CERT EMPIRE

AWS Cloud Practitioner CLF-C02 Exam Questions

Total Questions: 770+
Demo Questions: 29
Version: Updated for 2025

Prepared and Verified by Cert Empire – Your Trusted IT Certification Partner

For Access to the full set of Updated Questions – Visit: CLF-C02 Exam Dumps by Cert Empire

Question: 1

A retail company is migrating its IT infrastructure applications from on premises to the AWS Cloud. Which costs will the company eliminate with this migration? (Select TWO.)

A: Cost of data center operations

B: Cost of application licensing

C: Cost of marketing campaigns

D: Cost of physical server hardware

E: Cost of network management

Correct Answer:

A, D

Explanation:

Migrating from an on-premises infrastructure to the AWS Cloud allows a company to eliminate the capital expenditure (CapEx) and operational costs associated with managing physical hardware and data centers. The company no longer needs to purchase, maintain, or replace physical servers. Similarly, the operational costs of running a data center—such as power, cooling, physical security, and facility maintenance—are eliminated, as AWS manages the physical infrastructure. This shifts spending from a CapEx model to a variable operational expenditure (OpEx) model.

Why Incorrect Options are Wrong:

B: Application licensing costs are typically not eliminated; they may be restructured or transferred to the cloud using models like Bring Your Own License (BYOL).

C: Marketing campaign costs are a business expense entirely separate from IT infrastructure hosting and are unaffected by a cloud migration.

E: Network management is not eliminated; the responsibility shifts from managing physical network devices to configuring and managing virtual network resources within AWS (e.g., VPCs, subnets).

References:

1. AWS Cloud Adoption Framework (CAF) - Business Perspective: The Business Perspective of the AWS CAF focuses on ensuring that IT aligns with business needs. It

explicitly addresses the reduction of capital expenditures and the financial benefits of moving away from on-premises data centers. See the "Cost Optimization" capability.

AWS. (n.d.). AWS Cloud Adoption Framework. AWS Whitepapers & Guides. Retrieved from https://docs.aws.amazon.com/whitepapers/latest/aws-cloud-adoption-framework/business-perspective.html

2. AWS Total Cost of Ownership (TCO) Concept: AWS documentation on TCO directly compares the costs of running infrastructure on-premises versus on AWS. It highlights the elimination of server hardware costs and data center operations costs (like power and cooling) as key areas of savings.

AWS. (n.d.). AWS Total Cost of Ownership (TCO) Calculators. AWS Economics. Retrieved from https://aws.amazon.com/tco/

3. Overview of Amazon Web Services Whitepaper: This foundational document explains the value proposition of AWS, including "Trade capital expense for variable expense," which directly refers to avoiding the large upfront costs of physical hardware and data center infrastructure.

AWS. (2023, September). Overview of Amazon Web Services. AWS Whitepapers & Guides, p. 6. Retrieved from https://d1.awsstatic.com/whitepapers/aws-overview.pdf

Question: 2

Which AWS services can use AWS WAF to protect against common web exploitations? (Select TWO.)

A: Amazon Route 53

B: Amazon CloudFront

C: AWS Transfer Family

D: AWS Site-to-Site VPN

E: Amazon API Gateway

Correct Answer:

B, E

Explanation:

AWS WAF is a web application firewall that helps protect web applications or APIs against common web exploits that may affect availability, compromise security, or consume excessive resources. AWS WAF can be deployed on specific AWS services that serve web content or APIs. According to official AWS documentation, AWS WAF integrates directly with Amazon CloudFront, Application Load Balancer (ALB), Amazon API Gateway, and AWS AppSync, among others. Therefore, CloudFront, a content delivery network, and API Gateway, a service for managing APIs, are the correct services from the list that can be protected by AWS WAF.

Why Incorrect Options are Wrong:

A: Amazon Route 53: This is a DNS service. It routes end users to internet applications but does not host the applications themselves, and WAF does not directly protect it.

C: AWS Transfer Family: This service manages file transfers using protocols like SFTP and FTPS, not the HTTP/S traffic that AWS WAF is designed to inspect.

D: AWS Site-to-Site VPN: This is a network connectivity service operating at the network layer (IPsec), whereas AWS WAF operates at the application layer (Layer 7).

References:

AWS WAF Developer Guide. (2024). What is AWS WAF? AWS Documentation. "You can deploy AWS WAF on AWS resources such as Amazon CloudFront, Application Load

Balancer, Amazon API Gateway, and AWS AppSync." URL: https://docs.aws.amazon.com/waf/latest/developerguide/what-is-aws-waf.html

AWS Cloud Practitioner Essentials. (2023). Module 6: Security and Compliance. AWS Skill Builder. This module describes AWS WAF as a service that protects applications from common web exploits and is associated with services like CloudFront and Application Load Balancer.

Question: 3

A company needs a centralized, secure way to create and manage cryptographic keys. The company will use the keys across a wide range of AWS services and applications. The company needs to track and document when the keys are created, used, and deleted. Which AWS service or feature will meet these requirements?

A: AWS Secrets Manager

B: AWS License Manager

C: AWS Systems Manager Parameter Store

D: AWS Key Management Service (AWS KMS)

Correct Answer:

D

Explanation:

AWS Key Management Service (AWS KMS) is a managed service designed specifically for creating and controlling cryptographic keys used to encrypt data. It provides a centralized, secure location for key management. AWS KMS is integrated with AWS CloudTrail, which provides logs of all API requests, including key creation, usage, and deletion. This directly meets the company's requirement for a centralized system to create, manage, and audit the lifecycle of cryptographic keys across various AWS services and applications.

Why Incorrect Options are Wrong:

A: AWS Secrets Manager: This service helps protect secrets needed to access applications, services, and IT resources. It manages secrets, not the underlying cryptographic keys themselves.

B: AWS License Manager: This service is used for managing software licenses from various vendors across AWS and on-premises environments, which is unrelated to cryptographic key management.

C: AWS Systems Manager Parameter Store: This service provides secure storage for configuration data and secrets. While it uses encryption, its primary function is not the management of the cryptographic keys.

References:

AWS Key Management Service (AWS KMS): AWS Documentation. (2024). What is AWS Key Management Service? AWS Key Management Service Developer Guide. Retrieved from https://docs.aws.amazon.com/kms/latest/developerguide/overview.html

Logging AWS KMS API calls with AWS CloudTrail: AWS Documentation. (2024). Logging AWS KMS API calls with AWS CloudTrail. AWS Key Management Service Developer Guide. Retrieved from https://docs.aws.amazon.com/kms/latest/developerguide/logging-using-cloudtrail.html

AWS Secrets Manager: AWS Documentation. (2024). What is AWS Secrets Manager? AWS Secrets Manager User Guide. Retrieved from https://docs.aws.amazon.com/secretsmanager/latest/userguide/intro.html

AWS Systems Manager Parameter Store: AWS Documentation. (2024). AWS Systems Manager Parameter Store. AWS Systems Manager User Guide. Retrieved from https://docs.aws.amazon.com/systems-manager/latest/userguide/systems-manager-parameter-store.html

Question: 4

A company needs steady and predictable performance from its Amazon EC2 instances at the lowest possible cost. The company also needs the ability to scale resources to ensure that it has the right resources available at the right time. Which AWS service or resource will meet these requirements?

A: Amazon CloudWatch

B: Application Load Balancer

C: AWS Batch

D: Amazon EC2 Auto Scaling

Correct Answer:

D

Explanation:

Amazon EC2 Auto Scaling is the service designed to maintain application availability and allow you to automatically add or remove Amazon EC2 instances according to conditions you define. It ensures predictable performance by scaling out during demand spikes and achieves lower costs by scaling in during lulls, directly meeting the requirements for performance and cost-effective scaling. It ensures the right number of resources are available at the right time.

Why Incorrect Options are Wrong:

A: Amazon CloudWatch: This is a monitoring and observability service. While it can provide the metrics that trigger an Auto Scaling policy, it does not perform the scaling action itself.

B: Application Load Balancer: This service distributes incoming traffic across multiple EC2 instances. It works with Auto Scaling to ensure traffic reaches the scaled fleet but does not manage the scaling.

C: AWS Batch: This is a specialized service for running large-scale batch computing jobs. It is not intended for the general-purpose, demand-based scaling of a typical web application.

References:

Amazon EC2 Auto Scaling: AWS Documentation states, "Amazon EC2 Auto Scaling helps you ensure that you have the correct number of Amazon EC2 instances available to handle the load for your application."

Source: AWS Documentation, "What is Amazon EC2 Auto Scaling?", https://docs.aws.amazon.com/autoscaling/ec2/userguide/what-is-amazon-ec2-auto-scaling.html

AWS Cloud Practitioner Essentials: The official courseware explains that Auto Scaling "adds or replaces EC2 instances automatically across Availability Zones... based on demand." This aligns with cost savings and performance.

Source: AWS Skill Builder, "AWS Cloud Practitioner Essentials" course, Module 6: Compute.

Application Load Balancer: AWS Documentation clarifies its role: "An Application Load Balancer functions at the application layer... to route traffic to targets... within Amazon Virtual Private Cloud (Amazon VPC)."

Source: AWS Documentation, "What is an Application Load Balancer?", https://docs.aws.amazon.com/elasticloadbalancing/latest/application/introduction.html

Question: 5

A company runs a web application on Amazon EC2 instances. The application must run constantly and is expected to run indefinitely without interruption. Which EC2 instance purchasing options will meet these requirements MOST cost-effectively? (Select TWO.)

A: On-Demand Instances

B: Spot Instances.

C: Reserved Instances

D: Savings Plans

E: Dedicated Hosts

Correct Answer:

C, D

Explanation:

For workloads that are expected to run constantly and indefinitely, the most cost-effective purchasing options involve a long-term commitment in exchange for a lower price. Reserved Instances (RIs) provide a significant discount (up to 72%) over On-Demand pricing when you commit to a specific instance family in a region for a 1- or 3-year term. Similarly, Savings Plans offer comparable discounts for a commitment to a consistent amount of compute usage (measured in $/hour) over a 1- or 3-year term. Both models are designed specifically for steady-state, predictable workloads, directly matching the scenario's requirements for constant operation and cost optimization.

Why Incorrect Options are Wrong:

A: On-Demand Instances: This option is not the most cost-effective for long-term, constant workloads. It offers flexibility but at the highest price point.

B: Spot Instances: These instances can be terminated by AWS with a two-minute warning, which violates the requirement that the application must run constantly and without interruption.

E: Dedicated Hosts: This is a physical server dedicated to a single customer, primarily used for compliance or licensing needs. It is generally a more expensive option than RIs or Savings Plans.

References:

1. AWS Documentation - EC2 Pricing: "Savings Plans are a flexible pricing model that offer low prices on EC2...usage, in exchange for a commitment to a consistent amount of usage...for a 1 or 3 year term." and "Amazon EC2 Reserved Instances (RIs) provide a significant discount...compared to On-Demand pricing...for applications with steady state or predictable usage."

URL: https://aws.amazon.com/ec2/pricing/

2. AWS Documentation - Savings Plans: "Savings Plans is a flexible pricing model that provides savings of up to 72% on your AWS compute usage. This pricing model offers lower prices on Amazon EC2 instances usage, regardless of instance family, size, OS, tenancy or AWS Region..."

URL: https://aws.amazon.com/savingsplans/

3. AWS Documentation - Reserved Instances: "Amazon EC2 Reserved Instances (RI) provide a significant discount (up to 72%) compared to On-Demand pricing and provide a capacity reservation when used in a specific Availability Zone. RIs are recommended for applications with steady-state usage."

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

Question: 6

A company plans to launch an application that will run in multiple locations within the United States. The company needs to identify the two AWS Regions where the application can operate at the lowest price. Which AWS service or feature should the company use to determine the Regions that offer the lowest price?

A: Cost Explorer

B: AWS Budgets

C: AWS Trusted Advisor

D: AWS Pricing Calculator

Correct Answer:

D

Explanation:

The AWS Pricing Calculator is a web-based planning tool used to create cost estimates for AWS use cases. A company can model its planned application, specifying services like Amazon EC2 and Amazon S3, and configure them for different AWS Regions. By creating separate estimates for each target Region in the United States, the company can directly compare the projected monthly costs and identify the two most price-effective locations for their launch. This tool is specifically designed for pre-deployment cost estimation and comparison.

Why Incorrect Options are Wrong:

A: Cost Explorer: This tool is used to analyze existing or historical AWS cost and usage data. It is not designed for estimating costs for a new, un-deployed application.

B: AWS Budgets: This service is for setting spending limits and receiving alerts when costs approach or exceed a defined threshold. It helps manage costs, not estimate them initially.

C: AWS Trusted Advisor: This tool provides recommendations to optimize existing AWS environments across cost, performance, and security. It does not provide a feature to estimate and compare costs for new workloads.

References:

AWS Pricing Calculator: "AWS Pricing Calculator is a web service that you can use to create estimates for the cost of your AWS use cases... It is useful for both people who have

never used AWS and for those who want to reorganize or expand their usage." - AWS Documentation, What is AWS Pricing Calculator?

Cost Explorer: "AWS Cost Explorer has an easy-to-use interface that lets you visualize, understand, and manage your AWS costs and usage over time." - AWS Documentation, Using Cost Explorer to analyze your costs.

AWS Budgets: "You can use AWS Budgets to set custom budgets to track your cost and usage from the simplest to the most complex use cases." - AWS Documentation, Managing your costs with AWS Budgets.

AWS Trusted Advisor: "AWS Trusted Advisor helps you optimize your AWS infrastructure, improve security and performance, reduce your overall costs, and monitor service limits." - AWS Documentation, AWS Trusted Advisor.

Question: 7

Which AWS benefit enables users to deploy cloud infrastructure that consists of multiple geographic regions connected by a network with low latency, high throughput, and redundancy?

A: Economies of scale

B: Security

C: Elasticity

D: Global reach

Correct Answer:

D

Explanation:

The AWS benefit of global reach allows users to deploy applications and infrastructure across multiple AWS Regions worldwide. The AWS global network is a private, purpose-built infrastructure that connects these regions with low latency, high throughput, and redundant fiber optic cables. This enables customers to serve end-users with lower latency, build fault-tolerant applications with geographic separation, and meet data sovereignty requirements by placing data in specific geographic locations.

Why Incorrect Options are Wrong:

A: Economies of scale: This refers to the cost savings AWS passes to customers due to its massive operational scale, not the geographic distribution of its infrastructure.

B: Security: While the AWS global infrastructure is secure, "Security" is a broad benefit and does not specifically describe the capability to deploy across multiple geographic regions.

C: Elasticity: This is the ability to acquire and release resources as needed to match demand, which is about scaling capacity, not geographic deployment.

References:

AWS Documentation - AWS Global Infrastructure: "The AWS Global Cloud Infrastructure is the most secure, extensive, and reliable cloud platform, offering over 200 fully featured services from data centers globally. ... This global network provides the high availability and low latency that our customers need." (Direct URL: https://aws.amazon.com/about-aws/global-infrastructure/)

AWS Whitepaper - Overview of Amazon Web Services: In the section "The Six Advantages of Cloud Computing," the benefit "Go global in minutes" is detailed: "Easily deploy your application in multiple regions around the world with just a few clicks. This means you can provide lower latency and a better experience for your customers at minimal cost." (Page 6, Direct URL: https://docs.aws.amazon.com/whitepapers/latest/aws-overview/six-advantages-of-cloud-computing.html)

AWS Whitepaper - AWS Cloud Adoption Framework: Cloud-Ready Organization: The Business Perspective of the AWS CAF highlights "Global focus" as a capability, enabling organizations to "serve customers in their preferred localities and languages." (Page 10, Direct URL: https://d1.awsstatic.com/whitepapers/aws-caf-cloud-ready-organization.pdf)

Question: 8

A company needs to use machine learning and pattern matching to identify and protect sensitive data that the company stores in the AWS Cloud. Which AWS service will meet these requirements?

A: Amazon Inspector

B: Amazon Macie

C: Amazon GuardDuty

D: AWS Audit Manager

Correct Answer:

B

Explanation:

Amazon Macie is a fully managed data security and data privacy service that uses machine learning (ML) and pattern matching to discover and protect sensitive data in Amazon S3. It automatically provides an inventory of S3 buckets and evaluates them for security controls. Macie identifies sensitive data such as personally identifiable information (PII) and financial data, providing dashboards and alerts that give visibility into how this data is being accessed or moved. This directly addresses the company's need to use ML and pattern matching to identify and protect sensitive data.

Why Incorrect Options are Wrong:

A: Amazon Inspector: This is an automated vulnerability management service that scans AWS workloads (like EC2 instances and container images) for software vulnerabilities and unintended network exposure, not for sensitive data discovery.

C: Amazon GuardDuty: This is a threat detection service that monitors for malicious activity and unauthorized behavior. It analyzes logs but does not scan data content to classify sensitive information.

D: AWS Audit Manager: This service helps you continuously audit your AWS usage to simplify risk and compliance assessment. It is for auditing and reporting, not for data discovery and classification.

References:

Amazon Macie: AWS Documentation. (n.d.). What is Amazon Macie? Retrieved from https://docs.aws.amazon.com/macie/latest/user/what-is-macie.html

Amazon Inspector: AWS Documentation. (n.d.). What is Amazon Inspector? Retrieved from https://docs.aws.amazon.com/inspector/latest/user/what-is-inspector.html

Amazon GuardDuty: AWS Documentation. (n.d.). What is Amazon GuardDuty? Retrieved from https://docs.aws.amazon.com/guardduty/latest/ug/what-is-guardduty.html

AWS Audit Manager: AWS Documentation. (n.d.). What is AWS Audit Manager? Retrieved from https://docs.aws.amazon.com/audit-manager/latest/userguide/what-is.html

Question: 9

A company wants to expand from one AWS Region into a second AWS Region. What does the company need to do to expand into the second Region?

A: Contact an AWS account manager to sign a new contract.

B: Move an Availability Zone to the second Region.

C: Begin to deploy resources in the second Region.

D: Download the AWS Management Console for the second Region.

Correct Answer:

C

Explanation:

An AWS account is global and provides access to multiple AWS Regions. To expand operations into a new Region, a user simply needs to select the target Region in the AWS Management Console, AWS Command Line Interface (CLI), or via an SDK, and then begin provisioning resources. This action does not require a new contract or a separate console. AWS is designed for self-service, allowing users to deploy infrastructure globally on demand.

Why Incorrect Options are Wrong:

A: An AWS account is global. A single account and its associated contract allow resource deployment in any standard AWS Region without requiring new agreements.

B: Availability Zones are distinct, isolated locations within an AWS Region. They are fixed physical infrastructure and cannot be moved from one Region to another.

D: The AWS Management Console is a unified, web-based interface. It is not a downloadable application, and users switch between Regions using a selector within the single console.

References:

1. AWS Documentation - Regions and Zones: "An AWS account provides you with the ability to use any of the AWS Regions... When you work with a resource, you must specify the Region where the resource will be used." This confirms that an existing account can be used to deploy resources by simply specifying the new Region.

Source: AWS Documentation, "Regions and Zones," AWS General Reference. https://docs.aws.amazon.com/general/latest/gr/rande.html

2. AWS Documentation - Getting Started with the AWS Management Console: "You can select an AWS Region to work in. The list of Regions is on the navigation bar... The resources you view and manage in that Region are displayed." This demonstrates that expanding to a new Region is a matter of selecting it in the console.

Source: AWS Documentation, "Getting Started with the AWS Management Console," AWS Management Console User Guide. https://docs.aws.amazon.com/awsconsolehelpdocs/latest/gsg/getting-started.html#select-region

3. AWS Documentation - AWS Global Infrastructure: This resource explains that Regions are physically separate geographic areas and Availability Zones are contained within them, making option B logically impossible.

Source: AWS Official Website, "AWS Global Infrastructure." https://aws.amazon.com/about-aws/global-infrastructure/

Question: 10

Which of the following are AWS best practice recommendations for the use of AWS Identity and Access Management (IAM)? (Select TWO.)

A: Use the AWS account root user for daily access.

B: Use access keys and secret access keys on Amazon EC2.

C: Rotate credentials on a regular basis.

D: Create a shared set of access keys for system administrators.

E: Configure multi-factor authentication (MFA).

Correct Answer:

C, E

Explanation:

AWS Identity and Access Management (IAM) best practices are designed to enhance the security of an AWS account. A core security principle is to regularly rotate credentials, such as passwords and access keys, to limit the potential damage if they are compromised. Another fundamental best practice is to enable multi-factor authentication (MFA), which adds a crucial second layer of security to user sign-ins and resource access. MFA is especially critical for the root user and any highly privileged IAM users, protecting against unauthorized access even if a password is stolen.

Why Incorrect Options are Wrong:

A: AWS explicitly advises against using the root user for daily tasks. The root user should be secured and used only for tasks that require root-level permissions.

B: The best practice for EC2 instances is to use IAM roles, which provide temporary credentials, rather than storing long-term access keys on the instance.

D: Sharing credentials violates the principle of individual accountability and least privilege. Each user should have their own unique credentials.

References:

AWS IAM User Guide - Security best practices in IAM: This official guide explicitly states, "Rotate your credentials regularly" (confirming option C) and "Enable MFA for privileged users" (confirming option E).

URL: https://docs.aws.amazon.com/IAM/latest/UserGuide/best-practices.html

AWS IAM User Guide - Tasks that require root user credentials: This page details the limited set of tasks for which the root user is necessary, reinforcing that it should not be used for daily access (refuting option A).

URL: https://docs.aws.amazon.com/IAM/latest/UserGuide/root-user-tasks.html

AWS IAM User Guide - IAM roles for Amazon EC2: This document explains that using IAM roles is the secure way to grant permissions to applications running on EC2 instances, avoiding the need for long-term access keys (refuting option B).

URL: https://docs.aws.amazon.com/IAM/latest/UserGuide/idrolesuseswitch-role-ec2.html

Question: 11

Which of the following are aspects of the AWS shared responsibility model? (Select TWO.)

A: Configuration management of infrastructure devices is the customer’s responsibility.

B: For Amazon S3, AWS operates the infrastructure layer, the operating systems, and the platforms.

C: AWS is responsible for protecting the physical cloud infrastructure.

D: AWS is responsible for training the customer’s employees on AWS products and services.

E: For Amazon EC2, AWS is responsible for maintaining the guest operating system.

Correct Answer:

B, C

Explanation:

The AWS shared responsibility model divides security responsibilities between AWS and the customer. AWS is responsible for the security OF the cloud, while the customer is responsible for security IN the cloud.

C is correct because AWS's responsibility includes protecting the physical infrastructure that runs all AWS services, such as hardware, software, networking, and the facilities themselves.

B is correct because for managed services like Amazon S3, AWS manages the underlying infrastructure, operating system, and platform, relieving the customer of these operational burdens. The customer remains responsible for data and access management.

Why Incorrect Options are Wrong:

A: This statement is imprecise. AWS manages the physical "infrastructure devices." While customers are responsible for configuring their virtual resources (like EC2 instances), the term used is ambiguous and not standard.

D: Customers are responsible for training their own employees. While AWS provides training resources, it is not an AWS responsibility under the shared responsibility model.

E: For Infrastructure as a Service (IaaS) like Amazon EC2, the customer is responsible for managing and patching the guest operating system, not AWS.

References:

1. AWS Shared Responsibility Model Documentation: "AWS is responsible for protecting the infrastructure that runs all of the services offered in the AWS Cloud. This infrastructure is composed of the hardware, software, networking, and facilities that run AWS Cloud services." This supports option The same page shows a diagram where for managed services like S3, AWS manages the "Software" (OS, Platform) and "Hardware" (Infrastructure). This supports option B.

URL: https://aws.amazon.com/compliance/shared-responsibility-model/

2. AWS Well-Architected Framework, Security Pillar: "For IaaS services, such as Amazon EC2, you are responsible for managing the guest operating system (including security updates and patches)..." This directly refutes option E.

URL: https://docs.aws.amazon.com/wellarchitected/latest/security-pillar/shared-responsibility-model.html (Section: Differentiating between the responsibilities)

3. AWS Cloud Practitioner Essentials (Digital Training): This courseware clarifies that AWS manages the components "from the virtualization layer down to the physical security of the facilities," while the customer manages the guest OS (for EC2) and configures their resources. This supports the correctness of C and the incorrectness of E.

URL: https://explore.skillbuilder.aws/learn/course/external/view/elearning/134/aws-cloud-practitioner-essentials (Module 3: Cloud Security)

Question: 12

Which characteristic of the AWS Cloud helps users eliminate underutilized CPU capacity?

A: Agility

B: Elasticity

C: Reliability

D: Durability

Correct Answer:

B

Explanation:

Elasticity is the core AWS Cloud concept that allows users to acquire resources as they are needed and release them when they are no longer required. This capability directly addresses the problem of underutilized capacity. By automatically or manually scaling down computing resources (like EC2 instances) during periods of low demand, organizations can stop paying for idle or underutilized CPU, thus optimizing costs and eliminating waste. This aligns with the pay-as-you-go pricing model of the cloud.

Why Incorrect Options are Wrong:

A: Agility: Refers to the speed at which users can innovate and deploy services, not the dynamic scaling of resource capacity to match load.

C: Reliability: Pertains to a system's ability to function correctly and recover from infrastructure or service disruptions, not resource utilization.

D: Durability: Relates to the long-term protection and persistence of data, ensuring it is not lost, which is distinct from managing compute capacity.

References:

1. AWS Official Documentation - "What is Cloud Computing?": Under the "Benefits of cloud computing" section, AWS states: "With cloud computing, you don’t have to over-provision resources up front to handle peak levels of business activity in the future. Instead, you provision the amount of resources that you actually need... You can scale these resources up or down to instantly grow and shrink capacity as your business needs change." This describes elasticity.

URL: https://aws.amazon.com/what-is-cloud-computing/

2. AWS Cloud Adoption Framework (CAF) Whitepaper: The whitepaper discusses elasticity as a key capability. The "Platform Perspective" focuses on principles to "implement a hybrid architecture that is elastic and scalable," directly linking elasticity to scalable capacity.

URL: https://docs.aws.amazon.com/whitepapers/latest/aws-cloud-adoption-framework/platform-perspective.html

3. AWS Well-Architected Framework - Performance Efficiency Pillar: This document emphasizes "right-sizing" workloads to eliminate waste. It states, "Over-provisioning is a common cause of waste... By using elasticity, you can ensure that you are only paying for the resources you need."

URL: https://docs.aws.amazon.com/wellarchitected/latest/performance-efficiency-pillar/cost-2.html

Question: 13

A company is moving its office and must establish an encrypted connection to AWS. Which AWS service will help meet this requirement?

A: AWS VPN

B: Amazon Route 53

C: Amazon API Gateway

D: Amazon Connect

Correct Answer:

A

Explanation:

AWS VPN provides services to establish secure, encrypted connections. Specifically, AWS Site-to-Site VPN is designed to create an encrypted tunnel between an on-premises location, such as a company office, and an Amazon Virtual Private Cloud (VPC). This service uses IPsec (Internet Protocol Security) to secure the traffic in transit over the internet, directly fulfilling the requirement for an encrypted connection to AWS from a physical office.

Why Incorrect Options are Wrong:

B: Amazon Route 53: This is a scalable Domain Name System (DNS) web service. It translates domain names to IP addresses and does not establish network-level encrypted connections.

C: Amazon API Gateway: This is a managed service for creating and managing APIs. While it secures API calls, it does not provide a general-purpose encrypted network connection for an entire office.

D: Amazon Connect: This is a cloud-based contact center service for managing customer interactions. It is unrelated to establishing private network connectivity between an office and AWS.

References:

AWS Site-to-Site VPN Documentation: "AWS Site-to-Site VPN creates a secure connection between your data center or branch office and your AWS cloud resources." AWS

Documentation, What is AWS Site-to-Site VPN?, Retrieved from: https://docs.aws.amazon.com/vpn/latest/s2svpn/VPCVPN.html

Amazon Route 53 Documentation: "Amazon Route 53 is a highly available and scalable cloud Domain Name System (DNS) web service." AWS Documentation, What is Amazon Route 53?, Retrieved from: https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/Welcome.html

Overview of Amazon Web Services Whitepaper: This document describes AWS VPN as a service for extending on-premises networks to the AWS Cloud. AWS Whitepapers, Overview of Amazon Web Services, Page 41, Retrieved from: https://d1.awsstatic.com/whitepapers/aws-overview.pdf

Question: 14

How does consolidated billing help reduce costs for a company that has multiple AWS accounts?

A: It aggregates usage across accounts so that the company can reach volume discount thresholds sooner.

B: It offers an additional 5% discount on purchases of All Upfront Reserved Instances.

C: It provides a simplified billing invoice that the company can process more quickly than a standard invoice.

D: It gives AWS resellers the ability to bill their customers for usage.

Correct Answer:

A

Explanation:

Consolidated billing, a feature of AWS Organizations, combines usage from all member accounts into a single bill for the management account. For services with volume pricing tiers, such as Amazon S3 and data transfer, AWS aggregates the usage across all accounts. This combined usage allows the organization to reach higher-usage, lower-priced tiers more quickly than individual accounts would, directly reducing the overall cost. This is a primary financial benefit of using AWS Organizations.

Why Incorrect Options are Wrong:

B: Consolidated billing does not offer a specific, additional discount on Reserved Instances. It facilitates the sharing of RI discounts across accounts, but it does not add a new percentage discount.

C: While a single invoice is simpler and may reduce administrative overhead, the most significant and direct cost reduction comes from qualifying for volume discounts, not faster processing.

D: This describes a functionality for AWS Partners in the reseller program, not a direct cost-saving mechanism for a standard company managing its own multiple accounts.

References:

AWS Billing and Cost Management User Guide: "For pricing, AWS treats all the accounts in an organization as if they were one account. Some services, such as Amazon EC2 and

Amazon S3, have volume pricing tiers that give you lower prices the more you use the service. With consolidated billing, AWS combines the usage from all accounts to determine which volume pricing tier to apply, giving you a lower overall price whenever possible."

URL: https://docs.aws.amazon.com/awsaccountbilling/latest/aboutv2/consolidated-billing.html

AWS Organizations User Guide: "You can combine the usage from all accounts in the organization to qualify for volume pricing discounts. You then get the benefit of a single bill for all of your member accounts."

URL: https://docs.aws.amazon.com/organizations/latest/userguide/orgsintroduction.html

Question: 15

What is an AWS Region?

A: A broad set of global, cloud-based products that include compute, storage, and databases

B: A physical location around the world where data centers are clustered

C: One or more discrete data centers with redundant power, networking, and connectivity

D: A service that developers use to build applications that deliver latencies of single-digit milliseconds to users

Correct Answer:

B

Explanation:

An AWS Region is a physical, geographic location anywhere in the world where AWS clusters data centers. Each Region is designed to be completely isolated from other AWS Regions. This isolation helps achieve the greatest possible fault tolerance and stability. A Region consists of multiple, isolated, and physically separate Availability Zones within that geographic area. This structure is a core part of how AWS provides resilient and highly available cloud services.

Why Incorrect Options are Wrong:

A: This describes the broad portfolio of AWS services, not the physical infrastructure concept of a Region.

C: This is the definition of an Availability Zone (AZ), which is a component within an AWS Region, not the Region itself.

D: This describes a benefit or a specific service like AWS Local Zones, which are extensions of a Region, but it does not define the Region itself.

References:

AWS Global Infrastructure: "An AWS Region is a physical location in the world where we have multiple Availability Zones." AWS Global Infrastructure, Regions and Availability Zones. https://aws.amazon.com/about-aws/global-infrastructure/regionsaz/

AWS Documentation - Regions and Zones: "Amazon Web Services has the concept of a Region, which is a physical location around the world where we cluster data centers. We call each group of logical data centers an Availability Zone." AWS Regions and Zones. https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/using-regions-availability-zones.html

AWS Cloud Practitioner Essentials Courseware: "An AWS Region is a geographical area... Each Region consists of multiple, isolated, and physically separate AZs within a geographic area." AWS Skill Builder, AWS Cloud Practitioner Essentials, Module 3: Global Infrastructure and Reliability. https://explore.skillbuilder.aws/learn/course/external/view/elearning/13286/aws-cloud-practitioner-essentials

Question: 16

A company needs an AWS Support plan that provides programmatic case management through the AWS Support API. Which support plan will meet this requirement MOST cost-effectively?

A: AWS Business Support

B: AWS Basic Support

C: AWS Developer Support

D: AWS Enterprise Support

Correct Answer:

A

Explanation:

The AWS Support API provides programmatic access to create, manage, and resolve support cases. According to the official AWS Support plan comparison, this feature is first available at the AWS Business Support tier. While the AWS Enterprise Support plan also includes this feature, the Business Support plan is the most cost-effective option that fulfills the stated requirement. The Developer and Basic plans do not offer access to the AWS Support API for case management.

Why Incorrect Options are Wrong:

B: AWS Basic Support: This plan does not provide access to the AWS Support API for programmatic case management.

C: AWS Developer Support: This plan is designed for early development and testing and does not include access to the AWS Support API.

D: AWS Enterprise Support: Although this plan includes the AWS Support API, it is more expensive than the Business plan and therefore not the MOST cost-effective choice.

References:

AWS Support Plans: The official AWS documentation provides a detailed feature comparison table for all support plans. The "AWS Support API" feature is explicitly listed as available for Business and Enterprise plans, but not for Developer or Basic plans.

AWS. (n.d.). Compare AWS Support Plans. AWS Premium Support. Retrieved from https://aws.amazon.com/premiumsupport/plans/

AWS Support API Documentation: This documentation confirms that access to the Support API is a feature of the Business, Enterprise On-Ramp, and Enterprise Support plans.

AWS. (n.d.). AWS Support API Reference. Retrieved from https://docs.aws.amazon.com/awssupport/latest/APIReference/Welcome.html (See the introductory section on access).

Question: 17

Which of the following are AWS compute services? (Select TWO.)

A: Amazon Lightsail

B: AWS Systems Manager

C: AWS CloudFormation

D: AWS Batch

E: Amazon Inspector

Correct Answer:

A, D

Explanation:

AWS compute services provide the processing power and infrastructure required to run applications and workloads. Amazon Lightsail is a compute service that offers easy-to-use virtual private servers, storage, and networking, simplifying the process of launching and managing a web application. AWS Batch is a fully managed service that enables users to run batch computing jobs at any scale. It dynamically provisions the optimal quantity and type of compute resources (such as Amazon EC2 or AWS Fargate) based on job requirements, placing it firmly in the compute category.

Why Incorrect Options are Wrong:

B: AWS Systems Manager: This is a management and governance service used for operational visibility and control over AWS infrastructure, not a core compute service.

C: AWS CloudFormation: This is an Infrastructure as Code (IaC) service for modeling and provisioning AWS resources, categorized under management and governance.

E: Amazon Inspector: This is a security service that provides automated vulnerability management and security assessments for AWS workloads.

References:

AWS Compute Services: The official AWS page for compute services lists both Amazon Lightsail and AWS Batch. AWS. (n.d.). Cloud-based Compute Services. Retrieved from https://aws.amazon.com/products/compute/

Amazon Lightsail: "Amazon Lightsail is an easy-to-use cloud platform that offers you everything you need to build an application or website... Lightsail provides... a virtual private server (instance)." AWS. (n.d.). What is Amazon Lightsail?. Retrieved from https://aws.amazon.com/lightsail/

AWS Batch: "AWS Batch enables developers, scientists, and engineers to easily and efficiently run hundreds of thousands of batch computing jobs on AWS." AWS. (n.d.). AWS Batch. Retrieved from https://aws.amazon.com/batch/

AWS Cloud Practitioner Essentials Courseware: This official AWS training course categorizes services. Compute services include EC2, Lambda, Lightsail, and Batch. Management and Governance services include CloudFormation and Systems Manager. Security services include Inspector. AWS Training and Certification. (2023). AWS Cloud Practitioner Essentials. Module 3: AWS Global Infrastructure and Reliability & Module 4: AWS Cloud Security. Retrieved from https://explore.skillbuilder.aws/learn/course/external/view/elearning/134/aws-cloud-practitioner-essentials

Question: 18

A company has existing software licenses that it wants to bring to AWS, but the licensing model requires licensing physical cores. How can the company meet this requirement in the AWS Cloud?

A: Launch an Amazon EC2 instance with default tenancy.

B: Launch an Amazon EC2 instance on a Dedicated Host.

C: Create an On-Demand Capacity Reservation.

D: Purchase Dedicated Reserved Instances.

Correct Answer:

B

Explanation:

An Amazon EC2 Dedicated Host is a physical server with EC2 instance capacity that is fully dedicated to a single customer's use. This service is specifically designed for scenarios where customers need to use their existing server-bound software licenses, such as those based on physical cores or sockets. By provisioning a Dedicated Host, the company gains visibility and control over the underlying physical server, allowing them to place instances strategically to meet the compliance requirements of their per-core licensing model.

Why Incorrect Options are Wrong:

A: Launch an Amazon EC2 instance with default tenancy. Default tenancy instances run on shared hardware, which does not provide the necessary visibility or control over physical cores for license compliance.

C: Create an On-Demand Capacity Reservation. This reserves capacity in an Availability Zone but does not dedicate a physical server, which is the core requirement for per-core licensing.

D: Purchase Dedicated Reserved Instances. Dedicated Instances run on hardware dedicated to a single customer, but Dedicated Hosts provide greater control and visibility over the physical server needed for specific BYOL licensing.

References:

AWS Documentation - Amazon EC2 Dedicated Hosts: "Dedicated Hosts are physical servers with EC2 instance capacity fully dedicated to your use. Dedicated Hosts can help

you reduce costs by allowing you to use your existing server-bound software licenses, including Windows Server, SQL Server, and SUSE Linux Enterprise Server... and can also help you meet compliance requirements." (URL: https://aws.amazon.com/ec2/dedicated-hosts/)

AWS Documentation - Comparing Dedicated Hosts and Dedicated Instances: "With Dedicated Hosts, you get an entire physical server, and you are billed per host... This is ideal for BYOL scenarios where you are licensed per-socket or per-core." (This concept is detailed in FAQs and comparison pages). (URL: https://aws.amazon.com/ec2/dedicated-hosts/faqs/ - See "Q: How are Dedicated Hosts different from Dedicated Instances?")

AWS Documentation - Amazon EC2 User Guide for Linux Instances: "A Dedicated Host is a physical server with EC2 instance capacity fully dedicated to your use. Dedicated Hosts allow you to use your existing per-socket, per-core, or per-VM software licenses..." (Section: Dedicated Hosts). (URL: https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/dedicated-hosts-overview.html)

Question: 19

Which action will provide documentation to help a company evaluate whether its use of the AWS Cloud is compliant with local regulatory standards?

A: Running Amazon GuardDuty

B: Using AWS Artifact

C: Creating an AWS Support ticket

D: AWS Cloud Trail logs

Correct Answer:

B

Explanation:

AWS Artifact is a self-service portal that provides on-demand access to AWS's security and compliance reports. These documents, such as Service Organization Control (SOC) reports, Payment Card Industry (PCI) reports, and certifications from various accreditation bodies, are essential for customers to perform their due diligence. Companies can use these official reports to evaluate and validate that AWS's infrastructure and services meet their specific local regulatory and compliance requirements. This directly addresses the need for documentation for compliance evaluation.

Why Incorrect Options are Wrong:

A: Running Amazon GuardDuty: This is a threat detection service that monitors for malicious activity. It helps with operational security but does not provide formal compliance attestation reports about the AWS platform itself.

C: Creating an AWS Support ticket: While AWS Support can provide guidance, it is not the direct source for compliance documentation. The support team would likely redirect the user to AWS Artifact for these reports.

D: AWS CloudTrail logs: CloudTrail provides audit logs of API calls and actions within a customer's account. It is used to audit the customer's own activities, not to obtain documentation on AWS's compliance.

References:

1. AWS Artifact: AWS Official Documentation. "AWS Artifact is your go-to, central resource for compliance-related information that matters to you. It provides on-demand access to AWS's security and compliance reports and select online agreements."

URL: https://aws.amazon.com/artifact/

2. AWS Compliance: AWS Official Documentation. "AWS Artifact Reports provides on-demand access to AWS security and compliance documents, such as Service Organization Control (SOC) reports, Payment Card Industry (PCI) reports..."

URL: https://docs.aws.amazon.com/aws-technical-content/latest/aws-overview/governance-in-the-cloud.html#aws-compliance (Section: AWS Compliance)

3. Amazon GuardDuty: AWS Official Documentation. "Amazon GuardDuty is a threat detection service that continuously monitors your AWS accounts and workloads for malicious activity and delivers detailed security findings for visibility and remediation."

URL: https://aws.amazon.com/guardduty/

Question: 20

A company has a centralized group of users with large file storage requirements that have exceeded the space available on premises. The company wants to extend its file storage capabilities for this group while retaining the performance benefit of sharing content locally. What is the MOST operationally efficient AWS solution for this scenario?

A: Create an Amazon S3 bucket for each user. Mount each bucket by using an S3 file system mounting utility.

B: Configure and deploy an AWS Storage Gateway file gateway. Connect each user's workstation to the file gateway.

C: Move each user's working environment to Amazon WorkSpaces. Set up an Amazon WorkDocs account for each user.

D: Deploy an Amazon EC2 instance and attach an Amazon Elastic Block Store (Amazon EBS) Provisioned IOPS volume. Share the EBS volume directly with the users.

Correct Answer:

B

Explanation:

The AWS Storage Gateway, specifically the File Gateway type, is the most operationally efficient solution for this use case. It provides a seamless way to extend on-premises file storage to the virtually unlimited capacity of Amazon S3. The File Gateway is deployed on-premises as a virtual appliance and presents a standard file share using NFS or SMB protocols. It maintains a local cache of the most frequently accessed data, which directly addresses the requirement to retain the performance benefits of local access. This managed service simplifies administration compared to building and maintaining a custom solution.

Why Incorrect Options are Wrong:

Create an Amazon S3 bucket for each user. Mount each bucket by using an S3 file system mounting utility.

Move each user's working environment to Amazon WorkSpaces. Set up an Amazon WorkDocs account for each user.

Deploy an Amazon EC2 instance and attach an Amazon Elastic Block Store (Amazon EBS) Provisioned IOPS volume. Share the EBS volume directly with the users.

References:

1. AWS Storage Gateway Documentation: "File Gateway provides a virtual on-premises file server, which enables you to store and retrieve files as objects in Amazon S3... File Gateway offers SMB or NFS-based access to data in Amazon S3 with local caching. The cache provides low-latency access to the recently used data."

URL: https://docs.aws.amazon.com/storagegateway/latest/userguide/WhatIsStorageGateway.html#file-gateway-concepts

2. AWS Cloud Practitioner Essentials - Digital Training: The "Storage" module describes AWS Storage Gateway as a hybrid cloud storage service that enables on-premises applications to seamlessly use AWS cloud storage. It highlights the File Gateway's role in providing file-based, cached access to S3.

URL: https://explore.skillbuilder.aws/learn/course/external/view/elearning/11458/aws-cloud-practitioner-essentials (Module 5: Storage)

3. Amazon EBS Documentation: "An Amazon EBS volume is a durable, block-level storage device that you can attach to a single EC2 instance." This confirms that EBS volumes are not designed for direct sharing with multiple on-premises users.

URL: https://docs.aws.amazon.com/ebs/latest/userguide/ebs-volumes.html

Question: 21

Which option is a customer responsibility when using Amazon DynamoDB under the AWS Shared Responsibility Model?

A: Physical security of DynamoDB

B: Patching of DynamoDB

C: Access to DynamoDB tables

D: Encryption of data at rest in DynamoDB

Correct Answer:

Access to DynamoDB tables

Explanation:

Under the AWS Shared Responsibility Model, AWS is responsible for the security of the cloud, while the customer is responsible for security in the cloud. For a managed service like Amazon DynamoDB, AWS manages the underlying infrastructure, including physical security, hardware, and software patching. The customer is responsible for managing their data and configuring access to it. This includes using AWS Identity and Access Management (IAM) to create policies that define which users and roles can perform specific actions on their DynamoDB tables. This control over data access is a core customer responsibility.

Why Incorrect Options are Wrong:

References:

1. AWS Shared Responsibility Model: AWS Documentation. This page outlines the division of responsibilities. It states customers are responsible for "Identity and access management" and their "Data". AWS is responsible for the "Compute, Storage, Database, Networking" infrastructure and the "Software" that runs the services.

URL: https://aws.amazon.com/compliance/shared-responsibility-model/

2. Security in Amazon DynamoDB: AWS Documentation. This document specifies security configurations. Section "Identity and Access Management for Amazon DynamoDB" details how customers use IAM to control access, which is their responsibility.

URL: https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/security.html

3. AWS Well-Architected Framework - Security Pillar: AWS Whitepaper. Page 6 discusses the Shared Responsibility Model, clarifying that customers are responsible for managing access to their data and systems.

URL: https://docs.aws.amazon.com/wellarchitected/latest/security-pillar/the-shared-responsibility-model.html

Question: 22

A company wants to run a NoSQL database on Amazon EC2 instances. Which task is the responsibility of AWS in this scenario?

A: Update the guest operating system of the EC2 instances.

B: Maintain high availability at the database layer.

C: Patch the physical infrastructure that hosts the EC2 instances.

D: Configure the security group firewall.

Correct Answer:

C

Explanation:

This question relates to the AWS Shared Responsibility Model. For Infrastructure as a Service (IaaS) products like Amazon EC2, AWS is responsible for the security of the cloud. This includes managing and patching the underlying physical infrastructure, such as the compute, storage, and networking hardware, as well as the virtualization layer (hypervisor). The customer is responsible for security in the cloud, which includes the guest operating system, applications (like the NoSQL database), and network configurations such as security groups. Therefore, patching the physical infrastructure is an AWS responsibility.

Why Incorrect Options are Wrong:

Update the guest operating system of the EC2 instances: The customer is responsible for patching and maintaining the guest operating system installed on their EC2 instances.

Maintain high availability at the database layer: The customer is responsible for architecting and managing their applications, including configuring the NoSQL database for high availability.

Configure the security group firewall: While AWS provides the security group service, the customer is responsible for defining and managing the specific firewall rules.

References:

AWS Documentation - Shared Responsibility Model: This official page clearly outlines the division of responsibilities. It states, "AWS is responsible for protecting the infrastructure that runs all of the services offered in the AWS Cloud. This infrastructure is composed of the

hardware, software, networking, and facilities that run AWS Cloud services." This directly supports that AWS patches the physical infrastructure.

URL: https://aws.amazon.com/compliance/shared-responsibility-model/

AWS Well-Architected Framework - Security Pillar Whitepaper: This document details security best practices and reiterates the shared model. On page 6, it shows a diagram where "Hardware/AWS Global Infrastructure" and "Software" (Compute, Storage, Database, Networking) are AWS responsibilities, while "Platform, Applications, Identity & Access Management" and "Operating System, Network & Firewall Configuration" are customer responsibilities for IaaS.

URL: https://docs.aws.amazon.com/wellarchitected/latest/security-pillar/the-shared-responsibility-model.html

Question: 23

Which of the following are benefits of using AWS Trusted Advisor? (Choose two.)

A: Providing high-performance container orchestration

B: Creating and rotating encryption keys

C: Detecting underutilized resources to save costs

D: Improving security by proactively monitoring the AWS environment

E: Implementing enforced tagging across AWS resources

Correct Answer:

C, D

Explanation:

AWS Trusted Advisor is an online tool that provides real-time guidance to help you provision your resources following AWS best practices. It inspects your AWS environment and makes recommendations in five categories. Two of these categories are Cost Optimization, which includes checks for underutilized resources like idle Amazon EC2 instances or unassociated Elastic IP addresses, and Security, which includes checks for permissions, MFA on the root account, and other security configurations to help improve your security posture.

Why Incorrect Options are Wrong:

Providing high-performance container orchestration: This is a function of services like Amazon Elastic Container Service (ECS) or Amazon Elastic Kubernetes Service (EKS), not Trusted Advisor.

Creating and rotating encryption keys: This is the primary function of AWS Key Management Service (KMS), which manages cryptographic keys.

Implementing enforced tagging across AWS resources: This is typically managed using AWS Organizations Service Control Policies (SCPs) or AWS Config rules, not Trusted Advisor.

References:

AWS Trusted Advisor. (n.d.). How Trusted Advisor works. AWS Documentation. Retrieved from https://aws.amazon.com/premiumsupport/technology/trusted-advisor/

AWS Trusted Advisor. (n.d.). AWS Trusted Advisor check reference. AWS Support. Retrieved from https://docs.aws.amazon.com/awssupport/latest/user/trusted-advisor-check-reference.html (See "Cost Optimization" and "Security" sections).

Question: 24

A company is using a central data platform to manage multiple types of data for its customers. The company wants to use AWS services to discover, transform, and visualize the data. Which combination of AWS services should the company use to meet these requirements? (Choose two.)

A: AWS Glue

B: Amazon Elastic File System (Amazon EFS)

C: Amazon Redshift

D: Amazon QuickSight

D: Amazon Quantum Ledger Database (Amazon QLDB)

Correct Answer:

A, D

Explanation:

The question requires services for data discovery, transformation, and visualization. AWS Glue is a serverless data integration service that makes it easy to discover, prepare, and combine data. Its features, such as the AWS Glue Data Catalog and ETL (extract, transform, and load) jobs, directly address the discovery and transformation requirements. Amazon QuickSight is a scalable, serverless, embeddable, machine learning-powered business intelligence (BI) service. It is designed specifically to create and publish interactive dashboards and visualizations, fulfilling the visualization requirement. Together, these two services provide a comprehensive solution for the company's needs.

Why Incorrect Options are Wrong:

References:

1. AWS Glue: AWS Documentation. (n.d.). What is AWS Glue? "AWS Glue is a serverless data integration service... AWS Glue consists of a central metadata repository known as the AWS Glue Data Catalog, an ETL engine that automatically generates Python or Scala code, and a flexible scheduler that handles dependency resolution, job monitoring, and retries." Retrieved from https://docs.aws.amazon.com/glue/latest/dg/what-is-glue.html

2. Amazon QuickSight: AWS Documentation. (n.d.). What Is Amazon QuickSight? "Amazon QuickSight is a fast, cloud-powered business intelligence (BI) service that makes it easy for

you to deliver insights to everyone in your organization... you can create and publish interactive dashboards that can be accessed from browsers or mobile devices." Retrieved from https://docs.aws.amazon.com/quicksight/latest/user/what-is-quicksight.html

3. Analytics on AWS: AWS Whitepaper. (2022). Big Data Analytics Options on AWS. This whitepaper outlines the analytics pipeline, positioning AWS Glue for data processing/ETL (p. 21) and Amazon QuickSight for visualization (p. 29). Retrieved from https://d1.awsstatic.com/whitepapers/big-data-analytics-options-on-aws.pdf

Question: 25

A company is storing sensitive customer data in an Amazon S3 bucket. The company wants to protect the data from accidental deletion or overwriting. Which S3 feature should the company use to meet these requirements?

A: S3 Lifecycle rules

B: S3 Versioning

C: S3 bucket policies

D: S3 server-side encryption

Correct Answer:

S3 Versioning

Explanation:

S3 Versioning is a feature that keeps multiple variants of an object in the same bucket. When you enable S3 Versioning for a bucket, Amazon S3 preserves existing objects even if you overwrite or delete them. If an object is overwritten, a new version is created. If an object is deleted, a delete marker is added, but the previous versions are retained. This allows for the restoration of objects from unintended user actions, directly meeting the requirement to protect data from accidental deletion or overwriting.

Why Incorrect Options are Wrong:

References:

AWS Documentation - Using versioning in S3 buckets: "Versioning is a means of keeping multiple variants of an object in the same bucket. You can use the S3 Versioning feature to preserve, retrieve, and restore every version of every object stored in your buckets. With versioning, you can easily recover from both unintended user actions and application failures." (URL: https://docs.aws.amazon.com/AmazonS3/latest/userguide/Versioning.html)

AWS Documentation - Managing your storage lifecycle: "To manage your objects so that they are stored cost effectively throughout their lifecycle, configure their S3 Lifecycle." (URL: https://docs.aws.amazon.com/AmazonS3/latest/userguide/object-lifecycle-mgmt.html)

AWS Documentation - Protecting data using server-side encryption: "Server-side encryption protects data at rest. Amazon S3 encrypts your object as it writes it to disks in its data

centers and decrypts it for you when you access it." (URL:
https://docs.aws.amazon.com/AmazonS3/latest/userguide/serv-side-encryption.html)

Question: 26

An online gaming company needs to choose a purchasing option to run its Amazon EC2 instances for 1 year. The web traffic is consistent, and any increases in traffic are predictable. The EC2 instances must be online and available without any disruption. Which EC2 instance purchasing option will meet these requirements MOST cost-effectively?

A: On-Demand Instances

B: Reserved Instances

C: Spot Instances

D: Spot Fleet

Correct Answer:

B

Explanation:

Reserved Instances (RIs) are the most cost-effective option for workloads with consistent and predictable usage over a fixed term. By committing to a 1-year term, the company can achieve significant savings (up to 72%) compared to On-Demand pricing. RIs also provide a capacity reservation, which guarantees that the required EC2 instances will be available for the duration of the term, fulfilling the requirement for no disruptions. This model is specifically designed for steady-state workloads as described in the scenario.

Why Incorrect Options are Wrong:

On-Demand Instances: This option is not the most cost-effective for a predictable, 1-year workload. It is better suited for short-term, irregular workloads that cannot be interrupted.

Spot Instances: This option is unsuitable because Spot Instances can be terminated by AWS with a two-minute warning, which violates the requirement that instances must be available without disruption.

Spot Fleet: A Spot Fleet is a collection of Spot Instances and is therefore also subject to interruptions, making it inappropriate for a core workload that requires constant availability.

References:

AWS EC2 User Guide for Linux Instances, "Reserved Instances": "If you have a steady-state workload, using Reserved Instances can provide you with a significant discount compared to using On-Demand Instances." This directly aligns with the scenario's

consistent and predictable traffic. (URL: https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ec2-reserved-instances.html)

AWS EC2 Pricing Page: This page details the various purchasing options. It states that RIs are for "applications with steady state or predictable usage" and On-Demand is for "applications with short-term, spiky, or unpredictable workloads that cannot be interrupted." (URL: https://aws.amazon.com/ec2/pricing/)

AWS Well-Architected Framework, Cost Optimization Pillar, "Select the best pricing model": "For steady-state, predictable workloads, you can save money using Reserved Instances (RIs) or Savings Plans. For workloads that can be interrupted, you can save money using Spot Instances." (URL: https://docs.aws.amazon.com/wellarchitected/latest/cost-optimization-pillar/select-the-best-pricing-model.html)

Question: 27

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

Correct Answer:

C

Explanation:

AWS Cost Explorer is the correct tool for this purpose. It has an easy-to-use interface that lets you visualize, understand, and manage your AWS costs and usage over a specified time period. You can analyze your data at a high level (e.g., total costs and usage across all accounts) or dive deeper into your cost and usage data to identify trends, pinpoint cost drivers, and detect anomalies. It provides default reports and graphs to help visualize spending and usage patterns.

Why Incorrect Options are Wrong:

References:

AWS Cost Explorer: AWS Documentation. (n.d.). What is AWS Cost Explorer? AWS Billing and Cost Management. Retrieved from https://docs.aws.amazon.com/cost-management/latest/userguide/ce-what-is.html

AWS Organizations: AWS Documentation. (n.d.). What is AWS Organizations? AWS Organizations. Retrieved from https://docs.aws.amazon.com/organizations/latest/userguide/orgsintroduction.html

AWS Pricing Calculator: AWS. (n.d.). AWS Pricing Calculator. Retrieved from https://calculator.aws/

AWS Service Catalog: AWS Documentation. (n.d.). What is AWS Service Catalog? AWS Service Catalog. Retrieved from https://docs.aws.amazon.com/servicecatalog/latest/adminguide/introduction.html

Question: 28

A global company wants to migrate its third-party applications to the AWS Cloud. The company wants help from a global team of experts to complete the migration faster and more reliably in accordance with AWS internal best practices. Which AWS service or resource will meet these requirements?

A: AWS Support

B: AWS Professional Services

C: AWS Launch Wizard

D: AWS Managed Services (AMS)

Correct Answer:

B

Explanation:

AWS Professional Services is a global team of experts that helps customers realize their desired business outcomes when using the AWS Cloud. This organization assists enterprises with specific projects, such as large-scale application migrations. They provide guidance and hands-on keyboard help, leveraging AWS internal best practices, frameworks, and methodologies to ensure migrations are completed faster and more reliably. This directly aligns with the company's requirement for a global team of experts to assist with a migration project.

Why Incorrect Options are Wrong:

AWS Support: This service provides reactive technical assistance and troubleshooting for AWS services (break/fix), not proactive, project-based consulting for migrations.

AWS Launch Wizard: This is a self-service tool that guides the deployment of specific third-party applications, not a team of experts that manages a comprehensive migration project.

AWS Managed Services (AMS): This service focuses on the ongoing operational management of an AWS environment after migration, handling tasks like patching and monitoring, rather than the migration project itself.

References:

AWS Professional Services: AWS. (n.d.). AWS Professional Services. AWS Documentation. Retrieved from https://aws.amazon.com/professional-services/ ("Our global team of experts

can help you realize your desired business outcomes when using the AWS Cloud... We help you with the enterprise cloud adoption journey... including migrating legacy applications").

AWS Support: AWS. (n.d.). AWS Support. AWS Documentation. Retrieved from https://aws.amazon.com/premiumsupport/ (Describes support plans for technical assistance and troubleshooting).

AWS Launch Wizard: AWS. (n.d.). What is AWS Launch Wizard? AWS Documentation. Retrieved from https://docs.aws.amazon.com/launchwizard/latest/userguide/what-is-launch-wizard.html ("Launch Wizard offers a guided way of sizing, configuring, and deploying AWS resources for third-party applications").

AWS Managed Services (AMS): AWS. (n.d.). AWS Managed Services. AWS Documentation. Retrieved from https://aws.amazon.com/managed-services/ ("AWS Managed Services (AMS) helps you operate your AWS infrastructure more efficiently and securely... by providing ongoing management of your AWS infrastructure").

Question: 29

A developer wants to deploy an application quickly on AWS without manually creating the required resources. Which AWS service will meet these requirements?

A: Amazon EC2

B: AWS Elastic Beanstalk

C: AWS CodeBuild

D: Amazon Personalize

Correct Answer:

B

Explanation:

AWS Elastic Beanstalk is an orchestration service that automates the process of deploying and scaling applications. A developer can simply upload their application code, and Elastic Beanstalk automatically handles the deployment details, including capacity provisioning, load balancing, Auto Scaling, and application health monitoring. This allows for rapid application deployment without the need to manually create and configure the underlying AWS resources like Amazon EC2 instances, load balancers, or security groups.

Why Incorrect Options are Wrong:

Amazon EC2: This is an Infrastructure as a Service (IaaS) offering. It requires the developer to manually provision, configure, and manage the virtual servers and related resources, which contradicts the requirement.

AWS CodeBuild: This is a continuous integration service that compiles source code and runs tests. It is part of the development pipeline but does not deploy or manage the application infrastructure itself.

Amazon Personalize: This is a managed machine learning service for creating real-time personalized recommendations. It is not a service for general application deployment.

References:

AWS Elastic Beanstalk: AWS Documentation. (n.d.). What is AWS Elastic Beanstalk? AWS. Retrieved from https://docs.aws.amazon.com/elasticbeanstalk/latest/dg/Welcome.html

Amazon EC2: AWS Documentation. (n.d.). What is Amazon EC2? AWS. Retrieved from https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/concepts.html

AWS CodeBuild: AWS Documentation. (n.d.). What is AWS CodeBuild? AWS. Retrieved from https://docs.aws.amazon.com/codebuild/latest/userguide/welcome.html

Amazon Personalize: AWS Documentation. (n.d.). What is Amazon Personalize? AWS. Retrieved from https://docs.aws.amazon.com/personalize/latest/dg/what-is-personalize.html