



**Amazon**

## **Exam Questions AWS-Certified-Cloud-Practitioner**

Amazon AWS Certified Cloud Practitioner

### NEW QUESTION 1

- (Topic 2)

A company wants to migrate its on-premises application to the AWS Cloud. The company is legally obligated to retain certain data in its on-premises data center. Which AWS service or feature will support this requirement?

- A. AWS Wavelength
- B. AWS Local Zones
- C. VMware Cloud on AWS
- D. AWS Outposts

**Answer:** D

#### Explanation:

AWS Outposts is a fully managed service that extends AWS infrastructure, AWS services, APIs, and tools to virtually any datacenter, co-location space, or on-premises facility for a truly consistent hybrid experience. AWS Outposts enables you to run AWS services in your on-premises data center, which can support the requirement of retaining certain data on-premises due to legal obligations<sup>5</sup>.

### NEW QUESTION 2

- (Topic 2)

A company is collecting user behavior patterns to identify how to meet goals for sustainability impact. Which guidelines are best practices for the company to implement to meet these goals? (Select TWO.)

- A. Scale infrastructure with user load.
- B. Maximize the geographic distance between workloads and user locations.
- C. Eliminate creation and maintenance of unused assets.
- D. Scale resources with excess capacity and remove auto scaling.
- E. Scale infrastructure based on the number of users.

**Answer:** AC

#### Explanation:

To meet the goals for sustainability impact, the company should follow the best practices of scaling infrastructure with user load and eliminating creation and maintenance of unused assets. Scaling infrastructure with user load means adjusting the capacity of the infrastructure to match the demand of the users, which can reduce the energy consumption and carbon footprint of the system. Eliminating creation and maintenance of unused assets means avoiding the waste of resources and money on assets that are not needed or used, which can also improve the environmental and economic efficiency of the system<sup>3</sup>.

### NEW QUESTION 3

- (Topic 2)

Which AWS service or tool helps companies measure the environmental impact of their AWS usage?

- A. AWS customer carbon footprint tool
- B. AWS Compute Optimizer
- C. Sustainability pillar
- D. OS-Climate (Open Source Climate Data Commons)

**Answer:** A

#### Explanation:

AWS customer carbon footprint tool is an AWS service or tool that helps companies measure the environmental impact of their AWS usage. It allows users to estimate the carbon emissions associated with their AWS resources and services, such as EC2, S3, and Lambda. It also provides recommendations and best practices to reduce the carbon footprint and improve the sustainability of their AWS workloads<sup>4</sup>. AWS Compute Optimizer is an AWS service that helps users optimize the performance and cost of their EC2 instances and Auto Scaling groups. It provides recommendations for optimal instance types, sizes, and configurations based on the workload characteristics and utilization metrics. It does not help users measure the environmental impact of their AWS usage. Sustainability pillar is a concept that refers to the ability of a system to operate in an environmentally friendly and socially responsible manner. It is not an AWS service or tool that helps users measure the environmental impact of their AWS usage. OS-Climate (Open Source Climate Data Commons) is an initiative that aims to provide open source data, tools, and platforms to accelerate climate action and innovation. It is not an AWS service or tool that helps users measure the environmental impact of their AWS usage.

### NEW QUESTION 4

- (Topic 2)

A company wants to develop a shopping application that records customer orders. The application needs to use an AWS managed database service to store data. Which AWS service should the company use to meet these requirements?

- A. Amazon RDS
- B. Amazon Redshift
- C. Amazon ElastiCache
- D. Amazon Neptune

**Answer:** A

#### Explanation:

A is correct because Amazon RDS is the AWS service that provides a managed relational database service that supports various database engines, such as MySQL, PostgreSQL, Oracle, and SQL Server. B is incorrect because Amazon Redshift is the AWS service that provides a managed data warehouse service that is optimized for analytical queries. C is incorrect because Amazon ElastiCache is the AWS service that provides a managed in-memory data store service that supports Redis and Memcached. D is incorrect because Amazon Neptune is the AWS service that provides a managed graph database service that supports property graph and RDF models.

#### NEW QUESTION 5

- (Topic 2)

A company wants to migrate its Microsoft SQL Server database management system from on premises to the AWS Cloud. Which AWS service should the company use to reduce management overhead for this environment?

- A. Amazon Elastic Container Service (Amazon ECS)
- B. Amazon SageMaker
- C. Amazon RDS
- D. Amazon Athena

**Answer: C**

#### Explanation:

Amazon Relational Database Service (Amazon RDS) is the AWS service that the company should use to migrate its Microsoft SQL Server database management system from on premises to the AWS Cloud. Amazon RDS is a fully managed service that provides a scalable, secure, and high-performance relational database platform. Amazon RDS supports several database engines, including Microsoft SQL Server. Amazon RDS reduces the management overhead for the database environment by taking care of tasks such as provisioning, patching, backup, recovery, and monitoring. For more information, see [What is Amazon Relational Database Service \(Amazon RDS\)?](#) and [Amazon RDS for SQL Server](#).

#### NEW QUESTION 6

- (Topic 2)

A company is preparing to launch a redesigned website on AWS. Users from around the world will download digital handbooks from the website. Which AWS solution should the company use to provide these static files securely?

- A. Amazon Kinesis Data Streams
- B. Amazon CloudFront with Amazon S3
- C. Amazon EC2 instances with an Application Load Balancer
- D. Amazon Elastic File System (Amazon EFS)

**Answer: B**

#### Explanation:

Amazon CloudFront with Amazon S3 is a solution that allows you to provide static files securely to users from around the world. Amazon CloudFront is a fast content delivery network (CDN) service that securely delivers data, videos, applications, and APIs to customers globally with low latency, high transfer speeds, all within a developer-friendly environment. Amazon S3 is an object storage service that offers industry-leading scalability, data availability, security, and performance. You can use Amazon S3 to store and retrieve any amount of data from anywhere. You can also configure Amazon S3 to work with Amazon CloudFront to distribute your content to edge locations near your users for faster delivery and lower latency. Amazon Kinesis Data Streams is a service that enables you to build custom applications that process or analyze streaming data for specialized needs. This option is not relevant for providing static files securely. Amazon EC2 instances with an Application Load Balancer is a solution that allows you to distribute incoming traffic across multiple targets, such as EC2 instances, in multiple Availability Zones. This option is suitable for dynamic web applications, but not necessary for static files. Amazon Elastic File System (Amazon EFS) is a service that provides a simple, scalable, fully managed elastic NFS file system for use with AWS Cloud services and on-premises resources. This option is not relevant for providing static files securely.

#### NEW QUESTION 7

- (Topic 2)

A company wants to run its production workloads on AWS. The company needs concierge service, a designated AWS technical account manager (TAM), and technical support that is available 24 hours a day, 7 days a week. Which AWS Support plan will meet these requirements?

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

**Answer: B**

#### Explanation:

B is correct because AWS Enterprise Support is the AWS Support plan that provides concierge service, a designated AWS technical account manager (TAM), and technical support that is available 24 hours a day, 7 days a week. This plan is designed for customers who run mission-critical workloads on AWS and need the highest level of support. A is incorrect because AWS Basic Support is the AWS Support plan that provides customer service and support for billing and account issues, service limit increases, and technical support for a limited set of AWS services. It does not provide concierge service, a designated TAM, or 24/7 technical support. C is incorrect because AWS Business Support is the AWS Support plan that provides customer service and support for billing and account issues, service limit increases, and technical support for all AWS services, as well as access to AWS Trusted Advisor and AWS Support API. It does not provide concierge service or a designated TAM. D is incorrect because AWS Developer Support is the AWS Support plan that provides customer service and support for billing and account issues, service limit increases, and technical support for all AWS services, as well as access to AWS Trusted Advisor. It does not provide concierge service, a designated TAM, or 24/7 technical support.

#### NEW QUESTION 8

- (Topic 2)

A company wants to push VPC Flow Logs to an Amazon S3 bucket.

A company wants to optimize long-term compute costs of AWS Lambda functions and Amazon EC2 instances.

Which AWS purchasing option should the company choose to meet these requirements?

- A. Dedicated Hosts
- B. Compute Savings Plans
- C. Reserved Instances
- D. Spot Instances

**Answer: B**

**Explanation:**

Compute Savings Plans are a flexible and cost-effective way to optimize long-term compute costs of AWS Lambda functions and Amazon EC2 instances. With Compute Savings Plans, customers can commit to a consistent amount of compute usage (measured in \$/hour) for a 1-year or 3-year term and receive a discount of up to 66% compared to On-Demand prices<sup>3</sup>. Dedicated Hosts are physical servers with EC2 instance capacity fully dedicated to the customer's use. They are suitable for customers who have specific server-bound software licenses or compliance requirements<sup>4</sup>. Reserved Instances are a pricing model that provides a significant discount (up to 75%) compared to On-Demand pricing and a capacity reservation for EC2 instances. They are available in 1-year or 3-year terms and different payment options<sup>5</sup>. Spot Instances are spare EC2 instances that are available at up to 90% discount compared to On-Demand prices. They are suitable for customers who have flexible start and end times, can withstand interruptions, and can handle excess capacity.

**NEW QUESTION 9**

- (Topic 2)

Which AWS service requires the customer to patch the guest operating system?

- A. AWS Lambda
- B. Amazon OpenSearch Service
- C. Amazon EC2
- D. Amazon ElastiCache

**Answer: C**

**Explanation:**

The AWS service that requires the customer to patch the guest operating system is Amazon EC2. Amazon EC2 is a service that provides scalable compute capacity in the cloud, and allows customers to launch and run virtual servers, called instances, with a variety of operating systems, configurations, and specifications. The customer is responsible for patching and updating the guest operating system and any applications that run on the EC2 instances, as part of the security in the cloud. AWS Lambda, Amazon OpenSearch Service, and Amazon ElastiCache are not services that require the customer to patch the guest operating system. AWS Lambda is a serverless compute service that allows customers to run code without provisioning or managing servers. Amazon OpenSearch Service is a fully managed service that makes it easy to deploy, operate, and scale OpenSearch clusters in the AWS Cloud. Amazon ElastiCache is a fully managed service that provides in-memory data store and cache solutions, such as Redis and Memcached. These services are managed by AWS, and AWS is responsible for patching and updating the underlying infrastructure and software.

**NEW QUESTION 10**

- (Topic 1)

Who is responsible for decommissioning end-of-life underlying storage devices that are used to host data on AWS?

- A. Customer
- B. AWS
- C. Account creator
- D. Auditing team

**Answer: B**

**Explanation:**

AWS is responsible for decommissioning end-of-life underlying storage devices that are used to host data on AWS. AWS follows strict and audited data destruction processes to ensure that customer data is not exposed to unauthorized individuals or devices when an AWS storage device reaches the end of its useful life. AWS uses techniques detailed in DoD 5220.22-M ("National Industrial Security Program Operating Manual") or NIST 800-88 ("Guidelines for Media Sanitization") to destroy data as part of the decommissioning process<sup>3</sup>.

**NEW QUESTION 10**

- (Topic 1)

Which AWS service or tool does AWS Control Tower use to create resources?

- A. AWS CloudFormation
- B. AWS Trusted Advisor
- C. AWS Directory Service
- D. AWS Cost Explorer

**Answer: A**

**Explanation:**

AWS Control Tower uses AWS CloudFormation to create resources in your landing zone. AWS CloudFormation is a service that helps you model and set up your AWS resources using templates. AWS Control Tower supports creating AWS::ControlTower::EnabledControl resources in AWS CloudFormation. Therefore, the correct answer is A. You can learn more about AWS Control Tower and AWS CloudFormation from this page.

**NEW QUESTION 11**

- (Topic 1)

Which AWS service will help protect applications running on AWS from DDoS attacks?

- A. Amazon GuardDuty
- B. AWS WAF
- C. AWS Shield
- D. Amazon Inspector

**Answer: C**

**Explanation:**

AWS Shield is a managed Distributed Denial of Service (DDoS) protection service that safeguards applications running on AWS. AWS Shield 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<sup>3</sup>.

#### NEW QUESTION 14

- (Topic 1)

Which of the following are pillars of the AWS Well-Architected Framework? (Select TWO.)

- A. Availability
- B. Reliability
- C. Scalability
- D. Responsive design
- E. Operational excellence

**Answer:** BE

#### Explanation:

The correct answers to the questions are B and E because reliability and operational excellence are pillars of the AWS Well-Architected Framework. 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 AWS Well-Architected Framework consists of five pillars: operational excellence, security, reliability, performance efficiency, and cost optimization. Each pillar has a set of design principles that describe the characteristics of a well-architected system. Reliability is the pillar that focuses on the ability of a system to recover from failures and meet business and customer demand. Operational excellence is the pillar that focuses on the ability of a system to run and monitor processes that support business outcomes and continually improve. The other options are incorrect because they are not pillars of the AWS Well-Architected Framework. Availability, scalability, and responsive design are important aspects of cloud architecture, but they are not separate pillars in the framework. Availability and scalability are related to the reliability and performance efficiency pillars, while responsive design is related to the customer experience and user interface. Reference: AWS Well-Architected Framework

#### NEW QUESTION 15

- (Topic 1)

A company needs to run its existing custom, nonproduction workloads in the AWS Cloud quickly and cost-effectively. The workloads can recover from interruptions easily. Which pricing model should the company use?

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

**Answer:** C

#### Explanation:

The correct answer is C because Spot Instances are the pricing model that enables the company to run its existing custom, nonproduction workloads in the AWS Cloud quickly and cost-effectively. Spot Instances are spare Amazon EC2 instances that are available at up to 90% discount compared to On-Demand prices. Spot Instances are suitable for stateless, fault-tolerant, and flexible workloads that can recover from interruptions easily. The other options are incorrect because they are not the pricing model that enables the company to run its existing custom, nonproduction workloads in the AWS Cloud quickly and cost-effectively. Reserved Instances are Amazon EC2 instances that are reserved for a specific period of time (one or three years) in exchange for a lower hourly rate. Reserved Instances are suitable for steady-state or predictable workloads that run for a long duration. On-Demand Instances are Amazon EC2 instances that are launched and billed at a fixed hourly rate. On-Demand Instances are suitable for short-term, irregular, or unpredictable workloads that cannot be interrupted. Dedicated Hosts are physical servers that are dedicated to a single customer. Dedicated Hosts are suitable for workloads that require regulatory compliance or data isolation. Reference: Amazon EC2 Instance Purchasing Options

#### NEW QUESTION 18

- (Topic 1)

Which statement describes a characteristic of the AWS global infrastructure?

- A. Edge locations contain multiple AWS Regions.
- B. AWS Regions contain multiple Regional edge caches.
- C. Availability Zones contain multiple data centers.
- D. Each data center contains multiple edge locations.

**Answer:** C

#### Explanation:

Availability Zones contain multiple data centers. This is a characteristic of the AWS global infrastructure, which consists of AWS Regions, Availability Zones, and edge locations. AWS Regions are geographically isolated areas that contain multiple Availability Zones. Availability Zones are physically separate locations within an AWS Region that are engineered to be isolated from failures and connected by low-latency, high-throughput, and highly redundant networking. Each Availability Zone contains one or more data centers that house the servers and storage devices that run AWS services. Edge locations are sites that are located closer to the end users and provide caching and content delivery services. AWS Global Infrastructure AWS Certified Cloud Practitioner - aws.amazon.com

#### NEW QUESTION 20

- (Topic 1)

Which of the following are design principles for reliability in the AWS Cloud? (Select TWO.)

- A. Build architectures with tightly coupled resources.
- B. Use AWS Trusted Advisor to meet security best practices.
- C. Use automation to recover immediately from failure.
- D. Rightsize Amazon EC2 instances to ensure optimal performance.
- E. Simulate failures to test recovery processes.

**Answer:** CE

#### Explanation:

The design principles for reliability in the AWS Cloud are:

? Test recovery procedures. The best way to ensure that systems can recover from failures is to regularly test them using simulated scenarios. This can help

identify gaps and improve the recovery process.

? Automatically recover from failure. By using automation, systems can detect and correct failures without human intervention. This can reduce the impact and duration of failures and improve the availability of the system.

? Scale horizontally to increase aggregate system availability. By adding more redundant resources to the system, the impact of individual resource failures can be reduced. This can also improve the performance and scalability of the system.

? Stop guessing capacity. By using monitoring and automation, systems can adjust the capacity based on the demand and performance metrics. This can prevent failures due to insufficient or excessive capacity and optimize the cost and efficiency of the system.

? Manage change in automation. By using automation, changes to the system can be applied in a consistent and controlled manner. This can reduce the risk of human errors and configuration drifts that can cause failures. AWS Well- Architected Framework

#### NEW QUESTION 22

- (Topic 1)

What is an Availability Zone?

- A. A location where users can deploy compute, storage, database, and other select AWS services where no AWS Region currently exists
- B. One or more discrete data centers with redundant power, networking, and connectivity
- C. One or more clusters of servers where new workloads can be deployed
- D. A fast content delivery network (CDN) service that securely delivers data, videos, applications, and APIs to users globally

**Answer: B**

#### Explanation:

An Availability Zone is one or more discrete data centers with redundant power, networking, and connectivity. Availability Zones are part of the AWS global infrastructure, which consists of AWS Regions, Availability Zones, and edge locations. Availability Zones are physically separate locations within an AWS Region that are engineered to be isolated from failures and connected by low-latency, high-throughput, and highly redundant networking. Each Availability Zone contains one or more data centers that house the servers and storage devices that run AWS services. Availability Zones enable users to design and operate fault-tolerant and high-availability applications on AWS. AWS Global Infrastructure AWS Certified Cloud Practitioner - [aws.amazon.com](https://aws.amazon.com)

#### NEW QUESTION 26

- (Topic 1)

A company has an application with robust hardware requirements. The application must be accessed by students who are using lightweight, low-cost laptops. Which AWS service will help the company deploy the application without investing in backend infrastructure or high end client hardware?

- A. Amazon AppStream 2.0
- B. AWS AppSync
- C. Amazon WorkLink
- D. AWS Elastic Beanstalk

**Answer: A**

#### Explanation:

The correct answer is A because Amazon AppStream 2.0 is a service that will help the company deploy the application without investing in backend infrastructure or high end client hardware. Amazon AppStream 2.0 is a fully managed, secure application streaming service that allows customers to stream desktop applications from AWS to any device running a web browser. Amazon AppStream 2.0 handles the provisioning, scaling, patching, and maintenance of the backend infrastructure, and delivers high performance and responsive user experience. The other options are incorrect because they are not services that will help the company deploy the application without investing in backend infrastructure or high end client hardware. AWS AppSync is a service that enables customers to create flexible APIs for synchronizing data across multiple data sources. Amazon WorkLink is a service that enables customers to provide secure, one-click access to internal websites and web apps from mobile devices. AWS Elastic Beanstalk is a service that enables customers to deploy and manage web applications using popular platforms such as Java, .NET, PHP, and Node.js. Reference: [Amazon AppStream 2.0 FAQs]

#### NEW QUESTION 27

- (Topic 1)

Which of the following is available to a company that has an AWS Business Support plan?

- A. AWS Support concierge
- B. AWS DDoS Response Team (DRT)
- C. AWS technical account manager (TAM)
- D. AWS Health API

**Answer: D**

#### Explanation:

AWS Health API is available to a company that has an AWS Business Support plan. The AWS Health API provides programmatic access to the AWS Health information that is presented in the AWS Personal Health Dashboard. The AWS Health API can help users get timely and personalized information about events that can affect the availability and performance of their AWS resources, such as scheduled maintenance, network issues, or service disruptions. The AWS Health API can also integrate with other AWS services, such as Amazon CloudWatch Events and AWS Lambda, to enable automated actions and notifications. AWS Health API Overview AWS Support Plans

#### NEW QUESTION 28

- (Topic 1)

Which option is an advantage of AWS Cloud computing that minimizes variable costs?

- A. High availability
- B. Economies of scale
- C. Global reach
- D. Agility

**Answer: B**

**Explanation:**

One of the advantages of AWS Cloud computing is that it minimizes variable costs by leveraging economies of scale. This means that AWS can achieve lower costs per unit of computing resources by spreading the fixed costs of building and maintaining data centers over a large number of customers. As a result, AWS can offer lower and more predictable prices to its customers, who only pay for the resources they consume. Therefore, the correct answer is B. You can learn more about AWS pricing and economies of scale from this page.

**NEW QUESTION 33**

- (Topic 1)

A cloud engineer needs to download AWS security and compliance documents for an upcoming audit. Which AWS service can provide the documents?

- A. AWS Trusted Advisor
- B. AWS Artifact
- C. AWS Well-Architected Tool
- D. AWS Systems Manager

**Answer: B**

**Explanation:**

AWS Artifact is the AWS service that can provide security and compliance documents for an upcoming audit. AWS Artifact is a self-service portal that allows users to access and download AWS compliance reports and agreements. These documents provide evidence of AWS's compliance with global, regional, and industry-specific security standards and regulations

**NEW QUESTION 35**

- (Topic 1)

A company's application stores data in an Amazon S3 bucket. The company has an AWS Lambda function that processes data in the S3 bucket. The company needs to invoke the function once a day at a specific time. Which AWS service should the company use to meet this requirement?

- A. AWS Managed Services (AMS)
- B. AWS CodeStar
- C. Amazon EventBridge
- D. AWS Step Functions

**Answer: C**

**Explanation:**

Amazon EventBridge is the service that the company should use to meet the requirement of invoking the Lambda function once a day at a specific time. Amazon EventBridge is a serverless event bus service that allows you to easily connect your applications with data from AWS services, SaaS applications, and your own applications. You can use Amazon EventBridge to create rules that match events and route them to targets such as AWS Lambda functions, Amazon SNS topics, Amazon SQS queues, or other AWS services. You can also use Amazon EventBridge to create scheduled rules that trigger your targets at a specific time or interval, such as once a day. AWS Managed Services (AMS), AWS CodeStar, and AWS Step Functions are not services that the company should use to meet this requirement. AMS is a service that provides operational management for your AWS infrastructure and applications. AWS CodeStar is a service that provides a unified user interface for managing software development projects on AWS. AWS Step Functions is a service that coordinates multiple AWS services into serverless workflows.

**NEW QUESTION 37**

- (Topic 1)

A company hosts an application on an Amazon EC2 instance. The EC2 instance needs to access several AWS resources, including Amazon S3 and Amazon DynamoDB.

What is the MOST operationally efficient solution to delegate permissions?

- A. Create an IAM role with the required permission
- B. Attach the role to the EC2 instance.
- C. Create an IAM user and use its access key and secret access key in the application.
- D. Create an IAM user and use its access key and secret access key to create a CLI profile in the EC2 instance.
- E. Create an IAM role with the required permission
- F. Attach the role to the administrative IAM user.

**Answer: A**

**Explanation:**

Creating an IAM role with the required permissions and attaching the role to the EC2 instance is the most operationally efficient solution to delegate permissions. An IAM role is an entity that defines a set of permissions for making AWS service requests. An IAM role can be assumed by an EC2 instance to access other AWS resources, such as Amazon S3 and Amazon DynamoDB, without having to store any credentials on the instance. This solution is more secure and scalable than using IAM users and their access keys. For more information, see [IAM Roles for Amazon EC2] and [Using an IAM Role to Grant Permissions to Applications Running on Amazon EC2 Instances].

**NEW QUESTION 42**

- (Topic 1)

A company has a social media platform in which users upload and share photos with other users. The company wants to identify and remove inappropriate photos. The company has no machine learning (ML) scientists and must build this detection capability with no ML expertise. Which AWS service should the company use to build this capability?

- A. Amazon SageMaker
- B. Amazon Textract
- C. Amazon Rekognition
- D. Amazon Comprehend

**Answer: C**

**Explanation:**

Amazon Rekognition is the AWS service that the company should use to build the capability of identifying and removing inappropriate photos. Amazon Rekognition is a service that uses deep learning technology to analyze images and videos for various purposes, such as face detection, object recognition, text extraction, and content moderation. Amazon Rekognition can help users detect unsafe or inappropriate content in images and videos, such as nudity, violence, or drugs, and provide confidence scores for each label. Amazon Rekognition does not require any machine learning expertise, and users can easily integrate it with other AWS services

**NEW QUESTION 45**

- (Topic 1)

Which AWS service or feature is used to send both text and email messages from distributed applications?

- A. Amazon Simple Notification Service (Amazon SNS)
- B. Amazon Simple Email Service (Amazon SES)
- C. Amazon CloudWatch alerts
- D. Amazon Simple Queue Service (Amazon SQS)

**Answer:** A

**Explanation:**

Amazon Simple Notification Service (Amazon SNS) is the AWS service or feature that is used to send both text and email messages from distributed applications. Amazon SNS is a fully managed pub/sub messaging service that enables the user to send messages to multiple subscribers or endpoints, such as email addresses, phone numbers, HTTP endpoints, AWS Lambda functions, and more. Amazon SNS can be used to send notifications, alerts, confirmations, and reminders from applications to users or other applications.

**NEW QUESTION 49**

- (Topic 1)

According to the AWS shared responsibility model, which of the following are AWS responsibilities? (Select TWO.)

- A. Network infrastructure and virtualization of infrastructure
- B. Security of application data
- C. Guest operating systems
- D. Physical security of hardware
- E. Credentials and policies

**Answer:** AD

**Explanation:**

The correct answers are A and D because network infrastructure and virtualization of infrastructure and physical security of hardware are AWS responsibilities according to the AWS shared responsibility model. The AWS shared responsibility model is a framework that defines 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 global infrastructure, such as the regions, availability zones, and edge locations; the hardware, software, networking, and facilities that run the AWS services; and the virtualization layer that separates the customer instances and storage. The customer is responsible for the security in the cloud, which includes the customer data, the guest operating systems, the applications, the identity and access management, the firewall configuration, and the encryption. The other options are incorrect because they are not AWS responsibilities according to the AWS shared responsibility model. Security of application data, guest operating systems, and credentials and policies are customer responsibilities according to the AWS shared responsibility model. Reference: [AWS Shared Responsibility Model]

**NEW QUESTION 53**

- (Topic 1)

A company needs to run code in response to an event notification that occurs when objects are uploaded to an Amazon S3 bucket. Which AWS service will integrate directly with the event notification?

- A. AWS Lambda
- B. Amazon EC2
- C. Amazon Elastic Container Registry (Amazon ECR)
- D. AWS Elastic Beanstalk

**Answer:** A

**Explanation:**

AWS Lambda is a service that lets you run code without provisioning or managing servers. You can use Lambda to process event notifications from Amazon S3 when objects are uploaded or deleted. Lambda integrates directly with the event notification and invokes your code automatically. Therefore, the correct answer is A.

**NEW QUESTION 58**

- (Topic 1)

Which database engine is compatible with Amazon RDS?

- A. Apache Cassandra
- B. MongoDB
- C. Neo4j
- D. PostgreSQL

**Answer:** D

**Explanation:**

Amazon RDS supports six database engines: Amazon Aurora, MySQL, MariaDB, PostgreSQL, Oracle, and SQL Server. Apache Cassandra, MongoDB, and Neo4j are not compatible with Amazon RDS. Therefore, the correct answer is D. You can learn more about Amazon RDS and its supported database engines from this page.

#### NEW QUESTION 62

- (Topic 1)

Which AWS Well-Architected Framework concept represents a system's ability to remain functional when the system encounters operational problems?

- A. Consistency
- B. Elasticity
- C. Durability
- D. Latency

**Answer: B**

#### Explanation:

The AWS Well-Architected Framework is a set of best practices and guidelines for designing and operating systems in the cloud. The framework consists of five pillars: operational excellence, security, reliability, performance efficiency, and cost optimization. The concept of elasticity represents a system's ability to adapt to changes in demand by scaling resources up or down automatically. Therefore, the correct answer is B. You can learn more about the AWS Well-Architected Framework and its pillars from this page.

#### NEW QUESTION 64

- (Topic 1)

A company wants to use the AWS Cloud as an offsite backup location for its on-premises infrastructure.

Which AWS service will meet this requirement MOST cost-effectively?

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

**Answer: A**

#### Explanation:

Amazon S3 is the most cost-effective service for storing offsite backups of on-premises infrastructure. Amazon S3 offers low-cost, durable, and scalable storage that can be accessed from anywhere over the internet. Amazon S3 also supports lifecycle policies, versioning, encryption, and cross-region replication to optimize the backup and recovery process. Amazon EFS, Amazon FSx, and Amazon EBS are more suitable for storing data that requires high performance, low latency, and frequent access<sup>12</sup>

#### NEW QUESTION 65

- (Topic 1)

Which activity is a customer responsibility in the AWS Cloud according to the AWS shared responsibility model?

- A. Ensuring network connectivity from AWS to the internet
- B. Patching and fixing flaws within the AWS Cloud infrastructure
- C. Ensuring the physical security of cloud data centers
- D. Ensuring Amazon EBS volumes are backed up

**Answer: D**

#### Explanation:

The AWS shared responsibility model describes how AWS and the customer share responsibility for security and compliance of the AWS environment. AWS is responsible for the security of the cloud, which includes the physical security of AWS facilities, the infrastructure, hardware, software, and networking that run AWS services. The customer is responsible for security in the cloud, which includes the configuration of security groups, the encryption of customer data on AWS, the management of AWS Lambda infrastructure, and the management of network throughput of each AWS Region. One of the customer responsibilities is to ensure that Amazon EBS volumes are backed up.

#### NEW QUESTION 69

- (Topic 1)

What can a user accomplish using AWS CloudTrail?

- A. Generate an IAM user credentials report.
- B. Record API calls made to AWS services.
- C. Assess the compliance of AWS resource configurations with policies and guidelines.
- D. Ensure that Amazon EC2 instances are patched with the latest security update
- E. A company uses Amazon Workspaces.

**Answer: B**

#### Explanation:

AWS CloudTrail is an AWS service that enables users to accomplish the task of recording API calls made to AWS services. AWS CloudTrail is a service that tracks user activity and API usage across the AWS account. AWS CloudTrail records the details of every API call made to AWS services, such as the identity of the caller, the time of the call, the source IP address of the caller, the parameters and responses of the call, and more. Users can use AWS CloudTrail to audit, monitor, and troubleshoot their AWS resources and actions. The other options are incorrect because they are not tasks that users can accomplish using AWS CloudTrail. Generating an IAM user credentials report is a task that users can accomplish using IAM, which is an AWS service that enables users to manage access and permissions to AWS resources and services. Assessing the compliance of AWS resource configurations with policies and guidelines is a task that users can accomplish using AWS Config, which is an AWS service that enables users to assess, audit, and evaluate the configurations of their AWS resources. Ensuring that Amazon EC2 instances are patched with the latest security updates is a task that users can accomplish using AWS Systems Manager, which is an AWS service that enables users to automate operational tasks, manage configuration and compliance, and monitor system health and performance. Reference: AWS CloudTrail FAQs

#### NEW QUESTION 74

- (Topic 1)

Which benefit does Amazon Rekognition provide?

- A. The ability to place watermarks on images
- B. The ability to detect objects that appear in pictures
- C. The ability to resize millions of images automatically
- D. The ability to bid on object detection jobs

**Answer: B**

**Explanation:**

Amazon Rekognition is a service that provides deep learning-based image and video analysis. One of the benefits of Amazon Rekognition is the ability to detect objects that appear in pictures, such as faces, landmarks, animals, text, and scenes. This can enable applications to perform tasks such as face recognition, face verification, face comparison, face search, celebrity recognition, emotion detection, age range estimation, gender identification, facial analysis, facial expression recognition, and more. Amazon Rekognition Overview AWS Certified Cloud Practitioner - [aws.amazon.com](https://aws.amazon.com)

**NEW QUESTION 78**

- (Topic 1)

Which AWS benefit is demonstrated by on-demand technology services that enable companies to replace upfront fixed expenses with variable expenses?

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

**Answer: C**

**Explanation:**

Pay-as-you-go pricing is an AWS benefit that demonstrates the ability of users to replace upfront fixed expenses with variable expenses. With pay-as-you-go pricing, users only pay for the resources they consume, without any long-term contracts or commitments. This can lower the total cost of ownership and increase the return on investment. Pay-as-you-go pricing also provides flexibility and scalability, as users can adjust their resource usage according to their changing needs and demands. AWS Cloud Value Framework AWS Certified Cloud Practitioner - [aws.amazon.com](https://aws.amazon.com)

**NEW QUESTION 83**

- (Topic 1)

Which AWS services or features can control VPC traffic? (Select TWO.)

- A. Security groups
- B. AWS Direct Connect
- C. Amazon GuardDuty
- D. Network ACLs
- E. Amazon Connect

**Answer: AD**

**Explanation:**

The AWS services or features that can control VPC traffic are security groups and network ACLs. Security groups are stateful firewalls that control the inbound and outbound traffic at the instance level. You can assign one or more security groups to each instance in a VPC, and specify the rules that allow or deny traffic based on the protocol, port, and source or destination. Network ACLs are stateless firewalls that control the inbound and outbound traffic at the subnet level. You can associate one network ACL with each subnet in a VPC, and specify the rules that allow or deny traffic based on the protocol, port, and source or destination. AWS Direct Connect, Amazon GuardDuty, and Amazon Connect are not services or features that can control VPC traffic. AWS Direct Connect is a service that establishes a dedicated network connection between your premises and AWS. Amazon GuardDuty is a service that monitors your AWS account and workloads for malicious or unauthorized activity. Amazon Connect is a service that provides a cloud-based contact center solution.

**NEW QUESTION 84**

SIMULATION - (Topic 1)

A company runs thousands of simultaneous simulations using AWS Batch. Each simulation is stateless, is fault tolerant, and runs for up to 3 hours.

Which pricing model enables the company to optimize costs and meet these requirements?

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

**Answer: B**

**Explanation:**

The correct answer is B because Spot Instances enable the company to optimize costs and meet the requirements. Spot Instances are spare EC2 instances that are available at up to 90% discount compared to On-Demand prices. Spot Instances are suitable for stateless, fault-tolerant, and flexible applications that can run for any duration. The other options are incorrect because they do not enable the company to optimize costs and meet the requirements. Reserved Instances are EC2 instances that are reserved for a specific period of time (one or three years) in exchange for a lower hourly rate. Reserved Instances are suitable for steady-state or predictable workloads that run for a long duration. On-Demand Instances are EC2 instances that are launched and billed at a fixed hourly rate. On-Demand Instances are suitable for short-term, irregular, or unpredictable workloads that cannot be interrupted. Dedicated Instances are EC2 instances that run on hardware that is dedicated to a single customer. Dedicated Instances are suitable for workloads that require regulatory compliance or data isolation. Reference: [Amazon EC2 Instance Purchasing Options]

**NEW QUESTION 86**

- (Topic 1)

A company needs to use dashboards and charts to analyze insights from business data. Which AWS service will provide the dashboards and charts for these

insights?

- A. Amazon Macie
- B. Amazon Aurora
- C. Amazon QuickSight
- D. AWS CloudTrail

**Answer: C**

**Explanation:**

The correct answer is C because Amazon QuickSight is an AWS service that will provide the dashboards and charts for the insights from business data. Amazon QuickSight is a fully managed, scalable, and serverless business intelligence service that enables users to create and share interactive dashboards and charts. Amazon QuickSight can connect to various data sources, such as Amazon S3, Amazon RDS, Amazon Redshift, and more. Amazon QuickSight also provides users with machine learning insights, such as anomaly detection, forecasting, and natural language narratives. The other options are incorrect because they are not AWS services that will provide the dashboards and charts for the insights from business data. Amazon Macie is an AWS service that helps users discover, classify, and protect sensitive data stored in Amazon S3. Amazon Aurora is an AWS service that provides a relational database that is compatible with MySQL and PostgreSQL. AWS CloudTrail is an AWS service that enables users to track user activity and API usage across their AWS account. Reference: Amazon QuickSight FAQs

**NEW QUESTION 91**

- (Topic 1)

A company's user base needs to remotely access virtual desktop computers from the internet Which AWS service provides this functionality?

- A. Amazon Connect
- B. Amazon Cognito
- C. Amazon Workspaces
- D. Amazon Upstream 2.0

**Answer: C**

**Explanation:**

Amazon Workspaces is the AWS service that provides the functionality of remotely accessing virtual desktop computers from the internet. Amazon Workspaces is a fully managed, secure desktop-as-a-service (DaaS) solution that allows users to provision cloud-based virtual desktops and access them from anywhere, using any supported device. Amazon Workspaces helps users reduce the complexity and cost of managing and maintaining physical desktops, and provides a consistent and secure user experience

**NEW QUESTION 95**

- (Topic 1)

Which AWS service aggregates, organizes, and prioritizes security alerts and findings from multiple AWS services?

- A. Amazon Detective
- B. Amazon Inspector
- C. Amazon Macie
- D. AWS Security Hub

**Answer: D**

**Explanation:**

The correct answer is D because AWS Security Hub is a service that aggregates, organizes, and prioritizes security alerts and findings from multiple AWS services, such as Amazon GuardDuty, Amazon Inspector, Amazon Macie, AWS Firewall Manager, and AWS IAM Access Analyzer. The other options are incorrect because they are not services that aggregate security alerts and findings from multiple AWS services. Amazon Detective is a service that helps users analyze and visualize security data to investigate and remediate potential issues. Amazon Inspector is a service that helps users find security vulnerabilities and deviations from best practices in their Amazon EC2 instances. Amazon Macie is a service that helps users discover, classify, and protect sensitive data stored in Amazon S3. Reference: AWS Security Hub FAQs

**NEW QUESTION 99**

- (Topic 1)

Which task is a customer's responsibility, according to the AWS shared responsibility model?

- A. Management of the guest operating systems
  - B. Maintenance of the configuration of infrastructure devices
  - C. Management of the host operating systems and virtualization
  - D. Maintenance of the software that powers Availability Zones
- A company has refined its workload to use specific AWS services to improve efficiency and reduce cost.

**Answer: A**

**Explanation:**

Management of the guest operating systems is a customer's responsibility, according to the AWS shared responsibility model. The AWS shared responsibility model defines the different security and compliance responsibilities of AWS and the customer. AWS is responsible for the security of the cloud, which includes the physical infrastructure, hardware, software, and facilities that run the AWS Cloud. The customer is responsible for security in the cloud, which includes the configuration and management of the guest operating systems, applications, data, and network traffic protection

**NEW QUESTION 103**

- (Topic 1)

Which AWS service or tool helps to centrally manage billing and allow controlled access to resources across AWS accounts?

- A. AWS Identity and Access Management (IAM)
- B. AWS Organizations

- C. AWS Cost Explorer
- D. AWS Budgets

**Answer:** B

**Explanation:**

AWS Organizations helps to centrally manage billing and allow controlled access to resources across AWS accounts. AWS Organizations is a service that enables the user to consolidate multiple AWS accounts into an organization that can be managed as a single unit. AWS Organizations allows the user to create groups of accounts and apply policies to them, such as service control policies (SCPs) that specify the services and actions that users and roles can access in the accounts. AWS Organizations also enables the user to use consolidated billing, which combines the usage and charges from all the accounts in the organization into a single bill.

**NEW QUESTION 106**

- (Topic 1)

What are some advantages of using Amazon EC2 instances to host applications in the AWS Cloud instead of on premises? (Select TWO.)

- A. EC2 includes operating system patch management
- B. EC2 integrates with Amazon VPC
- C. AWS CloudTrail, and AWS Identity and Access Management (IAM)
- D. EC2 has a 100% service level agreement (SLA).
- E. EC2 has a flexible, pay-as-you-go pricing model.
- F. EC2 has automatic storage cost optimization.

**Answer:** BD

**Explanation:**

Some of the advantages of using Amazon EC2 instances to host applications in the AWS Cloud instead of on premises are:

? EC2 integrates with Amazon VPC, AWS CloudTrail, and AWS Identity and Access Management (IAM). Amazon VPC lets you provision a logically isolated section of the AWS Cloud where you can launch AWS resources in a virtual network that you define. AWS CloudTrail enables governance, compliance, operational auditing, and risk auditing of your AWS account. AWS IAM enables you to manage access to AWS services and resources securely. Therefore, the correct answer is B. You can learn more about Amazon EC2 and its integration with other AWS services from this page.

? EC2 has a flexible, pay-as-you-go pricing model. You only pay for the compute capacity you use, and you can scale up and down as needed. You can also choose from different pricing options, such as On-Demand, Savings Plans, Reserved Instances, and Spot Instances, to optimize your costs. Therefore, the correct answer is D. You can learn more about Amazon EC2 pricing from this page.

The other options are incorrect because:

? EC2 does not include operating system patch management. You are responsible for managing and maintaining your own operating systems on EC2 instances. You can use AWS Systems Manager to automate common maintenance tasks, such as applying patches, or use Amazon EC2 Image Builder to create and maintain secure images. Therefore, the incorrect answer is A.

? EC2 does not have a 100% service level agreement (SLA). The EC2 SLA guarantees 99.99% availability for each EC2 Region, not for each individual instance. Therefore, the incorrect answer is C.

? EC2 does not have automatic storage cost optimization. You are responsible for choosing the right storage option for your EC2 instances, such as Amazon Elastic Block Store (EBS) or Amazon Elastic File System (EFS), and monitoring and optimizing your storage costs. You can use AWS Cost Explorer or AWS Trusted Advisor to analyze and reduce your storage spending. Therefore, the incorrect answer is E.

**NEW QUESTION 110**

- (Topic 1)

Which feature of the AWS Cloud gives users the ability to pay based on current needs rather than forecasted needs?

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

**Answer:** B

**Explanation:**

Pay-as-you-go pricing is the feature of the AWS Cloud that gives users the ability to pay based on current needs rather than forecasted needs. Pay-as-you-go pricing means that users only pay for the AWS services and resources they use, without any upfront or long-term commitments. This allows users to scale up or down their usage depending on their changing business requirements, and avoid paying for idle or unused capacity. Pay-as-you-go pricing also enables users to benefit from the economies of scale and lower costs of AWS as they grow their business.

**NEW QUESTION 113**

- (Topic 1)

A company has an application that uses AWS services. During scaling events, the company wants to keep application usage within AWS service quotas.

Which AWS services or tools can report on the quotas so that the company can improve the reliability of the application? (Select TWO.)

- A. Service Quotas console
- B. AWS Trusted Advisor
- C. AWS Systems Manager
- D. AWS Shield
- E. AWS Cost Explorer

**Answer:** AB

**Explanation:**

The correct answers are A and B because Service Quotas console and AWS Trusted Advisor are AWS services or tools that can report on the quotas so that the company can improve the reliability of the application. Service Quotas console is an AWS tool that enables users to view and manage their quotas for AWS services from a central location. Users can use Service Quotas console to request quota increases, track quota usage, and set up alarms for approaching quota limits. AWS Trusted Advisor is an AWS service that provides real-time guidance to help users follow AWS best practices for security, performance, cost

optimization, and fault tolerance. One of the categories of checks that AWS Trusted Advisor performs is service limits, which monitors the usage of each AWS service and alerts users when they are close to reaching the default limit. The other options are incorrect because they are not AWS services or tools that can report on the quotas so that the company can improve the reliability of the application. AWS Systems Manager is an AWS service that enables users to automate operational tasks, manage configuration and compliance, and monitor system health and performance. AWS Shield is an AWS service that protects users from distributed denial of service (DDoS) attacks. AWS Cost Explorer is an AWS tool that enables users to visualize, understand, and manage their AWS costs and usage. Reference: Service Quotas, AWS Trusted Advisor FAQs

#### NEW QUESTION 114

- (Topic 1)

A company has been storing monthly reports in an Amazon S3 bucket. The company exports the report data into comma-separated values (.csv) files. A developer wants to write a simple query that can read all of these files and generate a summary report.

Which AWS service or feature should the developer use to meet these requirements with the LEAST amount of operational overhead?

- A. Amazon S3 Select
- B. Amazon Athena
- C. Amazon Redshift
- D. Amazon EC2

**Answer:** B

#### Explanation:

Amazon Athena is the AWS service that the developer should use to write a simple query that can read all of the .csv files stored in an Amazon S3 bucket and generate a summary report. Amazon Athena is an interactive query service that allows users to analyze data in Amazon S3 using standard SQL. Amazon Athena does not require any server setup or management, and users only pay for the queries they run. Amazon Athena can handle various data formats, including .csv, and can integrate with other AWS services such as Amazon QuickSight for data visualization

#### NEW QUESTION 118

- (Topic 1)

Which AWS services and features are provided to all customers at no charge? (Select TWO.)

- A. Amazon Aurora
- B. VPC
- C. Amazon SageMaker
- D. AWS Identity and Access Management (IAM)
- E. Amazon Polly

**Answer:** BD

#### Explanation:

The AWS services and features that are provided to all customers at no charge are VPC and AWS Identity and Access Management (IAM). VPC is a service that allows you to launch AWS resources in a logically isolated virtual network that you define. You can create and use a VPC at no additional charge, and you only pay for the resources that you launch in the VPC, such as EC2 instances or EBS volumes. IAM is a service that allows you to manage access and permissions to AWS resources. You can create and use IAM users, groups, roles, and policies at no additional charge, and you only pay for the AWS resources that the IAM entities access. Amazon Aurora, Amazon SageMaker, and Amazon Polly are not free services, and they charge based on the usage and features that you choose

#### NEW QUESTION 121

- (Topic 1)

Which AWS feature or resource is a deployable Amazon EC2 instance template that is prepackaged with software and security requirements?

- A. Amazon Elastic Block Store (Amazon EBS) volume
- B. AWS CloudFormation template
- C. Amazon Elastic Block Store (Amazon EBS) snapshot
- D. Amazon Machine Image (AMI)

**Answer:** D

#### Explanation:

An Amazon Machine Image (AMI) is a deployable Amazon EC2 instance template that is prepackaged with software and security requirements. It provides the information required to launch an instance, which is a virtual server in the cloud. You can use an AMI to launch as many instances as you need. You can also create your own custom AMIs or use AMIs shared by other AWS users

#### NEW QUESTION 125

- (Topic 1)

A company is migrating an application that includes an Oracle database to AWS. The company cannot rewrite the application. To which AWS service could the company migrate the database?

- A. Amazon Athena
- B. Amazon DynamoDB
- C. Amazon RDS
- D. Amazon DocumentDB (with MongoDB compatibility)

**Answer:** C

#### Explanation:

Amazon Relational Database Service (Amazon RDS) is a service that provides fully managed relational database engines. Amazon RDS supports several database engines, including Oracle, MySQL, PostgreSQL, MariaDB, SQL Server, and Amazon Aurora. Amazon RDS can be used to migrate an application that includes an Oracle database to AWS without rewriting the application, as long as the application is compatible with the Oracle version and edition supported by Amazon RDS. Amazon RDS can also provide benefits such as high availability, scalability, security, backup and restore, and performance optimization. [Amazon

### NEW QUESTION 130

- (Topic 1)

Which AWS service should a cloud practitioner use to receive real-time guidance for provisioning resources, based on AWS best practices related to security, cost optimization, and service limits?

- A. AWS Trusted Advisor
- B. AWS Config
- C. AWS Security Hub
- D. AWS Systems Manager

**Answer:** A

#### Explanation:

AWS Trusted Advisor is the AWS service that provides real-time guidance for provisioning resources, based on AWS best practices related to security, cost optimization, and service limits. AWS Trusted Advisor inspects the user's AWS environment and provides recommendations for improving performance, security, and reliability, reducing costs, and following best practices. AWS Trusted Advisor also alerts the user when they are approaching or exceeding their service limits, and helps them request limit increases<sup>3</sup>.

### NEW QUESTION 134

- (Topic 1)

Which duties are the responsibility of a company that is using AWS Lambda? (Select TWO.)

- A. Security inside of code
- B. Selection of CPU resources
- C. Patching of operating system
- D. Writing and updating of code
- E. Security of underlying infrastructure

**Answer:** AD

#### Explanation:

The duties that are the responsibility of a company that is using AWS Lambda are security inside of code and writing and updating of code. AWS Lambda is a serverless compute service that allows you to run code without provisioning or managing servers, scaling, or patching. AWS Lambda takes care of the security of the underlying infrastructure, such as the operating system, the network, and the firewall. However, the company is still responsible for the security of the code itself, such as encrypting sensitive data, validating input, and handling errors. The company is also responsible for writing and updating the code that defines the Lambda function, and choosing the runtime environment, such as Node.js, Python, or Java. AWS Lambda does not require the selection of CPU resources, as it automatically allocates them based on the memory configuration<sup>34</sup>.

### NEW QUESTION 136

- (Topic 1)

A company wants to manage access and permissions for its third-party software as a service (SaaS) applications. The company wants to use a portal where end users can access assigned AWS accounts and AWS Cloud applications. Which AWS service should the company use to meet these requirements?

- A. Amazon Cognito
- B. AWS IAM Identity Center (AWS Single Sign-On)
- C. AWS Identity and Access Management (IAM)
- D. AWS Directory Service for Microsoft Active Directory

**Answer:** B

#### Explanation:

AWS IAM Identity Center (AWS Single Sign-On) is the AWS service that the company should use to meet the requirements of managing access and permissions for its third-party SaaS applications. 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 Single Sign-On to enable your users to sign in to a user portal with their existing corporate credentials and access all of their assigned accounts and applications from one place<sup>4</sup>.

### NEW QUESTION 139

- (Topic 1)

Which AWS service is a highly available and scalable DNS web service?

- A. Amazon VPC
- B. Amazon CloudFront
- C. Amazon Route 53
- D. Amazon Connect

**Answer:** C

#### Explanation:

Amazon Route 53 is a highly available and scalable DNS web service. It is designed to give developers and businesses an extremely reliable and cost-effective way to route end users to Internet applications by translating domain names into the numeric IP addresses that computers use to connect to each other<sup>2</sup>. Amazon Route 53 also offers other features such as health checks, traffic management, domain name registration, and DNSSEC<sup>3</sup>.

### NEW QUESTION 141

- (Topic 1)

A user wants to identify any security group that is allowing unrestricted incoming SSH traffic. Which AWS service can be used to accomplish this goal?

- A. Amazon Cognito
- B. AWS Shield
- C. Amazon Macie
- D. AWS Trusted Advisor

**Answer:** D

**Explanation:**

The correct answer to the question is D because AWS Trusted Advisor is an AWS service that can be used to accomplish the goal of identifying any security group that is allowing unrestricted incoming SSH traffic. AWS Trusted Advisor is a service that provides customers with recommendations that help them follow AWS best practices. Trusted Advisor evaluates the customer's AWS environment and identifies ways to optimize their AWS infrastructure, improve security and performance, reduce costs, and monitor service quotas. One of the checks that Trusted Advisor performs is the Security Groups - Specific Ports Unrestricted check, which flags security groups that allow unrestricted access to specific ports, such as port 22 for SSH. Customers can use this check to review and modify their security group rules to restrict SSH access to only authorized sources. Reference: Security Groups - Specific Ports Unrestricted

**NEW QUESTION 144**

- (Topic 1)

A company is launching a new application in the AWS Cloud. The application will run on an Amazon EC2 instance. More EC2 instances will be needed when the workload increases.

Which AWS service or tool can the company use to launch the number of EC2 instances that will be needed to handle the workload?

- A. Elastic Load Balancing
- B. Amazon EC2 Auto Scaling
- C. AWS App2Container (A2C)
- D. AWS Systems Manager

**Answer:** B

**Explanation:**

Amazon EC2 Auto Scaling is the AWS service or tool that can help the company launch the number of EC2 instances that will be needed to handle the workload. Amazon EC2 Auto Scaling automatically adjusts the capacity of the EC2 instances based on the demand and the predefined scaling policies. Amazon EC2 Auto Scaling also helps to improve availability and reduce costs by scaling in and out as needed. For more information, see [What is Amazon EC2 Auto Scaling?](#) and [\[Getting Started with Amazon EC2 Auto Scaling\]](#).

**NEW QUESTION 146**

- (Topic 1)

Which of the following is an advantage that users experience when they move on-premises workloads to the AWS Cloud?

- A. Elimination of expenses for running and maintaining data centers
- B. Price discounts that are identical to discounts from hardware providers
- C. Distribution of all operational controls to AWS
- D. Elimination of operational expenses

**Answer:** A

**Explanation:**

The advantage that users experience when they move on-premises workloads to the AWS Cloud is: elimination of expenses for running and maintaining data centers. By moving on-premises workloads to the AWS Cloud, users can reduce or eliminate the costs associated with owning and operating physical servers, storage, network equipment, and facilities. These costs include hardware purchase, maintenance, repair, power, cooling, security, and staff. Users can also benefit from the pay-as-you-go pricing model of AWS, which allows them to pay only for the resources they use, and scale up or down as needed.

**NEW QUESTION 149**

- (Topic 1)

Which AWS service provides the ability to host a NoSQL database in the AWS Cloud?

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

**Answer:** B

**Explanation:**

Amazon DynamoDB is a fully managed NoSQL database service that provides fast and predictable performance with seamless scalability. It supports both key-value and document data models, and allows you to create tables that can store and retrieve any amount of data, and serve any level of request traffic. You can also use DynamoDB Streams to capture data modification events in DynamoDB tables.

**NEW QUESTION 154**

- (Topic 1)

A company deploys its application to multiple AWS Regions and configures automatic failover between those Regions.

Which cloud concept does this architecture represent?

- A. Security
- B. Reliability
- C. Scalability
- D. Cost optimization

**Answer:** B

**Explanation:**

Reliability is the cloud concept that this architecture represents. Reliability is 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. Deploying an application to multiple AWS Regions and configuring automatic failover between those Regions enhances the reliability of the application by reducing the impact of regional failures and increasing the availability of the application<sup>4</sup>

**NEW QUESTION 158**

- (Topic 3)

Which AWS service or tool helps users visualize, understand, and manage spending and usage over time?

- A. AWS Organizations
- B. AWS Pricing Calculator
- C. AWS Cost Explorer
- D. AWS Service Catalog

**Answer:** C

**Explanation:**

AWS Cost Explorer is the AWS service or tool that helps users visualize, understand, and manage spending and usage over time. AWS Cost Explorer is a web-based interface that allows users to access interactive graphs and tables that display their AWS costs and usage data. Users can create custom reports that analyze cost and usage data by various dimensions, such as service, region, account, tag, and more. Users can also view historical data for up to the last 12 months, forecast future costs for up to the next 12 months, and get recommendations for cost optimization. AWS Cost Explorer also provides preconfigured views that show common cost and usage scenarios, such as monthly spend by service, daily spend by linked account, and Reserved Instance utilization. Users can use AWS Cost Explorer to monitor their AWS spending and usage trends, identify cost drivers and anomalies, and optimize their resource allocation and budget planning. References: Cloud Cost Analysis - AWS Cost Explorer - AWS, Analyzing your costs with AWS Cost Explorer

**NEW QUESTION 161**

- (Topic 3)

A company wants to migrate a database from an on-premises environment to Amazon RDS. After the migration is complete, which management task will the company still be responsible for?

- A. Hardware lifecycle management
- B. Application optimization
- C. Server maintenance
- D. Power, network, and cooling provisioning

**Answer:** B

**Explanation:**

Amazon RDS is a managed database service that handles most of the common database administration tasks, such as hardware provisioning, server maintenance, backup and recovery, patching, scaling, and replication. However, Amazon RDS does not optimize the application that interacts with the database. The company is still responsible for tuning the performance, security, and availability of the application according to its business requirements and best practices<sup>12</sup>. References:

- ? What is Amazon Relational Database Service (Amazon RDS)?
- ? Perform common DBA tasks for Amazon RDS DB instances

**NEW QUESTION 163**

- (Topic 3)

Which of the following is an AWS Well-Architected Framework design principle for operational excellence in the AWS Cloud?

- A. Go global in minutes
- B. Make frequent, small, reversible changes
- C. Implement a strong foundation of identity and access management
- D. Stop spending money on hardware infrastructure for data center operations

**Answer:** B

**Explanation:**

Making frequent, small, reversible changes is one of the design principles for operational excellence in the AWS Cloud, as defined by the AWS Well-Architected Framework. This principle means that you should design your workloads to allow for rapid and safe changes, such as deploying updates, rolling back failures, and experimenting with new features. By making small and reversible changes, you can reduce the risk of errors, minimize the impact of failures, and increase the speed of recovery<sup>2</sup>. References: <sup>2</sup>: AWS Documentation - AWS Well-Architected Framework - Operational Excellence Pillar

**NEW QUESTION 166**

- (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 169

- (Topic 3)

A company is running and managing its own Docker environment on Amazon EC2 instances. The company wants an alternative to help manage cluster size, scheduling, and environment maintenance.

Which AWS service meets these requirements?

- A. AWS Lambda
- B. Amazon RDS
- C. AWS Fargate
- D. Amazon Athena

**Answer: C**

#### Explanation:

AWS Fargate is a serverless compute engine for containers that works with both Amazon Elastic Container Service (Amazon ECS) and Amazon Elastic Kubernetes Service (Amazon EKS). AWS Fargate allows you to run containers without having to manage servers or clusters of Amazon EC2 instances. With AWS Fargate, you only pay for the compute resources you use to run your containers, and you don't need to worry about scaling, patching, securing, or maintaining the underlying infrastructure. AWS Fargate simplifies the deployment and management of containerized applications, and enables you to focus on building and running your applications instead of managing the infrastructure. References: AWS Fargate, What is AWS Fargate?

#### NEW QUESTION 170

- (Topic 2)

Which option is a pillar of the AWS Well-Architected Framework?

- A. Patch management
- B. Cost optimization
- C. Business technology strategy
- D. Physical and environmental controls

**Answer: B**

#### Explanation:

The AWS Well-Architected Framework helps you understand the pros and cons of decisions you make while building systems on AWS. By using the Framework, you will learn architectural best practices for designing and operating reliable, secure, efficient, and cost-effective systems in the cloud. The Framework consists of five pillars: operational excellence, security, reliability, performance efficiency, and cost optimization<sup>2</sup>.

#### NEW QUESTION 173

- (Topic 2)

A company wants to securely store Amazon RDS database credentials and automatically rotate user passwords periodically.

Which AWS service or capability will meet these requirements?

- A. Amazon S3
- B. AWS Systems Manager Parameter Store
- C. AWS Secrets Manager
- D. AWS CloudTrail

**Answer: C**

#### Explanation:

AWS Secrets Manager is a service that helps you protect access to your applications, services, and IT resources. This service enables you to easily rotate, manage, and retrieve database credentials, API keys, and other secrets throughout their lifecycle<sup>1</sup>. Amazon S3 is a storage service that does not offer automatic rotation of credentials. AWS Systems Manager Parameter Store is a service that provides secure, hierarchical storage for configuration data management and secrets management<sup>2</sup>, but it does not offer automatic rotation of credentials. AWS CloudTrail is a service that enables governance, compliance, operational auditing, and risk auditing of your AWS account<sup>3</sup>, but it does not store or rotate credentials.

#### NEW QUESTION 174

- (Topic 2)

Which AWS service can a company use to securely store and encrypt passwords for a database?

- A. AWS Shield
- B. AWS Secrets Manager
- C. AWS Identity and Access Management (IAM)
- D. Amazon Cognito

**Answer: B**

**Explanation:**

AWS Secrets Manager is an AWS service that can be used to securely store and encrypt passwords for a database. It allows users to manage secrets, such as database credentials, API keys, and tokens, in a centralized and secure way. It also provides features such as automatic rotation, fine-grained access control, and auditing. AWS Shield is an AWS service that provides protection against Distributed Denial of Service (DDoS) attacks for AWS resources and services. It does not store or encrypt passwords for a database. AWS Identity and Access Management (IAM) is an AWS service that allows users to manage access to AWS resources and services. It can be used to create users, groups, roles, and policies that control who can do what in AWS. It does not store or encrypt passwords for a database. Amazon Cognito is an AWS service that provides user identity and data synchronization for web and mobile applications. It can be used to authenticate and authorize users, manage user profiles, and sync user data across devices. It does not store or encrypt passwords for a database.

**NEW QUESTION 177**

- (Topic 2)

Which design principles should a company apply to AWS Cloud workloads to maximize sustainability and minimize environmental impact? (Select TWO.)

- A. Maximize utilization of Amazon EC2 instances.
- B. Minimize utilization of Amazon EC2 instances.
- C. Minimize usage of managed services.
- D. Force frequent application reinstallations by users.
- E. Reduce the need for users to reinstall applications.

**Answer:** AE

**Explanation:**

To maximize sustainability and minimize environmental impact, a company should apply the following design principles to AWS Cloud workloads: maximize utilization of Amazon EC2 instances and reduce the need for users to reinstall applications. Maximizing utilization of Amazon EC2 instances means that the company can optimize the performance and efficiency of their compute resources, and avoid wasting energy and money on idle or underutilized instances. The company can use features such as Amazon EC2 Auto Scaling, Amazon EC2 Spot Instances, and AWS Compute Optimizer to automatically adjust the number and type of instances based on demand, cost, and performance. Reducing the need for users to reinstall applications means that the company can minimize the amount of data and bandwidth required to deliver their applications to users, and avoid unnecessary downloads and updates that consume energy and resources. The company can use services such as Amazon CloudFront, AWS AppStream 2.0, and AWS Amplify to deliver their applications faster, more securely, and more efficiently to users across the globe. Minimizing utilization of Amazon EC2 instances, minimizing usage of managed services, and forcing frequent application reinstallations by users are not design principles that would maximize sustainability and minimize environmental impact. Minimizing utilization of Amazon EC2 instances would reduce the performance and efficiency of the compute resources, and potentially increase the costs and complexity of the cloud workloads. Minimizing usage of managed services would increase the operational overhead and responsibility of the company, and potentially expose them to more security and reliability risks. Forcing frequent application reinstallations by users would increase the amount of data and bandwidth required to deliver the applications to users, and potentially degrade the user experience and satisfaction.

**NEW QUESTION 182**

- (Topic 2)

Which AWS solution provides the ability for a company to run AWS services in the company's on-premises data center?

- A. AWS Direct Connect
- B. AWS Outposts
- C. AWS Systems Manager hybrid activations
- D. AWS Storage Gateway

**Answer:** B

**Explanation:**

AWS Outposts is a fully managed service that extends AWS infrastructure, AWS services, APIs, and tools to virtually any datacenter, co-location space, or on-premises facility for a truly consistent hybrid experience. AWS Outposts enables you to run AWS services in your on-premises data center1.

**NEW QUESTION 187**

- (Topic 2)

A retail company has recently migrated its website to AWS. The company wants to ensure that it is protected from SQL injection attacks. The website uses an Application Load Balancer to distribute traffic to multiple Amazon EC2 instances.

Which AWS service or feature can be used to create a custom rule that blocks SQL injection attacks?

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

**Answer:** B

**Explanation:**

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. AWS WAF gives you control over how traffic reaches your applications by enabling you to create security rules that block common attack patterns, such as SQL injection or cross-site scripting, and rules that filter out specific traffic patterns you define2. You can use AWS WAF to create a custom rule that blocks SQL injection attacks on your website.

**NEW QUESTION 191**

- (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 193**

- (Topic 2)

Which AWS services allow users to monitor and retain records of account activities that include governance, compliance, and auditing?  
(Select TWO.)

- A. Amazon CloudWatch
- B. AWS CloudTrail
- C. Amazon GuardDuty
- D. AWS Shield
- E. AWS WAF

**Answer:** AB

**Explanation:**

Amazon CloudWatch and AWS CloudTrail are the AWS services that allow users to monitor and retain records of account activities that include governance, compliance, and auditing. Amazon CloudWatch is a service that collects and tracks metrics, collects and monitors log files, and sets alarms. AWS CloudTrail is a service that enables governance, compliance, operational auditing, and risk auditing of your AWS account. Amazon GuardDuty, AWS Shield, and AWS WAF are AWS services that provide security and protection for AWS resources, but they do not monitor and retain records of account activities. These concepts are explained in the AWS Cloud Practitioner Essentials course<sup>3</sup>.

**NEW QUESTION 194**

- (Topic 2)

A company wants to use Amazon EC2 instances to run a stateless and restartable process after business hours.  
Which AWS service provides DNS resolution?

- A. Amazon CloudFront
- B. Amazon VPC
- C. Amazon Route 53
- D. AWS Direct Connect

**Answer:** C

**Explanation:**

Amazon Route 53 is the AWS service that provides DNS resolution. DNS (Domain Name System) is a service that translates domain names into IP addresses. Amazon Route 53 is a highly available and scalable cloud DNS service that offers domain name registration, DNS routing, and health checking. Amazon Route 53 can route the traffic to various AWS services, such as Amazon EC2, Amazon S3, and Amazon CloudFront. Amazon Route 53 can also integrate with other AWS services, such as AWS Certificate Manager, AWS Shield, and AWS WAF. For more information, see [What is Amazon Route 53?] and [Amazon Route 53 Features].

**NEW QUESTION 199**

- (Topic 2)

A company is using AWS Organizations to configure AWS accounts.

A company is planning its migration to the AWS Cloud. The company is identifying its capability gaps by using the AWS Cloud Adoption Framework (AWS CAF) perspectives.

Which phase of the cloud transformation journey includes these identification activities?

- A. Envision
- B. Align
- C. Scale
- D. Launch

**Answer:** A

**Explanation:**

The Envision phase of the cloud transformation journey is where the company defines its vision, business drivers, and desired outcomes for the cloud adoption. The company also identifies its capability gaps by using the AWS Cloud Adoption Framework (AWS CAF) perspectives, which are business, people, governance, platform, security, and operations<sup>2</sup>.

**NEW QUESTION 203**

- (Topic 2)

Which AWS service provides the SIMPLEST way for the company to establish a website on AWS?

- A. Amazon Elastic File System (Amazon EFS)
- B. AWS Elastic Beanstalk
- C. AWS Lambda
- D. Amazon Lightsail

**Answer:** D

**Explanation:**

Amazon Lightsail is an easy-to-use cloud platform that offers you everything needed to build an application or website, plus a cost-effective, monthly plan. Whether you're new to the cloud or looking to get on the cloud quickly with AWS infrastructure you trust, we've got you covered. Lightsail provides the simplest way for the company to establish a website on AWS.

**NEW QUESTION 204**

- (Topic 2)

Which credential allows programmatic access to AWS resources for use from the AWS CLI or the AWS API?

- A. User name and password
- B. Access keys
- C. SSH public keys
- D. AWS Key Management Service (AWS KMS) keys

**Answer:** B

**Explanation:**

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 using the AWS CLI or AWS API<sup>1</sup>. User name and password are credentials that you use to sign in to the AWS Management Console or the AWS Management Console mobile app<sup>2</sup>. SSH public keys are credentials that you use to authenticate with EC2 instances that are launched from certain Linux AMIs<sup>3</sup>. AWS Key Management Service (AWS KMS) keys are customer master keys (CMKs) that you use to encrypt and decrypt your data and to control access to your data across AWS services and in your applications<sup>4</sup>.

**NEW QUESTION 205**

- (Topic 2)

Which AWS service or tool provides recommendations to help users get rightsized Amazon EC2 instances based on historical workload usage data?

- A. AWS Pricing Calculator
- B. AWS Compute Optimizer
- C. AWS App Runner
- D. AWS Systems Manager

**Answer:** B

**Explanation:**

The AWS service or tool that provides recommendations to help users get rightsized Amazon EC2 instances based on historical workload usage data is AWS Compute Optimizer. AWS Compute Optimizer is a service that analyzes the configuration and performance of the AWS resources, such as Amazon EC2 instances, and provides recommendations for optimal resource types and sizes based on the workload patterns and metrics. AWS Compute Optimizer helps users improve the performance, availability, and cost efficiency of their AWS resources. AWS Pricing Calculator, AWS App Runner, and AWS Systems Manager are not the best services or tools to use for this purpose. AWS Pricing Calculator is a tool that helps users estimate the cost of using AWS services based on their requirements and preferences. AWS App Runner is a service that helps users easily and quickly deploy web applications and APIs without managing any infrastructure. AWS Systems Manager is a service that helps users automate and manage the configuration and operation of their AWS resources and applications<sup>34</sup>.

**NEW QUESTION 206**

- (Topic 2)

A company needs to host a highly available application in the AWS Cloud. The application runs infrequently for short periods of time. Which AWS service will meet these requirements with the LEAST amount of operational overhead?

- A. Amazon EC2
- B. AWS Fargate
- C. AWS Lambda
- D. Amazon Aurora

**Answer:** C

**Explanation:**

The AWS service that will meet the requirements of the company that needs to host a highly available application in the AWS Cloud that runs infrequently for short periods of time with the least amount of operational overhead is AWS Lambda. AWS Lambda is a serverless compute service that allows customers to run code without provisioning or managing servers. The company can use AWS Lambda to create and deploy their application as functions that are triggered by events, such as API calls, messages, or schedules. AWS Lambda automatically scales the compute resources based on the demand, and customers only pay for the compute time they consume. AWS Lambda also simplifies the management and maintenance of the application, as customers do not need to worry about the underlying infrastructure, security, or availability. Amazon EC2, AWS Fargate, and Amazon Aurora are not the best services to use for this purpose. Amazon EC2 is a service that provides scalable compute capacity in the cloud, and allows customers to launch and run virtual servers, called instances, with a variety of operating systems, configurations, and specifications. Amazon EC2 requires customers to provision and manage the instances, and pay for the instance hours they use, regardless of the application usage. AWS Fargate is a serverless compute engine for containers that allows customers to run containerized applications without managing servers or clusters. AWS Fargate requires customers to specify the amount of CPU and memory resources for each container, and pay for the resources they allocate, regardless of the application usage.

Amazon Aurora is a fully managed relational database service that provides high performance, availability, and compatibility. Amazon Aurora is not a compute service, and it is not suitable for hosting an application that runs infrequently for short periods of time<sup>12</sup>.

**NEW QUESTION 207**

- (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 208**

- (Topic 2)

A company must store call recordings for 6 years. The storage system should be highly durable and cost-effective. Which AWS service meets these requirements?

- A. AWS Snowball
- B. Amazon S3
- C. AWS Storage Gateway
- D. Amazon Kinesis

**Answer:** B

**Explanation:**

Amazon S3 is a service that provides highly durable and cost-effective object storage for a variety of use cases, including backup and archive, big data analytics, disaster recovery, and cloud applications. Amazon S3 offers 99.999999999% (11 9's) of durability, meaning that data is designed to withstand the loss of two facilities concurrently. Amazon S3 also offers several storage classes with different price and performance characteristics, such as S3 Glacier and S3 Glacier Deep Archive, which are ideal for long-term archival of data that is rarely accessed. AWS Snowball, AWS Storage Gateway, and Amazon Kinesis are not designed to provide the same level of durability and cost-effectiveness as Amazon S3 for storing call recordings for 6 years. Source: Amazon S3

**NEW QUESTION 213**

- (Topic 2)

A user is moving a workload from a local data center to an architecture that is distributed between the local data center and the AWS Cloud. Which type of migration is this?

- A. On-premises to cloud native
- B. Hybrid to cloud native
- C. On-premises to hybrid
- D. Cloud native to hybrid

**Answer:** C

**Explanation:**

C is correct because moving a workload from a local data center to an architecture that is distributed between the local data center and the AWS Cloud is an example of an on-premises to hybrid migration. A 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. A is incorrect because on-premises to cloud native migration is the process of moving a workload from a local data center to an architecture that is fully hosted and managed on the AWS Cloud. B is incorrect because hybrid to cloud native migration is the process of moving a workload from an architecture that is distributed between the local data center and the AWS Cloud to an architecture that is fully hosted and managed on the AWS Cloud. D is incorrect because cloud native to hybrid migration is the process of moving a workload from an architecture that is fully hosted and managed on the AWS Cloud to an architecture that is distributed between the local data center and the AWS Cloud.

**NEW QUESTION 214**

- (Topic 2)

A user discovered that an Amazon EC2 instance is missing an Amazon Elastic Block Store (Amazon EBS) data volume. The user wants to determine when the EBS volume was removed. Which AWS service will provide this information?

- A. AWS Config
- B. AWS Trusted Advisor
- C. Amazon Timestream
- D. Amazon QuickSight

**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. AWS Config can help you determine when an EBS volume was removed from an EC2 instance by providing a timeline of configuration changes and compliance status. AWS Trusted Advisor, Amazon Timestream, and Amazon QuickSight do not provide the same level of configuration tracking and auditing as AWS Config. Source: AWS Config

**NEW QUESTION 216**

- (Topic 2)

A company needs to launch an Amazon EC2 instance. Which of the following can the company use during the launch process to configure the root volume of the EC2 instance?

- A. Amazon EC2 Auto Scaling
- B. Amazon Data Lifecycle Manager (Amazon DLM)

- C. Amazon Machine Image (AMI)
- D. Amazon Elastic Block Store (Amazon EBS) volume

**Answer:** C

**Explanation:**

Amazon Machine Image (AMI) is the option that the company can use during the launch process to configure the root volume of the EC2 instance. An AMI is a template that contains the software configuration, such as the operating system, applications, and settings, required to launch an EC2 instance. An AMI also specifies the volume size and type of the root device for the instance. The company can choose an AMI provided by AWS, the AWS Marketplace, or the AWS community, or create a custom AMI. For more information, see [Amazon Machine Images (AMI)] and [Launching an Instance Using the Launch Instance Wizard].

**NEW QUESTION 218**

- (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 220**

- (Topic 2)

A developer wants to use an Amazon S3 bucket to store application logs that contain sensitive data.

Which AWS service or feature should the developer use to restrict read and write access to the S3 bucket?

- A. Security groups
- B. Amazon CloudWatch
- C. AWS CloudTrail
- D. ACLs

**Answer:** D

**Explanation:**

ACLs are an AWS service or feature that the developer can use to restrict read and write access to the S3 bucket. ACLs are access control lists that grant basic permissions to other AWS accounts or predefined groups. They can be used to grant read or write access to an S3 bucket or an object3. Security groups are virtual firewalls that control the inbound and outbound traffic for Amazon EC2 instances. They are not a service or feature that can be used to restrict access to an S3 bucket. Amazon CloudWatch is a service that provides monitoring and observability for AWS resources and applications. It can be used to collect and analyze metrics, logs, events, and alarms. It is not a service or feature that can be used to restrict access to an S3 bucket. AWS CloudTrail is a service that provides governance, compliance, and audit for AWS accounts and resources. It can be used to track and record the API calls and user activity in AWS. It is not a service or feature that can be used to restrict access to an S3 bucket.

**NEW QUESTION 221**

- (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 workloads34

**NEW QUESTION 223**

- (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. 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 228**

- (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 229**

- (Topic 2)

A company is running an order processing system on Amazon EC2 instances. The company wants to migrate microservices-based application. Which combination of AWS services can the application use to meet these requirements? (Select TWO.)

- A. Amazon Simple Queue Service (Amazon SQS)
- B. AWS Lambda
- C. AWS Migration Hub
- D. AWS AppSync
- E. AWS Application Migration Service

**Answer: AB**

**Explanation:**

The combination of AWS services that the application can use to migrate to a microservices-based application are Amazon Simple Queue Service (Amazon SQS) and AWS Lambda. Amazon SQS is a fully managed message queuing service that enables customers to decouple and scale microservices, distributed systems, and serverless applications. The application can use Amazon SQS to send, store, and receive messages between the microservices, ensuring that each message is processed only once and in the right order. AWS Lambda is a serverless compute service that allows customers to run code without provisioning or managing servers. The application can use AWS Lambda to create and deploy microservices as functions that are triggered by events, such as messages from Amazon SQS. AWS Migration Hub, AWS AppSync, and AWS Application Migration Service are not the best services to use for migrating to a microservices-based application. AWS Migration Hub is a service that provides a single location to track the progress of application migrations across multiple AWS and partner solutions. AWS AppSync is a service that simplifies the development of GraphQL APIs for real-time and offline data synchronization. AWS Application Migration Service is a service that enables customers to migrate their on-premises applications to AWS without making any changes to the applications, servers, or databases.

**NEW QUESTION 233**

- (Topic 2)

A new AWS user who has little cloud experience wants to build an application by using AWS services. The user wants to learn how to implement specific AWS services from other customer examples. The user also wants to ask questions to AWS experts. Which AWS service or resource will meet these requirements?

- A. AWS Online Tech Talks
- B. AWS documentation
- C. AWS Marketplace
- D. AWS Health Dashboard

**Answer: A**

**Explanation:**

AWS Online Tech Talks are online presentations that cover a broad range of topics at varying technical levels and provide a live Q&A session with AWS experts. They are a great resource for new AWS users who want to learn how to implement specific AWS services from other customer examples and ask questions to AWS experts. AWS documentation, AWS Marketplace, and AWS Health Dashboard do not offer the same level of interactivity and guidance as AWS Online Tech Talks. Source: AWS Online Tech Talks

**NEW QUESTION 234**

- (Topic 2)

Which options are common stakeholders for the AWS Cloud Adoption Framework (AWS CAF) platform perspective? (Select TWO.)

- A. Chief financial officers (CFOs)
- B. IT architects
- C. Chief information officers (CIOs)
- D. Chief data officers (CDOs)
- E. Engineers

**Answer:** BE

**Explanation:**

The common stakeholders for the AWS Cloud Adoption Framework (AWS CAF) platform perspective are IT architects and engineers. The AWS CAF is a guidance that helps organizations design and travel an accelerated path to successful cloud adoption. The AWS CAF organizes the cloud adoption process into six areas of focus, called perspectives, which are business, people, governance, platform, security, and operations. Each perspective is divided into capabilities, which are further divided into skills and responsibilities. The platform perspective focuses on the provisioning and management of the cloud infrastructure and services that support the business applications. The platform perspective capabilities are design, implementation, and optimization. The stakeholders for the platform perspective are the IT architects and engineers who are responsible for designing, implementing, and optimizing the cloud platform. Chief financial officers (CFOs), chief information officers (CIOs), and chief data officers (CDOs) are not the common stakeholders for the AWS CAF platform perspective. CFOs are the common stakeholders for the AWS CAF business perspective, which focuses on the value realization of the cloud adoption. CIOs are the common stakeholders for the AWS CAF governance perspective, which focuses on the alignment of the IT strategy and processes with the business strategy and goals. CDOs are the common stakeholders for the AWS CAF security perspective, which focuses on the protection of the information assets and systems in the cloud.

**NEW QUESTION 236**

- (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 239**

- (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 244**

- (Topic 2)

How should the company deploy the application to meet these requirements?

- A. In a single Availability Zone
- B. On AWS Direct Connect
- C. On Reserved Instances
- D. In multiple Availability Zones

**Answer:** D

**Explanation:**

Deploying the application in multiple Availability Zones is the best way to ensure high availability for the application. Availability Zones are isolated locations within an AWS Region that are engineered to be fault-tolerant from failures in other Availability Zones. By deploying the application in multiple Availability Zones, the company can reduce the impact of outages and increase the resilience of the application. Deploying the application in a single Availability Zone, on AWS Direct Connect, or on Reserved Instances does not provide the same level of high availability as deploying the application in multiple Availability Zones. Source: Availability Zones

**NEW QUESTION 247**

- (Topic 2)

A company is running workloads for multiple departments within a single VPC. The company needs to be able to bill each department for its resource usage. Which action should the company take to accomplish this goal with the LEAST operational overhead?

- A. Add a department tag to each resource and configure cost allocation tags.
- B. Move each department resource to its own VPC.
- C. Move each department resource to its own AWS account.
- D. Use AWS Organizations to get a billing report for each department.

**Answer:** A

**Explanation:**

Adding a department tag to each resource and configuring cost allocation tags is an action that can help you accomplish the goal of billing each department for its resource usage with the least operational overhead. Tags are simple labels consisting of a key and an optional value that you can assign to AWS resources. You can use tags to organize your resources and track your AWS costs on a detailed level. Cost allocation tags enable you to track your AWS costs on a detailed level. After you activate cost allocation tags, AWS uses the cost allocation tags to organize your resource costs on your cost allocation report, to make it easier for you to categorize and track your AWS costs<sup>2</sup>. Moving each department resource to its own VPC or its own AWS account is an action that can help you isolate and control the resources for each department, but it would incur more operational overhead than using tags. Using AWS Organizations to get a billing report for each department is an action that can help you consolidate billing and payment across multiple AWS accounts, but it would not help you bill each department for its resource usage within a single VPC.

**NEW QUESTION 249**

- (Topic 2)

A company plans to migrate its on-premises workload to AWS. Before the migration, the company needs to estimate its future AWS service costs. Which AWS service or tool should the company use to meet this requirement?

- A. AWS Trusted Advisor
- B. AWS Budgets
- C. AWS Pricing Calculator
- D. AWS Cost Explorer

**Answer:** C

**Explanation:**

AWS Pricing Calculator is the AWS service or tool that the company should use to estimate its future AWS service costs before the migration. AWS Pricing Calculator is a web-based tool that allows the company to create cost estimates for various AWS services and scenarios. AWS Pricing Calculator helps the company to compare the costs of running the workload on premises versus on AWS, and to optimize the costs by choosing the best options for the workload. AWS Pricing Calculator also provides a detailed breakdown of the cost components and a downloadable report. For more information, see [AWS Pricing Calculator] and [Getting Started with AWS Pricing Calculator].

**NEW QUESTION 253**

- (Topic 2)

Which AWS services can a company use to host and run a MySQL database? (Select TWO.)

- A. Amazon RDS
- B. Amazon DynamoDB
- C. Amazon S3
- D. Amazon EC2
- E. Amazon MQ

**Answer:** AD

**Explanation:**

Amazon RDS and Amazon EC2 are two AWS services that you can use to host and run a MySQL database. Amazon RDS is a service that makes it easy to set up, operate, and scale a relational database in the cloud. You can use Amazon RDS to launch a MySQL database instance and let Amazon RDS manage common database tasks such as backups, patching, scaling, and replication<sup>6</sup>. Amazon EC2 is a service that provides secure, resizable compute capacity in the cloud. You can use Amazon EC2 to launch a virtual server and install MySQL software on it. You have complete control over your database configuration, but you are responsible for managing and maintaining the database software and the underlying infrastructure<sup>7</sup>. Amazon DynamoDB is a key-value and document database that delivers single-digit millisecond performance at any scale. Amazon S3 is an object storage service that offers industry-leading scalability, data availability, security, and performance. Amazon MQ is a managed message broker service for Apache ActiveMQ. None of these services can help you host and run a MySQL database.

**NEW QUESTION 256**

- (Topic 2)

Which AWS service is designed to help users build conversational interfaces into applications using voice and text?

- A. Amazon Lex
- B. Amazon Transcribe
- C. Amazon Comprehend

D. Amazon Timestream

**Answer:** A

**Explanation:**

A is correct because Amazon Lex is the AWS service that helps users build conversational interfaces into applications using voice and text. B is incorrect because Amazon Transcribe is the AWS service that helps users convert speech to text. C is incorrect because Amazon Comprehend is the AWS service that helps users analyze text using natural language processing. D is incorrect because Amazon Timestream is the AWS service that helps users collect, store, and process time series data.

**NEW QUESTION 261**

- (Topic 2)

What is a characteristic of Convertible Reserved Instances (RIs)?

- A. Users can exchange Convertible RIs for other Convertible RIs from a different instance family.
- B. Users can exchange Convertible RIs for other Convertible RIs in different AWS Regions.
- C. Users can sell and buy Convertible RIs on the AWS Marketplace.
- D. Users can shorten the term of their Convertible RIs by merging them with other Convertible RIs.

**Answer:** A

**Explanation:**

Convertible Reserved Instances (RIs) are a type of Reserved Instance that allow you to change the attributes of the RI as long as the exchange results in the creation of Reserved Instances of equal or greater value. You can exchange Convertible RIs for other Convertible RIs from a different instance family, size, platform, tenancy, or scope (Region or Availability Zone).

**NEW QUESTION 264**

- (Topic 2)

A company is hosting a web application on Amazon EC2 instances. The company wants to implement custom conditions to filter and control inbound web traffic. Which AWS service will meet these requirements?

- A. Amazon GuardDuty
- B. AWS WAF
- C. Amazon Macie
- D. AWS Shield

**Answer:** B

**Explanation:**

The AWS service that will meet the requirements of the company that is hosting a web application on Amazon EC2 instances and wants to implement custom conditions to filter and control inbound web traffic is AWS WAF. AWS WAF is a web application firewall that helps protect web applications from common web exploits that could affect availability, compromise security, or consume excessive resources. The company can use AWS WAF to create custom rules that block malicious requests that match certain patterns, such as SQL injection or cross-site scripting. AWS WAF can be applied to web applications that are behind an Application Load Balancer, Amazon CloudFront, or Amazon API Gateway. Amazon GuardDuty, Amazon Macie, and AWS Shield 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. Amazon Macie is a data security and data privacy service that uses machine learning and pattern matching to discover, classify, and protect sensitive data stored in Amazon S3. AWS Shield is a managed distributed denial of service (DDoS) protection service that safeguards web applications running on AWS. These services are more useful for detecting and preventing different types of threats and attacks, rather than filtering and controlling inbound web traffic based on custom conditions.

**NEW QUESTION 266**

- (Topic 2)

Which AWS service or tool offers consolidated billing?

- A. AWS Artifact
  - B. AWS Budgets
  - C. AWS Organizations
  - D. AWS Trusted Advisor
- A company wants to limit its employees' AWS access to a portfolio of predefined AWS resources.

**Answer:** C

**Explanation:**

AWS Organizations is a service that enables you to consolidate multiple AWS accounts into an organization that you create and centrally manage. With AWS Organizations, you can create a single payment method for all the AWS accounts in your organization through consolidated billing. Consolidated billing enables you to see a combined view of AWS charges incurred by all accounts in your organization, as well as get a detailed cost report for each individual AWS account associated with your organization. AWS Artifact is a service that provides on-demand access to AWS' security and compliance reports and select online agreements. AWS Budgets is a service that enables you to plan your service usage, service costs, and instance reservations. AWS Trusted Advisor is a service that provides real-time guidance to help you provision your resources following AWS best practices. None of these services or tools offer consolidated billing.

**NEW QUESTION 271**

- (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 274**

- (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 279**

- (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 280**

- (Topic 2)

A company is building an application that will receive millions of database queries each second. The company needs the data store for the application to scale to meet these needs.

Which AWS service will meet this requirement?

- A. Amazon DynamoDB
- B. AWS Cloud9
- C. Amazon ElastiCache for Memcached
- D. Amazon Neptune

**Answer:** A

**Explanation:**

Amazon DynamoDB is the AWS service that will meet the requirement of building an application that will receive millions of database queries each second. Amazon DynamoDB is a fully managed NoSQL database service that provides fast and consistent performance, scalability, and durability. Amazon DynamoDB can handle any level of request traffic and automatically scale up or down the capacity based on the demand. Amazon DynamoDB also supports in-memory caching with Amazon DynamoDB Accelerator (DAX) to improve the response time and reduce the cost. For more information, see [What is Amazon DynamoDB?](#) and [Amazon DynamoDB Features](#).

**NEW QUESTION 283**

.....

## 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

### AWS-Certified-Cloud-Practitioner Practice Exam Features:

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

**100% Actual & Verified — Instant Download, Please Click**  
[Order The AWS-Certified-Cloud-Practitioner Practice Test Here](#)