instance volume at launch.
- C. Create a Provisioned IOPS SSD (io2) Amazon Elastic Block Store (Amazon EBS) volume with Multi-Attach set to enabled. Attach the EBS volume to the ECS EC2 instance. Configure ECS task definitions to mount the EBS instance volume at launch.
- D. Launch an EC2 instance with several Provisioned IOPS SSD (io2) Amazon Elastic Block Store (Amazon EBS) volumes attached in a RAID 0 configuration. Configure the EC2 instance as an NFS storage server. Configure ECS task definitions to mount the volumes at launch.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 626

A company has two AWS accounts in the same AWS Region. One account is a publisher account, and the other account is a subscriber account. Each account has its own Amazon S3 bucket.

An application puts media objects into the publisher account's S3 bucket. The objects are encrypted with server-side encryption with customer-provided encryption keys (SSE-C). The company needs a solution that will automatically copy the objects to the subscriber's account's S3 bucket.

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

- A. Enable S3 Versioning on the publisher account's S3 bucket. Configure S3 Same-Region Replication of the objects to the subscriber account's S3 bucket.
- B. Create an AWS Lambda function that is invoked when objects are published in the publisher account's S3 bucket. Configure the Lambda function to copy the objects to the subscriber account's S3 bucket.
- C. Configure Amazon EventBridge (Amazon CloudWatch Events) to invoke an AWS Lambda function when objects are published in the publisher account's S3 bucket. Configure the Lambda function to copy the objects to the subscriber account's S3 bucket.
- D. Configure Amazon EventBridge (Amazon CloudWatch Events) to publish Amazon Simple Notification Service (Amazon SNS) notifications when objects are published in the publisher account's S3 bucket. When notifications are received, use the S3 console to copy the objects to the subscriber account's S3 bucket.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 627

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

Which solution will meet this requirement?

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

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 628

A company's workload is running in an on-premises data center and on AWS in the ap-south-1 Region. The company wants to use the ap-southeast-1 Region as its secondary site for disaster recovery. The company needs to have a single, dedicated connection with 4 Gbps of bandwidth throughput from its data center to both AWS Regions.

What should a solutions architect do to meet this requirement?

- A. Create one AWS Site-to-Site VPN connection in ap-south-1. Connect the company's data center to the virtual private gateway.
- B. Create two AWS Site-to-Site VPN connections: one in ap-south-1 and one in ap-southeast-1. Connect the company's data center to both virtual private gateways.
- C. Create one AWS Direct Connect gateway. Provision a 4 Gbps Direct Connect hosted connection from an AWS Direct Connect Partner. Associate the connection with the Direct Connect gateway.
- D. Create two AWS Direct Connect gateways: one in ap-south-1 and one in ap-southeast-1. Provision a 4 Gbps Direct Connect hosted connection from an AWS Direct Connect Partner. Associate the connection with both Direct Connect gateways.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 629

A solutions architect launches an Amazon EC2 instance inside a new VPC. The solutions architect configures network ACL rules and security group rules that allow the appropriate traffic to flow to and from the instance. An Elastic IP address is associated with the instance. The solutions architect needs to be able to access the instance from the internet.

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

- A. Create an internet gateway. Attach the internet gateway to the VPC.
- B. Create an internet gateway. Attach the internet gateway to the instance's subnet.
- C. Create an internet gateway. Attach the internet gateway to the instance.
- D. Add a route to the route table of the instance's subnet. Route traffic from the instance's subnet to the internet gateway.
- E. Add a route to the route table of the instance's subnet. Route traffic from the internet gateway to the instance's subnet.

The logo for Udumps, featuring a stylized orange 'U' followed by the word 'dumps' in a grey, lowercase, sans-serif font.

Correct Answer: CE

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 630

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:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": ["ec2:TerminateInstances"],
      "Resource": ["*"]
    },
    {
      "Effect": "Deny",
      "Action": ["ec2:TerminateInstances"],
      "Condition": {
        "NotIpAddress": {
          "aws:SourceIp": [
            "192.0.2.0/24",
            "203.0.113.0/24"
          ]
        }
      },
      "Resource": ["*"]
    }
  ]
}
```



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.
- D. 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.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 631

A company has deployed a serverless application that invokes an AWS Lambda function when new documents are uploaded to an Amazon S3 bucket. The application uses the Lambda function to process the documents. After a recent marketing campaign, the company noticed that the application did not process many of the documents. What should a solutions architect do to improve the architecture of this application?

- A. Set the Lambda function's runtime timeout value to 15 minutes.
- B. Configure an S3 bucket replication policy. Stage the documents in the S3 bucket for later processing.
- C. Deploy an additional Lambda function. Load balance the processing of the documents across the two Lambda functions.
- D. Create an Amazon Simple Queue Service (Amazon SQS) queue. Send the requests to the queue. Configure the queue as an event source for Lambda.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 632

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 instance.
- B. Create an Amazon S3 bucket. Mount the S3 bucket on each of the EC2 instances.
- C. Ensure that the EC2 instances are 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. Mount the file system on each of the EC2 instances.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 633

A company is expanding a secure, on-premises network to AWS. The on-premises network has no direct internet access. The company is setting up an AWS Direct Connect connection between the on-premises network and AWS. An application that runs in the on-premises network needs to use the AWS software development kits (SDKs).

A solutions architect must design a solution that supports this connectivity. However, the solution cannot incur additional cost beyond the cost of the Direct Connect connection.

Which solution will meet these requirements?

- A. Create a public virtual interface (VIF). Route the AWS traffic over the public VIF.
- B. Create a VPC and a NAT gateway. Route the AWS traffic from on premises to the NAT gateway.
- C. Create a VPC and an Amazon S3 interface endpoint. Route the AWS traffic from on premises to the S3 interface endpoint.
- D. Create a VPC peering connection between the on-premises network and Direct Connect. Route the AWS traffic over the peering connection.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 634

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

Which solution meets these requirements?

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

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 635

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

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

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

delay.
Which solution meets these requirements?

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

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:

QUESTION 636

The DNS provider that hosts a company's domain name records is experiencing outages that cause service disruption for a website running on AWS. The company needs to migrate to a more resilient managed DNS service and wants the service to run on AWS.
What should a solutions architect do to rapidly migrate the DNS hosting service?

- A. Create an Amazon Route 53 public hosted zone for the domain name. Import the zone file containing the domain records hosted by the previous provider.
- B. Create an Amazon Route 53 private hosted zone for the domain name. Import the zone file containing the domain records hosted by the previous provider.
- C. Create a Simple AD directory in AWS. Enable zone transfer between the DNS provider and AWS Directory Service for Microsoft Active Directory for the domain records.
- D. Create an Amazon Route 53 Resolver inbound endpoint in the VPC. Specify the IP addresses that the provider's DNS will forward DNS queries to. Configure the provider's DNS to forward DNS queries for the domain to the IP addresses that are specified in the inbound endpoint.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:



QUESTION 637

A company is designing a cloud communications platform that is driven by APIs. The application is hosted on Amazon EC2 instances behind a Network Load Balancer (NLB). The company uses Amazon API Gateway to provide external users with access to the application through APIs. The company wants to protect the platform against web exploits like SQL injection and also wants to detect and mitigate large, sophisticated DDoS attacks.
Which combination of solutions provides the MOST protection? (Choose two.)

- A. Use AWS WAF to protect the NLB.
- B. Use AWS Shield Advanced with the NLB.
- C. Use AWS WAF to protect Amazon API Gateway.
- D. Use Amazon GuardDuty with AWS Shield Standard.
- E. Use AWS Shield Standard with Amazon API Gateway.

Correct Answer: AD
Section: (none)
Explanation

Explanation/Reference:
Explanation:

QUESTION 638

A survey company has gathered data for several years from areas in the United States. The company hosts the data in an Amazon S3 bucket that is 3 TB in size and growing. The company has started to share the data with a European marketing firm that has S3 buckets. The company wants to ensure that its data transfer costs remain as low as possible.
Which solution will meet these requirements?

- A. Configure the Requester Pays feature on the company's S3 bucket.
- B. Configure S3 Cross-Region Replication from the company's S3 bucket to one of the marketing firm's S3 buckets.

- C. Configure cross-account access for the marketing firm so that the marketing firm has access to the company's S3 bucket.
- D. Configure the company's S3 bucket to use S3 Intelligent-Tiering. Sync the S3 bucket to one of the marketing firm's S3 buckets.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 639

A company develops applications in separate AWS accounts that are all part of an organization in AWS Organizations. An operations team creates an IAM user for each developer for a given application. As the company has grown, the number of applications has increased. Developers now work on several applications and need to view and access all their project accounts.

A solutions architect must design a solution that minimizes the operational overhead for the operations team.

What should the solutions architect do to meet these requirements?

- A. Implement AWS Single Sign-On for the organization.
- B. Consolidate all the AWS accounts into a single account for all users and applications.
- C. Use AWS CloudFormation StackSets to programmatically create IAM users in each account.
- D. Create a shared services account. Create all the IAM users in the shared services account. Configure cross-account access roles with appropriate access to each account.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 640

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

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 641

A company has a stateless, asynchronous application that runs in an Apache Hadoop cluster. The application is invoked on demand to run extract, transform, and load (ETL) jobs several times a day.

A solutions architect needs to migrate this application to the AWS Cloud by designing an Amazon EMR cluster for the workload. The cluster must be available immediately to process jobs.

Which implementation meets these requirements MOST cost-effectively?

- A. Use zonal Reserved Instances for the master nodes and the core nodes. Use a Spot Fleet for the task nodes.
- B. Use zonal Reserved Instances for the master nodes. Use Spot Instances for the core nodes and the task nodes.
- C. Use regional Reserved Instances for the master nodes. Use a Spot Fleet for the core nodes and the task nodes.
- D. Use regional Reserved Instances for the master nodes. Use On-Demand Capacity Reservations for the core nodes and the task nodes.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 642

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 EFS) Standard-Infrequent Access (Standard-IA) to store data and provide shared access to the instances.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 643

A company runs several websites on AWS for its different brands. Each website generates tens of gigabytes of web traffic logs each day. A solutions architect needs to design a scalable solution to give the company's developers the ability to analyze traffic patterns across all the company's websites. This analysis by the developers will occur on demand once a week over the course of several months. The solution must support queries with standard SQL.

Which solution will meet these requirements MOST cost-effectively?

- A. Store the logs in Amazon S3. Use Amazon Athena for analysis.
- B. Store the logs in Amazon RDS. Use a database client for analysis.
- C. Store the logs in Amazon OpenSearch Service (Amazon Elasticsearch Service). Use Amazon OpenSearch Service (Amazon Elasticsearch Service) for analysis.
- D. Store the logs in an Amazon EMR cluster. Use a supported open-source framework for SQLbased analysis.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 644

A company is deploying a new application to Amazon Elastic Kubernetes Service (Amazon EKS) with an AWS Fargate cluster. The application needs a storage solution for data persistence. The solution must be highly available and fault tolerant. The solution also must be shared between multiple application containers.

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

- A. Create Amazon Elastic Block Store (Amazon EBS) volumes in the same Availability Zones where EKS worker nodes are placed. Register the volumes in a StorageClass object on an EKS cluster. Use EBS Multi-Attach to share the data between containers.
- B. Create an Amazon Elastic File System (Amazon EFS) file system. Register the file system in a StorageClass object on an EKS cluster. Use the same file system for all containers.
- C. Create an Amazon Elastic Block Store (Amazon EBS) volume. Register the volume in a StorageClass object on an EKS cluster. Use the same volume for all containers.
- D. Create Amazon Elastic File System (Amazon EFS) file systems in the same Availability Zones where EKS worker nodes are placed. Register the file systems in a StorageClass object on an EKS cluster. Create an AWS Lambda function to synchronize the data between file systems.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 645

A manufacturing company has machine sensors that upload .csv files to an Amazon S3 bucket.

These .csv files must be converted into images and must be made available as soon as possible for the automatic generation of graphical reports. The images become irrelevant after 1 month, but the .csv files must be kept to train machine learning (ML) models twice a year. The ML trainings and audits are planned weeks in advance. Which combination of steps will meet these requirements MOST cost-effectively? (Choose two.)

- A. Launch an Amazon EC2 Spot Instance that downloads the .csv files every hour, generates the image files, and uploads the images to the S3 bucket.
- B. Design an AWS Lambda function that converts the .csv files into images and stores the images in the S3 bucket. Invoke the Lambda function when a .csv file is uploaded.
- C. Create S3 Lifecycle rules for .csv files and image files in the S3 bucket. Transition the .csv files from S3 Standard to S3 Glacier 1 day after they are uploaded. Expire the image files after 30 days.
- D. Create S3 Lifecycle rules for .csv files and image files in the S3 bucket. Transition the .csv files from S3 Standard to S3 One Zone-Infrequent Access (S3 One Zone-IA) 1 day after they are uploaded. Expire the image files after 30 days.
- E. Create S3 Lifecycle rules for .csv files and image files in the S3 bucket. Transition the .csv files from S3 Standard to S3 Standard-Infrequent Access (S3 Standard-IA) 1 day after they are uploaded. Keep the image files in Reduced Redundancy Storage (RRS).

Correct Answer: BD

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 646

A company has 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.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 647

A company is creating a new web application for its subscribers. The application will consist of a static single page and a persistent database layer. The application will have millions of users for 4 hours in the morning, but the application will have only a few thousand users during the rest of the day. The company's data architects have requested the ability to rapidly evolve their schema. Which solutions will meet these requirements and provide the MOST scalability? (Choose two.)

- A. Deploy Amazon DynamoDB as the database solution. Provision on-demand capacity.
- B. Deploy Amazon Aurora as the database solution. Choose the serverless DB engine mode.
- C. Deploy Amazon DynamoDB as the database solution. Ensure that DynamoDB auto scaling is enabled.
- D. Deploy the static content into an Amazon S3 bucket. Provision an Amazon CloudFront distribution with the S3 bucket as the origin.
- E. Deploy the web servers for static content across a fleet of Amazon EC2 instances in Auto Scaling groups. Configure the instances to periodically refresh the content from an Amazon Elastic File System (Amazon EFS) volume.

Correct Answer: AB

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 648

A company has multiple Windows file servers on premises. The company wants to migrate and consolidate its files into an Amazon FSx for Windows File Server file system. File permissions must be preserved to ensure that access rights do not change. Which solutions will meet these requirements? (Choose two.)

- A. Deploy AWS DataSync agents on premises. Schedule DataSync tasks to transfer the data to the FSx for Windows File Server file system.



- B. Copy the shares on each file server into Amazon S3 buckets by using the AWS CLI. Schedule AWS DataSync tasks to transfer the data to the FSx for Windows File Server file system.
- C. Remove the drives from each file server. Ship the drives to AWS for import into Amazon S3. Schedule AWS DataSync tasks to transfer the data to the FSx for Windows File Server file system.
- D. Order an AWS Snowcone device. Connect the device to the on-premises network. Launch AWS DataSync agents on the device. Schedule DataSync tasks to transfer the data to the FSx for Windows File Server file system.
- E. Order an AWS Snowball Edge Storage Optimized device. Connect the device to the on-premises network. Copy data to the device by using the AWS CLI. Ship the device back to AWS for import into Amazon S3. Schedule AWS DataSync tasks to transfer the data to the FSx for Windows File Server file system.

Correct Answer: AD

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 649

A company has deployed a business-critical application in the AWS Cloud. The application uses Amazon EC2 instances that run in the us-east-1 Region. The application uses Amazon S3 for storage of all critical data.

To meet compliance requirements, the company must create a disaster recovery (DR) plan that provides the capability of a full failover to another AWS Region.

What should a solutions architect recommend for this DR plan?

- A. Deploy the application to multiple Availability Zones in us-east-1. Create a resource group in AWS Resource Groups. Turn on automatic failover for the application to use a predefined recovery Region.
- B. Perform a virtual machine (VM) export by using AWS Import/Export on the existing EC2 instances. Copy the exported instances to the destination Region. In the event of a disaster, provision new EC2 instances from the exported EC2 instances.
- C. Create snapshots of all Amazon Elastic Block Store (Amazon EBS) volumes that are attached to the EC2 instances in us-east-1. Copy the snapshots to the destination Region. In the event of a disaster, provision new EC2 instances from the EBS snapshots.
- D. Use S3 Cross-Region Replication for the data that is stored in Amazon S3. Create an AWS CloudFormation template for the application with an S3 bucket parameter. In the event of a disaster, deploy the template to the destination Region and specify the local S3 bucket as the parameter.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:



QUESTION 650

A company wants to migrate a Windows-based application from on premises to the AWS Cloud.

The application has three tiers: an application tier, 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 processing between the tiers.

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

- A. Host all three tiers on Amazon EC2 instances. Use Amazon FSx File Gateway for file sharing between the tiers.
- B. Host all three tiers on Amazon EC2 instances. Use Amazon FSx for Windows File Server for file sharing between the tiers.
- C. Host the application tier and the business tier on Amazon EC2 instances. Host the database tier on Amazon RDS. Use Amazon Elastic File System (Amazon EFS) for file sharing between the tiers.
- D. Host the application tier and the business tier on Amazon EC2 instances. Host the database tier on Amazon RDS. Use a Provisioned IOPS SSD (io2) Amazon Elastic Block Store (Amazon EBS) volume for file sharing between the tiers.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 651

A company has three AWS accounts: Management, Development, and Production. These accounts use AWS services only in the us-east-1 Region. All accounts have a VPC with VPC Flow Logs configured to publish data to an Amazon S3 bucket in each separate account. For compliance reasons, the company needs an ongoing method to aggregate all the VPC flow logs across all accounts into one destination S3 bucket in the Management account.

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

- A. Add S3 Same-Region Replication rules in each S3 bucket that stores VPC flow logs to replicate objects to the destination S3 bucket. Configure the destination S3 bucket to allow objects to be received from the S3 buckets in other accounts.
- B. Set up an IAM user in the Management account. Grant permissions to the IAM user to access the S3 buckets that contain the VPC flow logs. Run the aws s3 sync command in the AWS CLI to copy the objects to the destination S3

bucket.

C. Use an S3 inventory report to specify which objects in the S3 buckets to copy. Perform an S3 batch operation to copy the objects into the destination S3 bucket in the Management account with a single request.

D. Create an AWS Lambda function in the Management account. Grant S3 GET permissions on the source S3 buckets. Grant S3 PUT permissions on the destination S3 bucket. Configure the function to invoke when objects are loaded in the source S3 buckets.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 652

A company is concerned about the security of its public web application due to recent web attacks.

The application uses an Application Load Balancer (ALB). A solutions architect must reduce the risk of DDoS attacks against the application.

What should the solutions architect do to meet this requirement?

A. Add an Amazon Inspector agent to the ALB

B. Configure Amazon Macie to prevent attacks

C. Enable AWS Shield Advanced to prevent attacks

D. Configure Amazon GuardDuty to monitor the ALB

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 653

A company built an application with Docker containers and needs to run the application in the AWS Cloud. The company wants to use a managed service to host the application.

The solution must scale in and out appropriately according to demand on the individual container services. The solution also must not result in additional operational overhead or infrastructure to manage.

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

A. Use Amazon Elastic Container Service (Amazon ECS) with AWS Fargate

B. Use Amazon Elastic Kubernetes Service (Amazon EKS) with AWS Fargate

C. Provision an Amazon API Gateway API. Connect the API to AWS Lambda to run the containers

D. Use Amazon Elastic Container Service (Amazon ECS) with Amazon EC2 worker nodes

E. Use Amazon Elastic Kubernetes Service (Amazon EKS) with Amazon EC2 worker nodes

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 654

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. Configure the web server in the Lightsail instance. Upload website content by using an SFTP client.

B. Create an AWS Auto Scaling group for Amazon EC2 instances. Use an Application Load Balancer. Upload website content by using an SFTP client.

C. Create a private Amazon S3 bucket. Use an S3 bucket policy to allow access from a CloudFront origin access identity (OAI). Upload website content by using the AWS CLI.

D. Create a public Amazon S3 bucket. Configure AWS Transfer for SFTP. Configure the S3 bucket for website hosting. Upload website content by using the SFTP client.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 655

A company is building an application that gives users the ability to upload video files. The company needs to store the video files securely in Amazon S3 for future access. The total storage requirement for the video files will grow to many petabytes. Users will rarely access the stored video files, but the stored video files must be available within minutes when users request access.

The company needs to minimize storage costs.

Which S3 storage class should a solutions architect choose to meet these requirements?

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

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 656

A company is running a batch application on Amazon EC2 instances. The application consists of a backend with multiple Amazon RDS databases. The application is causing a high number of reads on the databases. A solutions architect must reduce the number of database reads while ensuring high availability.

What should the solutions architect do to meet this requirement?

- A. Add Amazon RDS read replicas
- B. Use Amazon ElastiCache for Redis
- C. Use Amazon Route 53 DNS caching
- D. Use Amazon ElastiCache for Memcached



Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 657

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

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

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

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

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 658

A company has an application that serves clients that are deployed in more than 20,000 retail storefront locations around the world. The application consists of backend web services that are exposed over HTTPS on port 443. The application is hosted on Amazon EC2 instances behind an Application Load Balancer (ALB). The retail locations communicate with the web application over the public internet. The company allows each retail location to register the IP address that the retail location has been allocated by its local ISP.

The company's security team recommends to increase the security of the application endpoint by restricting access to only the IP addresses registered by the retail locations. What should a solutions architect do to meet these requirements?

- A. Associate an AWS WAF web ACL with the ALB. Use IP rule sets on the ALB to filter traffic.
Update the IP addresses in the rule to include the registered IP addresses.
- B. Deploy AWS Firewall Manager to manage the ALB. Configure firewall rules to restrict traffic to the ALB. Modify the firewall rules to include the registered IP addresses.
- C. Store the IP addresses in an Amazon DynamoDB table. Configure an AWS Lambda authorization function on the ALB to validate that incoming requests are from the registered IP addresses.
- D. Configure the network ACL on the subnet that contains the public interface of the ALB. Update the ingress rules on the network ACL with entries for each of the registered IP addresses.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 659

A social media company allows users to upload images to its website. The website runs on Amazon EC2 instances. During upload requests, the website resizes the images to a standard size and stores the resized images in Amazon S3. Users are experiencing slow upload requests to the website.

The company needs to reduce coupling within the application and improve website performance.

A solutions architect must design the most operationally efficient process for image uploads.

Which combination of actions should the solutions architect take to meet these requirements?

(Choose two.)

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

Correct Answer: CD

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 660

A photo-sharing company makes a personalized collage of each user's photographs available online. The company's custom software builds each collage by using 10-20 available images for each user. The application stores the original images and the collages in Amazon S3 Standard storage. After a collage is created, the original files are needed only if the company must recreate the collage.

The company has thousands of users and wants to reduce storage costs as much as possible. A solutions architect must recommend which S3 storage classes the company should use to store the original images and the collages.

What should the solutions architect recommend to meet these requirements?

- A. Move the original images to S3 Glacier. Move the collages to S3 Standard-Infrequent Access (S3 Standard-IA)
- B. Move the original images to S3 Glacier. Move the collages to S3 One Zone-Infrequent Access (S3 One Zone-IA)
- C. Move the original images to S3 Standard-Infrequent Access (S3 Standard-IA). Move the collages to S3 One Zone-Infrequent Access (S3 One Zone-IA)
- D. Move the original images to S3 One Zone-Infrequent Access (S3 One Zone-IA). Move the collages to S3 Standard-Infrequent Access (S3 Standard-IA)

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 661

A company is expecting rapid growth in the near future. A solutions architect needs to configure existing users and grant permissions to new users on AWS. The solutions architect has decided to create IAM groups. The solutions architect will add the new users to IAM groups based on department.

Which additional action is the MOST secure way to grant permissions to the new users?

- A. Apply service control policies (SCPs) to manage access permissions
- B. Create IAM roles that have least privilege permission. Attach the roles to the IAM groups
- C. Create an IAM policy that grants least privilege permission. Attach the policy to the IAM groups
- D. Create IAM roles. Associate the roles with a permissions boundary that defines the maximum permissions

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 662

A company hosts an application in its own data center. The application uses Amazon S3 for data storage. The application transfers several hundred terabytes of data each month to and from Amazon S3. The company wants to minimize the cost of this data transfer.

Which solution will meet these requirements?

- A. Create an FTPS server by using AWS Transfer Family. Configure the application to use the FTPS server to store and retrieve files.
- B. Enable S3 Transfer Acceleration on the S3 bucket. Configure the application to use the S3 Transfer Acceleration endpoint to upload and download the data.
- C. Set up an AWS Direct Connect connection between the AWS Region in use and the company's data center. Configure the application to use the Direct Connect connection to upload and download the data.
- D. Create an Amazon S3 File Gateway. Deploy the software appliance in the company's data center. Configure the application to use the S3 File Gateway to store and retrieve files.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 663

A solutions architect is designing a two-tiered architecture that includes a public subnet and a database subnet. The web servers in the public subnet must be open to the internet on port 443. The Amazon RDS for MySQL DB instance in the database subnet must be accessible only to the web servers on port 3306.

Which combination of steps should the solutions architect take to meet these requirements?

(Choose two.)

- A. Create a network ACL for the public subnet. Add a rule to deny outbound traffic to 0.0.0.0/0 on port 3306.
- B. Create a security group for the DB instance. Add a rule to allow traffic from the public subnet CIDR block on port 3306.
- C. Create a security group for the web servers in the public subnet. Add a rule to allow traffic from 0.0.0.0/0 on port 443.
- D. Create a security group for the DB instance. Add a rule to allow traffic from the web servers' security group on port 3306.
- E. Create a security group for the DB instance. Add a rule to deny all traffic except traffic from the web servers' security group on port 3306.

Correct Answer: CD

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 664

A company has hired an external vendor to perform work in the company's AWS account. The vendor uses an automated tool that is hosted in an AWS account that the vendor owns. The vendor does not have IAM access to the company's AWS account.

How should a solutions architect grant this access to the vendor?



- A. Create a IAM role in the company's account to delegate access to the vendor's IAM role. Attach the appropriate IAM policies to the role for the permissions that the vendor requires.
- B. Create an IAM user in the company's account with a password that meets the password complexity requirements. Attach the appropriate IAM policies to the user for the permissions that the vendor requires.
- C. Create an IAM group in the company's account. Add the tool's IAM user from the vendor account to the group for the permissions that the vendor requires.
- D. Create a new identity provider by choosing "AWS account" as the provider type in the IAM console. Supply the vendor's AWS account ID and user name. Attach the appropriate IAM policies to the new provider for the permissions that the vendor requires.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 665

A company wants to migrate two DNS servers to AWS. The servers host a total of approximately 200 zones and receive 1 million requests each day on average. The company wants to maximize availability while minimizing the operational overhead that is related to the management of the two servers.

What should a solutions architect recommend to meet these requirements?

- A. Create 200 new hosted zones in the Amazon Route 53 console. Import zone files.
- B. Launch a single large Amazon EC2 instance. Import zone files. Configure Amazon CloudWatch alarms and notifications to alert the company about any downtime.
- C. Migrate the servers to AWS by using AWS Server Migration Service (AWS SMS). Configure Amazon CloudWatch alarms and notifications to alert the company about any downtime.
- D. Launch an Amazon EC2 instance in an Auto Scaling group across two Availability Zones. Import zone files. Set the desired capacity to 1 and the maximum capacity to 3 for the Auto Scaling group. Configure scaling alarms to scale based on CPU utilization.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:



QUESTION 666

A company runs a stateless web application in production on a group of Amazon EC2 On-Demand Instances behind an Application Load Balancer. The application experiences heavy usage during an 8-hour period each business day.

Application usage is moderate and steady overnight.

Application usage is low during weekends.

The company wants to minimize its EC2 costs without affecting the availability of the application.

Which solution will meet these requirements?

- A. Use Spot Instances for the entire workload
- B. Use Reserved Instances for the baseline level of usage. Use Spot Instances for any additional capacity that the application needs.
- C. Use On-Demand Instances for the baseline level of usage. Use Spot Instances for any additional capacity that the application needs.
- D. Use Dedicated Instances for the baseline level of usage. Use On-Demand Instances for any additional capacity that the application needs.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 667

A company runs a latency-sensitive gaming service in the AWS Cloud. The gaming service runs on a fleet of Amazon EC2 instances behind an Application Load Balancer (ALB). An Amazon DynamoDB table stores the gaming data. All the infrastructure is in a single AWS Region. The main user base is in that same Region.

A solutions architect needs to update the architecture to support a global expansion of the gaming service. The gaming service must operate with the least possible latency.

Which solution will meet these requirements?

- A. Create an Amazon CloudFront distribution in front of the ALB
- B. Deploy an Amazon API Gateway regional API endpoint. Integrate the API endpoint with the ALB
- C. Create an accelerator in AWS Global Accelerator. Add a listener. Configure the endpoint to point to the ALB.

D. Deploy the ALB and the fleet of EC2 instances to another Region. Use Amazon Route 53 with geolocation routing.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:

QUESTION 668

A company hosts a two-tier website that runs on Amazon EC2 instances. The website has a database that runs on Amazon RDS for MySQL. All users are required to log in to the website to see their own customized pages. The website typically experiences low traffic. Occasionally, the website experiences sudden increases in traffic and becomes unresponsive. During these increases in traffic, the database experiences a heavy write load. A solutions architect must improve the website's availability without changing the application code. What should the solutions architect do to meet these requirements?

- A. Create an Amazon ElastiCache for Redis cluster. Configure the application to cache common database queries in the ElastiCache cluster.
- B. Create an Auto Scaling group. Configure Amazon CloudWatch alarms to scale the number of EC2 instances based on the percentage of CPU in use during the traffic increases.
- C. Create an Amazon CloudFront distribution that points to the EC2 instances as the origin. Enable caching of dynamic content, and configure a write throttle from the EC2 instances to the RDS database.
- D. Migrate the database to an Amazon Aurora Serverless cluster. Set the maximum Aurora capacity units (ACUs) to a value high enough to respond to the traffic increases. Configure the EC2 instances to connect to the Aurora database.

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Explanation:

QUESTION 669

A developer is creating an AWS Lambda function to perform dynamic updates to a database when an item is added to an Amazon Simple Queue Service (Amazon SQS) queue. A solutions architect must recommend a solution that tracks any usage of database credentials in AWS CloudTrail. The solution also must provide auditing capabilities. Which solution will meet these requirements?

- A. Store the encrypted credentials in a Lambda environment variable
- B. Create an Amazon DynamoDB table to store the credentials. Encrypt the table
- C. Store the credentials as a secure string in AWS Systems Manager Parameter Store
- D. Use an AWS Key Management Service (AWS KMS) key store to store the credentials

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:

QUESTION 670

A company is designing an application that will run on an AWS Lambda function within a VPC. An Amazon API Gateway API will invoke the Lambda function. A solutions architect needs to recommend an Amazon CloudWatch solution that developers can use to identify the users who are generating the most network traffic. Which solution will meet these requirements?

- A. Configure CloudWatch Lambda Insights. Examine the network usage graph by using the multifunction view in the performance dashboard.
- B. Create a canary in CloudWatch Synthetics. Turn on active tracing. Review the network usage graph in the Monitoring tab of the canary.
- C. Configure VPC flow logs to stream to CloudWatch Logs. Create a CloudWatch Contributor Insights rule from the sample blueprint.
- D. Add the application to CloudWatch Application Insights. View the graph for top network users in the dashboard that Application Insights creates automatically.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:

QUESTION 671

A company has users all around the world accessing its HTTP-based application deployed on Amazon EC2 instances in multiple AWS Regions. The company wants to improve the availability and performance of the application. The company also wants to protect the application against common web exploits that may affect availability, compromise security, or consume excessive resources. Static IP addresses are required. What should a solutions architect recommend to accomplish this?

- A. Put the EC2 instances behind Network Load Balancers (NLBs) in each Region. Deploy AWS WAF on the NLBs. Create an accelerator using AWS Global Accelerator and register the NLBs as endpoints.
- B. Put the EC2 instances behind Application Load Balancers (ALBs) in each Region. Deploy AWS WAF on the ALBs. Create an accelerator using AWS Global Accelerator and register the ALBs as endpoints.
- C. Put the EC2 instances behind Network Load Balancers (NLBs) in each Region. Deploy AWS WAF on the NLBs. Create an Amazon CloudFront distribution with an origin that uses Amazon Route 53 latency-based routing to route requests to the NLBs.
- D. Put the EC2 instances behind Application Load Balancers (ALBs) in each Region. Create an Amazon CloudFront distribution with an origin that uses Amazon Route 53 latency-based routing to route requests to the ALBs. Deploy AWS WAF on the CloudFront distribution.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 672

A company has more than 5 TB of file data on Windows file servers that run on premises. Users and applications interact with the data each day. The company is moving its Windows workloads to AWS. As the company continues this process, the company requires access to AWS and on-premises file storage with minimum latency. The company needs a solution that minimizes operational overhead and requires no significant changes to the existing file access patterns. The company uses an AWS Site-to-Site VPN connection for connectivity to AWS. What should a solutions architect do to meet these requirements?

- A. Deploy and configure Amazon FSx for Windows File Server on AWS. Move the on-premises file data to FSx for Windows File Server. Reconfigure the workloads to use FSx for Windows File Server on AWS.
- B. Deploy and configure an Amazon S3 File Gateway on premises. Move the on-premises file data to the S3 File Gateway. Reconfigure the on-premises workloads and the cloud workloads to use the S3 File Gateway.
- C. Deploy and configure an Amazon S3 File Gateway on premises. Move the on-premises file data to Amazon S3. Reconfigure the workloads to use either Amazon S3 directly or the S3 File Gateway, depending on each workload's location.
- D. Deploy and configure Amazon FSx for Windows File Server on AWS. Deploy and configure an Amazon FSx File Gateway on premises. Move the on-premises file data to the FSx File Gateway. Configure the cloud workloads to use FSx for Windows File Server on AWS. Configure the on-premises workloads to use the FSx File Gateway.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 673

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

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

Correct Answer: CD

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 674

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

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 675

A company wants its public web application to run on Amazon EC2 instances in an Auto Scaling group behind an Application Load Balancer (ALB). The application must use a publicly trusted SSL certificate. Which solution will meet these requirements MOST cost-effectively?

- A. Provision a public SSL/TLS certificate through AWS Certificate Manager (ACM). Configure the new certificate on the HTTPS listener for the ALB.
- B. Use AWS Certificate Manager Private Certificate Authority to issue an SSL/TLS certificate.
Configure the new certificate on the HTTPS listener for the ALB.
- C. Create a self-signed certificate on one of the EC2 instances in the Auto Scaling group. Export the certificate, and configure it on the HTTPS listener for the ALB.
- D. Deploy an EC2-hosted certificate authority (CA). Import a trusted root certificate. Issue a new SSL/TLS certificate. Configure the new certificate on the HTTPS listener for the ALB.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:



QUESTION 676

A solutions architect is tasked with transferring 750 TB of data from a network-attached file system located at a branch office to Amazon S3 Glacier. The solution must avoid saturating the branch office's low-bandwidth internet connection. What is the MOST cost-effective solution?

- A. Create a site-to-site VPN tunnel to an Amazon S3 bucket and transfer the files directly. Create a bucket policy to enforce a VPC endpoint.
- B. Order 10 AWS Snowball appliances and select an S3 Glacier vault as the destination. Create a bucket policy to enforce a VPC endpoint.
- C. Mount the network-attached file system to Amazon S3 and copy the files directly. Create a lifecycle policy to transition the S3 objects to Amazon S3 Glacier.
- D. Order 10 AWS Snowball appliances and select an Amazon S3 bucket as the destination. Create a lifecycle policy to transition the S3 objects to Amazon S3 Glacier.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 677

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

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

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 678

A company is building a solution that will report Amazon EC2 Auto Scaling events across all the applications in an AWS account. The company needs to use a serverless solution to store the EC2 Auto Scaling status data in Amazon S3. The company then will use the data in Amazon S3 to provide near-real-time updates in a dashboard. The solution must not affect the speed of EC2 instance launches. How should the company move the data to Amazon S3 to meet these requirements?

- A. Use an Amazon CloudWatch metric stream to send the EC2 Auto Scaling status data to Amazon Kinesis Data Firehose. Store the data in Amazon S3.
- B. Launch an Amazon EMR cluster to collect the EC2 Auto Scaling status data and send the data to Amazon Kinesis Data Firehose. Store the data in Amazon S3.
- C. Create an Amazon EventBridge (Amazon CloudWatch Events) rule to invoke an AWS Lambda function on a schedule. Configure the Lambda function to send the EC2 Auto Scaling status data directly to Amazon S3.
- D. Use a bootstrap script during the launch of an EC2 instance to install Amazon Kinesis Agent. Configure Kinesis Agent to collect the EC2 Auto Scaling status data and send the data to Amazon Kinesis Data Firehose. Store the data in Amazon S3.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 679

A company is testing an application that runs on an Amazon EC2 Linux instance. The instance contains a data volume of 500 GB that consists of a single Amazon Elastic Block Store (Amazon EBS) General Purpose SSD (gp2) volume. The application is now ready for production use and will be installed on multiple EC2 instances that run in an Auto Scaling group. All instances need access to the data that was stored on the 500 GB volume. The company needs a highly available and fault-tolerant solution that does not introduce any significant changes to the application's code. Which solution meets these requirements?

- A. Provision an EC2 instance with NFS server software that is configured with a single 500 GB gp2 volume.
- B. Use an Amazon FSx for Windows File Server file system that is configured as an SMB file store within a single Availability Zone.
- C. Migrate the data into an Amazon S3 bucket. Use an EC2 instance profile to access the contents of the bucket.
- D. Use an Amazon Elastic File System (Amazon EFS) file system that is configured with the General Purpose performance mode.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 680

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

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

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 681

A company is architecting a shared storage solution for a gaming application that is hosted in the AWS Cloud. The company needs the ability to use Lustre clients to access data. The solution must be fully managed. Which solution meets these requirements?

- A. Create an AWS DataSync task that shares the data as a mountable file system. Mount the file system to the application server.
- B. Create an AWS Storage Gateway file gateway. Create a file share that uses the required client protocol. Connect the application server to the file share.
- C. Create an Amazon Elastic File System (Amazon EFS) file system, and configure it to support Lustre. Attach the file system to the origin server. Connect the application server to the file system.
- D. Create an Amazon FSx for Lustre file system. Attach the file system to the origin server. Connect the application server to the file system.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 682

A company is creating an application that runs on containers in a VPC. The application stores and accesses data in an Amazon S3 bucket. During the development phase, the application will store and access 1 TB of data in Amazon S3 each day. The company wants to minimize costs and wants to prevent traffic from traversing the internet whenever possible. Which solution will meet these requirements?

- A. Enable S3 Intelligent-Tiering for the S3 bucket.
- B. Enable S3 Transfer Acceleration for the S3 bucket.
- C. Create a gateway VPC endpoint for Amazon S3. Associate this endpoint with all route tables in the VPC.
- D. Create an interface endpoint for Amazon S3 in the VPC. Associate this endpoint with all route tables in the VPC.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 683

A company is migrating a Linux-based web server group to AWS. The web servers must access files in a shared file store for some content. The company must not make any changes to the application. What should a solutions architect do to meet these requirements?

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

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 684

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. Configure Route 53 to route traffic to the CloudFront distribution.
- C. 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.
- D. 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. Create two domain names. 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.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 685

A company hosts a serverless application on AWS. The application uses Amazon API Gateway AWS Lambda, and an Amazon RDS for PostgreSQL database. The company notices an increase in application errors that result from database connection timeouts during times of peak traffic or unpredictable traffic. The company needs a solution that reduces the application failures with the least amount of change to the code. What should a solutions architect do to meet these requirements?

- A. Reduce the Lambda concurrency rate.
- B. Enable RDS Proxy on the RDS DB instance.
- C. Resize the RDS DB instance class to accept more connections.
- D. Migrate the database to Amazon DynamoDB with on-demand scaling.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 686

A company manages and runs a critical data management application in containers that are hosted on Amazon Elastic Container Service (Amazon ECS). The application has endpoints that are exposed through Application Load Balancers (ALBs). The application uses an Amazon Elastic File System (Amazon EFS) file system mount for persistent data storage. The company has configured Amazon ECS to use a minimal IAM instance role. Which combination of actions should a solutions architect take to improve the overall security posture of the application? (Choose two.)

- A. Decompose the Amazon ECS IAM instance role Use only ECS task roles.
- B. Enable EFS encryption in transit to protect data that is being written to Amazon EFS.
- C. Use AWS Config to define patch management policies on the container instances.
- D. Use Amazon Macie integration with Amazon EFS to monitor and protect sensitive information in the file system.
- E. Use Amazon GuardDuty to authenticate data access between the ALBs and the container instances.

Correct Answer: BC

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 687

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

The application must be secure and accessible for global customers that have dynamic IP addresses.

How should a solutions architect configure the security groups to meet these requirements?

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

D. Configure the security group for the web servers to allow inbound traffic on port 443 from 0.0.0.0/0. Configure the security group for the DB instance to allow inbound traffic on port 3306 from 0.0.0.0/0.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 688

A company needs to connect its on-premises data center network to a new VPC. The data center network has a 100 Mbps symmetrical internet connection. An application that is running on premises will transfer multiple gigabytes of data each day. The application will use an Amazon Kinesis Data Firehose delivery stream for processing. What should a solutions architect recommend for maximum performance?

- A. Create a VPC peering connection between the on-premises network and the VPC. Configure routing for the on-premises network to use the VPC peering connection.
- B. Procure an AWS Snowball Edge Storage Optimized device. After several days' worth of data has accumulated, copy the data to the device and ship the device to AWS for expedited transfer to Kinesis Data Firehose Repeat as needed.
- C. Create an AWS Site-to-Site VPN connection between the on-premises network and the VPC. Configure BGP routing between the customer gateway and the virtual private gateway. Use the VPN connection to send the data from on premises to Kinesis Data Firehose.
- D. Use AWS PrivateLink to create an interface VPC endpoint for Kinesis Data Firehose in the VPC. Set up a 1 Gbps AWS Direct Connect connection between the on-premises network and AWS. Use the PrivateLink endpoint to send the data from on premises to Kinesis Data Firehose.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 689

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

- A. Configure the web application to send an order message to Amazon Kinesis Data Firehose. Set the payment service to retrieve the message from Kinesis Data Firehose and process the order.
- B. Create a rule in AWS CloudTrail to invoke an AWS Lambda function based on the logged application path request. Use Lambda to query the database, call the payment service, and pass in the order information.
- C. Store the order in the database. Send a message that includes the order number to Amazon Simple Notification Service (Amazon SNS). Set the payment service to poll Amazon SNS retrieve the message, and process the order.
- D. Store the order in the database. Send a message that includes the order number to an Amazon Simple Queue Service (Amazon SQS) FIFO queue. Set the payment service to retrieve the message and process the order. Delete the message from the queue.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 690

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.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 691

A solutions architect is designing a multi-tier application for a company. The application's users upload images from a mobile device. The application generates a thumbnail of each image and returns a message to the user to confirm that the image was uploaded successfully.

The thumbnail generation can take up to 60 seconds, but the company wants to provide a faster response time to its users to notify them that the original image was received. The solutions architect must design the application to asynchronously dispatch requests to the different application tiers.

What should the solutions architect do to meet these requirements?

- A. Write a custom AWS Lambda function to generate the thumbnail and alert the user. Use the image upload process as an event source to invoke the Lambda function.
- B. Create an AWS Step Functions workflow. Configure Step Functions to handle the orchestration between the application tiers and alert the user when thumbnail generation is complete.
- C. Create an Amazon Simple Queue Service (Amazon SQS) message queue. As images are uploaded, place a message on the SQS queue for thumbnail generation. Alert the through an application message that the image was received.
- D. Create Amazon Simple Notification Service (Amazon SNS) notification topics and subscriptions.
Use one subscription with the application to generate the thumbnail after the image upload is complete. Use a second subscription to message the user's mobile app by way of a push notification after thumbnail generation is complete.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 692

A company copies 200 TB of data from a recent ocean survey onto AWS Snowball Edge Storage Optimized devices. The company has a high performance computing (HPC) cluster that is hosted on AWS to look for oil and gas deposits. A solutions architect must provide the cluster with consistent sub-millisecond latency and high-throughput access to the data on the Snowball Edge Storage Optimized devices. The company is sending the devices back to AWS.

Which solution will meet these requirements?

- A. Create an Amazon S3 bucket. Import the data into the S3 bucket. Configure an AWS Storage Gateway file gateway to use the S3 bucket. Access the file gateway from the HPC cluster instances.
- B. Create an Amazon S3 bucket. Import the data into the S3 bucket. Configure an Amazon FSx for Lustre file system, and integrate it with the S3 bucket. Access the FSx for Lustre file system from the HPC cluster instances.
- C. Create an Amazon S3 bucket and an Amazon Elastic File System (Amazon EFS) file system.
Import the data into the S3 bucket. Copy the data from the S3 bucket to the EFS file system. Access the EFS file system from the HPC cluster instances.
- D. Create an Amazon FSx for Lustre file system. Import the data directly into the FSx for Lustre file system. Access the FSx for Lustre file system from the HPC cluster instances.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 693

A company stores millions of objects in Amazon S3. The data is in JSON format and Apache Parquet format. The data is partitioned, and new objects are added daily. A solutions architect needs to create a solution so that employees can use SQL to perform one-time queries against all the data. The solution must avoid code changes and must minimize operational overhead.

Which solution will meet these requirements?

- A. Use S3 Select to perform queries against all the S3 objects.
- B. Create an AWS Glue table and an AWS Glue crawler. Schedule the crawler to run daily. Perform queries with Amazon Athena.
- C. Create an Amazon EMR cluster. Set up EMR File System (EMRFS) to access the S3 bucket.
Perform queries with Apache Spark.
- D. Create an Amazon Redshift cluster. Schedule an AWS Lambda function to perform the COPY command on the Redshift cluster to load the S3 data. Perform queries on the Redshift cluster.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 694

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. Run a cron job script on the EC2 instances to upload files to the S3 data lake.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 695

A company hosts a multi-tier web application on Amazon Linux Amazon EC2 instances behind an Application Load Balancer. The instances run in an Auto Scaling group across multiple Availability Zones. The company observes that the Auto Scaling group launches more On-Demand Instances when the application's end users access high volumes of static web content. The company wants to optimize cost. What should a solutions architect do to redesign the application MOST cost-effectively'?

- A. Update the Auto Scaling group to use Reserved Instances instead of On-Demand Instances.
- B. Update the Auto Scaling group to scale by launching Spot Instances instead of On-Demand Instances.
- C. Create an Amazon CloudFront distribution to host the static web contents from an Amazon S3 bucket.
- D. Create an AWS Lambda function behind an Amazon API Gateway API to host the static website contents.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 696

A company needs to develop a repeatable solution to process time-ordered information from websites around the world. The company collects the data from the websites by using Amazon Kinesis Data Streams and stores the data in Amazon S3. The processing logic needs to collect events and handle data from the last 5 years. The processing logic also must generate results in an S3 bucket so that a business intelligence application can analyze and compare the results. The processing must be repeated multiple times.

What should a solutions architect do to meet these requirements?

- A. Use Amazon S3 to collect events. Create an AWS Lambda function to process the events. Create different Lambda functions to handle repeated processing.
- B. Use Amazon EventBridge (Amazon CloudWatch Events) to collect events. Set AWS Lambda as an event target. Use EventBridge (CloudWatch Events) to create an archive for the events and to replay the events.
- C. Use an Amazon Simple Queue Service (Amazon SQS) FIFO queue to collect events. Process the events by using Amazon EC2. Use AWS Step Functions to create an archive for the events and to replay the events.
- D. Use Amazon Managed Streaming for Apache Kafka (Amazon MSK) to collect events. Process the events by using Amazon Elastic Kubernetes Service (Amazon EKS). Use Amazon MSK to create an archive for the events and to replay the events.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 697

A company runs a public three-tier web application in a VPC. The application runs on Amazon EC2 instances across multiple Availability Zones. The EC2 instances that run in private subnets need to communicate with a license server over the internet. The company needs a managed solution that minimizes operational maintenance.

Which solution meets these requirements?

- A. Provision a NAT instance in a public subnet. Modify each private subnet's route table with a default route that points to the NAT instance.



- B. Provision a NAT instance in a private subnet. Modify each private subnet's route table with a default route that points to the NAT instance.
- C. Provision a NAT gateway in a public subnet. Modify each private subnet's route table with a default route that points to the NAT gateway.
- D. Provision a NAT gateway in a private subnet. Modify each private subnet's route table with a default route that points to the NAT gateway.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 698

A company needs to store data from its healthcare application. The application's data frequently changes. A new regulation requires audit access at all levels of the stored data. The company hosts the application on an on-premises infrastructure that is running out of storage capacity. A solutions architect must securely migrate the existing data to AWS while satisfying the new regulation. Which solution will meet these requirements?

- A. Use AWS DataSync to move the existing data to Amazon S3. Use AWS CloudTrail to log data events.
- B. Use AWS Snowcone to move the existing data to Amazon S3. Use AWS CloudTrail to log management events.
- C. Use Amazon S3 Transfer Acceleration to move the existing data to Amazon S3. Use AWS CloudTrail to log data events.
- D. Use AWS Storage Gateway to move the existing data to Amazon S3. Use AWS CloudTrail to log management events.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 699

A company stores confidential data in an Amazon Aurora PostgreSQL database in the ap-southeast-3 Region. The database is encrypted with an AWS Key Management Service (AWS KMS) customer managed key. The company was recently acquired and must securely share a backup of the database with the acquiring company's AWS account in ap-southeast-3. What should a solutions architect do to meet these requirements?

- A. Create a database snapshot. Copy the snapshot to a new unencrypted snapshot. Share the new snapshot with the acquiring company's AWS account.
- B. Create a database snapshot. Add the acquiring company's AWS account to the KMS key policy. Share the snapshot with the acquiring company's AWS account.
- C. Create a database snapshot that uses a different AWS managed KMS key. Add the acquiring company's AWS account to the KMS key alias. Share the snapshot with the acquiring company's AWS account.
- D. Create a database snapshot. Download the database snapshot. Upload the database snapshot to an Amazon S3 bucket. Update the S3 bucket policy to allow access from the acquiring company's AWS account.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 700

A solutions architect must secure a VPC network that hosts Amazon EC2 instances. The EC2 instances contain highly sensitive data and run in a private subnet. According to company policy, the EC2 instances that run in the VPC can access only approved third-party software repositories on the internet for software product updates that use the third party's URL. Other internet traffic must be blocked. Which solution meets these requirements?

- A. Update the route table for the private subnet to route the outbound traffic to an AWS Network Firewall firewall. Configure domain list rule groups.
- B. Set up an AWS WAF web ACL. Create a custom set of rules that filter traffic requests based on source and destination IP address range sets.
- C. Implement strict inbound security group rules. Configure an outbound rule that allows traffic only to the authorized software repositories on the internet by specifying the URLs.
- D. Configure an Application Load Balancer (ALB) in front of the EC2 instances. Direct all outbound traffic to the ALB. Use a URL-based rule listener in the ALB's target group for outbound access to the internet.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 701

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

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

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

Correct Answer: AE

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 702

A company has a web application hosted over 10 Amazon EC2 instances with traffic directed by Amazon Route 53. The company occasionally experiences a timeout error when attempting to browse the application. The networking team finds that some DNS queries return IP addresses of unhealthy instances, resulting in the timeout error.

What should a solutions architect implement to overcome these timeout errors?

- A. Create a Route 53 simple routing policy record for each EC2 instance. Associate a health check with each record.
- B. Create a Route 53 failover routing policy record for each EC2 instance. Associate a health check with each record.
- C. Create an Amazon CloudFront distribution with EC2 instances as its origin. Associate a health check with the EC2 instances.
- D. Create an Application Load Balancer (ALB) with a health check in front of the EC2 instances. Route to the ALB from Route 53.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 703

A hospital is designing a new application that gathers symptoms from patients. The hospital has decided to use Amazon Simple Queue Service (Amazon SQS) and Amazon Simple Notification Service (Amazon SNS) in the architecture. A solutions architect is reviewing the infrastructure design. Data must be encrypted at rest and in transit. Only authorized personnel of the hospital should be able to access the data.

Which combination of steps should the solutions architect take to meet these requirements?

(Choose two.)

- A. Turn on server-side encryption on the SQS components. Update the default key policy to restrict key usage to a set of authorized principals.
- B. Turn on server-side encryption on the SNS components by using an AWS Key Management Service (AWS KMS) customer managed key. Apply a key policy to restrict key usage to a set of authorized principals.
- C. Turn on encryption on the SNS components. Update the default key policy to restrict key usage to a set of authorized principals. Set a condition in the topic policy to allow only encrypted connections over TLS.
- D. Turn on server-side encryption on the SQS components by using an AWS Key Management Service (AWS KMS) customer managed key. Apply a key policy to restrict key usage to a set of authorized principals. Set a condition in the queue policy to allow only encrypted connections over TLS.
- E. Turn on server-side encryption on the SQS components by using an AWS Key Management Service (AWS KMS) customer managed key. Apply an IAM policy to restrict key usage to a set of authorized principals. Set a condition in the queue policy to allow only encrypted connections over TLS.

Correct Answer: CD

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 704

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

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

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 705

A company has a multi-tier application deployed on several Amazon EC2 instances in an Auto Scaling group. An Amazon RDS for Oracle instance is the application's data layer that uses Oracle-specific PL/SQL functions. Traffic to the application has been steadily increasing. This is causing the EC2 instances to become overloaded and the RDS instance to run out of storage. The Auto Scaling group does not have any scaling metrics and defines the minimum healthy instance count only. The company predicts that traffic will continue to increase at a steady but unpredictable rate before leveling off. What should a solutions architect do to ensure the system can automatically scale for the increased traffic? (Choose two.)

- A. Configure storage Auto Scaling on the RDS for Oracle instance.
- B. Migrate the database to Amazon Aurora to use Auto Scaling storage.
- C. Configure an alarm on the RDS for Oracle instance for low free storage space.
- D. Configure the Auto Scaling group to use the average CPU as the scaling metric.
- E. Configure the Auto Scaling group to use the average free memory as the scaling metric.



Correct Answer: AC

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 706

A company runs an AWS Lambda function in private subnets in a VPC. The subnets have a default route to the internet through an Amazon EC2 NAT instance. The Lambda function processes input data and saves its output as an object to Amazon S3. Intermittently, the Lambda function times out while trying to upload the object because of saturated traffic on the NAT instance's network. The company wants to access Amazon S3 without traversing the internet. Which solution will meet these requirements?

- A. Replace the EC2 NAT instance with an AWS managed NAT gateway.
- B. Increase the size of the EC2 NAT instance in the VPC to a network optimized instance type.
- C. Provision a gateway endpoint for Amazon S3 in the VPC. Update the route tables of the subnets accordingly.
- D. Provision a transit gateway. Place transit gateway attachments in the private subnets where the Lambda function is running.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 707

A solutions architect is helping a company design storage for a high performance computing (HPC) environment that is based on Amazon Linux. The workload stores and processes a large number of engineering drawings that require shared storage. The company needs a solution that can support hundreds of Amazon EC2 instances. The solution also must provide sub-millisecond latencies.

Which solution meets these requirements?

- A. Amazon EC2 instance store
- B. Amazon Elastic Block Store (Amazon EBS) Provisioned IOPS SSD (io2)
- C. Amazon Elastic File System (Amazon EFS)
- D. Amazon FSx for Lustre

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 708

A company is running a critical business application on an Amazon EC2 instance. The EC2 instance is hosting an Apache web server and a MySQL database server. The application serves static content and dynamic content to end users. The application is experiencing severe availability issues because of heavy user demand. The company needs a solution that resolves the availability issues with the least operational effort and the least change to the application. What should a solutions architect do to meet these requirements?

- A. Deploy the application and the web server on AWS Fargate. Use a Network Load Balancer to route traffic. Migrate the database to Amazon DynamoDB.
- B. Create an Amazon Machine Image (AMI) from the current EC2 instance. Create an Auto Scaling group to provide more capacity as needed. Use a Network Load Balancer to route traffic.
- C. Host static content on Amazon S3. Deploy the application and the web server on AWS Fargate. Use an Application Load Balancer to route traffic. Migrate the database to Amazon Aurora Serverless.
- D. Host static content on Amazon S3. Deploy the application on EC2 instances that are configured in an Auto Scaling group. Use an Application Load Balancer to route traffic. Migrate the database to Amazon DynamoDB.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:



QUESTION 709

A gaming company uses Amazon DynamoDB to store user information such as geographic location, player data, and leaderboards. The company needs to configure continuous backups to an Amazon S3 bucket with a minimal amount of coding. The backups must not affect availability of the application and must not affect the read capacity units (RCUs) that are defined for the table. Which solution meets these requirements?

- A. Use an Amazon EMR cluster. Create an Apache Hive job to back up the data to Amazon S3.
- B. Export the data directly from DynamoDB to Amazon S3 with continuous backups. Turn on point-in-time recovery for the table.
- C. Configure Amazon DynamoDB Streams. Create an AWS Lambda function to consume the stream and export the data to an Amazon S3 bucket.
- D. Create an AWS Lambda function to export the data from the database tables to Amazon S3 on a regular basis. Turn on point-in-time recovery for the table.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 710

A company is building a new dynamic ordering website. The company wants to minimize server maintenance and patching. The website must be highly available and must scale read and write capacity as quickly as possible to meet changes in user demand.

Which solution will meet these requirements?

- A. Host static content in Amazon S3. Host dynamic content by using Amazon API Gateway and AWS Lambda. Use Amazon DynamoDB with on-demand capacity for the database. Configure Amazon CloudFront to deliver the website content.
- B. Host static content in Amazon S3. Host dynamic content by using Amazon API Gateway and AWS Lambda. Use Amazon Aurora with Aurora Auto Scaling for the database. Configure Amazon CloudFront to deliver the website content.
- C. Host all the website content on Amazon EC2 instances. Create an Auto Scaling group to scale the EC2 instances. Use an Application Load Balancer to distribute traffic. Use Amazon DynamoDB with provisioned write capacity for the database.

D. Host all the website content on Amazon EC2 instances. Create an Auto Scaling group to scale the EC2 instances. Use an Application Load Balancer to distribute traffic. Use Amazon Aurora with Aurora Auto Scaling for the database.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 711

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

Which combination of actions will meet this requirement with the LEAST operational overhead?

(Choose two.)

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

Correct Answer: AD

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 712

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

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 713

A company has a web application for travel ticketing. The application is based on a database that runs in a single data center in North America. The company wants to expand the application to serve a global user base. The company needs to deploy the application to multiple AWS Regions.

Average latency must be less than 1 second on updates to the reservation database.

The company wants to have separate deployments of its web platform across multiple Regions.

However, the company must maintain a single primary reservation database that is globally consistent.

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

- A. Convert the application to use Amazon DynamoDB. Use a global table for the center reservation table. Use the correct Regional endpoint in each Regional deployment.
- B. Migrate the database to an Amazon Aurora MySQL database. Deploy Aurora Read Replicas in each Region. Use the correct Regional endpoint in each Regional deployment for access to the database.
- C. Migrate the database to an Amazon RDS for MySQL database. Deploy MySQL read replicas in each Region. Use the correct Regional endpoint in each Regional deployment for access to the database.
- D. Migrate the application to an Amazon Aurora Serverless database. Deploy instances of the database to each Region. Use the correct Regional endpoint in each Regional deployment to access the database. Use AWS Lambda functions to process event streams in each Region to synchronize the databases.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 714

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

- A. Migrate the queue to a redundant pair (active/standby) of RabbitMQ instances on Amazon MQ.
Create a Multi-AZ Auto Scaling group for EC2 instances that host the application. Create another Multi-AZ Auto Scaling group for EC2 instances that host the PostgreSQL database.
- B. Migrate the queue to a redundant pair (active/standby) of RabbitMQ instances on Amazon MQ.
Create a Multi-AZ Auto Scaling group for EC2 instances that host the application. Migrate the database to run on a Multi-AZ deployment of Amazon RDS for PostgreSQL.
- C. Create a Multi-AZ Auto Scaling group for EC2 instances that host the RabbitMQ queue. Create another Multi-AZ Auto Scaling group for EC2 instances that host the application. Migrate the database to run on a Multi-AZ deployment of Amazon RDS for PostgreSQL.
- D. Create a Multi-AZ Auto Scaling group for EC2 instances that host the RabbitMQ queue. Create another Multi-AZ Auto Scaling group for EC2 instances that host the application. Create a third Multi-AZ Auto Scaling group for EC2 instances that host the PostgreSQL database.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 715

A company is building an application in the AWS Cloud. The application will store data in Amazon S3 buckets in two AWS Regions. The company must use an AWS Key Management Service (AWS KMS) customer managed key to encrypt all data that is stored in the S3 buckets. The data in both S3 buckets must be encrypted and decrypted with the same KMS key. The data and the key must be stored in each of the two Regions. Which solution will meet these requirements with the LEAST operational overhead?

- A. Create an S3 bucket in each Region. Configure the S3 buckets to use server-side encryption with Amazon S3 managed encryption keys (SSE-S3). Configure replication between the S3 buckets.
- B. Create a customer managed multi-Region KMS key. Create an S3 bucket in each Region.
Configure replication between the S3 buckets. Configure the application to use the KMS key with client-side encryption.
- C. Create a customer managed KMS key and an S3 bucket in each Region. Configure the S3 buckets to use server-side encryption with Amazon S3 managed encryption keys (SSE-S3).
Configure replication between the S3 buckets.
- D. Create a customer managed KMS key and an S3 bucket in each Region. Configure the S3 buckets to use server-side encryption with AWS KMS keys (SSE-KMS). Configure replication between the S3 buckets.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 716

A company has an API that receives real-time data from a fleet of monitoring devices. The API stores this data in an Amazon RDS DB instance for later analysis. The amount of data that the monitoring devices send to the API fluctuates. During periods of heavy traffic, the API often returns timeout errors. After an inspection of the logs, the company determines that the database is not capable of processing the volume of write traffic that comes from the API. A solutions architect must minimize the number of connections to the database and must ensure that data is not lost during periods of heavy traffic. Which solution will meet these requirements?

- A. Increase the size of the DB instance to an instance type that has more available memory.
- B. Modify the DB instance to be a Multi-AZ DB instance. Configure the application to write to all active RDS DB instances.
- C. Modify the API to write incoming data to an Amazon Simple Queue Service (Amazon SQS) queue.
Use an AWS Lambda function that Amazon SQS invokes to write data from the queue to the database.

D. Modify the API to write incoming data to an Amazon Simple Notification Service (Amazon SNS) topic. Use an AWS Lambda function that Amazon SNS invokes to write data from the topic to the database.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 717

A company is building a shopping application on AWS. The application offers a catalog that changes once each month and needs to scale with traffic volume. The company wants the lowest possible latency from the application. Data from each user's shopping cart needs to be highly available. User session data must be available even if the user is disconnected and reconnects.

What should a solutions architect do to ensure that the shopping cart data is preserved at all times?

- A. Configure an Application Load Balancer to enable the sticky sessions feature (session affinity) for access to the catalog in Amazon Aurora.
- B. Configure Amazon ElastiCache for Redis to cache catalog data from Amazon DynamoDB and shopping cart data from the user's session.
- C. Configure Amazon OpenSearch Service (Amazon Elasticsearch Service) to cache catalog data from Amazon DynamoDB and shopping cart data from the user's session.
- D. Configure an Amazon EC2 instance with Amazon Elastic Block Store (Amazon EBS) storage for the catalog and shopping cart. Configure automated snapshots.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 718

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

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

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

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 719

A company hosts a popular website in the AWS Cloud. A solutions architect needs to provide reports about user click behavior in near-real time as users navigate the website.

Which solution will meet this requirement?

- A. Store the clickstream data in Amazon DynamoDB. Deploy an application that runs on AWS Elastic Beanstalk to process and analyze the data.
- B. Push the clickstream data from each session to an Amazon Kinesis data stream. Analyze the data by using Amazon Kinesis Data Analytics.
- C. Store the clickstream data in an Amazon S3 bucket. Order the data by timestamp. Process the data with an AWS Lambda function that is subscribed to object creation events on the S3 bucket.
- D. Forward the clickstream data to Amazon Simple Queue Service (Amazon SQS). Store the data in an Amazon RDS for MySQL DB instance. Deploy Amazon EC2 instances to process and analyze the data.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 720

A company runs analytics software on Amazon EC2 instances. The software accepts job requests from users to process data that has been uploaded to Amazon S3. Users report that some submitted data is not being processed. Amazon CloudWatch reveals that the EC2 instances have a consistent CPU utilization at or near 100%. The company wants to improve system performance and scale the system based on user load. What should a solutions architect do to meet these requirements?

- A. Create a copy of the instance. Place all instances behind an Application Load Balancer.
- B. Create an S3 VPC endpoint for Amazon S3. Update the software to reference the endpoint.
- C. Stop the EC2 instances. Modify the instance type to one with a more powerful CPU and more memory. Restart the instances.
- D. Route incoming requests to Amazon Simple Queue Service (Amazon SQS). Configure an EC2 Auto Scaling group based on queue size. Update the software to read from the queue.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 721

A company is running a web-based game in two Availability Zones in the us-west-2 Region. The web servers use an Application Load Balancer (ALB) in public subnets. The ALB has an SSL certificate from AWS Certificate Manager (ACM) with a custom domain name. The game is written in JavaScript and runs entirely in a user's web browser.

The game is increasing in popularity in many countries around the world. The company wants to update the application architecture and optimize costs without compromising performance.

What should a solutions architect do to meet these requirements?

- A. Use Amazon CloudFront and create a global distribution that points to the ALB. Reuse the existing certificate from ACM for the CloudFront distribution. Use Amazon Route 53 to update the application alias to point to the distribution.
- B. Use AWS CloudFormation to deploy the application stack to AWS Regions near countries where the game is popular. Use ACM to create a new certificate for each application instance. Use Amazon Route 53 with a geolocation routing policy to direct traffic to the local application instance.
- C. Use Amazon S3 and create an S3 bucket in AWS Regions near countries where the game is popular. Deploy the HTML and JavaScript files to each S3 bucket. Use ACM to create a new certificate for each S3 bucket. Use Amazon Route 53 with a geolocation routing policy to direct traffic to the local S3 bucket.
- D. Use Amazon S3 and create an S3 bucket in us-west-2. Deploy the HTML and JavaScript files to the S3 bucket. Use Amazon CloudFront and create a global distribution with the S3 bucket as the origin. Use ACM to create a new certificate for the distribution. Use Amazon Route 53 to update the application alias to point to the distribution.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 722

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

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

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

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

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 723

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.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 724

A company wants to run an in-memory database for a latency-sensitive application that runs on Amazon EC2 instances. The application processes more than 100,000 transactions each minute and requires high network throughput. A solutions architect needs to provide a cost-effective network design that minimizes data transfer charges.

Which solution meets these requirements?

- A. Launch all EC2 instances in the same Availability Zone within the same AWS Region. Specify a placement group with cluster strategy when launching EC2 instances.
- B. Launch all EC2 instances in different Availability Zones within the same AWS Region. Specify a placement group with partition strategy when launching EC2 instances.
- C. Deploy an Auto Scaling group to launch EC2 instances in different Availability Zones based on a network utilization target.
- D. Deploy an Auto Scaling group with a step scaling policy to launch EC2 instances in different Availability Zones.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 725

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

Which combination of steps should a solutions architect take to meet these requirements?

(Choose two.)

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

Correct Answer: AE

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 726

A company has designed an application where users provide small sets of textual data by calling a public API. The application runs on AWS and includes a public Amazon API Gateway API that forwards requests to an AWS Lambda function for processing. The Lambda function then writes the data to an Amazon Aurora Serverless database for consumption.

The company is concerned that it could lose some user data if a Lambda function fails to process the request properly or reaches a concurrency limit.

What should a solutions architect recommend to resolve this concern?

- A. Split the existing Lambda function into two Lambda functions. Configure one function to receive API Gateway requests and put relevant items into Amazon Simple Queue Service (Amazon SQS). Configure the other function to read items from Amazon SQS and save the data into Aurora.
- B. Configure the Lambda function to receive API Gateway requests and write relevant items to Amazon ElastiCache. Configure ElastiCache to save the data into Aurora.
- C. Increase the memory for the Lambda function. Configure Aurora to use the Multi-AZ feature.

D. Split the existing Lambda function into two Lambda functions. Configure one function to receive API Gateway requests and put relevant items into Amazon Simple Notification Service (Amazon SNS). Configure the other function to read items from Amazon SNS and save the data into Aurora.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 727

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

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

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

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 728

A company is making a prototype of the infrastructure for its new website by manually provisioning the necessary infrastructure. This infrastructure includes an Auto Scaling group, an Application Load Balancer, and an Amazon RDS database. After the configuration has been thoroughly validated, the company wants the capability to immediately deploy the infrastructure for development and production use in two Availability Zones in an automated fashion.

What should a solutions architect recommend to meet these requirements?

- A. Use AWS Systems Manager to replicate and provision the prototype infrastructure in two Availability Zones.
- B. Define the infrastructure as a template by using the prototype infrastructure as a guide. Deploy the infrastructure with AWS CloudFormation.
- C. Use AWS Config to record the inventory of resources that are used in the prototype infrastructure.
Use AWS Config to deploy the prototype infrastructure into two Availability Zones.
- D. Use AWS Elastic Beanstalk and configure it to use an automated reference to the prototype infrastructure to automatically deploy new environments in two Availability Zones.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 729

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. Copy the data to the device. Create a custom transformation job by using AWS Glue.
- D. 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.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 730

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

What should a solutions architect do to grant the permissions?

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

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 731

A payment processing company records all voice communication with its customers and stores the audio files in an Amazon S3 bucket. The company needs to capture the text from the audio files. The company must remove from the text any personally identifiable information (PII) that belongs to customers.

What should a solutions architect do to meet these requirements?

- A. Process the audio files by using Amazon Kinesis Video Streams. Use an AWS Lambda function to scan for known PII patterns.
- B. When an audio file is uploaded to the S3 bucket, invoke an AWS Lambda function to start an Amazon Textract task to analyze the call recordings.
- C. Configure an Amazon Transcribe transcription job with PII redaction turned on. When an audio file is uploaded to the S3 bucket, invoke an AWS Lambda function to start the transcription job. Store the output in a separate S3 bucket.
- D. Create an Amazon Connect contact flow that ingests the audio files with transcription turned on. Embed an AWS Lambda function to scan for known PII patterns. Use Amazon EventBridge (Amazon CloudWatch Events) to start the contact flow when an audio file is uploaded to the S3 bucket.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 732

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.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 733

A company has a small Python application that processes JSON documents and outputs the results to an on-premises SQL database. The application runs thousands of times each day. The company wants to move the application to the AWS Cloud. The company needs a highly available solution that maximizes scalability and minimizes operational overhead.

Which solution will meet these requirements?

- A. Place the JSON documents in an Amazon S3 bucket. Run the Python code on multiple Amazon EC2 instances to process the documents. Store the results in an Amazon Aurora DB cluster.
- B. Place the JSON documents in an Amazon S3 bucket. Create an AWS Lambda function that runs the Python code to process the documents as they arrive in the S3 bucket. Store the results in an Amazon Aurora DB cluster.
- C. Place the JSON documents in an Amazon Elastic Block Store (Amazon EBS) volume. Use the EBS Multi-Attach feature to attach the volume to multiple Amazon EC2 instances. Run the Python code on the EC2 instances to process the documents. Store the results on an Amazon RDS DB instance.
- D. Place the JSON documents in an Amazon Simple Queue Service (Amazon SQS) queue as messages. Deploy the Python code as a container on an Amazon Elastic Container Service (Amazon ECS) cluster that is configured with the Amazon EC2 launch type. Use the container to process the SQS messages. Store the results on an Amazon RDS DB instance.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 734

A hospital recently deployed a RESTful API with Amazon API Gateway and AWS Lambda. The hospital uses API Gateway and Lambda to upload reports that are in PDF format and JPEG format. The hospital needs to modify the Lambda code to identify protected health information (PHI) in the reports.

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

- A. Use existing Python libraries to extract the text from the reports and to identify the PHI from the extracted text.
- B. Use Amazon Textract to extract the text from the reports. Use Amazon SageMaker to identify the PHI from the extracted text.
- C. Use Amazon Textract to extract the text from the reports. Use Amazon Comprehend Medical to identify the PHI from the extracted text.
- D. Use Amazon Rekognition to extract the text from the reports. Use Amazon Comprehend Medical to identify the PHI from the extracted text.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:



QUESTION 735

A solutions architect is investigating AWS file storage solutions that can be used with a company's on-premises Linux servers and applications. The company has an existing VPN connection set up between the company's VPC and its on-premises network.

Which AWS services should the solutions architect use? (Choose two.)

- A. AWS Backup
- B. AWS DataSync
- C. AWS Snowball Edge
- D. AWS Storage Gateway
- E. Amazon Elastic File System (Amazon EFS)

Correct Answer: BE

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 736

A company hosts a two-tier application on Amazon EC2 instances and Amazon RDS. The application's demand varies based on the time of day. The load is minimal after work hours and on weekends. The EC2 instances run in an EC2 Auto Scaling group that is configured with a minimum of two instances and a maximum of five instances. The application must be available at all times, but the company is concerned about overall cost.

Which solution meets the availability requirement MOST cost-effectively?

- A. Use all EC2 Spot Instances. Stop the RDS database when it is not in use.
- B. Purchase EC2 Instance Savings Plans to cover five EC2 instances. Purchase an RDS Reserved DB Instance.

- C. Purchase two EC2 Reserved Instances. Use up to three additional EC2 Spot Instances as needed. Stop the RDS database when it is not in use.
- D. Purchase EC2 Instance Savings Plans to cover two EC2 instances. Use up to three additional EC2 On-Demand Instances as needed. Purchase an RDS Reserved DB Instance.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 737

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? (Choose 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.

Correct Answer: AC

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 738

A company is building a microservices-based application to process information. The data that is transferred between the microservices is in CSV format. Communication between the microservices requires a latency response of less than 100 ms.

Which solution will meet these requirements?

- A. Run the microservices by using AWS Batch. Use an Amazon API Gateway API for communication between the microservices.
- B. Create an extract, transform, and load (ETL) data pipeline by using AWS Glue. Feed the data pipeline by using AWS Lambda functions.
- C. Build the microservices to run on AWS Lambda functions. Use an Amazon Simple Notification Service (Amazon SNS) topic for communication between the microservices.
- D. Build the microservices to run on containers in an Amazon Elastic Kubernetes Service (Amazon EKS) cluster. Configure Amazon EventBridge (Amazon CloudWatch Events) to receive events and route the events to the appropriate microservices.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 739

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

Which solution meets these requirements?

- A. Use Amazon DynamoDB with auto scaling Use on-demand backups and Amazon DynamoDB Streams
- B. Use Amazon Redshift. Configure concurrency scaling. Activate audit logging. Perform database snapshots every 4 hours
- C. Use Amazon RDS with Provisioned IOPS Activate the database auditing parameter. Perform database snapshots every 5 hours
- D. Use Amazon Aurora MySQL with auto scaling. Activate the database auditing parameter

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 740

A company is running a web application on Amazon EC2 instances in an Auto Scaling group. The application uses a database that runs on an Amazon RDS for PostgreSQL DB instance. The application performs slowly as traffic increases, and the database experiences a heavy read load during periods of high traffic.

Which actions should a solutions architect take to resolve these performance issues? (Choose two.)

- A. Enable auto scaling for the DB instance.
- B. Create a read replica for the DB instance. Configure the application to send read traffic to the read replica.
- C. Enable Multi-AZ for the DB instance. Configure the application to send read traffic to the standby DB instance.
- D. Create an Amazon ElastiCache cluster. Configure the application to cache query results in the ElastiCache cluster.
- E. Configure the Auto Scaling group subnets to ensure that the EC2 instances are provisioned in the same Availability Zone as the DB instance.

Correct Answer: BD

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 741

A company runs demonstration environments for its customers on Amazon EC2 instances. Each environment is isolated in its own VPC. The company's operations team needs to be notified when RDP or SSH access to an environment has been established.

What should a solutions architect recommend to meet these requirements?

- A. Configure Amazon CloudWatch Application Insights to create AWS Systems Manager OpsItems when RDP or SSH access is detected.
- B. Configure the EC2 instances with an IAM instance profile that has an IAM role with the AmazonSSMManagedInstanceCore policy attached.
- C. Publish VPC flow logs to Amazon CloudWatch Logs. Create required metric filters. Create an Amazon CloudWatch metric alarm with a notification action for when the alarm is in the ALARM state.
- D. Configure an Amazon EventBridge (Amazon CloudWatch Events) rule to listen for events of type EC2 Instance State-change Notification. Configure an Amazon Simple Notification Service (Amazon SNS) topic as a target. Subscribe the operations team to the topic.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 742

A company recently created a disaster recovery site in a different AWS Region. The company needs to transfer large amounts of data back and forth between NFS file systems in the two Regions on a periodic basis.

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

- A. Use AWS DataSync.
- B. Use AWS Snowball devices.
- C. Set up an SFTP server on Amazon EC2.
- D. Use AWS Database Migration Service (AWS DMS).

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 743

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

What should a solutions architect do to meet these requirements?

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

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 744

A company uses an Amazon S3 bucket as its data lake storage platform. The S3 bucket contains a massive amount of data that is accessed randomly by multiple teams and hundreds of applications. The company wants to reduce the S3 storage costs and provide immediate availability for frequently accessed objects.

What is the MOST operationally efficient solution that meets these requirements?

- A. Create an S3 Lifecycle rule to transition objects to the S3 Intelligent-Tiering storage class.
- B. Store objects in Amazon S3 Glacier. Use S3 Select to provide applications with access to the data.
- C. Use data from S3 storage class analysis to create S3 Lifecycle rules to automatically transition objects to the S3 Standard-Infrequent Access (S3 Standard-IA) storage class.
- D. Transition objects to the S3 Standard-Infrequent Access (S3 Standard-IA) storage class. Create an AWS Lambda function to transition objects to the S3 Standard storage class when they are accessed by an application.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 745

A company is using a SQL database to store movie data that is publicly accessible. The database runs on an Amazon RDS Single-AZ DB instance. A script runs queries at random intervals each day to record the number of new movies that have been added to the database. The script must report a final total during business hours.

The company's development team notices that the database performance is inadequate for development tasks when the script is running. A solutions architect must recommend a solution to resolve this issue.

Which solution will meet this requirement with the LEAST operational overhead?

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

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 746

An application that is hosted on Amazon EC2 instances needs to access an Amazon S3 bucket.

Traffic must not traverse the internet.

How should a solutions architect configure access to meet these requirements?

- A. Create a private hosted zone by using Amazon Route 53.
- B. Set up a gateway VPC endpoint for Amazon S3 in the VPC.
- C. Configure the EC2 instances to use a NAT gateway to access the S3 bucket.
- D. Establish an AWS Site-to-Site VPN connection between the VPC and the S3 bucket.

Correct Answer: B

Section: (none)



Explanation

Explanation/Reference:

Explanation:

QUESTION 747

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.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 748

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

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

What should a solutions architect recommend to fix the application?

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

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 749

A company uses an application to present metrics from sporting events to the public. The application must scale quickly during live events and must store these metrics for long-term reporting purposes. The company's architecture includes the following:

Amazon EC2 instances that run in an Auto Scaling group in private subnets
A Network Load Balancer that runs in public subnets
A MongoDB database cluster that runs across multiple EC2 instances
A solutions architect must implement a solution that minimizes operational overhead. The solution also must be able to scale automatically.

What should the solutions architect set up to meet these requirements?

- A. An Amazon DynamoDB database
- B. An Amazon RDS for MySQL DB instance
- C. EC2 instances that run MySQL
- D. Amazon Redshift

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 750

A company wants to implement a disaster recovery plan for its primary on-premises file storage volume. The file storage volume is mounted from an Internet Small Computer Systems Interface (iSCSI) device on a local storage server. The file storage volume holds hundreds of terabytes (TB) of data. The company wants to ensure that end users retain immediate access to all file types from the on-premises systems without experiencing latency. Which solution will meet these requirements with the LEAST amount of change to the company's existing infrastructure?

- A. Provision an Amazon S3 File Gateway as a virtual machine (VM) that is hosted on premises. Set the local cache to 10 TB. Modify existing applications to access the files through the NFS protocol. To recover from a disaster, provision an Amazon EC2 instance and mount the S3 bucket that contains the files.
- B. Provision an AWS Storage Gateway tape gateway. Use a data backup solution to back up all existing data to a virtual tape library. Configure the data backup solution to run nightly after the initial backup is complete. To recover from a disaster, provision an Amazon EC2 instance and restore the data to an Amazon Elastic Block Store (Amazon EBS) volume from the volumes in the virtual tape library.
- C. Provision an AWS Storage Gateway Volume Gateway cached volume. Set the local cache to 10 TB. Mount the Volume Gateway cached volume to the existing file server by using iSCSI, and copy all files to the storage volume. Configure scheduled snapshots of the storage volume. To recover from a disaster, restore a snapshot to an Amazon Elastic Block Store (Amazon EBS) volume and attach the EBS volume to an Amazon EC2 instance.
- D. Provision an AWS Storage Gateway Volume Gateway stored volume with the same amount of disk space as the existing file storage volume. Mount the Volume Gateway stored volume to the existing file server by using iSCSI, and copy all files to the storage volume. Configure scheduled snapshots of the storage volume. To recover from a disaster, restore a snapshot to an Amazon Elastic Block Store (Amazon EBS) volume and attach the EBS volume to an Amazon EC2 instance.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:

QUESTION 751

A company has an Amazon S3 data lake that is governed by AWS Lake Formation. The company wants to create a visualization in Amazon QuickSight by joining the data in the data lake with operational data that is stored in an Amazon Aurora MySQL database. The company wants to enforce column-level authorization so that the company's marketing team can access only a subset of columns in the database. Which solution will meet these requirements with the LEAST operational overhead?

- A. Use Amazon EMR to ingest the data directly from the database to the QuickSight SPICE engine. Include only the required columns.
- B. Use AWS Glue Studio to ingest the data from the database to the S3 data lake. Attach an IAM policy to the QuickSight users to enforce column-level access control. Use Amazon S3 as the data source in QuickSight.
- C. Use AWS Glue Elastic Views to create a materialized view for the database in Amazon S3. Create an S3 bucket policy to enforce column-level access control for the QuickSight users. Use Amazon S3 as the data source in QuickSight.
- D. Use a Lake Formation blueprint to ingest the data from the database to the S3 data lake. Use Lake Formation to enforce column-level access control for the QuickSight users. Use Amazon Athena as the data source in QuickSight.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:

QUESTION 752

A company's e-commerce website has unpredictable traffic and uses AWS Lambda functions to directly access a private Amazon RDS for PostgreSQL DB instance. The company wants to maintain predictable database performance and ensure that the Lambda invocations do not overload the database with too many connections. What should a solutions architect do to meet these requirements?

- A. Point the client driver at an RDS custom endpoint. Deploy the Lambda functions inside a VPC.
- B. Point the client driver at an RDS proxy endpoint. Deploy the Lambda functions inside a VPC.
- C. Point the client driver at an RDS custom endpoint. Deploy the Lambda functions outside a VPC.
- D. Point the client driver at an RDS proxy endpoint. Deploy the Lambda functions outside a VPC.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:

QUESTION 753

A company runs an application that receives data from thousands of geographically dispersed remote devices that use UDP. The application processes the data immediately and sends a message back to the device if necessary. No data is stored. The company needs a solution that minimizes latency for the data transmission from the devices. The solution also must provide rapid failover to another AWS Region. Which solution will meet these requirements?

- A. Configure an Amazon Route 53 failover routing policy Create a Network Load Balancer (NLB) in each of the two Regions Configure the NLB to invoke an AWS Lambda function to process the data
- B. Use AWS Global Accelerator Create a Network Load Balancer (NLB) in each of the two Regions as an endpoint. Create an Amazon Elastic Container Service (Amazon ECS) cluster with the Fargate launch type Create an ECS service on the cluster Set the ECS service as the target for the NLB Process the data in Amazon ECS.
- C. Use AWS Global Accelerator Create an Application Load Balancer (ALB) in each of the two Regions as an endpoint Create an Amazon Elastic Container Service (Amazon ECS) cluster with the Fargate launch type Create an ECS service on the cluster. Set the ECS service as the target for the ALB Process the data in Amazon ECS
- D. Configure an Amazon Route 53 failover routing policy Create an Application Load Balancer (ALB) in each of the two Regions Create an Amazon Elastic Container Service (Amazon ECS) cluster with the Fargate launch type Create an ECS service on the cluster Set the ECS service as the target for the ALB Process the data in Amazon ECS

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

QUESTION 754

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

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:



QUESTION 755

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

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

QUESTION 756

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.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

QUESTION 757

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

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

QUESTION 758

A company has a web application that is based on PHP and MySQL The company needs 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

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

QUESTION 759

A company stores confidential data in an Amazon Aurora PostgreSQL database in the ap-southeast-3 Region The database is encrypted with an AWS Key Management Service (AWS KMS) customer managed key The company was recently acquired and must securely share a backup of the database with the acquiring company's AWS account in ap-southeast-3. What should a solutions architect do to meet these requirements?

- A. Create a database snapshot Copy the snapshot to a new unencrypted snapshot Share the new snapshot with the acquiring company's AWS account
- B. Create a database snapshot Add the acquiring company's AWS account to the KMS key policy Share the snapshot with the acquiring company's AWS account
- C. Create a database snapshot that uses a different AWS managed KMS key Add the acquiring company's AWS account to the KMS key alias. Share the snapshot with the acquiring company's AWS account.
- D. Create a database snapshot Download the database snapshot Upload the database snapshot to an Amazon S3 bucket Update the S3 bucket policy to allow access from the acquiring company's AWS account

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

QUESTION 760

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.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

QUESTION 761

A company wants to migrate its on-premises application to AWS. The application produces output files that vary in size from tens of gigabytes to hundreds of terabytes. The application data must be stored in a standard file system structure. The company wants a solution that scales automatically, is highly available, and requires minimum operational overhead. Which solution will meet these requirements?

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

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

QUESTION 762

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

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

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

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

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

QUESTION 763

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. Run a cron job script on the EC2 instances to upload files to the S3 data lake.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

QUESTION 764

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.

Correct Answer: AC

Section: (none)

Explanation

Explanation/Reference:



QUESTION 765

A company's reporting system delivers hundreds of CSV files to an Amazon S3 bucket each day. The company must convert these files to Apache Parquet format and must store the files in a transformed data bucket.

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

- A. Create an Amazon EMR cluster with Apache Spark installed. Write a Spark application to transform the data. Use EMR File System (EMRFS) to write files to the transformed data bucket.
- B. Create an AWS Glue crawler to discover the data. Create an AWS Glue extract, transform, and load (ETL) job to transform the data. Specify the transformed data bucket in the output step.
- C. Use AWS Batch to create a job definition with Bash syntax to transform the data and output the data to the transformed data bucket. Use the job definition to submit a job. Specify an array job as the job type.
- D. Create an AWS Lambda function to transform the data and output the data to the transformed data bucket. Configure an event notification for the S3 bucket. Specify the Lambda function as the destination for the event notification.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

QUESTION 766

A company has on-premises servers that run a relational database. The database serves high-read traffic for users in different locations. The company wants to migrate the database to AWS with the least amount of effort. The database solution must support high availability and must not affect the company's current traffic flow.

Which solution meets these requirements?

- A. Use a database in Amazon RDS with Multi-AZ and at least one read replica.
- B. Use a database in Amazon RDS with Multi-AZ and at least one standby replica.
- C. Use databases that are hosted on multiple Amazon EC2 instances in different AWS Regions.

D. Use databases that are hosted on Amazon EC2 instances behind an Application Load Balancer in different Availability Zones

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

<https://aws.amazon.com/blogs/database/implementing-a-disaster-recovery-strategy-with-amazonrds/>

QUESTION 767

A company has more than 5 TB of file data on Windows file servers that run on premises. Users and applications interact with the data each day.

The company is moving its Windows workloads to AWS. As the company continues this process, the company requires access to AWS and on-premises file storage with minimum latency. The company needs a solution that minimizes operational overhead and requires no significant changes to the existing file access patterns. The company uses an AWS Site-to-Site VPN connection for connectivity to AWS.

What should a solutions architect do to meet these requirements?

- A. Deploy and configure Amazon FSx for Windows File Server on AWS. Move the on-premises file data to FSx for Windows File Server. Reconfigure the workloads to use FSx for Windows File Server on AWS.
- B. Deploy and configure an Amazon S3 File Gateway on premises. Move the on-premises file data to the S3 File Gateway. Reconfigure the on-premises workloads and the cloud workloads to use the S3 File Gateway.
- C. Deploy and configure an Amazon S3 File Gateway on premises. Move the on-premises file data to Amazon S3. Reconfigure the workloads to use either Amazon S3 directly or the S3 File Gateway, depending on each workload's location.
- D. Deploy and configure Amazon FSx for Windows File Server on AWS. Deploy and configure an Amazon FSx File Gateway on premises. Move the on-premises file data to the FSx File Gateway. Configure the cloud workloads to use FSx for Windows File Server on AWS. Configure the on-premises workloads to use the FSx File Gateway.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

QUESTION 768

A hospital recently deployed a RESTful API with Amazon API Gateway and AWS Lambda. The hospital uses API Gateway and Lambda to upload reports that are in PDF format and JPEG format. The hospital needs to modify the Lambda code to identify protected health information (PHI) in the reports. Which solution will meet these requirements with the LEAST operational overhead?

- A. Use existing Python libraries to extract the text from the reports and to identify the PHI from the extracted text.
- B. Use Amazon Textract to extract the text from the reports. Use Amazon SageMaker to identify the PHI from the extracted text.
- C. Use Amazon Textract to extract the text from the reports. Use Amazon Comprehend Medical to identify the PHI from the extracted text.
- D. Use Amazon Rekognition to extract the text from the reports. Use Amazon Comprehend Medical to identify the PHI from the extracted text.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

QUESTION 769

A company wants to create a mobile app that allows users to stream slow-motion video clips on their mobile devices. Currently, the app captures video clips and uploads the video clips in raw format into an Amazon S3 bucket. The app retrieves these video clips directly from the S3 bucket. However, the videos are large in their raw format.

Users are experiencing issues with buffering and playback on mobile devices. The company wants to implement solutions to maximize the performance and scalability of the app while minimizing operational overhead. Which combination of solutions will meet these requirements? (Select TWO.)

- A. Deploy Amazon CloudFront for content delivery and caching.
- B. Use AWS DataSync to replicate the video files across AWS Regions in other S3 buckets.
- C. Use Amazon Elastic Transcoder to convert the video files to more appropriate formats.
- D. Deploy an Auto Scaling group of Amazon EC2 instances in Local Zones for content delivery and caching.
- E. Deploy an Auto Scaling group of Amazon EC2 instances to convert the video files to more appropriate formats.

Correct Answer: CD

Section: (none)

Explanation

Explanation/Reference:

QUESTION 770

An online retail company needs to run near-real-time analytics on website traffic to analyze topselling 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-realtime 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).

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

QUESTION 771

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.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

<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." Q640

QUESTION 772

A company deploys Amazon EC2 instances that run in a VPC. The EC2 instances load source data into Amazon S3 buckets so that the data can be processed in the future. According to compliance laws, the data must not be transmitted over the public internet. Servers in the company's on-premises data center will consume the output from an application that runs on the EC2 instances. Which solution will meet these requirements?

- A. Deploy an interface VPC endpoint for Amazon EC2. Create an AWS Site-to-Site VPN connection between the company and the VPC.
- B. Deploy a gateway VPC endpoint for Amazon S3. Set up an AWS Direct Connect connection between the on-premises network and the VPC.
- C. Set up an AWS Transit Gateway connection from the VPC to the S3 buckets. Create an AWS Site-to-Site VPN connection between the company and the VPC.
- D. Set up proxy EC2 instances that have routes to NAT gateways. Configure the proxy EC2 instances to fetch S3 data and feed the application instances.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

QUESTION 773

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.

Correct Answer: CD

Section: (none)

Explanation

Explanation/Reference:

QUESTION 774

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.



Correct Answer: CD

Section: (none)

Explanation

Explanation/Reference:

QUESTION 775

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

Which solution will meet these requirements?

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

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

QUESTION 776

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

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

QUESTION 777

A company collects data from thousands of remote devices by using a RESTful web services application that runs on an Amazon EC2 instance. The EC2 instance receives the raw data, transforms the raw data, and stores all the data in an Amazon S3 bucket. The number of remote devices will increase into the millions soon. The company needs a highly scalable solution that minimizes operational overhead. Which combination of steps should a solutions architect take to meet these requirements? (Select TWO.)

- A. Use AWS Glue to process the raw data in Amazon S3.
- B. Use Amazon Route 53 to route traffic to different EC2 instances.
- C. Add more EC2 instances to accommodate the increasing amount of incoming data.
- D. Send the raw data to Amazon Simple Queue Service (Amazon SQS). Use EC2 instances to process the data.
- E. Use Amazon API Gateway to send the raw data to an Amazon Kinesis data stream. Configure Amazon Kinesis Data Firehose to use the data stream as a source to deliver the data to Amazon S3.

Correct Answer: BE

Section: (none)

Explanation

Explanation/Reference:

QUESTION 778

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

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

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

QUESTION 779

A company is deploying a web portal. The company wants to ensure that only the web portion of the application is publicly accessible. To accomplish this, the VPC was designed with two public subnets and two private subnets. The application will run on several Amazon EC2 instances in an Auto Scaling group. SSL termination must be offloaded from the EC2 instances. What should a solutions architect do to ensure these requirements are met? Configure a Network Load Balancer in the public subnets. Configure the Auto Scaling



- A. group in the private subnets and associate it with an Application Load Balancer Configure a Network Load Balancer in the public subnets. Configure the Auto Scaling
- B. group in the public subnets and associate it with an Application Load Balancer.
- C. Configure an Application Load Balancer in the public subnets. Configure the Auto Scaling group in the private subnets and associate it with the Application Load
- D. Balancer, Configure an Application Load Balancer in the private subnets. Configure the Auto Scaling group in the private subnets and associate it with the Application Load Balancer.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

QUESTION 780

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 tables. Attach the policy to the DynamoDB tables.
- D. Create a gateway VPC endpoint for DynamoDB that is associated with the Lambda VPC. Ensure that the Lambda execution role can access the gateway VPC endpoint.
- E. Create an interface VPC endpoint for DynamoDB that is associated with the Lambda VPC. Ensure that the Lambda execution role can access the interface VPC endpoint.

Correct Answer: AE

Section: (none)

Explanation

Explanation/Reference:



QUESTION 781

An online retail company has more than 50 million active customers and receives more than 25,000 orders each day. The company collects purchase data for customers and stores this data in Amazon S3. Additional customer data is stored in Amazon RDS. The company wants to make all the data available to various teams so that the teams can perform analytics. The solution must provide the ability to manage fine-grained permissions for the data and must minimize operational overhead. Which solution will meet these requirements?

- A. Migrate the purchase data to write directly to Amazon RDS. Use RDS access controls to limit access.
- B. Schedule an AWS Lambda function to periodically copy data from Amazon RDS to Amazon S3. Create an AWS Glue crawler. Use Amazon Athena to query the data. Use S3 policies to limit access.
- C. Create a data lake by using AWS Lake Formation. Create an AWS Glue JOBC connection to Amazon RDS. Register the S3 bucket in Lake Formation. Use Lake
- D. Formation access controls to limit access. Create an Amazon Redshift cluster Schedule an AWS Lambda function to periodically copy data from Amazon S3 and Amazon RDS to Amazon Redshift. Use Amazon Redshift access controls to limit access.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

QUESTION 782

A company runs a web application on Amazon EC2 instances in multiple Availability Zones. The EC2 instances are in private subnets. A solutions architect implements an internet-facing Application Load Balancer (ALB) and specifies the EC2 instances as the target group. However, the internet traffic is not reaching the EC2 instances. How should the solutions architect reconfigure the architecture to resolve this issue?

- A. Replace the ALB with a Network Load Balancer. Configure a NAT gateway in a public subnet to allow internet traffic.
- B. Move the EC2 instances to public subnets. Add a rule to the EC2 instances' security groups to allow outbound traffic to 0.0.0.0/0.
- C. Update the route tables for the EC2 instances' subnets to send 0.0.0.0/0 traffic through the internet gateway route. Add a rule to the EC2 instances' security groups to allow outbound traffic to 0.0.0.0/0.
- D. Create public subnets in each Availability Zone. Associate the public subnets with the ALB. Update the route tables for the public subnets with a route to the private subnets.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:

QUESTION 783

A company has an AWS Glue extract, transform, and load (ETL) job that runs every day at the same time. The job processes XML data that is in an Amazon S3 bucket. New data is added to the S3 bucket every day. A solutions architect notices that AWS Glue is processing all the data during each run. What should the solutions architect do to prevent AWS Glue from reprocessing old data?

- A. Edit the job to use job bookmarks.
- B. Edit the job to delete data after the data is processed
- C. Edit the job by setting the NumberOfWorkers field to 1.
- D. Use a FindMatches machine learning (ML) transform.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:

QUESTION 784

A company has an ordering application that stores customer information in Amazon RDS for MySQL. During regular business hours, employees run one-time queries for reporting purposes. Timeouts are occurring during order processing because the reporting queries are taking a long time to run. The company needs to eliminate the timeouts without preventing employees from performing queries. What should a solutions architect do to meet those requirements?

- A. Create a read replica Move reporting queries to the read replica.
- B. Create a read replica. Distribute the ordering application to the primary DB instance and the read replica.
- C. Migrate the ordering application to Amazon DynamoDB with on-demand capacity.
- D. Schedule the reporting queries for non-peak hours.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:

QUESTION 785

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

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

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:

QUESTION 786

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

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

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

QUESTION 787

A company's web application consists of an Amazon API Gateway API in front of an AWS Lambda function and an Amazon DynamoDB database. The Lambda function handles the business logic, and the DynamoDB table hosts the data.

- A. The application uses Amazon Cognito user pools to identify the individual users of the application. A solutions architect needs to update the application so that only users who have a subscription can access premium content.
- B. Enable API caching and throttling on the API Gateway API.
- C. Set up AWS WAF on the API Gateway API. Create a rule to filter users who have a subscription.
- D. Apply fine-grained IAM permissions to the premium content in the DynamoDB table.
- E. Implement API usage plans and API keys to limit the access of users who do not have a subscription.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:



QUESTION 788

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

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

QUESTION 789

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. Configure the ETL job to process the .csv files and store the processed data in Amazon Redshift.
- B. 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.
- C. Create an AWS Lambda function and an Amazon DynamoDB table. Use an S3 event to invoke the Lambda function. 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.

D. Use Amazon EventBridge (Amazon CloudWatch Events) to launch an Amazon EMR cluster on a weekly schedule. 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.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

QUESTION 790

A company produces batch data that comes from different databases. The company also produces live stream data from network sensors and application APIs. The company needs to consolidate all the data into one place for business analytics. The company needs to process the incoming data and then stage the data in different Amazon S3 buckets. Teams will later run onetime queries and import the data into a business intelligence tool to show key performance indicators (KPIs).

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

A. Use Amazon Athena for one-time queries. Use Amazon QuickSight to create dashboards for KPIs.

B. Use Amazon Kinesis Data Analytics for one-time queries. Use Amazon QuickSight to create dashboards for KPIs.

C. Create custom AWS Lambda functions to move the individual records from the databases to an Amazon Redshift cluster.

D. Use an AWS Glue extract, transform, and load (ETL) job to convert the data into JSON format. Load the data into multiple Amazon OpenSearch Service (Amazon Elasticsearch Service) clusters.

E. Use blueprints in AWS Lake Formation to identify the data that can be ingested into a data lake. Use AWS Glue to crawl the source, extract the data, and load the data into Amazon S3 in Apache Parquet format.

Correct Answer: CD

Section: (none)

Explanation

Explanation/Reference:

QUESTION 791

A company is running a popular social media website. The website gives users the ability to upload images to share with other users. The company wants to make sure that the images do not contain inappropriate content. The company needs a solution that minimizes development effort.

What should a solutions architect do to meet these requirements?

A. Use Amazon Comprehend to detect inappropriate content. Use human review for low-confidence predictions.

B. Use Amazon Rekognition to detect inappropriate content. Use human review for low-confidence predictions.

C. Use Amazon SageMaker to detect inappropriate content. Use ground truth to label low-confidence predictions.

D. Use AWS Fargate to deploy a custom machine learning model to detect inappropriate content. Use ground truth to label low-confidence predictions.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

QUESTION 792

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

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

A. Use AWS Batch to run the tasks as jobs. Schedule the jobs by using Amazon EventBridge (Amazon CloudWatch Events).

B. Convert the EC2 instance to a container. Use AWS App Runner to create the container on demand to run the tasks as jobs.

C. Copy the tasks into AWS Lambda functions. Schedule the Lambda functions by using Amazon EventBridge (Amazon CloudWatch Events).

D. Create an Amazon Machine Image (AMI) of the EC2 instance that runs the tasks. Create an Auto Scaling group with the AMI to run multiple copies of the instance.

Correct Answer: C

Section: (none)



Explanation

Explanation/Reference:

QUESTION 793

A media company collects and analyzes user activity data on premises. The company wants to migrate this capability to AWS. The user activity data store will continue to grow and will be petabytes in size. The company needs to build a highly available data ingestion solution that facilitates on-demand analytics of existing data and new data with SQL.

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

- A. Send activity data to an Amazon Kinesis data stream. Configure the stream to deliver the data to an Amazon S3 bucket.
- B. Send activity data to an Amazon Kinesis Data Firehose delivery stream. Configure the stream to deliver the data to an Amazon Redshift cluster
- C. Place activity data in an Amazon S3 bucket. Configure Amazon S3 to run an AWS Lambda function on the data as the data arrives in the S3 bucket.
- D. Create an ingestion service on Amazon EC2 instances that are spread across multiple Availability Zones. Configure the service to forward data to an Amazon RDS Multi-AZ database.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

QUESTION 794

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 WellArchitected 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 1AM role to each existing instance and new instance. Use AWS Systems Manager Session Manager to establish a remote SSH session.
- C. Create an administrative SSH key pair. Load the public key into each EC2 instance. Deploy a bastion host in a public subnet to provide a tunnel for administration of each instance.
- D. Establish an AWS Site-to-Site VPN connection. Instruct administrators to use their local onpremises machines to connect directly to the instances by using SSH keys across the VPN tunnel.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

QUESTION 795

A company hosts a multiplayer gaming application on AWS. The company wants the application to read data with sub-millisecond latency and run one-time queries on historical data.

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

- A. Use Amazon RDS for data that is frequently accessed. Run a periodic custom script to export the data to an Amazon S3 bucket.
- B. Store the data directly in an Amazon S3 bucket. Implement an S3 Lifecycle policy to move older data to S3 Glacier Deep Archive for long-term storage. Run one-time queries on the data in Amazon S3 by using Amazon Athena
- C. Use Amazon DynamoDB with DynamoDB Accelerator (DAX) for data that is frequently accessed. Export the data to an Amazon S3 bucket by using DynamoDB table export. Run one-time queries on the data in Amazon S3 by using Amazon Athena.
- D. Use Amazon DynamoDB for data that is frequently accessed Turn on streaming to Amazon Kinesis Data Streams. Use Amazon Kinesis Data Firehose to read the data from Kinesis Data Streams. Store the records in an Amazon S3 bucket.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

QUESTION 796

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.

Correct Answer: AD

Section: (none)

Explanation

Explanation/Reference:

QUESTION 797

A company wants to create an audio version of its product manual. The product manual contains custom product names and abbreviations. The product manual is divided into sections . Which solution will meet these requirements with the LEAST operational overhead?

- A. Use Amazon Polly Build custom lexicons for the product names and abbreviations Use the StartSpeechSynthesJsTaskAPI operation for each section of the product manual
- B. Use Amazon Polly Build custom Speech Synthesis Markup Language (SSML) for the product names and abbreviations Use the StartDocumentTextDetection API operation for each section of the product manual
- C. Use Amazon Textract Build custom Speech Synthesis Markup Language (SSML) for the product names and abbreviations Use the StartDocumentTextDetection API operation for each section of the product manual
- D. Use Amazon Textract Build custom lexicons for the product names and abbreviations Use the StartTranscriptionJob API operation for each section of the product manual

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:



QUESTION 798

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

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

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

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

QUESTION 799

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

- A. Store the credentials as secrets in AWS Secrets Manager Use multi-Region secret replication for the required Regions Configure Secrets Manager to rotate the secrets on a schedule
- B. Store the credentials as secrets in AWS Systems Manager by creating a secure string parameter Use multi-Region secret replication for the required Regions Configure Systems
- C. Manager to rotate the secrets on a schedule
- D. Store the credentials in an Amazon S3 bucket that has server-side encryption (SSE) enabled Use Amazon EventBridge (Amazon CloudWatch Events) to invoke an AWS Lambda function to rotate the credentials

E. Encrypt the credentials as secrets by using AWS Key Management Service (AWS KMS) multi-Region customer managed keys Store the secrets in an Amazon DynamoDB global table Use an AWS Lambda function to retrieve the secrets from DynamoDB Use the RDS API to rotate the secrets.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

QUESTION 800

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

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

QUESTION 801

A company wants to analyze and troubleshoot Access Denied errors and unauthorized errors that are related to IAM permissions. The company has AWS CloudTrail turned on. Which solution will meet these requirements with the LEAST effort?

- A. Use AWS Glue and write custom scripts to query CloudTrail logs for the errors.
- B. Use AWS Batch and write custom scripts to query CloudTrail logs for the errors.
- C. Search CloudTrail logs with Amazon Athena queries to identify the errors
- D. Search CloudTrail logs with Amazon QuickSight Create a dashboard to identify the errors

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

QUESTION 802

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 of 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. Create a transit gateway, and attach each VPC to the transit gateway. Establish connectivity between the Direct Connect connection and the transit gateway.

Correct Answer: C

Section: (none)

Explanation



Explanation/Reference:

QUESTION 803

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

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

Correct Answer: CE

Section: (none)

Explanation

Explanation/Reference:

QUESTION 804

A company wants to establish connectivity between its on-premises data center and AWS (or an existing workload). The workload runs on Amazon EC2 instances in two VPCs in different AWS Regions. The VPCs need to communicate with each other. The company needs to provide connectivity from its data center to both VPCs. The solution must support a bandwidth of 600 Mbps to the data center. Which solution will meet these requirements?

- A. Set up an AWS Site-to-Site VPN connection between the data center and one VPC. Create a VPC peering connection between the VPCs.
- B. Set up an AWS Site-to-Site VPN connection between the data center and each VPC. Create a VPC peering connection between the VPCs.
- C. Set up an AWS Direct Connect connection between the data center and one VPC. Create a VPC peering connection between the VPCs.
- D. Create a transit gateway. Attach both VPCs to the transit gateway. Create an AWS Site-to-Site VPN tunnel to the transit gateway.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

QUESTION 805

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. Use the AWS CLI to copy all files locally to the S3 bucket.
- B. Create an AWS Snowball Edge job. Receive a Snowball Edge device on premises. Use the Snowball Edge client to transfer data to the device. Return the device so that AWS can import the data into Amazon S3.
- C. Deploy an S3 File Gateway on premises. 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. Transfer the data from the existing NFS file share to the S3 File Gateway.
- D. Set up an AWS Direct Connect connection between the on-premises network and AWS. Deploy an S3 File Gateway on premises. Create a public virtual interface (VIF) 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. Transfer the data from the existing NFS file share to the S3 File Gateway.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

QUESTION 806

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

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

QUESTION 807

A company is building a new dynamic ordering website. The company wants to minimize server maintenance and patching. The website must be highly available and must scale read and write capacity as quickly as possible to meet changes in user demand.

- A. Host static content in Amazon S3. Host dynamic content by using Amazon API Gateway and AWS Lambda. Use Amazon DynamoDB with on-demand capacity for the database Configure Amazon CloudFront to deliver the website content.
- B. Host static content in Amazon S3. Host dynamic content by using Amazon API Gateway and AWS Lambda. Use Amazon Aurora with Aurora Auto Scaling for the database Configure Amazon CloudFront to deliver the website content.
- C. Host all the website content on Amazon EC2 instances. Create an Auto Scaling group to scale the EC2 instances. Use an Application Load Balancer to distribute traffic Use Amazon DynamoDB with provisioned write capacity for the database.
- D. Host all the website content on Amazon EC2 instances. Create an Auto Scaling group to scale the EC2 instances Use an Application Load Balancer to distribute traffic Use Amazon Aurora with Aurora Auto Scaling for the database.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

QUESTION 808

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

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

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

QUESTION 809

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 1AM policies of users who need to delete the objects

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

QUESTION 810

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

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

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

QUESTION 811

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

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

QUESTION 812

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

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

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

QUESTION 813

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

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

QUESTION 814

A company is building a solution that will report Amazon EC2 Auto Scaling events across all the applications in an AWS account. The company needs to use a serverless solution to store the EC2 Auto Scaling status data in Amazon S3. The company then will use the data in Amazon S3 to provide near-real time updates in a dashboard. The solution must not affect the speed of EC2 instance launches. How should the company move the data to Amazon S3 to meet these requirements?

- A. Use an Amazon CloudWatch metric stream to send the EC2 Auto Scaling status data to Amazon Kinesis Data Firehose. Store the data in Amazon S3.
- B. Launch an Amazon EMR cluster to collect the EC2 Auto Scaling status data and send the data to Amazon Kinesis Data Firehose. Store the data in Amazon S3.
- C. Create an Amazon EventBridge (Amazon CloudWatch Events) rule to invoke an AWS Lambda function on a schedule. Configure the Lambda function to send the EC2 Auto Scaling status data directly to Amazon S3.
- D. Use a bootstrap script during the launch of an EC2 instance to install Amazon Kinesis Agent. Configure Kinesis Agent to collect the EC2 Auto Scaling status data and send the data to Amazon Kinesis Data Firehose. Store the data in Amazon S3.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

QUESTION 815

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

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

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

QUESTION 816

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. Use an Amazon Elastic Container Service (Amazon ECS) cluster with the AWS Fargate launch type to run the containers. Use target tracking to scale automatically based on demand.
- B. Store container images in an Amazon Elastic Container Registry (Amazon ECR) repository. Use an Amazon Elastic Container Service (Amazon ECS) cluster with the Amazon EC2 launch type to run the containers. Use target tracking to scale automatically based on demand.
- C. Store container images in a repository that runs on an Amazon EC2 instance. Run the containers on EC2 instances that are spread across multiple Availability Zones. Monitor the average CPU utilization in Amazon CloudWatch. Launch new EC2 instances as needed.
- D. 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. Use an Amazon CloudWatch alarm to scale out EC2 instances when the average CPU utilization threshold is breached.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

QUESTION 817

A company has an event-driven application that invokes AWS Lambda functions up to 800 times each minute with varying runtimes. The Lambda functions access data that is stored in an Amazon Aurora MySQL DB cluster. The company is noticing connection timeouts as user activity increases. The database shows no signs of being overloaded. CPU, memory, and disk access metrics are all low. Which solution will resolve this issue with the LEAST operational overhead?

- A. Adjust the size of the Aurora MySQL nodes to handle more connections. Configure retry logic in the Lambda functions for attempts to connect to the database.
- B. Set up Amazon ElastiCache for Redis to cache commonly read items from the database. Configure the Lambda functions to connect to ElastiCache for reads.
- C. Add an Aurora Replica as a reader node. Configure the Lambda functions to connect to the reader endpoint of the DB cluster rather than to the writer endpoint.
- D. Use Amazon RDS Proxy to create a proxy. Set the DB cluster as the target database. Configure the Lambda functions to connect to the proxy rather than to the DB cluster.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:



QUESTION 818

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 highperforming storage. Which solution will meet these requirements?

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

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference: