

SAA-C02.premium.exam.182q - DEMO

Number: SAA-C02
Passing Score: 800
Time Limit: 120 min
File Version: 3.0



SAA-C02

AWS Certified Solutions Architect – Associate



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 (nonexistent 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 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 onpremises 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 process 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: C

Section: (none)

Explanation

Explanation/Reference:

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-realtime 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:

Reference: <https://d0.awsstatic.com/whitepapers/whitepaper-streaming-data-solutions-on-aws-with-amazon-kinesis.pdf> (9)

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 product team is creating a new application that will store a large amount of data. The data will be analyzed hourly and modified by multiple Amazon EC2 Linux instances. 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 EBS volume. Mount the EBS volume on the application instances.
- B. Store the data in an Amazon EFS file system. Mount the file system on the application instances.
- C. Store the data in Amazon S3 Glacier. Update the vault policy to allow access to the application instances.
- D. Store the data in Amazon S3 Standard-Infrequent Access (S3 Standard-IA). Update the bucket policy to allow access to 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 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:

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. Amazon EBS
- B. Amazon EC2 instance store
- C. Amazon EFS
- D. Amazon S3

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

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

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 database is currently less than 1 GB with unpredictable future growth. Data can be queried using simple key-value 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:

Reference: <https://aws.amazon.com/ec2/faqs/> (unpredictable) <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: B

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 customization. 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 customization. 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 customization. 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 customization. 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-a-match-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: A

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.

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.

Cross-origin resource sharing (CORS)

Cross-origin resource sharing (CORS) defines a way for client web applications that are loaded in one domain to interact with resources in a different domain. With CORS support, you can build rich client-side web applications with Amazon S3 and selectively allow cross-origin access to your Amazon S3 resources.

Reference: <https://medium.com/@KerrySheldon/s3-exercise-2-4-adding-objects-to-an-s3-bucket-with-cross-region-replication-a78b332b7697>
<https://docs.aws.amazon.com/AmazonS3/latest/dev/ObjectVersioning.html> <https://docs.aws.amazon.com/AmazonS3/latest/dev/cors.html>

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: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

VPC endpoints

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.

An interface endpoint is an elastic network interface with a private IP address from the IP address range of your subnet that serves as an entry point for traffic destined to a supported service. Interface endpoints are powered by AWS PrivateLink, a technology that enables you to privately access services by using private IP addresses. AWS PrivateLink restricts all network traffic between your VPC and services to the Amazon network. You do not need an internet gateway, a NAT device, or a virtual private gateway.

Reference: <https://aws.amazon.com/blogs/aws/new-vpc-endpoint-for-amazon-s3/> <https://docs.aws.amazon.com/vpc/latest/userguide/vpc-endpoints.html>

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 on-premises to the AWS Cloud. The company uses tiered storage on premises with hot high-performance 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 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:

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 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 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 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: A

Section: (none)

Explanation

Explanation/Reference:

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:

Reference: <https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/routing-to-cloudfront-distribution.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.

Which 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> (geolocation routing)

QUESTION 33 A solution 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: BD

Section: (none)

Explanation

Explanation/Reference:**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 IntelligentTiering 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 99.999999999% 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:

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

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/FIFO-queues.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 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:

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/asg-purchase-options.html>

VCEUp

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: C

Section: (none)

Explanation

Explanation/Reference:

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 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 install 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 recover if versioning enable, also it provide a 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=vv45DQAAQBAJ&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=ACfU3U3opL9Jha6jM2EI8x7EcjK4rigQHQ&hl=en&sa=X&ved=2ahUKewiS9e3yy7vpAhVx73MBHZN0DnQQ6AEwAHoECBQQAQ#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:

Reference: <https://aws.amazon.com/blogs/database/best-storage-practices-for-running-production-workloads-on-hosted-databases-with-amazon-rds-or-amazon-ec2/>

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 meet 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: A

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 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 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 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 permission 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-permission-guardrails-across-accounts-in-your-aws-organization/#:~: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: D

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 EBS volume. For better scalability and availability, the company duplicated the architecture and created a second EC2 instance and EBS volume in another Availability Zone, placing both behind an Application Load Balancer. After completing this change, users reported that each time they refreshed the website, they could see one subset of their documents or the other, but never all of the documents at the same time.

What should a solutions architect propose to ensure users see all of their documents at once?

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

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:

QUESTION 55 A solutions architect is tasked with transferring 750 TB of data from a network-attached file system located at a branch office 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 VPC endpoint.
- B. Order 10 AWS Snowball appliances and select an S3 Glacier vault as the destination. Create a bucket policy to enforce VPC endpoint.
- C. Mount the network-attached file system to Amazon S3 and copy the files directly. Create a lifecycle policy to 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:

Regional Limitations for AWS Snowball

The AWS Snowball service has two device types, the standard Snowball and the Snowball Edge. The following table highlights which of these devices are available in which regions.

Region	Snowball Availability	Snowball Edge Availability
US East (Ohio)	50 TB and 80 TB	100 TB
US East (N. Virginia)	50 TB and 80 TB	100 TB
US West (N. California)	50 TB and 80 TB	100 TB
US West (Oregon)	50 TB and 80 TB	100 TB
Canada (Central)	80 TB only	100 TB
Asia Pacific (Mumbai)	80 TB only	100 TB
Asia Pacific (Singapore)	80 TB only	100 TB
Asia Pacific (Sydney)	80 TB only	100 TB
Asia Pacific (Tokyo)	80 TB only	100 TB
EU (Frankfurt)	80 TB only	100 TB
EU (Ireland)	80 TB only	100 TB
EU (London)	80 TB only	100 TB
South America (São Paulo)	80 TB only	100 TB

Limitations on Jobs in AWS Snowball

The following limitations exist for creating jobs in AWS Snowball:

For security purposes, data transfers must be completed within 90 days of the Snowball being prepared.

Currently, AWS Snowball Edge device doesn't support server-side encryption with customer-provided keys (SSE-C). AWS Snowball Edge device does support server-side encryption with Amazon S3–managed encryption keys (SSE-S3) and server-side encryption with AWS Key Management Service – managed keys (SSE-KMS). For more information, see Protecting Data Using Server-Side Encryption in the Amazon Simple Storage Service Developer Guide.

In the US regions, Snowballs come in two sizes: 50 TB and 80 TB. All other regions have the 80 TB Snowballs only. If you're using Snowball to import data, and you need to transfer more data than will fit on a single Snowball, create additional jobs. Each export job can use multiple Snowballs.

The default service limit for the number of Snowballs you can have at one time is 1. If you want to increase your service limit, contact AWS Support.

All objects transferred to the Snowball have their metadata changed. The only metadata that remains the same is filename and filesize. All other metadata is set as in the following example: -rw-rw-r-- 1 root root [filesize] Dec 31 1969 [path/filename].

Object lifecycle management

To manage your objects so that they are stored cost effectively throughout their lifecycle, configure their Amazon S3 Lifecycle. An S3 Lifecycle configuration is a set of rules that define actions that Amazon S3 applies to a group of objects. There are two types of actions:

Transition actions – Define when objects transition to another storage class. For example, you might choose to transition objects to the S3 Standard-IA storage class 30 days after you created them, or archive objects to the S3 Glacier storage class one year after creating them.

Expiration actions – Define when objects expire. Amazon S3 deletes expired objects on your behalf.

The lifecycle expiration costs depend on when you choose to expire objects.

Reference: <https://docs.aws.amazon.com/snowball/latest/ug/limits.html> <https://docs.aws.amazon.com/AmazonS3/latest/dev/object-lifecycle-mgmt.html>

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:

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

How should security groups be configured in this situation? (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.

Correct Answer: AB

Section: (none)

Explanation

Explanation/Reference:

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. The application writes log files using standard API calls. For compliance reasons, all log files must be retained indefinitely and will be analyzed by a reporting tool that must access all files concurrently.

Which storage service should a solutions architect use to provide the MOST cost-effective solution?

- A. Amazon EBS
- B. 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 99.999999999% (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/fags/> <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:

Reference: https://docs.aws.amazon.com/dms/latest/userguide/CHAP_Target.S3.html

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-amazon-rds/>

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:

Reference: <https://docs.aws.amazon.com/autoscaling/ec2/userguide/as-dg.pdf>

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:

Reference: <https://docs.aws.amazon.com/AWSSimpleQueueService/latest/SQSDeveloperGuide/FIFO-queues.html#FIFO-queues-exactly-once-processing>

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",
      "Resources": "*",
      "Condition": {
        "IpAddress": {
          "aws:SourceIp": "10.100.100.0/24"
        }
      }
    },
    {
      "Effect": "Deny",
      "Action": "ec2:*",
      "Resources": "*",
      "Condition": {
        "StringNotEquals": {
          "ec2:Region": "us-east-1"
        }
      }
    }
  ]
}
```

VCEUp

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:

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:

Reference: <https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/on-demand-streaming-video.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:

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:

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:

QUESTION 75

A solutions architect needs to design a managed storage solution for a company's application that includes high-performance machine learning. 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 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 EBS.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

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

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: C

Section: (none)

Explanation

Explanation/Reference:

Reference: https://d1.awsstatic.com/whitepapers/Security/DDoS_White_Paper.pdf

QUESTION 78

A company has an application that calls AWS Lambda functions. A recent code review found database credentials stored in the source code. The database credentials need to be removed from the Lambda source code. The credentials must then be securely stored and rotated on an ongoing basis to meet security policy requirements.

What should a solutions architect recommend to meet these requirements?

- A. Store the password in AWS CloudHSM. Associate the Lambda function with a role that can retrieve the password from CloudHSM given its key ID.
- B. Store the password in AWS Secrets Manager. Associate the Lambda function with a role that can retrieve the password from Secrets Manager given its secret ID.
- C. Move the database password to an environment variable associated with the Lambda function. Retrieve the password from the environment variable upon execution.
- D. Store the password in AWS Key Management Service (AWS KMS). Associate the Lambda function with a role that can retrieve the password from AWS KMS given its key ID.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Reference: <https://aws.amazon.com/blogs/security/how-to-use-aws-secrets-manager-rotate-credentials-amazon-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: D

Section: (none)

Explanation

Explanation/Reference:**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:**QUESTION 81**

A company's operations teams have an existing Amazon S3 bucket configured to notify an Amazon SQS queue when new object is 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 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 update this queue when a new object is created.
- C. Create an Amazon SNS topic and SQS queue for the Update. 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 subscription for both queue in the topic.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

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: C

Section: (none)

Explanation

Explanation/Reference:

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:

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:

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 is instance in the private subnet that use a NAT gateway to connect to the internet. In case of an AZ failure, the company wants to ensure that the instance is 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 now 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:

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:

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 execute a script remotely on all EC2 instances to send data to the audit system.
- B. Use EC2 Auto Scaling lifecycle hooks to execute 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 execute a custom script through user data to send data to the audit system when instances are launched and terminated.
- D. Execute 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:

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: C

Section: (none)

Explanation

Explanation/Reference:

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:

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:**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:

QUESTION 92

A company recently expanded globally and wants to make its application accessible to users in those geographic locations. The application is deploying on Amazon EC2 instances behind an Application Load balancer in an Auto Scaling group. The company needs the ability 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 fouling policy.
- D. Configure an Amazon Route 53 multivalue answer routing policy.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

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:

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: D

Section: (none)

Explanation

Explanation/Reference:

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 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 EBS Provisioned IOPS SSD (101) volume. Attach the volume to all the EC2 instances.
- D. Create file systems on Amazon EBS volumes attached to each EC2 instance. Synchronize the Amazon EBS volumes across the different EC2 instances.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

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:

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 URL's. 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:

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:

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:

QUESTION 100

A solution 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": "*"
    }
  ]
}
```

VCEUp

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:

QUESTION 101

A company has an Amazon EC2 instance running on a private subnet that needs to access a public websites to download patches and updates. The company does not want external websites to see the EC2 instance IP address or initiate connection to it.

How can a solution 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:

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:

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:

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:**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:**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: A

Section: (none)

Explanation

Explanation/Reference:

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:**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: C

Section: (none)

Explanation

Explanation/Reference:

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 solution 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: A

Section: (none)

Explanation

Explanation/Reference:

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:

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 Amazon DynamoDB transactions
- B. Create 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 EBS Throughput Optimized HDD storage.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

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:

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:

QUESTION 114

A company currently stores symmetric encryption keys in a hardware security module (HSM). A solution 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: B

Section: (none)

Explanation

Explanation/Reference:

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 EFS file system that connects with the backup applications using the NFS interface.
- C. Set up an 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: A

Section: (none)

Explanation

Explanation/Reference:

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

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

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:

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:

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 solution architect needs to solve the problem with minimal changes to the existing web application.

What should the solution 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:

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:

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 solution 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:**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: A

Section: (none)

Explanation

Explanation/Reference:**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 environment and with AWS.

Which services mod the business requirements? (Choose two.)

- A. Amazon EBS
- B. Amazon EFS
- C. Amazon FSx for Windows
- D. Amazon S3
- E. AWS Storage Gateway file gateway

Correct Answer: CE

Section: (none)

Explanation

Explanation/Reference:**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: C

Section: (none)

Explanation

Explanation/Reference:

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 solution architect do to accomplish this?

- A. Configure a volume using Amazon EFS. Mount the EPS 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 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:

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 solution 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:

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: D

Section: (none)

Explanation

Explanation/Reference:

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:

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"
    }
  ]
}
```

```
{
  "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/*"
    }
  ]
}
```

B.

```

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

{
  "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"
    }
  ]
}

```

C.

VCEUp

D.

Correct Answer: A
 Section: (none)
 Explanation

Explanation/Reference:**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 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:**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.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:**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 EBS for maximum performance, Amazon S3 for durable data storage, and Amazon S3 Glacier for archival storage.
- B. Amazon EBS for maximum performance. Amazon EFS for durable data storage, and Amazon S3 Glacier for archival storage
- C. Amazon EC2 instance store for maximum performance, 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:**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 bucket 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 (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:

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 archived 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: DE

Section: (none)

Explanation

Explanation/Reference:

VCEUp

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:

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 group with instances in a partition placement group.
- C. Create a Network Load Balancer backed by the existing serves in different Availability Zones as the target.
- D. Create a Network Load Balancer backed by an Auto Scaling with instances in multiple Availability zones as the target

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

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:

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 M, 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: B

Section: (none)

Explanation

Explanation/Reference:

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 an Nitro-based Amazon EC2 instance with an 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:

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: A

Section: (none)

Explanation

Explanation/Reference:

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:

VCEUp

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:

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: B

Section: (none)

Explanation

Explanation/Reference:

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:

VCEUp

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:

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 EBS
- B. Amazon EC2
- C. Amazon FSx
- D. Amazon S3

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

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:

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: A

Section: (none)

Explanation

Explanation/Reference:

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, triggered 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: C

Section: (none)

Explanation

Explanation/Reference:

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:

QUESTION 151

A solution 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:

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:

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 applications 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 replication and connect the business reports to it.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

QUESTION 154

A company is running a two-tier ecommerce website using 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 solution 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:

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 Balances

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

QUESTION 156

A company wants to deploy a shared file system for its .NET application servers and Microsoft SQL Server database running on Amazon EC2 instance with Windows Server 2016. The solution must be able to be integrated in to the corporate Active Directory domain, be highly durable, be managed by AWS, and provided 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: D

Section: (none)

Explanation

Explanation/Reference:

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 (or 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:

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:

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

Correct Answer: AB

Section: (none)

Explanation

Explanation/Reference:

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 solution architect. Have the solution architect use the root user for daily administration tasks.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

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: B

Section: (none)

Explanation

Explanation/Reference:

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:

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:

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:

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 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 and route all traffic through VPC peering.
- C. Configure an AWS Direct Connect between all VPCs and VPNs.
- D. Configure a transit gateway with AWS Transit Gateway and connected all VPCs and VPNs.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

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: A

Section: (none)

Explanation

Explanation/Reference:

QUESTION 167

A solution 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 network-attached storage (NAS). The solution architected has proposed migrating the IIS web servers

Which replacement to the on-premises file share is MOST resilient and durable?

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

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

QUESTION 168 A company needs to implement a relational database with a multi-Region disaster recovery Recovery Point Objective (RPO) of 1 second and an 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: C

Section: (none)

Explanation

Explanation/Reference:

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:

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:

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: C

Section: (none)

Explanation

Explanation/Reference:

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:

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"
        }
      }
    }
  ]
}
```

VCEUp

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:

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:

QUESTION 175

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 access the queue. Update Amazon S3 to update the 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 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 subscription for both queues in the topic.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

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:

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: BC

Section: (none)

Explanation

Explanation/Reference:

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:

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:

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:

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:

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: