



# Amazon-Web-Services

## Exam Questions CLF-C02

AWS Certified Cloud Practitioner

### NEW QUESTION 1

- (Topic 3)

A company wants to monitor for misconfigured security groups that are allowing unrestricted access to specific ports. Which AWS service will meet this requirement?

- A. AWS Trusted Advisor
- B. Amazon CloudWatch
- C. Amazon GuardDuty
- D. AWS Health Dashboard

**Answer:** A

#### Explanation:

AWS Trusted Advisor is an online tool that provides you real time guidance to help you provision your resources following AWS best practices, including security and performance. It can help you monitor for misconfigured security groups that are allowing unrestricted access to specific ports. Amazon CloudWatch is a service that monitors your AWS resources and the applications you run on AWS. Amazon GuardDuty is a threat detection service that continuously monitors for malicious activity and unauthorized behavior. AWS Health Dashboard provides relevant and timely information to help you manage events in progress, and provides proactive notification to help you plan for scheduled activities.

### NEW QUESTION 2

- (Topic 3)

Which abilities are benefits of the AWS Cloud? (Select TWO.)

- A. Trade variable expenses for capital expenses.
- B. Deploy globally in minutes.
- C. Plan capacity in advance of deployments.
- D. Take advantage of economies of scale.
- E. Reduce dependencies on network connectivity.

**Answer:** AB

#### Explanation:

The AWS Cloud offers many benefits, such as:

? Trade variable expenses for capital expenses: You can pay only for the resources you use, instead of investing in fixed costs upfront. This reduces the risk and complexity of planning and managing your IT infrastructure<sup>4</sup>

? Deploy globally in minutes: You can leverage the global infrastructure of AWS to deploy your applications and data in multiple regions and availability zones. This enables you to reach your customers faster, improve performance, and increase reliability<sup>5</sup>

### NEW QUESTION 3

- (Topic 3)

Which of the following services can be used to block network traffic to an instance? (Select TWO.)

- A. Security groups
- B. Amazon Virtual Private Cloud (Amazon VPC) flow logs
- C. Network ACLs
- D. Amazon CloudWatch
- E. AWS CloudTrail

**Answer:** AC

#### Explanation:

Security groups and network ACLs are two AWS services that can be used to block network traffic to an instance. Security groups are virtual firewalls that control the inbound and outbound traffic for your instances at the instance level. You can specify which protocols, ports, and source or destination IP addresses are allowed or denied for each instance. Security groups are stateful, which means that they automatically allow return traffic for any allowed inbound or outbound traffic<sup>123</sup>. Network ACLs are virtual firewalls that control the inbound and outbound traffic for your subnets at the subnet level. You can create rules to allow or deny traffic based on protocols, ports, and source or destination IP addresses. Network ACLs are stateless, which means that you have to explicitly allow return traffic for any allowed inbound or outbound traffic<sup>456</sup>. References: 1: Security groups for your VPC - Amazon Virtual Private Cloud, 2: Security Groups for Your VPC - Amazon Elastic Compute Cloud, 3: AWS Security Groups: Everything You Need to Know, 4: Network ACLs - Amazon Virtual Private Cloud, 5: Control traffic to subnets using network ACLs - Amazon Virtual Private Cloud, 6: AWS Network ACLs: Everything You Need to Know

### NEW QUESTION 4

- (Topic 3)

A company wants high levels of detection and near-real-time (NRT) mitigation against large and sophisticated distributed denial of service (DDoS) attacks on applications running on AWS.

Which AWS service should the company use?

- A. Amazon GuardDuty
- B. Amazon Inspector
- C. AWS Shield Advanced
- D. Amazon Macie

**Answer:** C

#### Explanation:

AWS Shield Advanced is a service that provides high levels of detection and near-real-time (NRT) mitigation against large and sophisticated distributed denial of service (DDoS) attacks on applications running on AWS. AWS Shield Advanced also provides you with 24x7 access to the AWS DDoS Response Team (DRT)

and protection against DDoS attacks of any size or duration<sup>1</sup>. Amazon GuardDuty is a service that provides threat detection for your AWS accounts and workloads, but it does not offer DDoS protection<sup>3</sup>. Amazon Inspector is a service that helps you improve the security and compliance of your applications deployed on AWS by automatically assessing them for vulnerabilities and deviations from best practices. Amazon Macie is a service that uses machine learning and pattern matching to discover and protect your sensitive data in AWS.

#### NEW QUESTION 5

- (Topic 3)

A company wants to migrate its on-premises workloads to the AWS Cloud. The company wants to separate workloads for chargeback to different departments. Which AWS services or features will meet these requirements? (Select TWO.)

- A. Placement groups
- B. Consolidated billing
- C. Edge locations
- D. AWS Config
- E. Multiple AWS accounts

**Answer:** BE

#### Explanation:

Consolidated billing is a feature of AWS Organizations that enables customers to consolidate billing and payment for multiple AWS accounts. With consolidated billing, customers can group multiple AWS accounts under one payer account, making it easier to manage billing and track costs across multiple accounts. Consolidated billing also offers benefits such as volume discounts, Reserved Instance discounts, and Savings Plans discounts. Consolidated billing is offered at no additional cost.

Multiple AWS accounts is a feature of AWS Organizations that enables customers to create and manage multiple AWS accounts from a central location. With multiple AWS accounts, customers can isolate workloads for different departments, projects, or environments, and apply granular access controls and policies to each account. Multiple AWS accounts also helps customers improve security, compliance, and governance of their AWS resources<sup>56</sup>. References: 5:

Consolidated billing for AWS Organizations - AWS

Billing, 6: Understanding Consolidated Bills - AWS Billing, 7: AWS Consolidated Billing: Tutorial & Best Practices, 8: Simplifying Your Bills With Consolidated Billing on AWS - Aimably, 9: AWS Consolidated Billing - W3Schools

#### NEW QUESTION 6

- (Topic 3)

A company needs to engage third-party consultants to help maintain and support its AWS environment and the company's business needs. Which AWS service or resource will meet these requirements?

- A. AWS Support
- B. AWS Organizations
- C. AWS Service Catalog
- D. AWS Partner Network (APN)

**Answer:** D

#### Explanation:

The AWS service or resource that will meet these requirements is D. AWS Partner Network (APN).

AWS Partner Network (APN) is a global community of consulting and technology partners that offer a wide range of services and solutions for AWS customers. APN partners can help customers design, architect, build, migrate, and manage their workloads and applications on AWS. APN partners have access to various resources, training, tools, and support to enhance their AWS expertise and deliver value to customers<sup>12</sup>.

AWS Support is a service that provides technical assistance and guidance for AWS customers. AWS Support offers different plans with varying levels of response time, access channels, and features. AWS Support does not directly engage third-party consultants, but rather connects customers with AWS experts and resources<sup>3</sup>.

AWS Organizations is a service that allows customers to manage multiple AWS accounts within a single organization. AWS Organizations enables customers to create groups of accounts, apply policies, automate account creation, and consolidate billing. AWS Organizations does not directly engage third-party consultants, but rather helps customers simplify and optimize their AWS account management<sup>4</sup>.

AWS Service Catalog is a service that allows customers to create and manage catalogs of IT services that are approved for use on AWS. AWS Service Catalog enables customers to control the configuration, deployment, and governance of their IT services. AWS Service Catalog does not directly engage third-party consultants, but rather helps customers standardize and streamline their IT service delivery<sup>5</sup>.

References:

1: AWS Partner Network (APN) - Amazon Web Services (AWS) 2: Find an APN Partner - Amazon Web Services (AWS) 3: AWS Support – Amazon Web Services 4: AWS Organizations – Amazon Web Services 5: AWS Service Catalog – Amazon Web Services

#### NEW QUESTION 7

- (Topic 3)

A company is planning to migrate to the AWS Cloud and wants to become more responsive to customer inquiries and feedback. The company wants to focus on organizational transformation.

A company wants to give its customers the ability to view specific data that is hosted in Amazon S3 buckets. The company wants to keep control over the full datasets that the company shares with the customers.

Which S3 feature will meet these requirements?

- A. S3 Storage Lens
- B. S3 Cross-Region Replication (CRR)
- C. S3 Versioning
- D. S3 Access Points

**Answer:** D

#### Explanation:

S3 Access Points are a feature of Amazon S3 that allows you to easily manage access to specific data that is hosted in S3 buckets. S3 Access Points are unique hostnames that customers can use to access data in S3 buckets. You can create multiple access points for a single bucket, each with its own name and permissions. You can use S3 Access Points to provide different levels of access to different groups of customers, such as read-only or write-only access. You can also use S3 Access Points to enforce encryption or logging requirements for specific data. S3 Access Points help you keep control over the full datasets that you

share with your customers, while simplifying the access management and improving the performance and scalability of your applications.

### NEW QUESTION 8

- (Topic 3)

A software engineer wants to launch a virtual machine (VM) and MySQL database on AWS. Which AWS service will meet these requirements with the LEAST operational effort?

- A. Amazon Elastic Container Service (Amazon ECS)
- B. AWS Elastic Beanstalk
- C. Amazon Lightsail
- D. Amazon EC2

**Answer:** B

#### Explanation:

AWS Elastic Beanstalk is a service that enables you to quickly deploy and manage applications in the AWS Cloud without worrying about the infrastructure that runs those applications. You simply upload your application, and Elastic Beanstalk automatically handles the details of capacity provisioning, load balancing, scaling, and application health monitoring. Elastic Beanstalk supports several platform configurations for Java, .NET, PHP, Node.js, Python, Ruby, Go, and Docker web applications that can run on familiar servers such as Apache, Nginx, Passenger, and IIS. You can also use Elastic Beanstalk to launch a virtual machine (VM) and MySQL database on AWS with the least operational effort. Amazon Elastic Container Service (Amazon ECS) is a fully managed container orchestration service that enables you to easily run, scale, and secure Docker containerized applications on AWS. However, it requires more operational effort than Elastic Beanstalk, as you need to define your application architecture and the specifications of the containers that run it. Amazon Lightsail is an easy-to-use cloud platform that offers everything you need to build an application or website, plus a cost-effective, monthly plan. It is designed for developers who have little or no prior cloud experience and want to launch and manage applications on AWS with minimal complexity. However, it does not support MySQL databases, and it requires more operational effort than Elastic Beanstalk, as you need to configure your VM and database settings. Amazon EC2 is a web service that provides secure, resizable compute capacity in the cloud. It allows you to launch a virtual machine (VM) and MySQL database on AWS, but it requires the most operational effort, as you need to provision, monitor, and manage your EC2 instances and database.

### NEW QUESTION 9

- (Topic 3)

Which cloud computing advantage is a company applying when it uses AWS Regions to increase application availability to users in different countries?

- A. Pay-as-you-go pricing
- B. Capacity forecasting
- C. Economies of scale
- D. Global reach

**Answer:** D

#### Explanation:

Global reach is a cloud computing advantage that a company can apply when it uses AWS Regions to increase application availability to users in different countries. Global reach refers to the ability to deploy applications and services in multiple geographic locations around the world, and to serve customers with low latency and high performance. AWS has the largest and most reliable global infrastructure of any cloud provider, with 25 Regions and 81 Availability Zones across the Americas, Europe, Asia Pacific, Africa, and the Middle East<sup>123</sup>. By using AWS Regions, a company can choose the best location for its application based on customer proximity, compliance requirements, and disaster recovery strategies<sup>23</sup>. References: 1: AWS Global Infrastructure - Amazon Web Services (AWS), 2: Regions and Availability Zones - Amazon Elastic Compute Cloud, 3: AWS Infrastructure: Regions and Availability Zones Explained

### NEW QUESTION 10

- (Topic 3)

Which tasks are the responsibility of the customer, according to the AWS shared responsibility model? (Select TWO.)

- A. Patch the Amazon RDS operating system.
- B. Upgrade the firmware of the network infrastructure.
- C. Manage data encryption.
- D. Maintain physical access control in an AWS Region.
- E. Grant least privilege access to IAM users.

**Answer:** CE

#### Explanation:

According to the AWS shared responsibility model, the customer is responsible for security in the cloud, which includes the tasks of managing data encryption and granting least privilege access to IAM users. Data encryption is the process of transforming data into an unreadable format that can only be accessed with a key or a password. The customer must decide whether to encrypt their data at rest (when it is stored on AWS) or in transit (when it is moving between AWS and the customer or between AWS services). The customer must also choose the encryption method, algorithm, and key management solution that best suit their needs. AWS provides various services and features that support data encryption, such as AWS Key Management Service (AWS KMS), AWS Certificate Manager (ACM), and AWS Encryption SDK<sup>5</sup> IAM users are entities that represent the people or applications that interact with AWS resources and services. The customer must grant the IAM users the minimum permissions that they need to perform their tasks, and avoid giving them unnecessary or excessive access. This is known as the principle of least privilege, and it helps reduce the risk of unauthorized or malicious actions. The customer can use IAM policies, roles, groups, and permissions boundaries to manage the access of IAM users.

### NEW QUESTION 10

- (Topic 3)

According to the AWS shared responsibility model, which task is the customer's responsibility?

- A. Maintaining the infrastructure needed to run AWS Lambda
- B. Updating the operating system of Amazon DynamoDB instances
- C. Maintaining Amazon S3 infrastructure
- D. Updating the guest operating system on Amazon EC2 instances

**Answer:** D

**Explanation:**

The AWS shared responsibility model describes the division of responsibilities between AWS and the customer for security and compliance. AWS is responsible for the security of the cloud, which includes the hardware, software, networking, and facilities that run AWS services. The customer is responsible for security in the cloud, which includes the customer data, applications, operating systems, and network and firewall configurations. Therefore, updating the guest operating system on Amazon EC2 instances is the customer's responsibility<sup>2</sup>

**NEW QUESTION 13**

- (Topic 3)

Which AWS service gives users the ability to discover and protect sensitive data that is stored in Amazon S3 buckets?

- A. Amazon Macie
- B. Amazon Detective
- C. Amazon GuardDuty
- D. AWS IAM Access Analyzer

**Answer:** A

**Explanation:**

Amazon Macie is a data security and privacy service offered by AWS that uses machine learning and pattern matching to discover the sensitive data stored within Amazon S3. You can define your own custom type of sensitive data category that might be unique to your business or use case. Macie also provides you with dashboards and alerts that give you visibility into how your data is being accessed or moved. Macie helps you protect your data by enabling you to apply data protection techniques such as encryption, deletion, access control, and auditing. References: Strengthen the security of sensitive data stored in Amazon S3 by using additional AWS services, Security best practices for Amazon S3, Sensitive Data Protection on AWS, Sensitive Data Protection on Amazon Web Services

**NEW QUESTION 18**

- (Topic 3)

Which AWS service can provide a dedicated network connection with consistent low latency from on premises to the AWS Cloud?

- A. Amazon VPC
- B. Amazon Kinesis Data Streams
- C. AWS Direct Connect
- D. Amazon OpenSearch Service

**Answer:** C

**Explanation:**

AWS Direct Connect is a service that provides a dedicated network connection from on premises to the AWS Cloud. It can reduce network costs, increase bandwidth throughput, and provide a more consistent network experience than internet-based connections. It can also provide low latency for applications that require real-time data transfer<sup>4</sup>. Amazon VPC is a service that provides a logically isolated section of the AWS Cloud where users can launch AWS resources in a virtual network that they define. Amazon Kinesis Data Streams is a service that provides a scalable and durable stream of data records for real-time data processing. Amazon OpenSearch Service is a service that provides a fully managed, scalable, and secure search and analytics solution that is compatible with Elasticsearch.

**NEW QUESTION 23**

- (Topic 3)

Which Amazon EC2 instance pricing model can provide discounts of up to 90%?

- A. Reserved Instances
- B. On-Demand
- C. Dedicated Hosts
- D. Spot Instances

**Answer:** D

**Explanation:**

Spot Instances are Amazon EC2 instances that are available at a discounted price compared to On-Demand pricing. Spot Instances use spare EC2 capacity that is not being used by other customers, and the price fluctuates based on supply and demand. Customers can request Spot Instances for their applications and specify the maximum price they are willing to pay per hour. If the Spot price is lower than the customer's bid, the Spot Instance is launched and the customer pays the current Spot price. However, if the Spot price rises above the customer's bid, the Spot Instance is terminated by AWS and the customer is charged for the partial hour of usage. Therefore, Spot Instances can provide discounts of up to 90% or more, but they are not suitable for applications that require continuous or predictable availability. Spot Instances are recommended for applications that are flexible, fault-tolerant, or have low priority, such as batch processing, data analysis, or testing and development.

**NEW QUESTION 25**

- (Topic 3)

A company wants to set AWS spending targets and track costs against those targets. Which AWS tool or feature should the company use to meet these requirements?

- A. AWS Cost Explorer
- B. AWS Budgets
- C. AWS Cost and Usage Report
- D. Savings Plans

**Answer:** B

**Explanation:**

AWS Budgets is a tool that allows users to set AWS spending targets and track costs against those targets. Users can create budgets for various dimensions, such as service, linked account, tag, and more. Users can also receive alerts when the actual or forecasted costs exceed or are projected to exceed the budgeted amount. AWS Cost Explorer, AWS Cost and Usage Report, and Savings Plans are other AWS tools or features that can help users manage and optimize their AWS costs, but they do not enable users to set and track spending targets .

#### NEW QUESTION 28

- (Topic 3)

A company has a physical tape library to store data backups. The tape library is running out of space. The company needs to extend the tape library's capacity to the AWS Cloud.

Which AWS service should the company use to meet this requirement?

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

**Answer:** D

#### Explanation:

AWS Storage Gateway is a hybrid cloud storage service that provides on- premises access to virtually unlimited cloud storage. You can use AWS Storage Gateway to simplify storage management and reduce costs for key hybrid cloud storage use cases. One of these use cases is tape-based backup, which allows you to store data backups on virtual tapes in the AWS Cloud. You can use the Tape Gateway feature of AWS Storage Gateway to extend your existing physical tape library to the AWS Cloud. Tape Gateway provides a virtual tape infrastructure that scales seamlessly with your backup needs and eliminates the operational burden of provisioning, scaling, and maintaining a physical tape infrastructure<sup>123</sup>. References: 1: Cloud Storage Appliances, Hybrid Device - AWS Storage Gateway - AWS, 2: AWS Storage Gateway Documentation, 3: AWS Storage Gateway Features | Amazon Web Services

#### NEW QUESTION 29

- (Topic 3)

Which AWS service provides encryption at rest for Amazon RDS and for Amazon Elastic Block Store (Amazon EBS) volumes?

- A. AWS Lambda
- B. AWS Key Management Service (AWS KMS)
- C. AWSWAF
- D. Amazon Rekognition

**Answer:** B

#### Explanation:

AWS Key Management Service (AWS KMS) is a managed service that enables you to easily encrypt your data. AWS KMS provides you with centralized control of the encryption keys used to protect your data. You can use AWS KMS to encrypt data in Amazon RDS and Amazon EBS volumes<sup>12</sup>

#### NEW QUESTION 31

- (Topic 3)

Which AWS services or features give users the ability to create a network connection between two VPCs? (Select TWO.)

- A. VPC endpoints
- B. Amazon Route 53
- C. VPC peering
- D. AWS Direct Connect
- E. AWS Transit Gateway

**Answer:** CE

#### Explanation:

VPC peering and AWS Transit Gateway are two AWS services or features that give users the ability to create a network connection between two VPCs. VPC peering is a networking connection between two VPCs that enables you to route traffic between them privately. You can create a VPC peering connection between your own VPCs, with a VPC in another AWS account, or with a VPC in a different AWS Region. Traffic between peered VPCs never traverses the public internet. VPC peering does not support transitive peering relationships, which means that if VPC A is peered with VPC B, and VPC B is peered with VPC C, then VPC A and VPC C are not automatically peered<sup>789</sup>. AWS Transit Gateway is a networking service that acts as a regional router for your VPCs and on- premises networks. You can attach up to 5,000 VPCs and VPN connections to a single transit gateway and route traffic between them. AWS Transit Gateway simplifies the management and scalability of your network architecture, as you only need to create and manage a single connection from the central transit gateway to each connected network. AWS Transit Gateway supports transitive routing, which means that any network that is attached to the transit gateway can communicate with any other network that is attached to the same transit gateway . References: 7: VPC peering - Amazon Virtual Private Cloud, 8: Connect VPCs using VPC peering - Amazon Virtual Private Cloud, 9: Amazon VPC-to-Amazon VPC connectivity options - Amazon Virtual Private Cloud, : [AWS Transit Gateway - Amazon Web Services], : [Connect VPCs using AWS Transit Gateway - Amazon Virtual Private Cloud], : [AWS Transit Gateway: Simplify Your Network Architecture]

#### NEW QUESTION 35

- (Topic 3)

Elasticity in the AWS Cloud refers to which of the following? (Select TWO.)

- A. How quickly an Amazon EC2 instance can be restarted
- B. The ability to rightsized resources as demand shifts
- C. The maximum amount of RAM an Amazon EC2 instance can use
- D. The pay-as-you-go billing model
- E. How easily resources can be procured when they are needed

**Answer:** BE

#### Explanation:

Elasticity in the AWS Cloud refers to the ability to acquire resources as you need them and release resources when you no longer need them. In the cloud, you want to do this automatically<sup>1</sup>. This means that you can rightsized resources as demand shifts, and you can easily procure resources when they are needed. Elasticity is not related to how quickly an Amazon EC2 instance can be restarted, the maximum amount of RAM an Amazon EC2 instance can use, or the pay-as-you-go billing model. These are aspects of scalability, performance, and cost, respectively<sup>2</sup>.

For more information on elasticity, you can refer to the following sources:

- ? Elasticity - AWS Well-Architected Framework
- ? Elastic - Reactive Systems on AWS
- ? What is the difference between scalability and elasticity?

### NEW QUESTION 37

- (Topic 3)

An IT engineer needs to access AWS services from an on-premises application. Which credentials or keys does the application need for authentication?

- A. AWS account user name and password
- B. IAM access key and secret
- C. Amazon EC2 key pairs
- D. AWS Key Management Service (AWS KMS) keys

**Answer:** B

#### Explanation:

IAM access keys are long-term credentials that consist of an access key ID and a secret access key. You use access keys to sign programmatic requests that you make to AWS. If you need to access AWS services from an on-premises application, you can use IAM access keys to authenticate your requests. AWS account user name and password are used to sign in to the AWS Management Console. Amazon EC2 key pairs are used to connect to your EC2 instances using SSH. AWS Key Management Service (AWS KMS) keys are used to encrypt and decrypt your data using the AWS Encryption SDK or the AWS CLI.

### NEW QUESTION 38

- (Topic 3)

Which AWS services can a company use to achieve a loosely coupled architecture? (Select TWO.)

- A. Amazon Workspaces
- B. Amazon Simple Queue Service (Amazon SQS)
- C. Amazon Connect
- D. AWS Trusted Advisor
- E. AWS Step Functions

**Answer:** BE

#### Explanation:

Amazon Simple Queue Service (Amazon SQS) and AWS Step Functions are AWS services that can be used to achieve a loosely coupled architecture. Amazon SQS is a fully managed message queuing service that enables you to decouple and scale microservices, distributed systems, and serverless applications. AWS Step Functions lets you coordinate multiple AWS services into serverless workflows so you can build and update apps quickly. Using Step Functions, you can design and run workflows that stitch together services such as AWS Lambda and Amazon SNS into feature-rich applications. References: Amazon SQS, AWS Step Functions

### NEW QUESTION 39

- (Topic 3)

A company needs to control inbound and outbound traffic for an Amazon EC2 instance.

Which AWS service or feature can the company associate with the EC2 instance to meet this requirement?

- A. Network ACL
- B. Security group
- C. AWS WAF
- D. VPC route tables

**Answer:** B

#### Explanation:

A security group is a virtual firewall that can be associated with an Amazon EC2 instance to control the inbound and outbound traffic for the instance. You can specify which protocols, ports, and source or destination IP ranges are allowed or denied by the security group. A network ACL is a stateless filter that can be associated with a subnet to control the traffic to and from the subnet, but it is not associated with an EC2 instance<sup>4</sup>. AWS WAF is a web application firewall that helps protect your web applications or APIs against common web exploits that may affect availability, compromise security, or consume excessive resources. VPC route tables are used to determine where network traffic is directed within a VPC or to an internet gateway, virtual private gateway, NAT device, VPC peering connection, or VPC endpoint.

### NEW QUESTION 42

- (Topic 3)

A development team wants to deploy multiple test environments for an application in a fast repeatable manner.

Which AWS service should the team use?

- A. Amazon EC2
- B. AWS CloudFormation
- C. Amazon QuickSight
- D. Amazon Elastic Container Service (Amazon ECS)

**Answer:** B

#### Explanation:

AWS CloudFormation is a service that allows you to model and provision your AWS resources using templates. You can define your infrastructure as code and

automate the creation and update of your resources. AWS CloudFormation also supports nested stacks, change sets, and rollback features to help you manage complex and dynamic environments<sup>34</sup>. References:

? AWS CloudFormation

? AWS Certified Cloud Practitioner Exam Guide

#### NEW QUESTION 44

- (Topic 3)

A company needs to implement identity management for a fleet of mobile apps that are running in the AWS Cloud. Which AWS service will meet this requirement?

- A. Amazon Cognito
- B. AWS Security Hub
- C. AWS Shield
- D. AWS WAF

**Answer:** A

#### Explanation:

Amazon Cognito is a service that provides identity management for mobile and web applications, allowing users to sign up, sign in, and access AWS resources with different identity providers. AWS Security Hub is a service that provides a comprehensive view of the security posture of AWS accounts and resources. AWS Shield is a service that provides protection against distributed denial of service (DDoS) attacks. AWS WAF is a web application firewall that helps protect web applications from common web exploits.

#### NEW QUESTION 47

- (Topic 3)

A company wants to query its server logs to gain insights about its customers' experiences. Which AWS service will store this data MOST cost-effectively?

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

**Answer:** D

#### Explanation:

Amazon S3 is an AWS service that provides scalable, durable, and cost-effective object storage in the cloud. Amazon S3 can store any amount and type of data, such as server logs, and offers various storage classes with different performance and pricing characteristics. Amazon S3 is the most cost-effective option for storing server logs, as it offers low-cost storage classes, such as S3 Standard-Infrequent Access (S3 Standard-IA) and S3 Intelligent-Tiering, that are suitable for infrequently accessed or changing access patterns data. Amazon S3 also integrates with other AWS services, such as Amazon Athena and Amazon OpenSearch Service, that can query the server logs directly from S3 without requiring any additional data loading or transformation. References: Amazon S3, Amazon S3 Storage Classes, Querying Data in Amazon S3

#### NEW QUESTION 50

- (Topic 3)

A company wants to generate a list of IAM users. The company also wants to view the status of various credentials that are associated with the users, such as password, access keys, and multi-factor authentication (MFA) devices. Which AWS service or feature will meet these requirements?

- A. IAM credential report
- B. AWS IAM Identity Center (AWS Single Sign-On)
- C. AWS Identity and Access Management Access Analyzer
- D. AWS Cost and Usage Report

**Answer:** A

#### Explanation:

An IAM credential report is a feature of AWS Identity and Access Management (IAM) that allows you to view and download a report that lists all IAM users in your account and the status of their various credentials, such as passwords, access keys, and MFA devices. You can use this report to audit the security status of your IAM users and ensure that they follow the best practices for credential management<sup>1</sup>. References: 1: AWS Documentation - IAM User Guide - Getting credential reports for your AWS account

#### NEW QUESTION 54

- (Topic 3)

Which AWS service uses AWS Compute Optimizer to provide sizing recommendations based on workload metrics?

- A. Amazon EC2
- B. Amazon RDS
- C. Amazon Lightsail
- D. AWS Step Functions

**Answer:** A

#### Explanation:

Amazon EC2 is a web service that provides secure, resizable compute capacity in the cloud. It allows you to launch virtual servers, called instances, with different configurations of CPU, memory, storage, and networking resources. AWS Compute Optimizer analyzes the specifications and utilization metrics of your Amazon EC2 instances and generates recommendations for optimal instance types that can reduce costs and improve performance. You can view the recommendations on the AWS Compute Optimizer console or the Amazon EC2 console<sup>12</sup>. Amazon RDS, Amazon Lightsail, and AWS Step Functions are not supported by AWS Compute Optimizer. Amazon RDS is a managed relational database service that lets you set up, operate, and scale a relational database in the cloud. Amazon Lightsail is an easy-to-use cloud platform that offers everything you need to

build an application or website, plus a cost-effective, monthly plan. AWS Step Functions lets you coordinate multiple AWS services into serverless workflows so you can build and update apps quickly3 .

#### NEW QUESTION 57

- (Topic 3)

A company is moving to the AWS Cloud to reduce operational overhead for its application infrastructure. Which IT operation will the company still be responsible for after the migration to AWS?

- A. Security patching of AWS Elastic Beanstalk
- B. Backups of data that is stored in Amazon Aurora
- C. Termination of Amazon EC2 instances that are managed by AWS Auto Scaling
- D. Configuration of IAM access controls

**Answer:** D

#### Explanation:

AWS Elastic Beanstalk, Amazon Aurora, and AWS Auto Scaling are managed services that reduce the operational overhead for the customers. AWS is responsible for security patching, backups, and termination of these services. However, the customers are still responsible for configuring IAM access controls to manage the permissions and policies for their AWS resources. This is part of the AWS shared responsibility model, which defines the security and compliance responsibilities of AWS and the customers. You can learn more about the AWS shared responsibility model from this whitepaper or this digital course.

#### NEW QUESTION 61

- (Topic 3)

Which cloud concept is demonstrated by using AWS Compute Optimizer?

- A. Security validation
- B. Rightsizing
- C. Elasticity
- D. Global reach

**Answer:** B

#### Explanation:

Rightsizing is the cloud concept that is demonstrated by using AWS Compute Optimizer. Rightsizing is the process of adjusting the type and size of your cloud resources to match the optimal performance and cost for your workloads. AWS Compute Optimizer is a service that analyzes the configuration and utilization metrics of your AWS resources, such as Amazon EC2 instances, Amazon EBS volumes, AWS Lambda functions, and Amazon ECS services on AWS Fargate. It reports whether your resources are optimal, and generates optimization recommendations to reduce the cost and improve the performance of your workloads. AWS Compute Optimizer uses machine learning to analyze your historical utilization data and compare it with the most cost-effective AWS alternatives. You can use the recommendations to evaluate the trade-offs between cost and performance, and decide when to move or resize your resources to achieve the best results. References: Workload Rightsizing - AWS Compute Optimizer - AWS, What is AWS Compute Optimizer? - AWS Compute Optimizer

#### NEW QUESTION 63

- (Topic 3)

Which AWS service can a company use to find security and compliance reports, including International Organization for Standardization (ISO) reports?

- A. AWS Artifact
- B. Amazon CloudWatch
- C. AWS Config
- D. AWS Audit Manager

**Answer:** A

#### Explanation:

AWS Artifact is a self-service portal that provides on-demand access to AWS security and compliance reports and select online agreements. You can use AWS Artifact to download AWS service audit reports, such as ISO, PCI, and SOC, and to accept and manage agreements with AWS, such as the Business Associate Addendum (BAA).

#### NEW QUESTION 66

- (Topic 3)

A company has migrated its workloads to AWS. The company wants to adopt AWS at scale and operate more efficiently and securely. Which AWS service or framework should the company use for operational support?

- A. AWS Support
- B. AWS Cloud Adoption Framework (AWS CAF)
- C. AWS Managed Services (AMS)
- D. AWS Well-Architected Framework

**Answer:** D

#### Explanation:

The AWS Well-Architected Framework is a set of best practices and guidelines for designing and operating workloads on AWS. It helps customers achieve operational excellence, security, reliability, performance efficiency, cost optimization, and sustainability. The framework is based on six pillars, each with its own design principles, best practices, and questions. Customers can use the framework to assess their current state, identify gaps, and implement improvements12. AWS Support is a service that provides technical assistance, guidance, and resources for AWS customers. It offers different plans with varying levels of access to AWS experts, response times, and features3. AWS Support does not provide a comprehensive framework for operational support. AWS Cloud Adoption Framework (AWS CAF) is a guidance tool that helps customers plan and execute their cloud migration journey. It provides a set of perspectives, capabilities, and best practices to align the business and technical aspects of cloud adoption4. AWS CAF does not focus on operational support for existing workloads on AWS. AWS Managed Services (AMS) is a service that operates AWS infrastructure on behalf of customers. It provides a secure and compliant environment, automates

common activities, and applies best practices for provisioning, patching, backup, recovery, and monitoring<sup>5</sup>. AWS does not provide a framework for customers to operate their own workloads on AWS.

#### NEW QUESTION 67

- (Topic 3)

Which option is AWS responsible for under the AWS shared responsibility model?

- A. Network and firewall configuration
- B. Client-side data encryption
- C. Management of user permissions
- D. Hardware and infrastructure

**Answer:** D

#### Explanation:

Hardware and infrastructure is the option that AWS is responsible for under the AWS shared responsibility model. The AWS shared responsibility model describes how AWS and customers share responsibilities for security and compliance in the cloud. AWS is responsible for security of the cloud, which means protecting the infrastructure that runs all the services offered in the AWS Cloud. This infrastructure is composed of the hardware, software, networking, and facilities that run AWS Cloud services. Customers are responsible for security in the cloud, which means taking care of the security of their own applications, data, and operating systems. This includes network and firewall configuration, client-side data encryption, management of user permissions, and more.

#### NEW QUESTION 69

- (Topic 3)

A company is hosting an application in the AWS Cloud. The company wants to verify that underlying AWS services and general AWS infrastructure are operating normally.

Which combination of AWS services can the company use to gather the required information? (Select TWO.)

- A. AWS Personal Health Dashboard
- B. AWS Systems Manager
- C. AWS Trusted Advisor
- D. AWS Service Health Dashboard
- E. AWS Service Catalog

**Answer:** AD

#### Explanation:

AWS Personal Health Dashboard and AWS Service Health Dashboard are two AWS services that can help the company to verify that underlying AWS services and general AWS infrastructure are operating normally. AWS Personal Health Dashboard provides a personalized view into the performance and availability of the AWS services you are using, as well as alerts that are automatically triggered by changes in the health of those services. In addition to event-based alerts, Personal Health Dashboard provides proactive notifications of scheduled activities, such as any changes to the infrastructure powering your resources, enabling you to better plan for events that may affect you. These notifications can be delivered to you via email or mobile for quick visibility, and can always be viewed from within the AWS Management Console. When you get an alert, it includes detailed information and guidance, enabling you to take immediate action to address AWS events impacting your resources<sup>3</sup>. AWS Service Health Dashboard provides a general status of AWS services, and the Service health view displays the current and historical status of all AWS services. This page shows reported service events for services across AWS Regions. You don't need to sign in or have an AWS account to access the AWS Service Health Dashboard – Service health page. You can also subscribe to RSS feeds for specific services or regions to receive notifications about service events<sup>4</sup>. References: Getting started with your AWS Health Dashboard – Your account health, Introducing AWS Personal Health Dashboard

#### NEW QUESTION 74

- (Topic 3)

A company wants to allow users to authenticate and authorize multiple AWS accounts by using a single set of credentials.

Which AWS service or resource will meet this requirement?

- A. AWS Organizations
- B. IAM user
- C. AWS IAM Identity Center (AWS Single Sign-On)
- D. AWS Control Tower

**Answer:** C

#### Explanation:

AWS IAM Identity Center (AWS Single Sign-On) is a cloud-based service that makes it easy to centrally manage single sign-on (SSO) access to multiple AWS accounts and business applications. You can use AWS SSO to enable your users to sign in to the AWS Management Console or the AWS Command Line Interface (AWS CLI) with their existing corporate credentials<sup>2</sup>. You can also manage SSO access and user permissions across all your AWS accounts in AWS Organizations<sup>3</sup>. References: AWS Single Sign-On - AWS Documentation, AWS Organizations - AWS Documentation

#### NEW QUESTION 79

- (Topic 3)

Which pillar of the AWS Well-Architected Framework includes the AWS shared responsibility model?

- A. Operational excellence
- B. Performance efficiency
- C. Reliability
- D. Security

**Answer:** D

#### Explanation:

The AWS Well-Architected Framework is a set of best practices and guidelines for designing and operating reliable, secure, efficient, and cost-effective systems in the cloud. The framework consists of five pillars: operational excellence, performance efficiency, reliability, security, and cost optimization. The security pillar covers the AWS shared responsibility model, which defines the security and compliance responsibilities of AWS and the customers. You can learn more about the AWS Well-Architected Framework from [this whitepaper] or [this digital course].

#### NEW QUESTION 80

- (Topic 3)

A company is considering migration to the AWS Cloud. The company wants a fully managed service or feature that can transfer streaming data from multiple sources to an Amazon S3 bucket.

Which AWS service or feature should the company use to meet these requirements?

- A. AWS DataSync
- B. Amazon Kinesis Data Firehose
- C. S3 Select
- D. AWS Transfer Family

**Answer:** B

#### Explanation:

Amazon Kinesis Data Firehose is a fully managed service that delivers real-time streaming data to destinations such as Amazon S3, Amazon Redshift, Amazon Elasticsearch Service, and Splunk. You can use Amazon Kinesis Data Firehose to capture, transform, and load streaming data from multiple sources, such as web applications, mobile devices, IoT sensors, and social media.

#### NEW QUESTION 85

- (Topic 3)

What does the concept of agility mean in AWS Cloud computing? (Select TWO.)

- A. The speed at which AWS resources are implemented
- B. The speed at which AWS creates new AWS Regions
- C. The ability to experiment quickly
- D. The elimination of wasted capacity
- E. The low cost of entry into cloud computing

**Answer:** AC

#### Explanation:

Agility in AWS Cloud computing means the ability to rapidly provision and deprovision AWS resources as needed, and the ability to experiment quickly with new ideas and solutions. Agility helps businesses to respond to changing customer demands, market opportunities, and competitive threats, and to innovate faster and cheaper. Agility also reduces the risk of failure, as businesses can test and validate their assumptions before committing to large-scale deployments. Some of the benefits of agility in AWS Cloud computing are:

? The speed at which AWS resources are implemented: AWS provides a variety of services and tools that allow you to create, configure, and launch AWS resources in minutes, using the AWS Management Console, the AWS Command Line Interface (AWS CLI), the AWS Software Development Kits (AWS SDKs), or the AWS CloudFormation templates. You can also use the AWS Cloud Development Kit (AWS CDK) to define your AWS resources as code using familiar programming languages, and synthesize them into AWS CloudFormation templates. You can also use the AWS Service Catalog to create and manage standardized portfolios of AWS resources that meet your organizational policies and best practices. AWS also offers on-demand, pay-as-you-go pricing models, so you only pay for the resources you use, and you can scale them up or down as your needs change<sup>12345</sup>

? The ability to experiment quickly: AWS enables you to experiment quickly with new ideas and solutions, without having to invest in upfront capital or long-term commitments. You can use AWS to create and test multiple prototypes, hypotheses, and minimum viable products (MVPs) in parallel, and measure their performance and feedback. You can also use AWS to leverage existing services and solutions, such as AWS Marketplace, AWS Solutions, and AWS Quick Starts, that can help you accelerate your innovation process. AWS also supports a culture of experimentation and learning, by providing tools and resources for continuous integration and delivery (CI/CD), testing, monitoring, and analytics.

References: Six advantages of cloud computing - Overview of Amazon Web Services, AWS Cloud Development Kit (AWS CDK), AWS Service Catalog, AWS Pricing, AWS CloudFormation, [Experimentation and Testing - AWS Well-Architected Framework], [AWS Marketplace], [AWS Solutions], [AWS Quick Starts], [AWS Developer Tools]

#### NEW QUESTION 88

- (Topic 3)

A company wants to use the AWS Cloud to deploy an application globally.

Which architecture deployment model should the company use to meet this requirement?

- A. Multi-Region
- B. Single-Region
- C. Multi-AZ
- D. Single-AZ

**Answer:** A

#### Explanation:

The architecture deployment model that the company should use to meet this requirement is A. Multi-Region.

A multi-region deployment model is a cloud computing architecture that distributes an application and its data across multiple geographic regions. A multi-region deployment model enables a company to achieve global reach, high availability, disaster recovery, and performance optimization. By deploying an application in multiple regions, a company can serve customers from the nearest region, reduce latency, increase redundancy, and comply with data sovereignty regulations<sup>12</sup>.

A single-region deployment model is a cloud computing architecture that runs an application and its data within a single geographic region. A single-region deployment model is simpler and cheaper than a multi-region deployment model, but it has limited scalability, availability, and performance. A single-region deployment model may not be suitable for a company that wants to deploy an application globally, as it may face challenges such as network latency, regional outages, or regulatory compliance<sup>12</sup>.

A multi-AZ (Availability Zone) deployment model is a cloud computing architecture that distributes an application and its data across multiple isolated locations within a single region. An Availability Zone is a physically separate location within an AWS Region that has independent power, cooling, and networking. A multi-AZ deployment model enhances the availability and durability of an application by providing redundancy and fault tolerance within a region<sup>34</sup>.

A single-AZ deployment model is a cloud computing architecture that runs an application and its data within a single Availability Zone. A single-AZ deployment

model is the simplest and most cost-effective option, but it has no redundancy or fault tolerance. A single-AZ deployment model may not be suitable for a company that wants to deploy an application globally, as it may face challenges such as network latency, regional outages, or regulatory compliance<sup>34</sup>.

References:

1: AWS Cloud Computing - W3Schools 2: Understand the Different Cloud Computing Deployment Models Unit - Trailhead 3: Regions and Availability Zones - Amazon Elastic Compute Cloud 4: AWS Reference Architecture Diagrams

#### NEW QUESTION 92

- (Topic 3)

Which AWS service or feature identifies whether an Amazon S3 bucket or an IAM role has been shared with an external entity?

- A. AWS Service Catalog
- B. AWS Systems Manager
- C. AWS IAM Access Analyzer
- D. AWS Organizations

**Answer: C**

#### Explanation:

AWS IAM Access Analyzer is a service that helps you identify the resources in your organization and accounts, such as Amazon S3 buckets or IAM roles, that are shared with an external entity. This lets you identify unintended access to your resources and data, which is a security risk. IAM Access Analyzer uses logic-based reasoning to analyze the resource-based policies in your AWS environment. For each instance of a resource shared outside of your account, IAM Access Analyzer generates a finding. Findings include information about the access and the external principal granted to it<sup>345</sup>. References: 3: Using AWS Identity and Access Management Access Analyzer, 4: IAM Access Analyzer - Amazon Web Services (AWS), 5: Welcome - IAM Access Analyzer

#### NEW QUESTION 95

- (Topic 3)

At what support level do users receive access to a support concierge?

- A. Basic Support
- B. Developer Support
- C. Business Support
- D. Enterprise Support

**Answer: D**

#### Explanation:

Users receive access to a support concierge at the Enterprise Support level. A support concierge is a team of AWS billing and account experts that specialize in working with enterprise accounts. They can help users with billing and account inquiries, cost optimization, FinOps support, cost analysis, and prioritized answers to billing questions. The support concierge is included as part of the Enterprise Support plan, which also provides access to a Technical Account Manager (TAM), Infrastructure Event Management, AWS Trusted Advisor, and 24/7 technical support. References: AWS Support Plan Comparison, AWS Enterprise Support Plan, AWS Support Concierge

#### NEW QUESTION 97

- (Topic 3)

A company's IT team is managing MySQL database server clusters. The IT team has to patch the database and take backup snapshots of the data in the clusters. The company wants to move this workload to AWS so that these tasks will be completed automatically.

What should the company do to meet these requirements?

- A. Deploy MySQL database server clusters on Amazon EC2 instances.
- B. Use Amazon RDS with a MySQL database.
- C. Use an AWS CloudFormation template to deploy MySQL database servers on Amazon EC2 instances.
- D. Migrate all the MySQL database data to Amazon S3.

**Answer: B**

#### Explanation:

Amazon RDS is a service that makes it easy to set up, operate, and scale a relational database in the cloud. Amazon RDS supports MySQL as one of the database engines. By using Amazon RDS with a MySQL database, the company can offload the tasks of patching the database and taking backup snapshots to AWS. Amazon RDS automatically patches the database software and operating system of the database instances. Amazon RDS also automatically backs up the database and retains the backups for a user-defined retention period. The company can also restore the database to any point in time within the retention period. Deploying MySQL database server clusters on Amazon EC2 instances, using an AWS CloudFormation template to deploy MySQL database servers on Amazon EC2 instances, or migrating all the MySQL database data to Amazon S3 are not the best options to meet the requirements. These options would not automate the tasks of patching the database and taking backup snapshots, and would require more operational overhead from the company<sup>3</sup>

#### NEW QUESTION 102

- (Topic 3)

Which of the following actions are controlled with AWS Identity and Access Management (IAM)? (Select TWO.)

- A. Control access to AWS service APIs and to other specific resources.
- B. Provide intelligent threat detection and continuous monitoring.
- C. Protect the AWS environment using multi-factor authentication (MFA).
- D. Grant users access to AWS data centers.
- E. Provide firewall protection for applications from common web attacks.

**Answer: AC**

#### Explanation:

AWS Identity and Access Management (IAM) is a service that enables you to manage access to AWS services and resources securely. You can use IAM to perform the following actions:

? Control access to AWS service APIs and to other specific resources: You can create users, groups, roles, and policies that define who can access which AWS resources and how. You can also use IAM to grant temporary access to users or applications that need to perform certain tasks on your behalf<sup>3</sup>

? Protect the AWS environment using multi-factor authentication (MFA): You can enable MFA for your IAM users and root user to add an extra layer of security to your AWS account. MFA requires users to provide a unique authentication code from an approved device or SMS text message, in addition to their user name and password, when they sign in to AWS<sup>4</sup>

#### NEW QUESTION 106

- (Topic 3)

Which options are AWS Cloud Adoption Framework (AWS CAF) people perspective capabilities? (Select TWO.)

- A. Organizational alignment
- B. Portfolio management
- C. Organization design
- D. Risk management
- E. Modern application development

**Answer:** AC

#### Explanation:

The AWS Cloud Adoption Framework (AWS CAF) people perspective capabilities are the organizational skills and processes that enable effective cloud adoption. According to the AWS CAF people perspective whitepaper<sup>1</sup>, there are seven capabilities in this perspective, two of which are:

? Organizational alignment: This capability helps you align your organizational structure, roles, and responsibilities to support your cloud transformation goals and objectives. It involves assessing your current and desired state of alignment, identifying gaps and misalignments, and designing and implementing changes to optimize your cloud performance<sup>1</sup>.

? Organization design: This capability helps you design and evolve your organization to enable agility, innovation, and collaboration in the cloud. It involves defining your cloud operating model, identifying the skills and competencies needed for cloud roles, and creating career paths and development plans for your cloud workforce<sup>1</sup>.

The other options are not capabilities in the AWS CAF people perspective. Portfolio management, risk management, and modern application development are capabilities in the AWS CAF business perspective, governance perspective, and platform perspective respectively<sup>2</sup>.

References:

? 1: AWS Cloud Adoption Framework: People Perspective - AWS Cloud Adoption Framework: People Perspective

? 2: AWS Cloud Adoption Framework - AWS Cloud Adoption Framework

#### NEW QUESTION 110

- (Topic 3)

A company is using Amazon DynamoDB for its application database.

Which tasks are the responsibility of AWS, according to the AWS shared responsibility model? (Select TWO.)

- A. Classify data.
- B. Configure access permissions.
- C. Manage encryption options.
- D. Provide public endpoints to store and retrieve data.
- E. Manage the infrastructure layer and the operating system.

**Answer:** DE

#### Explanation:

According to the AWS shared responsibility model, AWS is responsible for security of the cloud, while customers are responsible for security in the cloud. This means that AWS is responsible for protecting the infrastructure that runs AWS services, such as hardware, software, networking, and facilities. Customers are responsible for managing their data, classifying their assets, and using IAM tools to apply the appropriate permissions. For abstracted services, such as Amazon DynamoDB, AWS operates the infrastructure layer, the operating system, and platforms, and provides customers with public endpoints to store and retrieve data. Customers are responsible for classifying their data, managing their encryption options, and configuring their access permissions. References: Shared Responsibility Model, Security and compliance in Amazon DynamoDB, [AWS Cloud Practitioner Essentials: Module 2 - Security in the Cloud]

#### NEW QUESTION 114

- (Topic 3)

A company wants a customized assessment of its current on-premises environment. The company wants to understand its projected running costs in the AWS Cloud.

Which AWS service or tool will meet these requirements?

- A. AWS Trusted Advisor
- B. Amazon Inspector
- C. AWS Control Tower
- D. Migration Evaluator

**Answer:** D

#### Explanation:

Migration Evaluator is an AWS service that provides a customized assessment of your current on-premises environment and helps you build a data-driven business case for migration to AWS. Migration Evaluator collects and analyzes data from your on-premises servers, such as CPU, memory, disk, network, and utilization metrics, and compares them with the most cost-effective AWS alternatives. Migration Evaluator also helps you understand your existing software licenses and running costs, and provides recommendations for Bring Your Own License (BYOL) and License Included (LI) options in AWS. Migration Evaluator generates a detailed report that shows your projected running costs in the AWS Cloud, along with potential savings and benefits. You can use this report to support your decision-making and planning for cloud migration. References: Cloud Business Case & Migration Plan - Amazon Migration Evaluator - AWS, Getting started with Migration Evaluator

#### NEW QUESTION 117

- (Topic 3)

Which option is the default pricing model for Amazon EC2 instances?

- A. On-Demand Instances
- B. Savings Plans
- C. Spot Instances
- D. Reserved Instances

**Answer:** A

**Explanation:**

On-Demand Instances are the default pricing model for Amazon EC2 instances. They allow users to pay for compute capacity by the second, with no long-term commitments or upfront payments. They are suitable for applications with short-term, irregular, or unpredictable workloads that cannot be interrupted<sup>3</sup>. Savings Plans are a pricing model that offer significant savings on Amazon EC2 and AWS Fargate usage, in exchange for a commitment to a consistent amount of usage (measured in \$/hour) for a 1- year or 3-year term. Spot Instances are a pricing model that offer spare Amazon EC2 compute capacity at up to 90% discount compared to On-Demand prices, but they can be interrupted by AWS with a two-minute notice when the demand exceeds the supply. Reserved Instances are a pricing model that offer up to 75% discount compared to On- Demand prices, in exchange for a commitment to use a specific instance type and size in a specific region for a 1-year or 3-year term.

**NEW QUESTION 118**

- (Topic 3)

Which of the following is a benefit of operating in the AWS Cloud?

- A. The ability to migrate on-premises network devices to the AWS Cloud
- B. The ability to expand compute, storage, and memory when needed
- C. The ability to host custom hardware in the AWS Cloud
- D. The ability to customize the underlying hypervisor layer for Amazon EC2

**Answer:** B

**Explanation:**

One of the benefits of operating in the AWS Cloud is the ability to expand compute, storage, and memory when needed, which enables users to scale their applications and resources up or down based on demand. This also helps users optimize their costs and performance. The ability to migrate on-premises network devices to the AWS Cloud, the ability to host custom hardware in the AWS Cloud, and the ability to customize the underlying hypervisor layer for Amazon EC2 are not benefits of operating in the AWS Cloud, as they are either not possible or not recommended by AWS .

**NEW QUESTION 120**

- (Topic 3)

A company wants to create a globally accessible ecommerce platform for its customers. The company wants to use a highly available and scalable DNS web service to connect users to the platform.

Which AWS service will meet these requirements?

- A. Amazon EC2
- B. Amazon VPC
- C. Amazon Route 53
- D. Amazon RDS

**Answer:** C

**Explanation:**

Amazon Route 53 is a highly available and scalable Domain Name System (DNS) web service that can route internet traffic to the company's ecommerce platform<sup>1</sup>. Route 53 can also register domain names, check the health of resources, and provide global DNS features<sup>2</sup>. Route 53 can connect users to the platform by translating human-readable names like `www.example.com` into the numeric IP addresses that computers use to communicate with each other<sup>2</sup>.  
References: 1: Amazon Route 53 | DNS Service | AWS; 2: What is Amazon Route 53? - Amazon Route 53

**NEW QUESTION 123**

- (Topic 3)

According to the AWS shared responsibility model, who is responsible for the virtualization layer down to the physical security of the facilities in which AWS services operate?

- A. It is the sole responsibility of the customer.
- B. It is the sole responsibility of AWS.
- C. It is a shared responsibility between AWS and the customer.
- D. The customer's AWS Support plan tier determines who manages the configuration.

**Answer:** B

**Explanation:**

According to the AWS shared responsibility model, AWS is responsible for the security of the cloud, which includes the virtualization layer down to the physical security of the facilities in which AWS services operate<sup>1</sup>. The customer is responsible for the security in the cloud, which includes the configuration and management of the AWS resources and applications that they use<sup>1</sup>.

**NEW QUESTION 125**

- (Topic 3)

Which AWS service or feature will search for and identify AWS resources that are shared externally?

- A. Amazon OpenSearch Service
- B. AWS Control Tower
- C. AWS IAM Access Analyzer
- D. AWS Fargate

**Answer:** C

**Explanation:**

AWS IAM Access Analyzer is an AWS service that helps customers identify and review the resources in their AWS account that are shared with an external entity, such as another AWS account, a root user, an organization, or a public entity. AWS IAM Access Analyzer uses automated reasoning, a form of mathematical logic and inference, to analyze the resource-based policies in the account and generate comprehensive findings that show the access level, the source of the access, the affected resource, and the condition under which the access applies. Customers can use AWS IAM Access Analyzer to audit their shared resources, validate their access policies, and monitor any changes to the resource sharing status. References: AWS IAM Access Analyzer, Identify and review resources shared with external entities, How AWS IAM Access Analyzer works

**NEW QUESTION 130**

- (Topic 3)

A company wants to use the latest technologies and wants to minimize its capital investment. Instead of upgrading on-premises infrastructure, the company wants to move to the AWS Cloud.

Which AWS Cloud benefit does this scenario describe?

- A. Increased speed to market
- B. The trade of infrastructure expenses for operating expenses
- C. Massive economies of scale
- D. The ability to go global in minutes

**Answer: B**

**Explanation:**

The trade of infrastructure expenses for operating expenses is one of the benefits of the AWS Cloud. By moving to the AWS Cloud, the company can avoid the upfront costs of purchasing and maintaining on-premises infrastructure, such as servers, storage, network, and software. Instead, the company can pay only for the AWS resources and services that they use, as they use them. This reduces the risk and complexity of planning and managing IT infrastructure, and allows the company to focus on innovation and growth. Increased speed to market, massive economies of scale, and the ability to go global in minutes are also benefits of the AWS Cloud, but they are not the best ones to describe this scenario. Increased speed to market means that the company can launch new products and services faster by using AWS services and tools. Massive economies of scale means that the company can benefit from the lower costs and higher performance that AWS achieves by operating at a large scale. The ability to go global in minutes means that the company can deploy their applications and data in multiple regions and availability zones around the world to reach their customers faster and improve performance and reliability.

**NEW QUESTION 134**

- (Topic 3)

A company website is experiencing DDoS attacks.

Which AWS service can help protect the company website against these attacks?

- A. AWS Resource Access Manager
- B. AWS Amplify
- C. AWS Shield
- D. Amazon GuardDuty

**Answer: C**

**Explanation:**

AWS Shield is a managed DDoS protection service that safeguards applications running on AWS from distributed denial of service (DDoS) attacks. DDoS attacks are malicious attempts to disrupt the normal functioning of a website or application by overwhelming it with a large volume of traffic from multiple sources. AWS Shield provides two tiers of protection: Standard and Advanced. AWS Shield Standard is automatically enabled for all AWS customers at no additional cost. It protects your AWS resources, such as Amazon CloudFront, AWS Global Accelerator, and Amazon Route 53, from the most common and frequently occurring network and transport layer DDoS attacks. AWS Shield Advanced is an optional paid service that provides additional protection for your AWS resources and applications, such as Amazon Elastic Compute Cloud (Amazon EC2), Elastic Load Balancing (ELB), Amazon Simple Storage Service (Amazon S3), Amazon Relational Database Service (Amazon RDS), and AWS Elastic Beanstalk. AWS Shield Advanced offers enhanced detection and mitigation capabilities, 24/7 access to the AWS DDoS Response Team (DRT), real-time visibility and reporting, and cost protection against DDoS-related spikes in your AWS bill. References: AWS Shield, What is a DDoS Attack & How to Protect Your Site Against One

**NEW QUESTION 136**

- (Topic 3)

AWS has the ability to achieve lower pay-as-you-go pricing by aggregating usage across hundreds of thousands of users.

This describes which advantage of the AWS Cloud?

- A. Launch globally in minutes
- B. Increase speed and agility
- C. High economies of scale
- D. No guessing about compute capacity

**Answer: C**

**Explanation:**

AWS has the ability to achieve lower pay-as-you-go pricing by aggregating usage across hundreds of thousands of users. This means that AWS can leverage its massive scale and purchasing power to reduce the costs of infrastructure, hardware, software, and operations. These savings are then passed on to the customers, who only pay for the resources they use. You can learn more about the AWS pricing model from [this webpage] or [this digital course].

**NEW QUESTION 137**

- (Topic 3)

A company wants to create a set of custom dashboards to collect metrics to monitor its applications.

Which AWS service will meet these requirements?

- A. Amazon CloudWatch
- B. AWS X-Ray
- C. AWS Systems Manager

D. AWS CloudTrail

**Answer:** A

**Explanation:**

Amazon CloudWatch is a service that provides monitoring and observability for AWS resources and applications. Users can create custom dashboards to collect and visualize metrics, logs, alarms, and events from different sources<sup>5</sup>. AWS X-Ray is a service that provides distributed tracing and analysis for applications. AWS Systems Manager is a service that provides operational management for AWS resources and applications. AWS CloudTrail is a service that provides governance, compliance, and auditing for AWS account activity.

**NEW QUESTION 139**

- (Topic 3)

A company is moving an on-premises data center to the AWS Cloud. The company must migrate 50 petabytes of file storage data to AWS with the least possible operational overhead.

Which AWS service or resource should the company use to meet these requirements?

- A. AWS Snowmobile
- B. AWS Snowball Edge
- C. AWS Data Exchange
- D. AWS Database Migration Service (AWS DMS)

**Answer:** A

**Explanation:**

The AWS service that the company should use to meet these requirements is A. AWS Snowmobile.

AWS Snowmobile is a service that allows you to migrate large amounts of data to AWS using a 45-foot long ruggedized shipping container that can store up to 100 petabytes of data. AWS Snowmobile is designed for situations where you need to move massive amounts of data to the cloud in a fast, secure, and cost-effective way. AWS Snowmobile has the least possible operational overhead because it eliminates the need to buy, configure, or manage hundreds or thousands of storage devices<sup>12</sup>.

AWS Snowball Edge is a service that allows you to migrate data to AWS using a physical device that can store up to 80 terabytes of data and has compute and storage capabilities to run applications on the device. AWS Snowball Edge is suitable for situations where you have limited or intermittent network connectivity, or where bandwidth costs are high. However, AWS Snowball Edge has more operational overhead than AWS Snowmobile because you need to request multiple devices and transfer your data onto them using the client<sup>3</sup>.

AWS Data Exchange is a service that allows you to find, subscribe to, and use third-party data in the cloud. AWS Data Exchange is not a data migration service, but rather a data marketplace that enables data providers and data consumers to exchange data sets securely and efficiently<sup>4</sup>.

AWS Database Migration Service (AWS DMS) is a service that helps migrate databases to AWS. AWS DMS does not migrate file storage data, but rather supports various database platforms and engines as sources and targets<sup>5</sup>.

References:

1: AWS Snowmobile – Move Exabytes of Data to the Cloud in Weeks 2: AWS Snowmobile

- Amazon Web Services 3: Automated Software Vulnerability Management - Amazon Inspector - AWS 4: AWS Data Exchange - Find, subscribe to, and use third-party data in ... 5: AWS Database Migration Service – Amazon Web Services

**NEW QUESTION 141**

- (Topic 3)

A company wants to store data with high availability, encrypt the data at rest, and have direct access to the data over the internet.

Which AWS service will meet these requirements MOST cost-effectively?

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

**Answer:** C

**Explanation:**

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 offers two storage classes: the Standard storage class, and the Infrequent Access storage class (EFS IA).

EFS IA provides price/performance that is cost-optimized for files not accessed every day. Amazon EFS encrypts data at rest and in transit, and supports direct access over the internet<sup>4</sup>.

**NEW QUESTION 144**

- (Topic 3)

A company seeks cost savings in exchange for a commitment to use a specific amount of an AWS service or category of AWS services for 1 year or 3 years.

Which AWS pricing model or offering will meet these requirements?

- A. Pay-as-you-go pricing
- B. Savings Plans
- C. AWS Free Tier
- D. Volume discounts

**Answer:** B

**Explanation:**

Savings Plans are an AWS pricing model or offering that can meet the requirements of seeking cost savings in exchange for a commitment to use a specific amount of an AWS service or category of AWS services for 1 year or 3 years. Savings Plans are flexible plans that offer significant discounts on AWS compute usage, such as EC2, Lambda, and Fargate. The company can choose from two types of Savings Plans: Compute Savings Plans and EC2 Instance Savings Plans. Compute Savings Plans provide the most flexibility and apply to any eligible compute usage, regardless of instance family, size, region, operating system, or tenancy. EC2 Instance Savings Plans provide more savings and apply to a specific instance family within a region. The company can select the amount of

compute usage per hour (e.g., \$10/hour) that they want to commit to for the duration of the plan (1 year or 3 years). The company will pay the discounted Savings Plan rate for the amount of usage that matches their commitment, and the regular on-demand rate for any usage beyond that

#### NEW QUESTION 148

- (Topic 3)

Which scenarios represent the concept of elasticity on AWS? (Select TWO.)

- A. Scaling the number of Amazon EC2 instances based on traffic
- B. Resizing Amazon RDS instances as business needs change
- C. Automatically directing traffic to less-utilized Amazon EC2 instances
- D. Using AWS compliance documents to accelerate the compliance process
- E. Having the ability to create and govern environments using code

**Answer:** AB

#### Explanation:

These are two scenarios that represent the concept of elasticity on AWS. Elasticity means the ability to adjust the resources and capacity of the system in response to changes in demand or environment. Scaling the number of Amazon EC2 instances based on traffic means using services such as AWS Auto Scaling or Elastic Load Balancing to add or remove instances as the traffic increases or decreases. Resizing Amazon RDS instances as business needs change means using the Amazon RDS console or API to modify the instance type, storage type, or storage size of the database as the workload grows or shrinks. You can learn more about the concept of elasticity on AWS from [this webpage] or [this digital course].

#### NEW QUESTION 153

- (Topic 3)

According to security best practices, how should an Amazon EC2 instance be given access to an Amazon S3 bucket?

- A. Hard code an IAM user's secret key and access key directly in the application, and upload the file.
- B. Store the IAM user's secret key and access key in a text file on the EC2 instance, read the keys, then upload the file.
- C. Have the EC2 instance assume a role to obtain the privileges to upload the file.
- D. Modify the S3 bucket policy so that any service can upload to it at any time.

**Answer:** C

#### Explanation:

According to security best practices, the best way to give an Amazon EC2 instance access to an Amazon S3 bucket is to have the EC2 instance assume a role to obtain the privileges to upload the file. A role is an AWS Identity and Access Management (IAM) entity that defines a set of permissions for making AWS service requests. You can use roles to delegate access to users, applications, or services that don't normally have access to your AWS resources. For example, you can create a role that allows EC2 instances to access S3 buckets, and then attach the role to the EC2 instance. This way, the EC2 instance can assume the role and obtain temporary security credentials to access the S3 bucket. This method is more secure and scalable than storing or hardcoding IAM user credentials on the EC2 instance, as it avoids the risk of exposing or compromising the credentials. It also allows you to manage the permissions centrally and dynamically, and to audit the access using AWS CloudTrail. For more information on how to create and use roles for EC2 instances, see [Using an IAM role to grant permissions to applications running on Amazon EC2 instances](#)<sup>1</sup>

The other options are not recommended for security reasons. Hardcoding or storing IAM user credentials on the EC2 instance is a bad practice, as it exposes the credentials to potential attackers or unauthorized users who can access the instance or the application code. It also makes it difficult to rotate or revoke the credentials, and to track the usage of the credentials. Modifying the S3 bucket policy to allow any service to upload to it at any time is also a bad practice, as it opens the bucket to potential data breaches, data loss, or data corruption. It also violates the principle of least privilege, which states that you should grant only the minimum permissions necessary for a task.

References: [Using an IAM role to grant permissions to applications running on Amazon EC2 instances](#)

#### NEW QUESTION 156

- (Topic 3)

Which Amazon S3 storage class is the MOST cost-effective for long-term storage?

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

**Answer:** A

#### Explanation:

Amazon S3 Glacier Deep Archive is the lowest-cost storage class in the cloud. It is designed for long-term data archiving that is rarely accessed. It offers a retrieval time of 12 hours and a durability of 99.999999999% (11 9's). It is ideal for data that must be retained for 7 years or longer to meet regulatory compliance requirements.

#### NEW QUESTION 159

- (Topic 3)

A company needs to apply security rules to specific Amazon EC2 instances. Which AWS service or feature provides this functionality?

- A. AWS Shield
- B. Network ACLs
- C. Security groups
- D. AWS Firewall Manager

**Answer:** C

#### Explanation:

Security groups act as a firewall for associated Amazon EC2 instances, controlling both inbound and outbound traffic at the instance level. You can use security groups to set rules that allow or deny traffic to or from your instances. You can modify the rules for a security group at any time; the new rules are automatically

applied to all instances that are associated with the security group.

#### NEW QUESTION 164

- (Topic 3)

Which option is a perspective that includes foundational capabilities of the AWS Cloud Adoption Framework (AWS CAF)?

- A. Sustainability
- B. Security
- C. Performance efficiency
- D. Reliability

**Answer: B**

#### Explanation:

The AWS Cloud Adoption Framework (AWS CAF) helps organizations understand how cloud adoption transforms the way they work, and it provides structure to identify and address gaps in skills and processes. The AWS CAF organizes guidance into six areas of focus, called perspectives. Each perspective reflects a different stakeholder viewpoint with its own distinct responsibilities, skills, and attributes. The Security Perspective helps you structure the selection and implementation of security controls that meet your organization's needs<sup>2</sup>.

#### NEW QUESTION 166

- (Topic 3)

Which AWS service provides protection against DDoS attacks for applications that run in the AWS Cloud?

- A. Amazon VPC
- B. AWS Shield
- C. AWS Audit Manager
- D. AWS Config

**Answer: B**

#### Explanation:

AWS Shield is an AWS service that provides protection against distributed denial of service (DDoS) attacks for applications that run in the AWS Cloud. DDoS attacks are attempts to make an online service unavailable by overwhelming it with traffic from multiple sources. AWS Shield provides two tiers of protection: AWS Shield Standard and AWS Shield Advanced. AWS Shield Standard is automatically enabled for all AWS customers at no additional charge. It provides protection against common and frequently occurring network and transport layer DDoS attacks. AWS Shield Advanced is an optional paid service that provides additional protection against larger and more sophisticated DDoS attacks. AWS Shield Advanced also provides access to 24/7 DDoS response team, cost protection, and enhanced detection and mitigation capabilities

#### NEW QUESTION 167

- (Topic 3)

A company is migrating its workloads to the AWS Cloud. The company must retain full control of patch management for the guest operating systems that host its applications.

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

- A. Amazon DynamoDB
- B. Amazon EC2
- C. AWS Lambda
- D. Amazon RDS

**Answer: B**

#### Explanation:

Amazon EC2 is the AWS service that the company should use to meet its requirements of retaining full control of patch management for the guest operating systems that host its applications. Amazon EC2 is a service that provides secure, resizable compute capacity in the cloud. Users can launch virtual servers, called instances, that run various operating systems, such as Linux, Windows, macOS, and more. Users have full administrative access to their instances and can install and configure any software, including patches and updates, on their instances. Users are responsible for managing the security and maintenance of their instances, including patching the guest operating system and applications. Users can also use AWS Systems Manager to automate and simplify the patching process for their EC2 instances. AWS Systems Manager is a service that helps users manage their AWS and on-premises resources at scale. Users can use AWS Systems Manager Patch Manager to scan their instances for missing patches, define patch baselines and maintenance windows, and apply patches automatically or manually across their instances. Users can also use AWS Systems Manager to monitor the patch compliance status and patching history of their instances. References: What is Amazon EC2?, AWS Systems Manager Patch Manager

#### NEW QUESTION 171

- (Topic 3)

A company uses AWS for its web application. The company wants to minimize latency and perform compute operations for the application as close to end users as possible.

Which AWS service or infrastructure component will provide this functionality?

- A. AWS Regions
- B. Availability Zones
- C. Edge locations
- D. AWS Direct Connect

**Answer: C**

#### Explanation:

Edge locations are sites that Amazon CloudFront uses to cache copies of your content for faster delivery to users at any location. You can use Amazon CloudFront to deliver your entire website, including dynamic, static, streaming, and interactive content using a global network of edge locations. Requests for your content are automatically routed to the nearest edge location, so content is delivered with the best possible performance<sup>3</sup>. Edge locations can also host AWS

Lambda functions to perform compute operations for your web application as close to end users as possible4.

#### NEW QUESTION 173

- (Topic 3)

A company wants to build a new web application by using AWS services. The application must meet the on-demand load for periods of heavy activity. Which AWS services or resources provide the necessary workload adjustments to meet these requirements? (Select TWO.)

- A. Amazon Machine Image (AMI)
- B. Amazon EC2 Auto Scaling
- C. Amazon EC2 instance
- D. AWS Lambda
- E. EC2 Image Builder

**Answer:** BD

#### Explanation:

Amazon EC2 Auto Scaling helps you ensure that you have the correct number of Amazon EC2 instances available to handle the load for your application. You create collections of EC2 instances, called Auto Scaling groups. You can specify the minimum number of instances in each Auto Scaling group, and Amazon EC2 Auto Scaling ensures that your group never goes below this size. You can specify the maximum number of instances in each Auto Scaling group, and Amazon EC2 Auto Scaling ensures that your group never goes above this size4. AWS Lambda lets you run code without provisioning or managing servers. You pay only for the compute time you consume. 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.

#### NEW QUESTION 175

- (Topic 3)

A company wants to grant users in one AWS account access to resources in another AWS account. The users do not currently have permission to access the resources.

Which AWS service will meet this requirement?

- A. IAM group
- B. IAM role
- C. IAM tag
- D. IAM Access Analyzer

**Answer:** B

#### Explanation:

IAM roles are a way to delegate access to resources in different AWS accounts. IAM roles allow users to assume a set of permissions for a limited time without having to create or share long-term credentials. IAM roles can be used to grant cross-account access by creating a trust relationship between the accounts and specifying the permissions that the role can perform. Users can then switch to the role and access the resources in the other account using temporary security credentials provided by the role. References: Cross account resource access in IAM, IAM tutorial: Delegate access across AWS accounts using IAM roles, How to Enable Cross-Account Access to the AWS Management Console

#### NEW QUESTION 178

- (Topic 3)

A company needs to securely store important credentials that an application uses to connect users to a database.

Which AWS service can meet this requirement with the MINIMAL amount of operational overhead?

- A. AWS Key Management Service (AWS KMS)
- B. AWS Config
- C. AWS Secrets Manager
- D. Amazon GuardDuty

**Answer:** C

#### Explanation:

AWS Secrets Manager is a service that helps you protect secrets needed to access your applications, services, and IT resources. You can use AWS Secrets Manager to store, rotate, and retrieve database credentials, API keys, and other secrets throughout their lifecycle. AWS Secrets Manager eliminates the need to hardcode sensitive information in plain text, and reduces the risk of unauthorized access or leakage. AWS Secrets Manager also integrates with other AWS services, such as AWS Lambda, Amazon RDS, and AWS CloudFormation, to simplify the management of secrets across your environment5

#### NEW QUESTION 182

- (Topic 3)

Which AWS Cloud deployment model uses AWS Outposts as part of the application deployment infrastructure?

- A. On-premises
- B. Serverless
- C. Cloud-native
- D. Hybrid

**Answer:** D

#### Explanation:

AWS Outposts is a fully managed service that extends AWS infrastructure, services, APIs, and tools to customer premises. By providing local access to AWS managed infrastructure, AWS Outposts enables customers to build and run applications on premises using the same programming interfaces as in AWS Regions, while using local compute and storage resources for lower latency and local data processing needs. An Outpost is a pool of AWS compute and storage capacity deployed at a customer site. AWS operates, monitors, and manages this capacity as part of an AWS Region. You can create subnets on your Outpost and specify

them when you create AWS resources such as EC2 instances, EBS volumes, ECS clusters, and RDS instances. Instances in Outpost subnets communicate with other instances in the AWS Region using private IP addresses, all within the same VPC. Outposts solutions allow you to extend and run native AWS services on premises, and is available in a variety of form factors, from 1U and 2U Outposts servers to 42U Outposts racks, and multiple rack deployments. With AWS Outposts, you can run some AWS services locally and connect to a broad range of services available in the local AWS Region<sup>2</sup>. AWS Outposts is a hybrid cloud deployment model that uses AWS Outposts as part of the application deployment infrastructure. Hybrid cloud is a cloud computing environment that uses a mix of on-premises, private cloud, and public cloud services with orchestration between the platforms. Hybrid cloud provides businesses with greater flexibility, more deployment options, and optimized costs. By using AWS Outposts, customers can benefit from the fully managed infrastructure, services, APIs, and tools of AWS on premises, while still having access to the full range of AWS services available in the Region for a truly consistent hybrid experience<sup>3</sup>. References: On-Premises Private Cloud - AWS Outposts Family - AWS, What is AWS Outposts? - AWS Outposts

#### NEW QUESTION 183

- (Topic 3)

A company deployed an Amazon EC2 instance last week. A developer realizes that the EC2 instance is no longer running. The developer reviews a list of provisioned EC2 instances, and the EC2 instance is no longer on the list.

What can the developer do to generate a recent history of the EC2 instance?

- A. Run Cost Explorer to identify the start time and end time of the EC2 instance.
- B. Use Amazon Inspector to find out when the EC2 instance was stopped.
- C. Perform a search in AWS CloudTrail to find all EC2 instance-related events.
- D. Use AWS Secrets Manager to display hidden termination logs of the EC2 instance.

**Answer:** C

#### Explanation:

AWS CloudTrail is a service that enables governance, compliance, operational auditing, and risk auditing of a customer's AWS account. AWS CloudTrail allows customers to track user activity and API usage across their AWS infrastructure. AWS CloudTrail can also provide a history of EC2 instance events, such as launch, stop, terminate, and reboot. Cost Explorer is a tool that enables customers to visualize, understand, and manage their AWS costs and usage over time. Amazon Inspector is an automated security assessment service that helps improve the security and compliance of applications deployed on AWS. AWS Secrets Manager helps customers protect secrets needed to access their applications, services, and IT resources.

#### NEW QUESTION 185

- (Topic 2)

A company wants to migrate its application to AWS. The company wants to replace upfront expenses with variable payment that is based on usage.

What should the company do to meet these requirements?

- A. Use pay-as-you-go pricing.
- B. Purchase Reserved Instances.
- C. Pay less by using more.
- D. Rightsize instances.

**Answer:** A

#### Explanation:

Pay-as-you-go pricing is one of the main benefits of AWS. With pay-as-you-go pricing, you pay only for what you use, when you use it. There are no long-term contracts, termination fees, or complex licensing. You replace upfront expenses with lower variable costs and pay only for the resources you consume.

#### NEW QUESTION 189

- (Topic 2)

A company needs help managing multiple AWS linked accounts that are reported on a consolidated bill.

Which AWS Support plan includes an AWS concierge whom the company can ask for assistance?

- A. AWS Developer Support
- B. AWS Enterprise Support
- C. AWS Business Support
- D. AWS Basic Support

**Answer:** B

#### Explanation:

AWS Enterprise Support is the AWS Support plan that includes an AWS concierge whom the company can ask for assistance. According to the AWS Support Plans page, AWS Enterprise Support provides "a dedicated Technical Account Manager (TAM) who provides advocacy and guidance to help plan and build solutions using best practices, coordinate access to subject matter experts, and proactively keep your AWS environment operationally healthy."<sup>2</sup> AWS Business Support, AWS Developer Support, and AWS Basic Support do not include a TAM or a concierge service.

#### NEW QUESTION 190

- (Topic 2)

A company wants to move its iOS application development and build activities to AWS. Which AWS service or resource should the company use for these activities?

- A. AWS CodeCommit
- B. Amazon EC2 M1 Mac instances
- C. AWS Amplify
- D. AWS App Runner

**Answer:** B

#### Explanation:

Amazon EC2 M1 Mac instances are the AWS service or resource that the company should use for its iOS application development and build activities, as they enable users to run macOS on AWS and access a broad and growing set of AWS services. AWS CodeCommit is a service that provides a fully managed source

control service that hosts secure Git-based repositories. AWS Amplify is a set of tools and services that enable developers to build full-stack web and mobile applications using AWS. AWS App Runner is a service that makes it easy for developers to quickly deploy containerized web applications and APIs. These concepts are explained in the AWS Developer Tools page4.

#### NEW QUESTION 191

- (Topic 2)

A company has multiple AWS accounts that include compute workloads that cannot be interrupted. The company wants to obtain billing discounts that are based on the company's use of AWS services.

Which AWS feature or purchasing option will meet these requirements?

- A. Resource tagging
- B. Consolidated billing
- C. Pay-as-you-go pricing
- D. Spot Instances

**Answer: B**

#### Explanation:

Consolidated billing is an AWS feature that allows users to combine the usage and costs of multiple AWS accounts into a single bill. This enables users to obtain billing discounts that are based on the company's use of AWS services, such as volume pricing tiers, Reserved Instance discounts, and Savings Plans discounts5. Resource tagging is an AWS feature that allows users to assign metadata to AWS resources, such as EC2 instances, S3 buckets, and Lambda functions. This enables users to organize, track, and manage their AWS resources, such as filtering, grouping, and reporting. Pay-as-you-go pricing is an AWS pricing model that allows users to pay only for the resources and services they use, without any upfront or long-term commitments. This enables users to lower their costs by scaling up or down as needed, and avoiding over-provisioning or under-utilization. Spot Instances are spare EC2 instances that are available at up to 90% discount compared to On-Demand prices. They are suitable for workloads that can tolerate interruptions, such as batch processing, data analysis, and testing. Spot Instances are allocated based on the current supply and demand, and can be reclaimed by AWS with a two-minute notice when the demand exceeds the supply.

#### NEW QUESTION 196

- (Topic 2)

Which AWS service is used to temporarily provide federated security credentials to a

- A. Amazon GuardDuty
- B. AWS Simple Token Service (AWS STS)
- C. AWS Secrets Manager
- D. AWS Certificate Manager

**Answer: B**

#### Explanation:

The AWS service that is used to temporarily provide federated security credentials to a user is AWS Security Token Service (AWS STS). AWS STS is a service that enables customers to request temporary, limited-privilege credentials for AWS Identity and Access Management (IAM) users or for users that they authenticate (federated users). The company can use AWS STS to grant federated users access to AWS resources without creating permanent IAM users or sharing long-term credentials. AWS STS helps customers manage and secure access to their AWS resources for federated users. Amazon GuardDuty, AWS Secrets Manager, and AWS Certificate Manager are not the best services to use for this purpose. Amazon GuardDuty is a threat detection service that monitors for malicious activity and unauthorized behavior across the AWS accounts and resources. AWS Secrets Manager is a service that helps customers manage and rotate secrets, such as database credentials, API keys, and passwords. AWS Certificate Manager is a service that helps customers provision, manage, and deploy public and private Secure Sockets Layer/Transport Layer Security (SSL/TLS) certificates for use with AWS services and internal connected resources. These services are more useful for different types of security and compliance tasks, rather than providing temporary federated security credentials to a user.

#### NEW QUESTION 198

- (Topic 2)

A company has a compliance requirement to record and evaluate configuration changes, as well as perform remediation actions on AWS resources.

Which AWS service should the company use?

- A. AWS Config
- B. AWS Secrets Manager
- C. AWS CloudTrail
- D. AWS Trusted Advisor

**Answer: A**

#### Explanation:

AWS Config is a service that enables you to assess, audit, and evaluate the configurations of your AWS resources. AWS Config continuously monitors and records your AWS resource configurations and allows you to automate the evaluation of recorded configurations against desired configurations. With AWS Config, you can review changes in configurations and relationships between AWS resources, dive into detailed resource configuration histories, and determine your overall compliance against the configurations specified in your internal guidelines. This can help you simplify compliance auditing, security analysis, change management, and operational troubleshooting1.

#### NEW QUESTION 200

- (Topic 2)

A company does not want to rely on elaborate forecasting to determine its usage of compute resources. Instead, the company wants to pay only for the resources that it uses. The company also needs the ability to increase or decrease its resource usage to meet business requirements.

Which pillar of the AWS Well-Architected Framework aligns with these requirements?

- A. Operational excellence
- B. Security
- C. Reliability
- D. Cost optimization

**Answer:** D

**Explanation:**

Cost optimization is the pillar of the AWS Well-Architected Framework that aligns with the requirements of not relying on elaborate forecasting and paying only for the resources that are used. The cost optimization pillar focuses on the ability of a system to deliver business value at the lowest price point. Cost optimization involves using the right AWS services and resources for the workload, measuring and monitoring the cost and usage, and continuously improving the cost efficiency. Cost optimization also leverages the benefits of the AWS Cloud, such as pay-as-you-go pricing, elasticity, and scalability. For more information, see [Cost Optimization Pillar] and [Cost Optimization].

**NEW QUESTION 203**

- (Topic 2)

An application runs on multiple Amazon EC2 instances that access a shared file system simultaneously. Which AWS storage service should be used?

- A. Amazon EBS
- B. Amazon EFS
- C. Amazon S3
- D. AWS Artifact

**Answer:** B

**Explanation:**

Amazon Elastic File System (Amazon EFS) is the AWS storage service that should be used for an application that runs on multiple Amazon EC2 instances that access a shared file system simultaneously. Amazon EFS is a fully managed service that provides a scalable, elastic, and highly available file system for Linux-based workloads. Amazon EFS supports the Network File System version 4 (NFSv4) protocol and allows multiple EC2 instances to read and write data to the same file system concurrently. Amazon EFS also integrates with other AWS services, such as AWS Backup, AWS CloudFormation, and AWS CloudTrail. For more information, see What is Amazon Elastic File System? and [Amazon EFS Use Cases].

**NEW QUESTION 205**

- (Topic 2)

A company is running an application on AWS. The company wants to identify and prevent the accidental Which AWS service or feature will meet these requirements?

- A. Amazon GuardDuty
- B. Network ACL
- C. AWS WAF
- D. AWS Network Firewall

**Answer:** A

**Explanation:**

Amazon GuardDuty is a threat detection service that continuously monitors for malicious activity and unauthorized behavior to protect your AWS accounts, workloads, and data stored in Amazon S3. With the cloud, the collection and aggregation of account and network activities is simplified, but it can be time consuming for security teams to continuously analyze event log data for potential threats. With GuardDuty, you can automate anomaly detection and get actionable findings to help you protect your AWS resources.

**NEW QUESTION 207**

- (Topic 2)

A company wants its workload to perform consistently and correctly. Which benefit of AWS Cloud computing does this goal represent?

- A. Security
- B. Elasticity
- C. Pay-as-you-go pricing
- D. Reliability

**Answer:** D

**Explanation:**

Reliability is the benefit of AWS Cloud computing that ensures the workload performs consistently and correctly. According to the AWS Cloud Practitioner Essentials course, reliability means "the ability of a system to recover from infrastructure or service disruptions, dynamically acquire computing resources to meet demand, and mitigate disruptions such as misconfigurations or transient network issues."<sup>1</sup> Elasticity, security, and pay-as-you-go pricing are also benefits of AWS Cloud computing, but they do not directly relate to the goal of consistent and correct performance.

**NEW QUESTION 212**

- (Topic 2)

Which service is an AWS in-memory data store service?

- A. Amazon Aurora
- B. Amazon RDS
- C. Amazon DynamoDB
- D. Amazon ElastiCache

**Answer:** D

**Explanation:**

Amazon ElastiCache is a service that offers fully managed in-memory data store and cache services that deliver sub-millisecond response times to applications. You can use Amazon ElastiCache to improve the performance of your applications by retrieving data from fast, managed, in-memory data stores, instead of relying entirely on slower disk-based databases. Amazon Aurora is a relational database service that combines the performance and availability of high-end commercial

databases with the simplicity and cost-effectiveness of open source databases. Amazon RDS is a service that makes it easy to set up, operate, and scale a relational database in the cloud. Amazon DynamoDB is a key-value and document database that delivers single-digit millisecond performance at any scale. None of these services are in-memory data store services.

#### NEW QUESTION 217

- (Topic 2)

Which benefit of AWS Cloud computing provides lower latency between users and applications?

- A. Agility
- B. Economies of scale
- C. Global reach
- D. Pay-as-you-go pricing

**Answer: C**

#### Explanation:

Global reach is the benefit of AWS Cloud computing that provides lower latency between users and applications. Global reach means that AWS customers can deploy their applications and data in multiple regions around the world, and deliver them to users with high performance and availability. AWS has the largest global infrastructure of any cloud provider, with 25 geographic regions and 81 Availability Zones, as well as 216 Points of Presence in 84 cities across 42 countries. Customers can choose the optimal locations for their applications and data based on their business requirements, such as compliance, data sovereignty, and customer proximity. Agility, economies of scale, and pay-as-you-go pricing are other benefits of AWS Cloud computing, but they do not directly provide lower latency between users and applications. Agility means that AWS customers can quickly and easily provision and scale up or down AWS resources as needed, without upfront costs or long-term commitments. Economies of scale means that AWS customers can benefit from the lower costs and higher efficiency that AWS achieves by operating at a massive scale and passing the savings to the customers. Pay-as-you-go pricing means that AWS customers only pay for the AWS resources they use, without any upfront costs or long-term contracts.

#### NEW QUESTION 220

- (Topic 2)

A company is planning a migration to the AWS Cloud and wants to examine the costs that are associated with different workloads. Which AWS tool will meet these requirements?

- A. AWS Budgets
- B. AWS Cost Explorer
- C. AWS Pricing Calculator
- D. AWS Cost and Usage Report

**Answer: C**

#### Explanation:

The AWS tool that will meet the requirements of the company that is planning a migration to the AWS Cloud and wants to examine the costs that are associated with different workloads is AWS Pricing Calculator. AWS Pricing Calculator is a tool that helps customers estimate the cost of using AWS services based on their requirements and preferences. The company can use AWS Pricing Calculator to compare the costs of different AWS services and configurations, such as Amazon EC2, Amazon S3, Amazon RDS, and more. AWS Pricing Calculator also provides detailed breakdowns of the cost components, such as compute, storage, network, and data transfer. AWS Pricing Calculator helps customers plan and optimize their cloud budget and migration strategy. AWS Budgets, AWS Cost Explorer, and AWS Cost and Usage Report are not the best tools to use for this purpose. AWS Budgets is a tool that helps customers monitor and manage their AWS spending and usage against predefined budget limits and thresholds. AWS Cost Explorer is a tool that helps customers analyze and visualize their AWS spending and usage trends over time. AWS Cost and Usage Report is a tool that helps customers access comprehensive and granular information about their AWS costs and usage in a CSV or Parquet file. These tools are more useful for tracking and optimizing the existing AWS costs and usage, rather than estimating the costs of different workloads<sup>34</sup>

#### NEW QUESTION 224

- (Topic 2)

A company has developed a distributed application that recovers gracefully from interruptions. The application periodically processes large volumes of data by using multiple Amazon EC2 instances. The application is sometimes idle for months.

Which EC2 instance purchasing option is MOST cost-effective for this use case?

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

**Answer: B**

#### Explanation:

Spot Instances are instances that use spare EC2 capacity that is available for up to 90% off the On-Demand price. Because Spot Instances can be interrupted by EC2 with two minutes of notification when EC2 needs the capacity back, you can use them for applications that have flexible start and end times, or that can withstand interruptions<sup>5</sup>. This option is most cost-effective for the use case described in the question. Reserved Instances are instances that you purchase for a one-year or three-year term, and pay a lower hourly rate compared to On-Demand Instances. This option is suitable for applications that have steady state or predictable usage. Dedicated Instances are instances that run on hardware that's dedicated to a single customer within an Amazon VPC. This option is suitable for applications that have stringent regulatory or compliance requirements. On-Demand Instances are instances that you pay for by the second, with no long-term commitments or upfront payments. This option is suitable for applications that have unpredictable or intermittent workloads.

#### NEW QUESTION 226

- (Topic 2)

A company wants to improve its security and audit posture by limiting Amazon EC2 inbound access. According to the AWS shared responsibility model, which task is the responsibility of the customer?

- A. Protect the global infrastructure that runs all of the services offered in the AWS Cloud.

- B. Configure logical access controls for resources, and protect account credentials.
- C. Configure the security used by managed services.
- D. Patch and back up Amazon Aurora.

**Answer:** B

**Explanation:**

According to the AWS shared responsibility model, the customer is responsible for configuring logical access controls for resources, and protecting account credentials. This includes managing IAM user permissions, security group rules, network ACLs, encryption keys, and other aspects of access management<sup>1</sup>. AWS is responsible for protecting the global infrastructure that runs all of the services offered in the AWS Cloud, such as the hardware, software, networking, and facilities. AWS is also responsible for configuring the security used by managed services, such as Amazon RDS, Amazon DynamoDB, and Amazon Aurora<sup>2</sup>.

**NEW QUESTION 229**

- (Topic 2)

Which of the following is the customer's responsibility, according to the AWS shared responsibility model?

- A. Identity and access management
- B. Hard drive initialization
- C. Protection of data center hardware
- D. Security of Availability Zones

**Answer:** A

**Explanation:**

Identity and access management is the customer's responsibility, according to the AWS shared responsibility model. This means that the customer is responsible for managing user access to the AWS resources, using tools such as AWS Identity and Access Management (IAM), AWS Single Sign-On (SSO), and AWS Organizations. The customer is also responsible for securing their data in transit and at rest, using encryption, key management, and other methods. Hard drive initialization, protection of data center hardware, and security of Availability Zones are AWS's responsibility, as they are part of the infrastructure, physical security, and network security that AWS provides to the customer<sup>12</sup>

**NEW QUESTION 231**

- (Topic 2)

A company wants guidance to optimize the cost and performance of its current AWS environment.

Which AWS service or tool should the company use to identify areas for optimization?

- A. Amazon QuickSight
- B. AWS Trusted Advisor
- C. AWS Organizations
- D. AWS Budgets

**Answer:** B

**Explanation:**

AWS Trusted Advisor is the AWS service or tool that the company should use to identify areas for optimization. According to the AWS Trusted Advisor User Guide, "AWS Trusted Advisor is an online tool that provides you real time guidance to help you provision your resources following AWS best practices. AWS Trusted Advisor checks help optimize your AWS infrastructure, increase security and performance, reduce your overall costs, and monitor service limits." Amazon QuickSight, AWS Organizations, and AWS Budgets are not designed to provide optimization recommendations for the current AWS environment.

**NEW QUESTION 232**

- (Topic 2)

A company has set up a VPC in its AWS account and has created a subnet in the VPC. The company wants to make the subnet public.

Which AWS features should the company use to meet this requirement? (Select TWO.)

- A. Amazon VPC internet gateway
- B. Amazon VPC NAT gateway
- C. Amazon VPC route tables
- D. Amazon VPC network ACL
- E. Amazon EC2 security groups

**Answer:** AC

**Explanation:**

To make a subnet public, the company should use an Amazon VPC internet gateway and an Amazon VPC route table. An internet gateway is a horizontally scaled, redundant, and highly available VPC component that allows communication between your VPC and the internet. A route table contains a set of rules, called routes, that are used to determine where network traffic from your subnet or gateway is directed. To enable internet access for a subnet, you need to attach an internet gateway to your VPC and add a route to the internet gateway in the route table associated with the subnet.

**NEW QUESTION 234**

- (Topic 2)

A company needs to host a web server on Amazon EC2 instances for at least 1 year. The web server cannot tolerate interruption.

Which EC2 instance purchasing option will meet these requirements MOST cost- effectively?

- A. On-Demand Instances
- B. Partial Upfront Reserved Instances
- C. Spot Instances
- D. No Upfront Reserved Instances

**Answer:** B

**Explanation:**

The most cost-effective EC2 instance purchasing option for the company that needs to host a web server on Amazon EC2 instances for at least 1 year and cannot tolerate interruption is Partial Upfront Reserved Instances. Reserved Instances are a pricing model that offer significant discounts compared to On-Demand Instances in exchange for a commitment to use a specific amount of compute capacity for a fixed period of time (1 or 3 years). Partial Upfront Reserved Instances require customers to pay a portion of the total cost upfront, and the remaining cost in monthly installments over the term. This option offers a lower effective hourly rate than No Upfront Reserved Instances, which require no upfront payment but have higher monthly payments. On-Demand Instances and Spot Instances are not the best options for the company. On-Demand Instances are a pricing model that offer the most flexibility and no long-term commitment, but have the highest hourly rate. Spot Instances are a pricing model that offer the lowest cost, but are subject to interruption based on supply and demand<sup>34</sup>

**NEW QUESTION 239**

- (Topic 2)

A company provides a software as a service (SaaS) application. The company has a new customer that is based in a different country.

The new customer's data needs to be hosted in that country.

Which AWS service or infrastructure component should the company use to meet this requirement?

- A. AWS Shield
- B. Amazon S3 Object Lock
- C. AWS Regions
- D. Placement groups

**Answer: C**

**Explanation:**

AWS Regions are geographic areas around the world where AWS has clusters of data centers. Each AWS Region consists of multiple, isolated, and physically separate AZ's within a geographic area. By hosting the customer's data in a specific AWS Region, the company can meet the requirement of hosting the data in the customer's country. AWS Shield is a service that provides always-on detection and automatic inline mitigations that minimize application downtime and latency, so there is no need to engage AWS Support to benefit from DDoS protection. Amazon S3 Object Lock is a feature that allows you to store objects using a write-once-read-many (WORM) model. You can use it to prevent an object from being deleted or overwritten for a fixed amount of time or indefinitely. Placement groups are logical grouping of instances within a single Availability Zone. Placement groups enable applications to participate in a low-latency, 10 Gbps network. None of these services or infrastructure components can help the company host the customer's data in a different country.

**NEW QUESTION 243**

- (Topic 2)

A company wants to migrate its applications to the AWS Cloud. The company plans to identify and prioritize any business transformation opportunities and evaluate its AWS Cloud readiness. Which AWS service or tool should the company use to meet these requirements?

- A. AWS Cloud Adoption Framework (AWS CAF)
- B. AWS Managed Services (AMS)
- C. AWS Well-Architected Framework
- D. AWS Migration Hub

**Answer: A**

**Explanation:**

AWS Cloud Adoption Framework (AWS CAF) is a service or tool that helps users migrate their applications to the AWS Cloud. It provides guidance and best practices to identify and prioritize any business transformation opportunities and evaluate their AWS Cloud readiness. It also helps users align their business and technical perspectives, create an actionable roadmap, and measure their progress. AWS Managed Services (AMS) is a service that provides operational services for AWS infrastructure and applications. It helps users reduce their operational overhead and risk, and focus on their core business. It does not help users identify and prioritize any business transformation opportunities and evaluate their AWS Cloud readiness. AWS Well-Architected Framework is a tool that helps users design and implement secure, high-performing, resilient, and efficient solutions on AWS. It provides a set of questions and best practices across five pillars: operational excellence, security, reliability, performance efficiency, and cost optimization. It does not help users identify and prioritize any business transformation opportunities and evaluate their AWS Cloud readiness. AWS Migration Hub is a service that provides a single location to track and manage the migration of applications to AWS. It helps users discover their on-premises servers, group them into applications, and choose the right migration tools. It does not help users identify and prioritize any business transformation opportunities and evaluate their AWS Cloud readiness.

**NEW QUESTION 247**

- (Topic 2)

An ecommerce company wants to design a highly available application that will be hosted on multiple Amazon EC2 instances.

How should the company deploy the EC2 instances to meet these requirements?

- A. Across multiple edge locations
- B. Across multiple VPCs
- C. Across multiple Availability Zones
- D. Across multiple AWS accounts

**Answer: C**

**Explanation:**

The company should deploy the EC2 instances across multiple Availability Zones to design a highly available application. Availability Zones are isolated locations within an AWS Region that are engineered to be fault-tolerant and operate independently of each other. By deploying the EC2 instances across multiple Availability Zones, the company can ensure that their application can withstand the failure of an entire Availability Zone and continue to operate with minimal disruption. Deploying the EC2 instances across multiple edge locations, VPCs, or AWS accounts will not provide the same level of availability and fault tolerance as Availability Zones. Edge locations are part of the Amazon CloudFront service, which is a content delivery network (CDN) that caches and serves web content to users. VPCs are virtual networks that isolate the AWS resources within an AWS Region. AWS accounts are the primary units of ownership and access control for AWS resources<sup>12</sup>

**NEW QUESTION 251**

- (Topic 2)

A company that is planning to migrate to the AWS Cloud is based in an isolated area that has limited internet connectivity. The company needs to perform local

data processing on premises. The company needs a solution that can operate without a stable internet connection. Which AWS service will meet these requirements?

- A. Amazon S3
- B. AWS Snowball Edge
- C. AWS StorageGateway
- D. AWS Backup

**Answer: B**

**Explanation:**

AWS Snowball Edge is a service that provides a physical device that can store up to 100 TB of data and perform local data processing on premises. It enables users to transfer data to and from the AWS Cloud in areas with limited or no internet connectivity. It also supports AWS Greengrass, which allows users to run AWS Lambda functions and other AWS services locally without a stable internet connection. Amazon S3 is a storage service that provides scalable, durable, and secure object storage. It requires a stable internet connection to transfer data to and from the AWS Cloud. AWS Storage Gateway is a service that provides a hybrid storage solution that connects on-premises applications to AWS Cloud storage services, such as Amazon S3, Amazon S3 Glacier, and Amazon EBS. It requires a stable internet connection to synchronize data between the on-premises and cloud storage. AWS Backup is a service that provides a centralized and automated solution to back up data across AWS services and on-premises resources. It requires a stable internet connection to transfer data to and from the AWS Cloud.

**NEW QUESTION 255**

- (Topic 2)

Which benefit of the AWS Cloud helps companies achieve lower usage costs because of the aggregate usage of all AWS users?

- A. No need to guess capacity
- B. Ability to go global in minutes
- C. Economies of scale
- D. Increased speed and agility

**Answer: C**

**Explanation:**

The benefit of the AWS Cloud that helps companies achieve lower usage costs because of the aggregate usage of all AWS users is economies of scale. Economies of scale means that AWS can achieve lower costs and higher efficiency by operating at a massive scale and passing the savings to the customers. AWS leverages the aggregate usage of all AWS users to negotiate better prices with hardware vendors, optimize power consumption, and improve operational processes. As a result, AWS can offer lower and more flexible pricing options to the customers, such as pay-as-you-go, reserved, and spot pricing models. No need to guess capacity, ability to go global in minutes, and increased speed and agility are other benefits of the AWS Cloud, but they are not directly related to the aggregate usage of all AWS users. No need to guess capacity means that AWS customers can avoid the risk of over-provisioning or under-provisioning resources, and scale up or down as needed. Ability to go global in minutes means that AWS customers can deploy their applications and data in multiple regions around the world, and deliver them to users with high performance and availability. Increased speed and agility means that AWS customers can quickly and easily provision and access AWS resources, and accelerate their innovation and time to market.

**NEW QUESTION 258**

- (Topic 2)

A company has an Amazon S3 bucket containing images of scanned financial invoices. The company is building an artificial intelligence (AI)-based application on AWS. The company wants the application to identify and read total balance amounts on the invoices.

Which AWS service will meet these requirements?

- A. Amazon Forecast
- B. Amazon Textract
- C. Amazon Rekognition
- D. Amazon Lex

**Answer: B**

**Explanation:**

Amazon Textract is a service that automatically extracts text and data from scanned documents. Amazon Textract goes beyond simple optical character recognition (OCR) to also identify the contents of fields in forms and information stored in tables. Amazon Textract can analyze images of scanned financial invoices and extract the total balance amounts, as well as other relevant information, such as invoice number, date, vendor name, etc5.

**NEW QUESTION 261**

- (Topic 2)

A company suspects that its AWS resources are being used for illegal activities. Which AWS group or team should the company notify?

- A. AWS Abuse team
- B. AWS Support team
- C. AWS technical account managers
- D. AWS Professional Services team

**Answer: A**

**Explanation:**

AWS Abuse team is the AWS group or team that the company should notify if it suspects that its AWS resources are being used for illegal activities. AWS Abuse team is a dedicated team that handles reports of abuse, such as spam, phishing, malware, denial-of-service attacks, and unauthorized access, involving AWS resources. The company can contact the AWS Abuse team by filling out the [Report Abuse of AWS Resources form] or sending an email to [abuse@amazonaws.com](mailto:abuse@amazonaws.com). The company should provide as much information as possible, such as the source and destination IP addresses, timestamps, log files, and screenshots, to help the AWS Abuse team investigate and take appropriate actions. For more information, see [Reporting Abuse] and [AWS Acceptable Use Policy].

**NEW QUESTION 263**

- (Topic 2)

Which AWS service is always free of charge for users?

- A. Amazon S3
- B. Amazon Aurora
- C. Amazon EC2
- D. AWS Identity and Access Management (IAM)

**Answer:** D

**Explanation:**

AWS Identity and Access Management (IAM) is a service that allows users to manage access to AWS resources and services. It enables users to create and manage users, groups, roles, and policies that control who can do what in AWS. IAM is always free of charge for users, as there is no additional cost for using IAM with any AWS service<sup>1</sup>. Amazon S3 is a storage service that provides scalable, durable, and secure object storage. Amazon S3 has a free tier that offers 5 GB of storage, 20,000 GET requests, and 2,000 PUT requests per month for one year. However, users are charged for any additional usage beyond the free tier limits<sup>2</sup>. Amazon Aurora is a relational database service that is compatible with MySQL and PostgreSQL. Amazon Aurora has a free tier that offers 750 hours of Aurora Single-AZ db.t2.small database usage and 20 GB of storage per month for one year. However, users are charged for any additional usage beyond the free tier limits<sup>3</sup>. Amazon EC2 is a compute service that provides resizable virtual servers. Amazon EC2 has a free tier that offers 750 hours of Linux and Windows t2.micro instances per month for one year. However, users are charged for any additional usage beyond the free tier limits<sup>4</sup>.

**NEW QUESTION 267**

.....

## Thank You for Trying Our Product

### We offer two products:

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

2nd - Questions and Answers in PDF Format

### CLF-C02 Practice Exam Features:

- \* CLF-C02 Questions and Answers Updated Frequently
- \* CLF-C02 Practice Questions Verified by Expert Senior Certified Staff
- \* CLF-C02 Most Realistic Questions that Guarantee you a Pass on Your FirstTry
- \* CLF-C02 Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year

**100% Actual & Verified — Instant Download, Please Click**  
**[Order The CLF-C02 Practice Test Here](#)**