

# Amazon-Web-Services

## Exam Questions CLF-C02

AWS Certified Cloud Practitioner



### NEW QUESTION 1

- (Topic 3)

A company plans to migrate to the AWS Cloud. The company is gathering information about its on-premises infrastructure and requires information such as the hostname, IP address, and MAC address.

Which AWS service will meet these requirements?

- A. AWS DataSync
- B. AWS Application Migration Service
- C. AWS Application Discovery Service
- D. AWS Database Migration Service (AWS DMS)

**Answer:** C

#### Explanation:

AWS Application Discovery Service is a service that helps you plan your migration to the AWS Cloud by collecting usage and configuration data about your on-premises servers and databases. This data includes information such as the hostname, IP address, and MAC address of each server, as well as the performance metrics, network connections, and processes running on them. You can use AWS Application Discovery Service to discover your on-premises inventory, map the dependencies between servers and applications, and estimate the cost and effort of migrating to AWS. You can also export the data to other AWS services, such as AWS Migration Hub and AWS Database Migration Service, to support your migration tasks. AWS Application Discovery Service offers two ways of performing discovery: agentless discovery and agent-based discovery. Agentless discovery uses a virtual appliance that you deploy on your VMware vCenter to collect data from your virtual machines and hosts. Agent-based discovery uses an agent that you install on each of your physical or virtual servers to collect data. You can choose the method that best suits your environment and needs. AWS DataSync is a service that helps you transfer data between your on-premises storage and AWS storage services, such as Amazon S3, Amazon EFS, and Amazon FSx for Windows File Server. AWS DataSync does not collect information about your on-premises infrastructure, but rather focuses on optimizing the data transfer speed, security, and reliability. AWS Application Migration Service is a service that helps you migrate your applications from your on-premises or cloud environment to AWS without making any changes to the applications, their architecture, or the migrated servers. AWS Application Migration Service does not collect information about your on-premises infrastructure, but rather uses a lightweight agent to replicate your servers as Amazon Machine Images (AMIs) and launch them as EC2 instances on AWS. AWS Database Migration Service is a service that helps you migrate your databases from your on-premises or cloud environment to AWS, either as a one-time migration or as a continuous replication. AWS Database Migration Service does not collect information about your on-premises infrastructure, but rather uses a source and a target endpoint to connect to your databases and transfer the data. References: AWS Application Discovery Service, AWS DataSync, AWS Application Migration Service, [AWS Database Migration Service]

### NEW QUESTION 2

- (Topic 3)

A company wants durable storage for static content and infinitely scalable data storage infrastructure at the lowest cost.

Which AWS service should the company choose?

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

**Answer:** B

#### Explanation:

Amazon S3 is a service that provides durable storage for static content and infinitely scalable data storage infrastructure at the lowest cost. Amazon S3 is an object storage service that allows you to store and retrieve any amount of data from anywhere on the internet. Amazon S3 offers industry-leading scalability, availability, and performance, as well as 99.999999999% (11 9s) of durability and multi-AZ resilience. Amazon S3 also provides various storage classes that offer different levels of performance and cost optimization, such as S3 Standard, S3 Intelligent-Tiering, S3 Standard-Infrequent Access (S3 Standard-IA), S3 One Zone-Infrequent Access (S3 One Zone-IA), and S3 Glacier<sup>456</sup>. Amazon S3 is ideal for storing static content, such as images, videos, documents, and web pages, as well as building data lakes, backup and archive solutions, big data analytics, and machine learning applications<sup>456</sup>. References: 4: Cloud Storage on AWS, 5: Object Storage - Amazon Simple Storage Service (S3) - AWS, 6: Amazon S3 Documentation

### NEW QUESTION 3

- (Topic 3)

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

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

**Answer:** AB

#### Explanation:

The AWS Cloud offers many benefits, such as:

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

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

### NEW QUESTION 4

- (Topic 3)

Which AWS service is a cloud security posture management (CSPM) service that aggregates alerts from various AWS services and partner products in a standardized format?

- A. AWS Security Hub
- B. AWS Trusted Advisor

- C. Amazon EventBridge
- D. Amazon GuardDuty

**Answer:** A

**Explanation:**

AWS Security Hub is a cloud security posture management (CSPM) service that performs security best practice checks, aggregates alerts, and enables automated remediation. Security Hub collects findings from the security services enabled across your AWS accounts, such as intrusion detection findings from Amazon GuardDuty, vulnerability scans from Amazon Inspector, and sensitive data identification findings from Amazon Macie. Security Hub also collects findings from partner security products using a standardized AWS Security Finding Format, eliminating the need for time-consuming data parsing and normalization efforts. Customers can designate an administrator account that can access all findings across their accounts. References: AWS Security Hub Overview, AWS Security Hub FAQs

**NEW QUESTION 5**

- (Topic 3)

What is the LEAST expensive AWS Support plan that provides the full set of AWS Trusted Advisor best practice checks for cost optimization?

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

**Answer:** B

**Explanation:**

AWS Business Support is the least expensive AWS Support plan that provides the full set of AWS Trusted Advisor best practice checks for cost optimization. AWS Trusted Advisor is a service that provides best practices and recommendations for cost optimization, performance, security, and fault tolerance. AWS Business Support also provides other benefits, such as 24/7 technical support, unlimited cases, and faster response times. AWS Enterprise Support is the most expensive AWS Support plan that provides the same benefits as AWS Business Support, plus additional benefits, such as a technical account manager and enterprise concierge support. AWS Developer Support and AWS Basic Support are cheaper AWS Support plans that provide only a limited set of AWS Trusted Advisor best practice checks for cost optimization .

**NEW QUESTION 6**

- (Topic 3)

A company is migrating its applications from on-premises to the AWS Cloud. The company wants to ensure that the applications are assigned only the minimum permissions that are needed to perform all operations. Which AWS service will meet these requirements'?

- A. AWS Identity and Access Management (IAM)
- B. Amazon CloudWatch
- C. Amazon Macie
- D. Amazon GuardDuty

**Answer:** A

**Explanation:**

AWS Identity and Access Management (IAM) is a service that helps you securely control access to AWS resources for your users. You use IAM to control who can use your AWS resources (authentication) and what resources they can use and in what ways (authorization). IAM also enables you to follow the principle of least privilege, which means granting only the permissions that are necessary to perform a task<sup>1</sup>. References: AWS Identity and Access Management (IAM) - AWS Documentation

**NEW QUESTION 7**

- (Topic 3)

A company uses AWS Organizations. The company wants to apply security best practices from the AWS Well-Architected Framework to all of its AWS accounts. Which AWS service will meet these requirements?

- A. Amazon Macie
- B. Amazon Detective
- C. AWS Control Tower
- D. AWS Secrets Manager

**Answer:** C

**Explanation:**

AWS Control Tower is the easiest way to set up and govern a secure, multi- account AWS environment based on best practices established through AWS's experience working with thousands of enterprises as they move to the cloud. With AWS Control Tower, builders can provision new AWS accounts in a few clicks, while you have peace of mind knowing your accounts conform to your organization's policies. AWS Control Tower automates the setup of a baseline environment, or landing zone, that is a secure, well- architected multi-account AWS environment<sup>1</sup>. AWS Control Tower helps you apply security best practices from the AWS Well-Architected Framework to all of your AWS accounts<sup>2</sup>.

**NEW QUESTION 8**

- (Topic 3)

A company wants to ensure that all of its Amazon EC2 instances have compliant operating system patches. Which AWS service will meet these requirements?

- A. AWS Compute Optimizer
- B. AWS Elastic Beanstalk
- C. AWS AppSync
- D. AWS Systems Manager

**Answer:** D

**Explanation:**

AWS Systems Manager gives you visibility and control of your infrastructure on AWS. Systems Manager provides a unified user interface so you can view operational data from multiple AWS services and allows you to automate operational tasks across your AWS resources. You can use Systems Manager to apply OS patches, create system images, configure Windows and Linux operating systems, and execute PowerShell commands<sup>5</sup>. Systems Manager can help you ensure that all of your Amazon EC2 instances have compliant operating system patches by using the Patch Manager feature.

**NEW QUESTION 9**

- (Topic 3)

A company needs to migrate a PostgreSQL database from on-premises to Amazon RDS. Which AWS service or tool should the company use to meet this requirement?

- A. Cloud Adoption Readiness Tool
- B. AWS Migration Hub
- C. AWS Database Migration Service (AWS DMS)
- D. AWS Application Migration Service

**Answer:** C

**Explanation:**

AWS Database Migration Service (AWS DMS) is a managed and automated service that helps you migrate your databases from your on-premises or cloud environment to AWS, either as a one-time migration or as a continuous replication. AWS DMS supports migration between 20-plus database and analytics engines, such as PostgreSQL, Oracle, MySQL, SQL Server, MongoDB, Amazon Aurora, Amazon RDS, Amazon Redshift, and Amazon S3. AWS DMS also provides schema conversion and validation tools, as well as monitoring and security features. AWS DMS is a cost-effective and reliable solution for database migration, as you only pay for the compute resources and additional log storage used during the migration process, and you can minimize the downtime and data loss with

Multi-AZ and ongoing replication<sup>12</sup>

To migrate a PostgreSQL database from on-premises to Amazon RDS using AWS DMS, you need to perform the following steps:

? Create an AWS DMS replication instance in the same AWS Region as your target Amazon RDS PostgreSQL DB instance. The replication instance is a server that runs the AWS DMS replication software and connects to your source and target endpoints. You can choose the instance type, storage, and network settings based on your migration requirements<sup>3</sup>

? Create a source endpoint that points to your on-premises PostgreSQL database.

You need to provide the connection details, such as the server name, port, database name, user name, and password. You also need to specify the engine name as postgres and the SSL mode as required<sup>4</sup>

? Create a target endpoint that points to your Amazon RDS PostgreSQL DB instance. You need to provide the connection details, such as the server name, port, database name, user name, and password. You also need to specify the engine name as postgres and the SSL mode as verify-full.

? Create a migration task that defines the migration settings and options, such as the replication instance, the source and target endpoints, the migration type (full load, full load and change data capture, or change data capture only), the table mappings, the task settings, and the task monitoring role. You can also use the AWS Schema Conversion Tool (AWS SCT) to convert your source schema to the target schema and apply it to the target endpoint before or after creating the migration task.

? Start the migration task and monitor its progress and status using the AWS DMS console, the AWS CLI, or the AWS DMS API. You can also use AWS CloudFormation to automate the creation and execution of the migration task.

The other options are not suitable for migrating a PostgreSQL database from on-premises to Amazon RDS. Cloud Adoption Readiness Tool is a tool that helps you assess your readiness for cloud adoption based on six dimensions: business, people, process, platform, operations, and security. It does not perform any database migration tasks. AWS Migration Hub is a service that helps you track and manage the progress of your application migrations across multiple AWS and partner services, such as AWS DMS, AWS Application Migration Service, AWS Server Migration Service, and CloudEndure Migration. It does not perform any database migration tasks itself, but rather integrates with other migration services. AWS Application Migration Service is a service that helps you migrate your applications from your on-premises or cloud environment to AWS without making any changes to the applications, their architecture, or the migrated servers. It does not support database migration, but rather replicates your servers as Amazon Machine Images (AMIs) and launches them as EC2 instances on AWS. References: AWS Database Migration Service, What is AWS Database Migration Service?, Working with an AWS DMS replication instance, Creating source and target endpoints for PostgreSQL, [Creating a target endpoint for Amazon RDS for PostgreSQL], [Creating a migration task for AWS DMS], [AWS Schema Conversion Tool], [Starting a migration task for AWS DMS], [AWS CloudFormation], [Cloud Adoption Readiness Tool], [AWS Migration Hub], [AWS Application Migration Service]

**NEW QUESTION 10**

- (Topic 3)

Which AWS service or feature offers security for a VPC by acting as a firewall to control traffic in and out of subnets?

- A. AWS Security Hub
- B. Security groups
- C. Network ACL
- D. AWSWAF

**Answer:** C

**Explanation:**

A network access control list (network ACL) is a feature that acts as a firewall for controlling traffic in and out of one or more subnets in a virtual private cloud (VPC). AWS Security Hub is a service that provides a comprehensive view of the security posture of AWS accounts and resources. Security groups are features that act as firewalls for controlling traffic at the instance level. AWS WAF is a web application firewall that helps protect web applications from common web exploits.

**NEW QUESTION 10**

- (Topic 3)

Which of the following is a benefit that AWS Professional Services provides?

- A. Management of the ongoing security of user data
- B. Advisory solutions for AWS adoption
- C. Technical support 24 hours a day, 7 days a week
- D. Monitoring of monthly billing costs in AWS accounts



**Answer:** B

**Explanation:**

AWS Professional Services is a team of experts that help customers achieve their desired outcomes using the AWS Cloud. One of the benefits that AWS Professional Services provides is advisory solutions for AWS adoption, which include guidance on cloud strategy, architecture, migration, and innovation<sup>2</sup>. Management of the ongoing security of user data, technical support 24 hours a day, 7 days a week, and monitoring of monthly billing costs in AWS accounts are not benefits that AWS Professional Services provides, as they are either the responsibility of the customer or the features of other AWS services or support plans<sup>3</sup>

**NEW QUESTION 12**

- (Topic 3)

A company wants to receive alerts to monitor its overall operating costs for its AWS public cloud infrastructure.

Which AWS offering will meet these requirements?

- A. Amazon EventBridge
- B. Compute Savings Plans
- C. AWS Budgets
- D. Migration Evaluator

**Answer:** C

**Explanation:**

AWS Budgets is a service that enables you to plan your service usage, service costs, and instance reservations. You can use AWS Budgets to create custom budgets that alert you when your costs or usage exceed (or are forecasted to exceed) your budgeted amount. You can also use AWS Budgets to monitor how close your usage and costs are to meeting your reservation purchases<sup>1</sup>

**NEW QUESTION 16**

- (Topic 3)

A company wants its AWS usage to be more sustainable. The company wants to track, measure, review, and forecast polluting emissions that result from its AWS applications.

Which AWS service or tool can the company use to meet these requirements?

- A. AWS Health Dashboard
- B. AWS customer carbon footprint tool
- C. AWS Support Center
- D. Amazon QuickSight

**Answer:** B

**Explanation:**

AWS customer carbon footprint tool is a tool that helps customers measure and manage their carbon emissions from their AWS usage. It provides data on the carbon intensity, energy consumption, and estimated emissions of AWS services across regions and time periods. It also enables customers to review and forecast their emissions, and compare them with industry benchmarks. AWS Health Dashboard is a service that provides personalized information about the health and performance of AWS services and resources. AWS Support Center is a service that provides access to AWS support resources, such as cases, forums, and documentation. Amazon QuickSight is a service that provides business intelligence and analytics for AWS data sources.

**NEW QUESTION 19**

- (Topic 3)

A company has a MySQL database running on a single Amazon EC2 instance. The company now requires higher availability in the event of an outage.

Which set of tasks would meet this requirement?

- A. Add an Application Load Balancer in front of the EC2 instance.
- B. Configure EC2 Auto Recovery to move the instance to another Availability Zone.
- C. Migrate to Amazon RDS and enable Multi-AZ.
- D. Enable termination protection for the EC2 instance to avoid outages.

**Answer:** C

**Explanation:**

The set of tasks that would meet the requirement of having higher availability for a MySQL database running on a single Amazon EC2 instance is to migrate to Amazon RDS and enable Multi-AZ. Amazon RDS is a fully managed relational database service that supports MySQL and other popular database engines. By enabling Multi-AZ, users can have a primary database in one Availability Zone and a synchronous standby replica in another Availability Zone. In case of a planned or unplanned outage of the primary database, Amazon RDS automatically fails over to the standby replica with minimal disruption<sup>3</sup>. Adding an Application Load Balancer in front of the EC2 instance, configuring EC2 Auto Recovery to move the instance to another Availability Zone, or enabling termination protection for the EC2 instance would not provide higher availability for the database, as they do not address the single point of failure or data replication issues.

**NEW QUESTION 22**

- (Topic 3)

A company needs to identify who accessed an AWS service and what action was performed for a given time period.

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

- A. Amazon CloudWatch
- B. AWS CloudTrail
- C. AWS Security Hub
- D. Amazon Inspector

**Answer:** B

**Explanation:**

AWS CloudTrail is a service that enables governance, compliance, operational auditing, and risk auditing of your AWS account. With CloudTrail, you can log, continuously monitor, and retain account activity related to actions across your AWS infrastructure. You can use CloudTrail to identify who accessed an AWS service and what action was performed for a given time period. Amazon CloudWatch, AWS Security Hub, and Amazon Inspector are AWS services that provide different types of monitoring and security capabilities.

**NEW QUESTION 23**

- (Topic 3)

A company needs to run a workload for several batch image rendering applications. It is acceptable for the workload to experience downtime. Which Amazon EC2 pricing model would be MOST cost-effective in this situation?

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

**Answer:** D

**Explanation:**

Amazon EC2 Spot Instances are instances that use spare EC2 capacity that is available at up to a 90% discount compared to On-Demand prices. You can use Spot Instances for various stateless, fault-tolerant, or flexible applications such as big data, containerized workloads, high-performance computing (HPC), and test & development workloads. Spot Instances are ideal for workloads that can be interrupted, such as batch image rendering applications<sup>1</sup>. On-Demand Instances are instances that let you pay for compute capacity by the hour or second (minimum of 60 seconds) with no long-term commitments. This frees you from the costs and complexities of planning, purchasing, and maintaining hardware and transforms what are commonly large fixed costs into much smaller variable costs<sup>2</sup>. Reserved Instances are instances that provide you with a significant discount (up to 75%) compared to On-Demand Instance pricing. In exchange, you select a term and make an upfront payment to reserve a certain amount of compute capacity for that term<sup>3</sup>. Dedicated Instances are instances that run in a VPC on hardware that's dedicated to a single customer. Your Dedicated Instances are physically isolated at the host hardware level from instances that belong to other AWS accounts<sup>4</sup>.

**NEW QUESTION 24**

- (Topic 3)

A company has deployed an application in the AWS Cloud. The company wants to ensure that the application is highly resilient. Which component of AWS infrastructure can the company use to meet this requirement?

- A. Content delivery network (CDN)
- B. Edge locations
- C. Wavelength Zones
- D. Availability Zones

**Answer:** D

**Explanation:**

Availability Zones are components of AWS infrastructure that can help the company ensure that the application is highly resilient. Availability Zones are multiple, isolated locations within each AWS Region. Each Availability Zone has independent power, cooling, and physical security, and is connected to the other Availability Zones in the same Region via low-latency, high-throughput, and highly redundant networking. Availability Zones allow you to operate production applications and databases that are more highly available, fault tolerant, and scalable than would be possible from a single data center.

**NEW QUESTION 29**

- (Topic 3)

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

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

**Answer:** CE

**Explanation:**

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

**NEW QUESTION 30**

- (Topic 3)

Which AWS service or feature is associated with a subnet in a VPC and is used to control inbound and outbound traffic?

- A. Amazon Inspector
- B. Network ACLs
- C. AWS Shield
- D. VPC Flow Logs

**Answer:** B

**Explanation:**

Network ACLs (network access control lists) are an optional layer of security for your VPC that act as a firewall for controlling traffic in and out of one or more subnets. You can use network ACLs to allow or deny traffic based on protocol, port, or source and destination IP address. Network ACLs are stateless, meaning that they do not track the traffic that flows through them. Therefore, you must create rules for both inbound and outbound traffic.

**NEW QUESTION 33**

- (Topic 3)

A company needs to store infrequently used data for data archives and long-term backups.

A company needs a history report about how its Amazon EC2 instances were modified last month.

Which AWS service can be used to meet this requirement?

- A. AWS Service Catalog
- B. AWS Config
- C. Amazon CloudWatch
- D. AWS Artifact

**Answer: B**

**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 also track changes to your EC2 instances over time and provide a history report of the modifications. AWS Service Catalog, Amazon CloudWatch, and AWS Artifact are not the best services to meet this requirement. AWS Service Catalog is a service that allows you to create and manage catalogs of IT services that are approved for use on AWS. Amazon CloudWatch is a service that monitors your AWS resources and applications and provides metrics, alarms, dashboards, and logs. AWS Artifact is a service that provides on-demand access to AWS security and compliance reports and online agreements

**NEW QUESTION 36**

- (Topic 3)

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

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

**Answer: A**

**Explanation:**

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

**NEW QUESTION 38**

- (Topic 3)

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

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

**Answer: C**

**Explanation:**

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

**NEW QUESTION 39**

- (Topic 3)

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

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

**Answer: D**

**Explanation:**

Spot Instances are Amazon EC2 instances that are available at a discounted price compared to On-Demand pricing. Spot Instances use spare EC2 capacity that is not being used by other customers, and the price fluctuates based on supply and demand. Customers can request Spot Instances for their applications and

specify the maximum price they are willing to pay per hour. If the Spot price is lower than the customer's bid, the Spot Instance is launched and the customer pays the current Spot price. However, if the Spot price rises above the customer's bid, the Spot Instance is terminated by AWS and the customer is charged for the partial hour of usage. Therefore, Spot Instances can provide discounts of up to 90% or more, but they are not suitable for applications that require continuous or predictable availability. Spot Instances are recommended for applications that are flexible, fault-tolerant, or have low priority, such as batch processing, data analysis, or testing and development.

**NEW QUESTION 43**

- (Topic 3)

A company is running a workload in the AWS Cloud.

Which AWS best practice ensures the MOST cost-effective architecture for the workload?

- A. Loose coupling
- B. Rightsizing
- C. Caching
- D. Redundancy

**Answer:** B

**Explanation:**

The AWS best practice that ensures the most cost-effective architecture for the workload is rightsizing. Rightsizing means selecting the most appropriate instance type or resource configuration that matches the needs of the workload. Rightsizing can help optimize performance and reduce costs by avoiding over-provisioning or under-provisioning of resources<sup>1</sup>. Loose coupling, caching, and redundancy are other AWS best practices that can improve the scalability, availability, and performance of the workload, but they do not necessarily ensure the most cost-effective architecture.

**NEW QUESTION 46**

- (Topic 3)

A company wants to launch its web application in a second AWS Region. The company needs to determine which services must be regionally configured for this launch.

Which AWS services can be configured at the Region level? (Select TWO.)

- A. Amazon EC2
- B. Amazon Route 53
- C. Amazon CloudFront
- D. AWS WAF
- E. Amazon DynamoDB

**Answer:** BD

**Explanation:**

Amazon Route 53 and AWS WAF are AWS services that can be configured at the Region level. Amazon Route 53 is a highly available and scalable cloud Domain Name System (DNS) web service that lets you register domain names, route traffic to resources, and check the health of your resources. 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. Amazon EC2, Amazon CloudFront, and Amazon DynamoDB are AWS services that can be configured at the global level or the Availability Zone level .

**NEW QUESTION 51**

- (Topic 3)

Which AWS service provides the ability to manage infrastructure as code?

- A. AWS CodePipeline
- B. AWS CodeDeploy
- C. AWS Direct Connect
- D. AWS CloudFormation

**Answer:** D

**Explanation:**

The AWS service that provides the ability to manage infrastructure as code is AWS CloudFormation. Infrastructure as code is a process of defining and provisioning AWS resources using code or templates, rather than manual actions or scripts. AWS CloudFormation allows you to create and update stacks of AWS resources based on predefined templates that describe the desired state and configuration of the resources. AWS CloudFormation automates and simplifies the deployment and management of AWS resources, and ensures consistency and repeatability across different environments and regions. AWS CloudFormation also supports rollback, change sets, drift detection, and nested stacks features that help you to monitor and control the changes to your infrastructure<sup>1</sup>.

**NEW QUESTION 54**

- (Topic 3)

A company wants to integrate natural language processing (NLP) into business intelligence (BI) dashboards. The company wants to ask questions and receive answers with relevant visualizations.

Which AWS service or tool will meet these requirements?

- A. Amazon Macie
- B. Amazon Rekognition
- C. Amazon QuickSight Q
- D. Amazon Lex

**Answer:** C

**Explanation:**

Amazon QuickSight Q is a natural language query feature that allows users to ask questions about their data and receive answers in the form of relevant visualizations<sup>1</sup>. Amazon Macie is a data security and data privacy service that uses machine learning and pattern matching to discover and protect sensitive data



in AWS2. Amazon Rekognition is a computer vision service that can analyze images and videos for faces, objects, scenes, text, and more3. Amazon Lex is a service for building conversational interfaces using voice and text4.

**NEW QUESTION 56**

- (Topic 3)

Which AWS service can identify when an Amazon EC2 instance was terminated?

- A. AWS Identity and Access Management (IAM)
- B. AWS CloudTrail
- C. AWS Compute Optimizer
- D. Amazon EventBridge

**Answer:** B

**Explanation:**

AWS CloudTrail is the AWS service that can identify when an Amazon EC2 instance was terminated. AWS CloudTrail is a service that records API calls and events for AWS accounts and resources. AWS CloudTrail can capture the `TerminateInstances` event, which is triggered when an EC2 instance is terminated by a user or an AWS service. The event contains information such as the instance ID, the user identity, the source IP address, the time, and the reason for the termination12. Customers can use the CloudTrail console, the AWS CLI, or the AWS SDKs to view and search for the `TerminateInstances` events in their event history or in their S3 buckets where they store their CloudTrail logs13.

**NEW QUESTION 57**

- (Topic 3)

A company is operating several factories where it builds products. The company needs the ability to process data, store data, and run applications with local system interdependencies that require low latency.

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

- A. AWS IoT Greengrass
- B. AWS Lambda
- C. AWS Outposts
- D. AWS Snowball Edge

**Answer:** C

**Explanation:**

AWS Outposts is a service that provides fully managed AWS infrastructure and services on premises. It allows users to run applications that require low latency and local data processing, while seamlessly connecting to the AWS Cloud for a consistent hybrid experience. AWS IoT Greengrass is a service that provides local compute, messaging, data caching, sync, and ML inference capabilities for connected devices. AWS Lambda is a service that allows users to run code without provisioning or managing servers. AWS Snowball Edge is a device that provides a petabyte-scale data transport and edge computing solution.

**NEW QUESTION 58**

- (Topic 3)

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

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

**Answer:** CE

**Explanation:**

VPC peering and AWS Transit Gateway are two AWS services or features that give users the ability to create a network connection between two VPCs. VPC peering is a networking connection between two VPCs that enables you to route traffic between them privately. You can create a VPC peering connection between your own VPCs, with a VPC in another AWS account, or with a VPC in a different AWS Region. Traffic between peered VPCs never traverses the public internet. VPC peering does not support transitive peering relationships, which means that if VPC A is peered with VPC B, and VPC B is peered with VPC C, then VPC A and VPC C are not automatically peered789. AWS Transit Gateway is a networking service that acts as a regional router for your VPCs and on- premises networks. You can attach up to 5,000 VPCs and VPN connections to a single transit gateway and route traffic between them. AWS Transit Gateway simplifies the management and scalability of your network architecture, as you only need to create and manage a single connection from the central transit gateway to each connected network. AWS Transit Gateway supports transitive routing, which means that any network that is attached to the transit gateway can communicate with any other network that is attached to the same transit gateway . References: 7: VPC peering - Amazon Virtual Private Cloud, 8: Connect VPCs using VPC peering - Amazon Virtual Private Cloud, 9: Amazon VPC-to-Amazon VPC connectivity options - Amazon Virtual Private Cloud, : [AWS Transit Gateway - Amazon Web Services], : [Connect VPCs using AWS Transit Gateway - Amazon Virtual Private Cloud], : [AWS Transit Gateway: Simplify Your Network Architecture]

**NEW QUESTION 63**

- (Topic 3)

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

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

**Answer:** BE

**Explanation:**

Elasticity in the AWS Cloud refers to the ability to acquire resources as you need them and release resources when you no longer need them. In the cloud, you

want to do this automatically<sup>1</sup>. This means that you can rightsized resources as demand shifts, and you can easily procure resources when they are needed. Elasticity is not related to how quickly an Amazon EC2 instance can be restarted, the maximum amount of RAM an Amazon EC2 instance can use, or the pay-as-you-go billing model. These are aspects of scalability, performance, and cost, respectively<sup>2</sup>. For more information on elasticity, you can refer to the following sources:

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

**NEW QUESTION 64**

- (Topic 3)

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

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

**Answer:** B

**Explanation:**

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

**NEW QUESTION 69**

- (Topic 3)

Which AWS service or feature enables users to encrypt data at rest in Amazon S3?

- A. IAM policies
- B. Server-side encryption
- C. Amazon GuardDuty
- D. Client-side encryption

**Answer:** B

**Explanation:**

Server-side encryption is an encryption option that Amazon S3 provides to encrypt data at rest in Amazon S3. With server-side encryption, Amazon S3 encrypts an object before saving it to disk in its data centers and decrypts it when you download the objects. You have three server-side encryption options to choose from: SSE-S3, SSE-C, and SSE-KMS. SSE-S3 uses keys that are managed by Amazon S3. SSE-C allows you to manage your own encryption keys. SSE-KMS uses keys that are managed by AWS Key Management Service (AWS KMS)<sup>5</sup>.

**NEW QUESTION 71**

- (Topic 3)

A company wants to launch multiple workloads on AWS. Each workload is related to a different business unit. The company wants to separate and track costs for each business unit.

Which solution will meet these requirements with the LEAST operational overhead?

- A. Use AWS Organizations and create one account for each business unit.
- B. Use a spreadsheet to control the owners and cost of each resource.
- C. Use an Amazon DynamoDB table to record costs for each business unit.
- D. Use the AWS Billing console to assign owners to resources and track costs.

**Answer:** A

**Explanation:**

AWS Organizations is a service that helps you centrally manage and govern your AWS environment. You can use AWS Organizations to create multiple accounts for different business units, and group them into organizational units (OUs) that reflect your organizational structure<sup>1</sup>. By doing so, you can separate and track costs for each business unit using the account ID as a cost allocation tag<sup>2</sup>. You can also use AWS Organizations to apply policies and controls to your accounts, such as service control policies (SCPs) and tag policies<sup>1</sup>.

The other options are not suitable for meeting the requirements with the least operational overhead. Using a spreadsheet or a DynamoDB table to control and record costs for each business unit would require manual data entry and maintenance, which is prone to errors and inconsistencies. Using the AWS Billing console to assign owners to resources and track costs would also require manual tagging of each resource, which is time-consuming and inefficient.

References:

? 1: What Is AWS Organizations? - AWS Organizations

? 2: Cost Tagging and Reporting with AWS Organizations | AWS Cloud Financial Management

**NEW QUESTION 74**

- (Topic 3)

A company is running a monolithic on-premises application that does not scale and is difficult to maintain. The company has a plan to migrate the application to AWS and divide the application into microservices.

Which best practice of the AWS Well-Architected Framework is the company following with this plan?

- A. Integrate functional testing as part of AWS deployment.
- B. Use automation to deploy changes.
- C. Deploy the application to multiple locations.
- D. Implement loosely coupled dependencies.

**Answer:** D

**Explanation:**

The company is following the best practice of implementing loosely coupled dependencies by migrating the application to AWS and dividing the application into microservices. Loosely coupled dependencies are a design principle of the AWS Well-Architected Framework that helps to reduce the interdependencies between components and improve the scalability, reliability, and performance of the system. By breaking down the monolithic application into smaller, independent, and modular services, the company can reduce the complexity and maintenance costs, increase the agility and flexibility, and enable faster and more frequent deployments. AWS CloudFormation is an AWS service that provides the ability to manage infrastructure as code. Infrastructure as code is a process of defining and provisioning AWS resources using code or templates, rather than manual actions or scripts. AWS CloudFormation allows users to create and update stacks of AWS resources based on predefined templates that describe the desired state and configuration of the resources. AWS CloudFormation automates and simplifies the deployment and management of AWS resources, and ensures consistency and repeatability across different environments and regions. AWS CloudFormation also supports rollback, change sets, drift detection, and nested stacks features that help users to monitor and control the changes to their infrastructure. References: Implementing Loosely Coupled Dependencies, What is AWS CloudFormation?

**NEW QUESTION 76**

- (Topic 3)

A company's application has high customer usage during certain times of the day. The company wants to reduce the number of Amazon EC2 instances that run when application usage is low.

Which AWS service or instance purchasing option should the company use to meet this requirement?

- A. EC2 Instance Savings Plans
- B. Spot Instances
- C. Reserved Instances
- D. Amazon EC2 Auto Scaling

**Answer: D**

**Explanation:**

Amazon EC2 Auto Scaling is an AWS service that can help users reduce the number of Amazon EC2 instances that run when application usage is low. Amazon EC2 Auto Scaling allows users to create scaling policies that automatically adjust the number of EC2 instances based on the demand or a schedule. EC2 Instance Savings Plans, Spot Instances, and Reserved Instances are instance purchasing options that can help users save money on EC2 usage, but they do not automatically scale the number of instances according to the application usage.

**NEW QUESTION 79**

- (Topic 3)

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

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

**Answer: BE**

**Explanation:**

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

**NEW QUESTION 83**

- (Topic 3)

What is a customer responsibility when using AWS Lambda according to the AWS shared responsibility model?

- A. Managing the code within the Lambda function
- B. Confirming that the hardware is working in the data center
- C. Patching the operating system
- D. Shutting down Lambda functions when they are no longer in use

**Answer: A**

**Explanation:**

According to the AWS shared responsibility model, AWS is responsible for the security of the cloud, while customers are responsible for the security in the cloud. This means that AWS is responsible for the physical servers, networking, and operating system that run Lambda functions, while customers are responsible for the security of their code and AWS IAM to the Lambda service and within their function. Customers need to manage the code within the Lambda function, such as writing, testing, debugging, deploying, and updating the code, as well as ensuring that the code does not contain any vulnerabilities or malicious code that could compromise the security or performance of the function. References: 2: AWS Lambda - Amazon Web Services (AWS), 3: AWS Lambda Documentation, 1: Amazon CLF-C02: What is customer responsibility under AWS ... - PUPUWEB

**NEW QUESTION 85**

- (Topic 3)

A company wants to design a reliable web application that is hosted on Amazon EC2. Which approach will achieve this goal?

- A. Launch large EC2 instances in the same Availability Zone.
- B. Spread EC2 instances across more than one security group.
- C. Spread EC2 instances across more than one Availability Zone.
- D. Use an Amazon Machine Image (AMI) from AWS Marketplace.

**Answer:** C

**Explanation:**

The approach that will achieve the goal of designing a reliable web application that is hosted on Amazon EC2 is to spread EC2 instances across more than one Availability Zone. An Availability Zone is a physically isolated location within an AWS Region that has its own power, cooling, and network connectivity. By spreading EC2 instances across multiple Availability Zones, users can increase the fault tolerance and availability of their web applications, as well as reduce latency for end users<sup>2</sup>. Launching large EC2 instances in the same Availability Zone, spreading EC2 instances across more than one security group, or using an Amazon Machine Image (AMI) from AWS Marketplace are not sufficient to ensure reliability, as they do not provide redundancy or resilience in case of an outage in one Availability Zone.

**NEW QUESTION 87**

- (Topic 3)

A company has teams that have different job roles and responsibilities. The company's employees often change teams. The company needs to manage permissions for the employees so that the permissions are appropriate for the job responsibilities.

Which IAM resource should the company use to meet this requirement with the LEAST operational overhead?

- A. IAM user groups
- B. IAM roles
- C. IAM instance profiles
- D. IAM policies for individual users

**Answer:** B

**Explanation:**

IAM roles are a way of granting temporary permissions to entities that need to access AWS resources, such as users, applications, or services. IAM roles allow customers to assign permissions to entities without having to create or manage IAM users or credentials for them. IAM roles can be assumed by different entities depending on the trust policy attached to the role. For example, IAM roles can be assumed by IAM users in the same or different AWS accounts, AWS services such as EC2 or Lambda, or external identities such as federated users or web identities. IAM roles can also be switched by IAM users to temporarily change their permissions. IAM roles are recommended for managing permissions for employees who often change teams, because they allow customers to define permissions based on job roles and responsibilities, and easily assign or revoke them as needed. IAM roles also reduce the operational overhead of creating, updating, or deleting IAM users or credentials for each employee or team change.

**NEW QUESTION 89**

- (Topic 3)

A company wants an AWS service to provide product recommendations based on its customer data.

Which AWS service will meet this requirement?

- A. Amazon Polly
- B. Amazon Personalize
- C. Amazon Comprehend
- D. Amazon Rekognition

**Answer:** B

**Explanation:**

Amazon Personalize is an AWS service that helps developers quickly build and deploy a custom recommendation engine with real-time personalization and user segmentation<sup>1</sup>. It uses machine learning (ML) to analyze customer data and provide relevant recommendations based on their preferences, behavior, and context. Amazon Personalize can be used for various use cases such as optimizing recommendations, targeting customers more accurately, maximizing the value of unstructured text, and promoting items using business rules<sup>1</sup>.

The other options are not suitable for providing product recommendations based on customer data. Amazon Polly is a service that converts text into lifelike speech. Amazon Comprehend is a service that uses natural language processing (NLP) to extract insights from text and documents. Amazon Rekognition is a service that uses computer vision (CV) to analyze images and videos for faces, objects, scenes, and activities.

References:

- ? 1: Cloud Products - Amazon Web Services (AWS)
- ? 2: Recommender System – Amazon Personalize – Amazon Web Services
- ? 3: Top 25 AWS Services List 2023 - GeeksforGeeks
- ? 4: AWS to Azure services comparison - Azure Architecture Center
- ? 5: The 25+ Best AWS Cost Optimization Tools (Updated 2023) - CloudZero
- ? 6: Amazon Polly – Text-to-Speech Service - AWS
- ? 7: Natural Language Processing - Amazon Comprehend - AWS
- ? 8: Image and Video Analysis - Amazon Rekognition - AWS

**NEW QUESTION 90**

- (Topic 3)

What is the purpose of having an internet gateway within a VPC?

- A. To create a VPN connection to the VPC
- B. To allow communication between the VPC and the internet
- C. To impose bandwidth constraints on internet traffic
- D. To load balance traffic from the internet across Amazon EC2 instances

**Answer:** B

**Explanation:**

An internet gateway is a service that allows for internet traffic to enter into a VPC. Otherwise, a VPC is completely segmented off and then the only way to get to it is potentially through a VPN connection rather than through internet connection. An internet gateway is a logical connection between an AWS VPC and the internet. It supports IPv4 and IPv6 traffic. It does not cause availability risks or bandwidth constraints on your network traffic<sup>1</sup>. An internet gateway enables resources in your public subnets (such as EC2 instances) to connect to the internet if the resource has a public IPv4 address or an IPv6 address. Similarly, resources on the internet can initiate a connection to resources in your subnet using the public IPv4 address or IPv6 address<sup>2</sup>. An internet gateway also provides a target in your VPC route tables for internet-routable traffic. For communication using IPv4, the internet gateway also performs network address translation (NAT).



For communication using IPv6, NAT is not needed because IPv6 addresses are public<sup>2</sup>. To enable access to or from the internet for instances in a subnet in a VPC using an internet gateway, you must create an internet gateway and attach it to your VPC, add a route to your subnet's route table that directs internet-bound traffic to the internet gateway, ensure that instances in your subnet have a public IPv4 address or an IPv6 address, and ensure that your network access control lists and security group rules allow the desired internet traffic to flow to and from your instance<sup>2</sup>. References: Connect to the internet using an internet gateway, AWS Internet Gateway and VPC Routing

**NEW QUESTION 91**

- (Topic 3)

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

- A. Configure the AWS provided security group firewall.
- B. Classify company assets in the AWS Cloud.
- C. Determine which Availability Zones to use for Amazon S3 buckets.
- D. Patch or upgrade Amazon DynamoDB.
- E. Select Amazon EC2 instances to run AWS Lambda on.
- F. AWS Config

**Answer:** AB

**Explanation:**

According to the AWS shared responsibility model, the customer is responsible for security in the cloud, which includes the tasks of configuring the AWS provided security group firewall and classifying company assets in the AWS Cloud. A security group is a virtual firewall that controls the inbound and outbound traffic for one or more EC2 instances. The customer must configure the security group rules to allow or deny traffic based on protocol, port, or source and destination IP address<sup>2</sup> Classifying company assets in the AWS Cloud means identifying the types, categories, and sensitivity levels of the data and resources that the customer stores and processes on AWS. The customer must also determine the applicable compliance requirements and regulations that apply to their assets, and implement the appropriate security controls and measures to protect them

**NEW QUESTION 92**

- (Topic 3)

A company needs an automated vulnerability management service that continually scans AWS workloads for software vulnerabilities. Which AWS service will meet these requirements?

- A. Amazon GuardDuty
- B. Amazon Inspector
- C. AWS Security Hub
- D. AWS Shield

**Answer:** B

**Explanation:**

The correct answer is B. Amazon Inspector.

Amazon Inspector is an automated vulnerability management service that continually scans AWS workloads for software vulnerabilities and unintended network exposure. Amazon Inspector automatically discovers workloads, such as Amazon EC2 instances, containers, and Lambda functions, and scans them for software vulnerabilities and unintended network exposure<sup>12</sup>.

Amazon GuardDuty is a threat detection service that monitors your AWS accounts and workloads for malicious or unauthorized activity. Amazon GuardDuty does not scan for software vulnerabilities, but rather analyzes AWS CloudTrail, Amazon VPC Flow Logs, and DNS logs to detect threats such as compromised credentials, backdoors, or crypto mining<sup>3</sup>.

AWS Security Hub is a security and compliance service that aggregates and prioritizes security findings from multiple AWS services and partner solutions. AWS Security Hub does not scan for software vulnerabilities, but rather provides a comprehensive view of your security posture across your AWS accounts<sup>4</sup>.

AWS Shield is a managed service that protects your web applications and network resources from distributed denial-of-service (DDoS) attacks. AWS Shield does not scan for software vulnerabilities, but rather provides detection and mitigation of DDoS attacks at the network and application layers<sup>5</sup>.

References:

1: Automated Software Vulnerability Management - Amazon Inspector - AWS 3: [Amazon GuardDuty – Intelligent Threat Detection Made Easy] 2: AWS Re-Launches Amazon Inspector with New Architecture and Features - InfoQ 4: [AWS Security Hub – Unified Security and Compliance Center] 5: [AWS Shield – Managed DDoS Protection]

**NEW QUESTION 97**

- (Topic 3)

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

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

**Answer:** A

**Explanation:**

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

**NEW QUESTION 98**

- (Topic 3)

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

- A. Amazon Aurora

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

**Answer:** D

**Explanation:**

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

**NEW QUESTION 102**

- (Topic 3)

Which capabilities are in the platform perspective of the AWS Cloud Adoption Framework (AWS CAF)? (Select TWO.)

- A. Performance and capacity management
- B. Data engineering
- C. Continuous integration and continuous delivery (CI/CD)
- D. Infrastructure protection
- E. Change and release management

**Answer:** BC

**Explanation:**

These are two of the seven capabilities that are in the platform perspective of the AWS Cloud Adoption Framework (AWS CAF). The platform perspective helps you build an enterprise-grade, scalable, hybrid cloud platform, modernize existing workloads, and implement new cloud-native solutions<sup>1</sup>. The other five capabilities are:

? Platform architecture – Establish and maintain guidelines, principles, patterns, and guardrails for your cloud environment.

? Platform engineering – Build a compliant multi-account cloud environment with enhanced security features, and packaged, reusable cloud products.

? Platform operations – Manage and optimize your cloud environment with automation, monitoring, and incident response.

? Application development – Develop and deploy cloud-native applications using modern architectures and best practices.

? Application migration – Migrate your existing applications to the cloud using proven methodologies and tools.

Performance and capacity management, infrastructure protection, and change and release management are not capabilities of the platform perspective. They are part of the operations perspective, which helps you achieve operational excellence in the cloud<sup>2</sup>. The operations perspective comprises six capabilities:

? Performance and capacity management – Monitor and optimize the performance and capacity of your cloud workloads.

? Infrastructure protection – Protect your cloud infrastructure from unauthorized access, malicious attacks, and data breaches.

? Change and release management – Manage changes and releases to your cloud workloads using automation and governance.

? Configuration management – Manage the configuration of your cloud resources and applications using automation and version control.

? Incident management – Respond to incidents affecting your cloud workloads using best practices and tools.

? Service continuity management – Ensure the availability and resilience of your cloud workloads using backup, recovery, and disaster recovery strategies.

**NEW QUESTION 104**

- (Topic 3)

A company wants to integrate natural language processing (NLP) into business intelligence (BI) dashboards. The company wants to ask questions and receive answers with relevant visualizations.

Which AWS service or tool will meet these requirements?

- A. Amazon Macie
- B. Amazon Rekognition
- C. Amazon QuickSight Q
- D. Amazon Lex

**Answer:** C

**Explanation:**

Amazon QuickSight Q is a natural language query feature that lets you ask questions about your data using everyday language and get answers in seconds. You can type questions such as “What are the total sales by region?” or “How did marketing campaign A perform?” and get answers in the form of relevant visualizations, such as charts or tables. You can also use Q to drill down into details, filter data, or perform calculations. Q uses machine learning to understand your data and your intent, and provides suggestions and feedback to help you refine your questions.

**NEW QUESTION 106**

- (Topic 3)

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

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

**Answer:** A

**Explanation:**

Amazon EC2 is a web service that provides secure, resizable compute capacity in the cloud. It allows you to launch virtual servers, called instances, with different configurations of CPU, memory, storage, and networking resources. AWS Compute Optimizer analyzes the specifications and utilization metrics of your Amazon EC2 instances and generates recommendations for optimal instance types that can reduce costs and improve performance. You can view the recommendations on the AWS Compute Optimizer console or the Amazon EC2 console<sup>12</sup>.

Amazon RDS, Amazon Lightsail, and AWS Step Functions are not supported by AWS Compute Optimizer. Amazon RDS is a managed relational database service

that lets you set up, operate, and scale a relational database in the cloud. Amazon Lightsail is an easy- to-use cloud platform that offers everything you need to build an application or website, plus a cost-effective, monthly plan. AWS Step Functions lets you coordinate multiple AWS services into serverless workflows so you can build and update apps quickly<sup>3</sup> .

**NEW QUESTION 109**

- (Topic 3)

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

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

**Answer: D**

**Explanation:**

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

**NEW QUESTION 114**

- (Topic 3)

A company hosts a large amount of data in AWS. The company wants to identify if any of the data should be considered sensitive. Which AWS service will meet the requirement?

- A. Amazon Inspector
- B. Amazon Macie
- C. AWS Identity and Access Management (IAM)
- D. Amazon CloudWatch

**Answer: B**

**Explanation:**

Amazon Macie is a fully managed service that uses machine learning and pattern matching to help you detect, classify, and better protect your sensitive data stored in the AWS Cloud<sup>1</sup>. Macie can automatically discover and scan your Amazon S3 buckets for sensitive data such as personally identifiable information (PII), financial information, healthcare information, intellectual property, and credentials<sup>1</sup>. Macie also provides you with a dashboard that shows the type, location, and volume of sensitive data in your AWS environment, as well as alerts and findings on potential security issues<sup>1</sup>.

The other options are not suitable for identifying sensitive data in AWS. Amazon Inspector is a service that helps you find security vulnerabilities and deviations from best practices in your Amazon EC2 instances<sup>2</sup>. AWS Identity and Access Management (IAM) is a service that helps you manage access to your AWS resources by creating users, groups, roles, and policies<sup>3</sup>. Amazon CloudWatch is a service that helps you monitor and troubleshoot your AWS resources and applications by collecting metrics, logs, events, and alarms<sup>4</sup>. References:

? 1: What Is Amazon Macie? - Amazon Macie

? 2: What Is Amazon Inspector? - Amazon Inspector

? 3: What Is IAM? - AWS Identity and Access Management

? 4: What Is Amazon CloudWatch? - Amazon CloudWatch

**NEW QUESTION 115**

- (Topic 3)

Which task must a user perform by using the AWS account root user credentials?

- A. Make changes to AWS production resources.
- B. Change AWS Support plans.
- C. Access AWS Cost and Usage Reports.
- D. Grant auditors' access to an AWS account for a compliance audit.

**Answer: B**

**Explanation:**

Changing AWS Support plans is a task that must be performed by using the AWS account root user credentials. The root user is the email address that you used to sign up for AWS. It has complete access to all AWS services and resources in the account. You should use the root user only to perform a few account and service management tasks, such as changing AWS Support plans, closing the account, or changing the account name or email address. Making changes to AWS production resources, accessing AWS Cost and Usage Reports, and granting auditors access to an AWS account for a compliance audit are tasks that can be performed by using IAM users or roles, which are entities that you create in AWS to delegate permissions to access AWS services and resources.

**NEW QUESTION 118**

- (Topic 3)

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

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

**Answer: D**

**Explanation:**

Hardware and infrastructure is the option that AWS is responsible for under the AWS shared responsibility model. The AWS shared responsibility model describes

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

**NEW QUESTION 122**

- (Topic 3)

An auditor is preparing for an annual security audit. The auditor requests certification details for a company's AWS hosted resources across multiple Availability Zones in the us-east-1 Region.

How should the company respond to the auditor's request?

- A. Open an AWS Support ticket to request that the AWS technical account manager (TAM) respond and help the auditor.
- B. Open an AWS Support ticket to request that the auditor receive approval to conduct an onsite assessment of the AWS data centers in which the company operates.
- C. Explain to the auditor that AWS does not need to be audited because the company's application is hosted in multiple Availability Zones.
- D. Use AWS Artifact to download the applicable report for AWS security control
- E. Provide the report to the auditor.

**Answer:** D

**Explanation:**

AWS Artifact is your go-to, central resource for compliance-related information that matters to you. It provides on-demand access to AWS' security and compliance reports and select online agreements. Reports available in AWS Artifact include our Service Organization Control (SOC) reports, Payment Card Industry (PCI) reports, and certifications from accreditation bodies across geographies and compliance verticals that validate the implementation and operating effectiveness of AWS security controls. Agreements available in AWS Artifact include the Business Associate Addendum (BAA) and the Nondisclosure Agreement (NDA). You can use AWS Artifact to download the applicable report for AWS security controls and provide it to the auditor.

**NEW QUESTION 126**

- (Topic 3)

Which of the following is a benefit of using an AWS managed service?

- A. Reduced operational overhead for a company's IT staff
- B. Increased fixed costs that can be predicted by a finance team
- C. Removal of the need to have a backup strategy
- D. Removal of the need to follow compliance standards

**Answer:** A

**Explanation:**

This is a benefit of using an AWS managed service, such as Amazon S3, Amazon DynamoDB, or AWS Lambda. AWS managed services are fully managed by AWS, which means that AWS handles the provisioning, scaling, patching, backup, and recovery of the underlying infrastructure and software. This reduces the operational overhead for the company's IT staff, who can focus on their core business logic and innovation. You can learn more about the AWS managed services from this webpage or this digital course.

**NEW QUESTION 127**

- (Topic 3)

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

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

**Answer:** D

**Explanation:**

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

**NEW QUESTION 128**

- (Topic 3)

A company plans to migrate to the AWS Cloud. The company wants to use the AWS Cloud Adoption Framework (AWS CAF) to define and track business outcomes as part of its cloud transformation journey.

Which AWS CAF governance perspective capability will meet these requirements?

- A. Benefits management
- B. Risk management
- C. Application portfolio management
- D. Cloud financial management

**Answer:** A

**Explanation:**

The correct answer is A. Benefits management.

Benefits management is the AWS CAF governance perspective capability that helps you define and track business outcomes as part of your cloud transformation journey. Benefits management helps you align your cloud initiatives with your business objectives, measure the value and impact of your cloud investments, and



communicate the benefits of cloud adoption to your stakeholders<sup>12</sup>.

Risk management is the AWS CAF governance perspective capability that helps you identify and mitigate the potential risks associated with cloud adoption, such as security, compliance, legal, and operational risks<sup>12</sup>.

Application portfolio management is the AWS CAF governance perspective capability that helps you assess and optimize your existing application portfolio for cloud migration or modernization. Application portfolio management helps you categorize your applications based on their business value and technical fit, prioritize them for cloud adoption, and select the best migration or modernization strategy for each application<sup>12</sup>.

Cloud financial management is the AWS CAF governance perspective capability that helps you manage and optimize the costs and value of your cloud resources. Cloud financial management helps you plan and budget for cloud adoption, track and allocate cloud costs, implement cost optimization strategies, and report on cloud financial performance<sup>12</sup>. References:

1: AWS Cloud Adoption Framework: Governance Perspective 2: All you need to know about AWS Cloud Adoption Framework — Governance Perspective

#### NEW QUESTION 131

- (Topic 3)

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

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

**Answer: D**

#### Explanation:

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

#### NEW QUESTION 133

- (Topic 3)

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

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

**Answer: B**

#### Explanation:

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

#### NEW QUESTION 134

- (Topic 3)

A company is launching a mobile app. The company wants customers to be able to use the app without upgrading their mobile devices. Which pillar of the AWS Well-Architected Framework does this goal represent?

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

**Answer: C**

#### Explanation:

Cost optimization is one of the five pillars of the AWS Well-Architected Framework. It focuses on avoiding unnecessary costs, understanding and controlling where money is being spent, selecting the most appropriate and right number of resource types, analyzing spend over time, and scaling to meet business needs without overspending.

#### NEW QUESTION 138

- (Topic 3)

Which AWS service is always provided at no charge?

- A. Amazon S3
- B. AWS Identity and Access Management (IAM)
- C. Elastic Load Balancers
- D. AWS WAF

**Answer: B**

**Explanation:**

AWS Identity and Access Management (IAM) is a web service that helps you securely control access to AWS resources. You can use IAM to create and manage AWS users and groups, and use permissions to allow and deny their access to AWS resources. IAM is always provided at no charge<sup>12</sup>. References: 1: AWS Identity and Access Management (IAM) - Amazon Web Services (AWS), 2: Which aws service is always provided at no charge? - Brainly.in

**NEW QUESTION 139**

- (Topic 3)

An ecommerce company wants to use Amazon EC2 Auto Scaling to add and remove EC2 instances based on CPU utilization. Which AWS service or feature can initiate an Amazon EC2 Auto Scaling action to achieve this goal?

- A. Amazon Simple Queue Service (Amazon SQS)
- B. Amazon Simple Notification Service (Amazon SNS)
- C. AWS Systems Manager
- D. Amazon CloudWatch alarm

**Answer:** D

**Explanation:**

Amazon CloudWatch alarm is an AWS service or feature that can initiate an Amazon EC2 Auto Scaling action based on CPU utilization. Amazon CloudWatch is a monitoring and observability service that collects and tracks metrics, logs, events, and alarms for your AWS resources and applications. Amazon CloudWatch alarms are actions that you can configure to send notifications or automatically make changes to the resources you are monitoring based on rules that you define<sup>67</sup>.

Amazon EC2 Auto Scaling is a service that helps you maintain application availability and allows you to automatically add or remove EC2 instances according to definable conditions. You can create dynamic scaling policies that track a specific CloudWatch metric, such as CPU utilization, and define what action to take when the associated CloudWatch alarm is in ALARM. When the policy is in effect, Amazon EC2 Auto Scaling adjusts the group's desired capacity up or down when the threshold of an alarm is

breached<sup>89</sup>. References: 6: Cloud Monitoring - Amazon CloudWatch - AWS, 7: Amazon

CloudWatch Documentation, 8: Dynamic scaling for Amazon EC2 Auto Scaling, 9: Amazon EC2 Auto Scaling Documentation

**NEW QUESTION 140**

- (Topic 3)

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

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

**Answer:** AC

**Explanation:**

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

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

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

**NEW QUESTION 141**

- (Topic 3)

A company wants to migrate its on-premises relational databases to the AWS Cloud. The company wants to use infrastructure as close to its current geographical location as possible.

Which AWS service or resource should the company use to select its Amazon RDS deployment area?

- A. Amazon Connect
- B. AWS Wavelength
- C. AWS Regions
- D. AWS Direct Connect

**Answer:** C

**Explanation:**

AWS Regions are the AWS service or resource that the company should use to select its Amazon RDS deployment area. AWS Regions are separate geographic areas where AWS clusters its data centers. Each AWS Region consists of multiple, isolated, and physically separate Availability Zones within a geographic area. Each AWS Region is designed to be isolated from the other AWS Regions to achieve the highest possible fault tolerance and stability. AWS provides a more extensive global footprint than any other cloud provider, and to support its global footprint and ensure customers are served across the world, AWS opens new Regions rapidly. AWS maintains multiple geographic Regions, including Regions in North America, South America, Europe, China, Asia Pacific, South Africa, and the Middle East. Amazon RDS is available in several AWS Regions worldwide. To create or work with an Amazon RDS DB instance in a specific AWS Region, you must use the corresponding regional service endpoint. You can choose the AWS Region that meets your latency or legal requirements. You can also use multiple AWS Regions to design a disaster recovery solution or to distribute your read workload. References: Global Infrastructure Regions & AZs - aws.amazon.com, Regions, Availability Zones, and Local Zones - Amazon Relational Database Service

**NEW QUESTION 145**

- (Topic 3)

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

- A. Organizational alignment

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

**Answer:** AC

**Explanation:**

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

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

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

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

References:

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

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

**NEW QUESTION 148**

- (Topic 3)

Which task must a user perform by using the AWS account root user credentials?

- A. Make changes to AWS production resources.
- B. Change AWS Support plans.
- C. Access AWS Cost and Usage Reports.
- D. Grant auditors' access to an AWS account for a compliance audit.

**Answer:** B

**Explanation:**

The AWS account root user is the email address that you used to sign up for AWS. The root user has complete access to all AWS services and resources in the account. You should use the root user only to perform a few account and service management tasks. One of these tasks is changing AWS Support plans, which requires root user credentials. For other tasks, you should create an IAM user or role with the appropriate permissions and use that instead of the root user.

**NEW QUESTION 152**

- (Topic 3)

A company is using Amazon DynamoDB for its application database.

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

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

**Answer:** DE

**Explanation:**

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

**NEW QUESTION 153**

- (Topic 3)

Which of the following is a software development framework that a company can use to define cloud resources as code and provision the resources through AWS CloudFormation?

- A. AWS CLI
- B. AWS Developer Center
- C. AWS Cloud Development Kit (AWS CDK)
- D. AWS CodeStar

**Answer:** C

**Explanation:**

AWS Cloud Development Kit (AWS CDK) is a software development framework that allows you to define cloud resources as code using familiar programming languages, such as TypeScript, Python, Java, .NET, and Go (in Developer Preview). You can use AWS CDK to model your application resources using high-level constructs that provide sensible defaults and best practices, or use low-level constructs that provide full access to the underlying AWS CloudFormation resources. AWS CDK synthesizes your code into AWS CloudFormation templates that you can deploy using the AWS CDK CLI or the AWS Management Console. AWS CDK also integrates with other AWS services, such as AWS CodeCommit, AWS CodeBuild, AWS CodePipeline, AWS Lambda, Amazon EC2, Amazon S3, and more, to help you automate your development and deployment processes. AWS CDK is an open-source framework that you can extend and contribute to. References:

Cloud Development Framework - AWS Cloud Development Kit -

AWS, AWS Cloud Development Kit Documentation, AWS Cloud Development Kit - Wikipedia, AWS CDK Intro Workshop | AWS CDK Workshop

**NEW QUESTION 158**

- (Topic 3)

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

- A. Redundancy
- B. Operational excellence
- C. Availability
- D. Multi-Region

**Answer: B**

**Explanation:**

The AWS Well-Architected Framework helps cloud architects build secure, high-performing, resilient, and efficient infrastructure for their applications and workloads. Based on five pillars — operational excellence, security, reliability, performance efficiency, and cost optimization — the Framework provides a consistent approach for customers and partners to evaluate architectures, and implement designs that can scale over time. Operational excellence is one of the pillars of the Framework, and it focuses on running and monitoring systems to deliver business value, and continually improving processes and procedures.

**NEW QUESTION 159**

- (Topic 3)

A company wants its Amazon EC2 instances to share the same geographic area but use multiple independent underlying power sources.

Which solution achieves this goal?

- A. Use EC2 instances in a single Availability Zone.
- B. Use EC2 instances in multiple AWS Regions.
- C. Use EC2 instances in multiple Availability Zones in the same AWS Region.
- D. Use EC2 instances in the same edge location and the same AWS Region.

**Answer: C**

**Explanation:**

The solution that achieves the goal of having Amazon EC2 instances share the same geographic area but use multiple independent underlying power sources is to use EC2 instances in multiple Availability Zones in the same AWS Region. An Availability Zone is a physically isolated location within an AWS Region that has its own power, cooling, and network connectivity. An AWS Region is a geographical area that consists of two or more Availability Zones. By using multiple Availability Zones, users can increase the fault tolerance and resilience of their applications, as well as reduce latency for end users<sup>3</sup>.

Using EC2 instances in a single Availability Zone, multiple AWS Regions, or the same edge location and the same AWS Region would not meet the requirement of having multiple independent power sources.

**NEW QUESTION 163**

- (Topic 3)

Which AWS services are connectivity services for a VPC? (Select TWO.)

- A. AWS Site-to-Site VPN
- B. AWS Direct Connect
- C. Amazon Connect
- D. AWS Key Management Service (AWS KMS)
- E. AWS Identity and Access Management (IAM)

**Answer: A**

**Explanation:**

AWS Site-to-Site VPN and AWS Direct Connect are AWS services that are connectivity services for a VPC. AWS Site-to-Site VPN is a service that enables you to securely connect your on-premises network or branch office site to your Amazon Virtual Private Cloud (Amazon VPC). You can establish VPN connections over the internet or over AWS Direct Connect<sup>1</sup>. AWS Direct Connect is a service that lets you establish a dedicated network connection between your network and one of the AWS Direct Connect locations. Using AWS Direct Connect, you can create a private connection between AWS and your datacenter, office, or colocation environment, which can reduce your network costs, increase bandwidth throughput, and provide a more consistent network experience than internet-based connections<sup>2</sup>. Amazon Connect is a service that lets you set up and manage a contact center in the cloud, but it does not provide network connectivity between the VPC and your on-premises network. AWS Key Management Service (AWS KMS) is a service that makes it easy for you to create and manage cryptographic keys and control their use across a wide range of AWS services and in your applications, but it does not provide network connectivity between the VPC and your on-premises network. AWS Identity and Access Management (IAM) is a service that enables you to manage access to AWS services and resources securely, but it does not provide network connectivity between the VPC and your on-premises network.

**NEW QUESTION 168**

- (Topic 3)

Which benefit does AWS offer exclusively to users who have an AWS Enterprise Support plan?

- A. Access to a technical project manager
  - B. Access to a technical account manager (TAM)
  - C. Access to a cloud support engineer
  - D. Access to a solutions architect
- A company wants to automatically set up and govern a multi-account AWS environment.

**Answer: B**

**Explanation:**

AWS Enterprise Support plan is the highest level of support that AWS offers to its customers. One of the exclusive benefits of this plan is the access to a technical account manager (TAM), who is a dedicated point of contact for guidance, advocacy, and support<sup>2</sup>. A technical project manager, a cloud support engineer, and a solutions architect are not exclusive benefits of the AWS Enterprise Support plan, as they are also available to customers with lower-tier support plans or through other AWS services or programs<sup>3,4,5</sup>.



**NEW QUESTION 171**

- (Topic 3)

A company needs to search for text in documents that are stored in Amazon S3. Which AWS service will meet these requirements?

- A. Amazon Kendra
- B. Amazon Rekognition
- C. Amazon Polly
- D. Amazon Lex

**Answer:** A

**Explanation:**

Amazon Kendra is a highly accurate and easy to use intelligent search service powered by machine learning. It enables users to easily find the content they are looking for, even when it is scattered across multiple locations and content repositories within their organization. Amazon Kendra supports natural language queries, and can search for text in documents stored in Amazon S3, as well as other sources such as SharePoint, OneDrive, Salesforce, ServiceNow, and more<sup>1</sup>. Amazon Rekognition is a computer vision service that makes it easy to add image and video analysis to applications. It can detect objects, faces, text, scenes, activities, and emotions in images and videos. However, it is not designed for searching for text in documents stored in Amazon S3<sup>2</sup>.

Amazon Polly is a text-to-speech service that turns text into lifelike speech. It can create audio versions of books, articles, podcasts, and more. However, it is not designed for searching for text in documents stored in Amazon S3<sup>3</sup>.

Amazon Lex is a service for building conversational interfaces using voice and text. It can create chatbots that can interact with users using natural language. However, it is not designed for searching for text in documents stored in Amazon S3<sup>4</sup>.

References:

? Amazon Kendra – Intelligent Search Service Powered by Machine Learning

? Amazon Rekognition – Video and Image - AWS

? Amazon Polly – Text-to-Speech Service - AWS

? Amazon Lex – Build Conversation Bots - AWS

**NEW QUESTION 176**

- (Topic 3)

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

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

**Answer:** A

**Explanation:**

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

**NEW QUESTION 178**

- (Topic 3)

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

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

**Answer:** B

**Explanation:**

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

**NEW QUESTION 180**

- (Topic 3)

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

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

**Answer:** C

**Explanation:**

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

their access policies, and monitor any changes to the resource sharing status. References: AWS IAM Access Analyzer, Identify and review resources shared with external entities, How AWS IAM Access Analyzer works

**NEW QUESTION 182**

- (Topic 3)

AWS has the ability to achieve lower pay-as-you-go pricing by aggregating usage across hundreds of thousands of users. This describes which advantage of the AWS Cloud?

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

**Answer:** C

**Explanation:**

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

**NEW QUESTION 183**

- (Topic 3)

Which mechanism allows developers to access AWS services from application code?

- A. AWS Software Development Kit
- B. AWS Management Console
- C. AWS CodePipeline
- D. AWS Config

**Answer:** A

**Explanation:**

AWS Software Development Kit (SDK) is a set of platform-specific building tools for developers. It allows developers to access AWS services from application code using familiar programming languages. It provides pre-built components and libraries that can be incorporated into applications, as well as tools to debug, monitor, and optimize performance<sup>2</sup>. References: What is SDK? - SDK Explained - AWS

**NEW QUESTION 187**

- (Topic 3)

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

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

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

**Answer:** A

**Explanation:**

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

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

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

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

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

References:

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

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

**NEW QUESTION 188**

- (Topic 3)

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

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

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

**Answer:** C

**Explanation:**

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

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

**NEW QUESTION 191**

- (Topic 3)

Which AWS service or feature can the company use to limit the access to AWS services for member accounts?

- A. AWS Identity and Access Management (IAM)
- B. Service control policies (SCPs)
- C. Organizational units (OUs)
- D. Access control lists (ACLs)

**Answer:** B

**Explanation:**

Service control policies (SCPs) are a type of organization policy that you can use to manage permissions in your organization. SCPs offer central control over the maximum available permissions for all accounts in your organization, allowing you to ensure your accounts stay within your organization's access control guidelines<sup>2</sup>. SCPs are available only in an organization that has all features enabled<sup>2</sup>.

**NEW QUESTION 194**

- (Topic 3)

A company wants to migrate its PostgreSQL database to AWS. The company does not use the database frequently.

Which AWS service or resource will meet these requirements with the LEAST management overhead?

- A. PostgreSQL on Amazon EC2
- B. Amazon RDS for PostgreSQL
- C. Amazon Aurora PostgreSQL-Compatible Edition
- D. Amazon Aurora Serverless

**Answer:** D

**Explanation:**

Amazon Aurora Serverless is an on-demand, auto-scaling configuration for Amazon Aurora PostgreSQL-Compatible Edition. It is a fully managed service that automatically scales up and down based on the application's actual needs. Amazon Aurora Serverless is suitable for applications that have infrequent, intermittent, or unpredictable database workloads, and that do not require the full power and range of options provided by provisioned Aurora clusters. Amazon Aurora Serverless eliminates the need to provision and manage database instances, and reduces the management overhead associated with database administration tasks such as scaling, patching, backup, and recovery. References: Amazon Aurora Serverless, Choosing between Aurora Serverless and provisioned Aurora DB clusters, [AWS Cloud Practitioner Essentials: Module 4 - Databases in the Cloud]

**NEW QUESTION 198**

- (Topic 3)

Which capabilities are in the platform perspective of the AWS Cloud Adoption Framework (AWS CAF)? (Select TWO.)

- A. Performance and capacity management
- B. Data engineering
- C. Continuous integration and continuous delivery (CI/CD)
- D. Infrastructure protection
- E. Change and release management

**Answer:** BC

**Explanation:**

The platform perspective of the AWS Cloud Adoption Framework (AWS CAF) helps you build an enterprise-grade, scalable, hybrid cloud platform, modernize existing workloads, and implement new cloud-native solutions<sup>1</sup>. It comprises seven capabilities, two of which are data engineering and CI/CD<sup>1</sup>.

? Data engineering: This capability helps you design and evolve a fit-for-purpose data and analytics architecture that can reduce complexity, cost, and technical debt while enabling you to gain actionable insights from exponentially growing data volumes<sup>1</sup>. It involves selecting key technologies for each of your architectural layers, such as ingestion, storage, catalog, processing, and consumption. It also involves supporting real-time data processing and adopting a Lake House architecture to facilitate data movements between data lakes and purpose-built data stores<sup>1</sup>.

? CI/CD: This capability helps you automate the delivery of your cloud solutions using a set of practices and tools that enable faster and more reliable deployments<sup>1</sup>. It involves establishing a pipeline that can build, test, and deploy your code across multiple environments. It also involves adopting a DevOps culture that fosters collaboration, feedback, and continuous improvement among your development and operations teams<sup>1</sup>.

References:

? 1: Platform perspective: infrastructure and applications - An Overview of the AWS Cloud Adoption Framework

**NEW QUESTION 202**

- (Topic 3)

Which actions are best practices for an AWS account root user? (Select TWO.)

- A. Share root user credentials with team members.
- B. Create multiple root users for the account, separated by environment.
- C. Enable multi-factor authentication (MFA) on the root user.
- D. Create an IAM user with administrator privileges for daily administrative tasks, instead of using the root user.
- E. Use programmatic access instead of the root user and password.

**Answer:** CD

**Explanation:**

The AWS account root user is the identity that has complete access to all AWS services and resources in the account. It is accessed by signing in with the email address and password that were used to create the account<sup>1</sup>. The root user should be protected and used only for a few account and service management tasks that require it<sup>1</sup>. Therefore, the following actions are best practices for an AWS account root user:

? Enable multi-factor authentication (MFA) on the root user. MFA is a security feature that requires users to provide two or more pieces of information to authenticate themselves, such as a password and a code from a device. MFA adds an extra layer of protection for the root user credentials, which can access sensitive information and perform critical operations in the account<sup>2</sup>.

? Create an IAM user with administrator privileges for daily administrative tasks, instead of using the root user. IAM is a service that helps customers manage access to AWS resources for users and groups. Customers can create IAM users and assign them permissions to perform specific tasks on specific resources. Customers can also create IAM roles and policies to delegate access to other AWS services or external entities<sup>3</sup>. By creating an IAM user with administrator privileges, customers can avoid using the root user for everyday tasks and reduce the risk of accidental or malicious changes to the account<sup>1</sup>.

**NEW QUESTION 204**

- (Topic 3)

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

- A. High availability
- B. Performance efficiency
- C. Cost optimization
- D. Going global in minutes
- E. Continuous development

**Answer:** BC

**Explanation:**

The AWS Well-Architected Framework is a set of six pillars and lenses that help cloud architects design and run workloads in the cloud. The six pillars are: operational excellence, security, reliability, performance efficiency, cost optimization, and sustainability. Each pillar has a set of design principles and best practices that guide the architectural decisions. High availability is not a separate pillar, but a quality that can be achieved by applying the principles of the reliability pillar. Going global in minutes and continuous development are not pillars of the framework, but possible benefits of using AWS services and following the framework's recommendations. References: AWS Well-Architected - Build secure, efficient cloud applications, AWS Well-Architected Framework, The 6 Pillars of the AWS Well-Architected Framework

**NEW QUESTION 205**

- (Topic 3)

Which option is an AWS Cloud Adoption Framework (AWS CAF) foundational capability for the operations perspective?

- A. Performance and capacity management
- B. Application portfolio management
- C. Identity and access management
- D. Product management

**Answer:** C

**Explanation:**

Identity and access management is one of the foundational capabilities for the operations perspective of the AWS Cloud Adoption Framework (AWS CAF). It involves managing the identities, roles, permissions, and credentials of users and systems that interact with AWS resources. Performance and capacity management is a capability for the platform perspective. Application portfolio management is a capability for the business perspective. Product management is a capability for the governance perspective.

**NEW QUESTION 210**

- (Topic 3)

Which of the following is a fully managed graph database service on AWS?

- A. Amazon Aurora
- B. Amazon FSx
- C. Amazon DynamoDB
- D. Amazon Neptune

**Answer:** D

**Explanation:**

Amazon Neptune is a fully managed graph database service on AWS. A graph database is a type of database that stores and queries data as a network of nodes and edges, representing entities and relationships. Graph databases are useful for applications that deal with highly connected data, such as social networks, recommendation engines, fraud detection, and knowledge graphs<sup>45</sup>. Amazon Neptune is a fast, reliable, and scalable graph database service that supports two popular graph models: property graphs and RDF. Amazon Neptune also supports two open standards for querying graphs: Apache TinkerPop Gremlin and SPARQL. Amazon Neptune handles the heavy lifting of managing the database, such as provisioning, patching, backup, recovery, encryption, and replication<sup>456</sup>. References: 4: Managed Graph Database - Amazon Neptune - AWS, 5: Amazon Neptune – A Fully Managed Graph Database Service, 6: Working with AWS Neptune. Neptune is a fully-managed graph ... - Medium

**NEW QUESTION 213**

- (Topic 3)

Which AWS service offers object storage?

- A. Amazon RDS
- B. Amazon Elastic File System (Amazon EFS)
- C. Amazon S3
- D. Amazon DynamoDB



**Answer:** C

**Explanation:**

Amazon S3 is the AWS service that offers object storage. Object storage is a technology that stores and manages data in an unstructured format called objects. Each object consists of the data, metadata, and a unique identifier. Object storage is ideal for storing large amounts of unstructured data, such as photos, videos, email, web pages, sensor data, and audio files<sup>1</sup>. Amazon S3 provides industry-leading scalability, data availability, security, and performance for object storage<sup>2</sup>. Amazon RDS is the AWS service that offers relational database storage. Relational database storage is a technology that stores and manages data in a structured format called tables. Each table consists of rows and columns that define the attributes and values of the data. Relational database storage is ideal for storing structured or semi-structured data, such as customer records, inventory, transactions, and analytics<sup>3</sup>.

Amazon Elastic File System (Amazon EFS) is the AWS service that offers file storage. File storage is a technology that stores and manages data in a hierarchical format called files and folders. Each file consists of the data and metadata, and each folder consists of files or subfolders. File storage is ideal for storing shared data that can be accessed by multiple users or applications, such as home directories, content repositories, media libraries, and configuration files<sup>4</sup>.

Amazon DynamoDB is the AWS service that offers NoSQL database storage. NoSQL database storage is a technology that stores and manages data in a flexible format called documents or key-value pairs. Each document or key-value pair consists of the data and metadata, and can have different attributes and values depending on the schema. NoSQL database storage is ideal for storing dynamic or unstructured data that requires high performance, scalability, and availability, such as web applications, social media, gaming, and IoT.

**NEW QUESTION 217**

- (Topic 3)

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

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

**Answer:** AB

**Explanation:**

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

**NEW QUESTION 218**

- (Topic 3)

What can a cloud practitioner use to retrieve AWS security and compliance documents and submit them as evidence to an auditor or regulator?

- A. AWS Certificate Manager
- B. AWS Systems Manager
- C. AWS Artifact
- D. Amazon Inspector

**Answer:** C

**Explanation:**

AWS Artifact is a service that provides on-demand access to AWS security and compliance documents, such as AWS ISO certifications, Payment Card Industry (PCI) reports, and Service Organization Control (SOC) reports. You can download these documents and submit them as evidence to your auditors or regulators to demonstrate the security and compliance of the AWS infrastructure and services that you use. AWS Artifact also allows you to review, accept, and manage AWS agreements, such as the Business Associate Addendum (BAA) for customers who are subject to the Health Insurance Portability and Accountability Act (HIPAA). References: AWS Artifact, What is AWS Artifact?

**NEW QUESTION 219**

- (Topic 3)

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

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

**Answer:** A

**Explanation:**

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

**NEW QUESTION 223**

- (Topic 3)

A company needs a bridge between technology and business to help evolve to a culture of continuous growth and learning. Which perspective in the AWS Cloud Adoption Framework (AWS CAF) serves as this bridge?

- A. People
- B. Governance
- C. Operations
- D. Security

**Answer:** A

**Explanation:**

The People perspective in the AWS Cloud Adoption Framework (AWS CAF) serves as a bridge between technology and business, accelerating the cloud journey to help organizations more rapidly evolve to a culture of continuous growth, learning, and where change becomes business-as-normal, with focus on culture, organizational structure, leadership, and workforce<sup>1</sup>. References: People Perspective - AWS Cloud Adoption Framework

**NEW QUESTION 224**

- (Topic 3)

A company's headquarters is located on a different continent from where the majority of the company's customers live. The company wants an AWS Cloud environment setup that will provide the lowest latency to the customers.

A company wants to automate the creation of new AWS accounts and automatically prevent all users from creating Amazon EC2 instances.

Which AWS service provides this functionality?

- A. AWS Service Catalog
- B. AWS Organizations
- C. EC2 Image Builder
- D. AWS Systems Manager

**Answer:** B

**Explanation:**

AWS Organizations is a service that enables you to create and manage multiple AWS accounts centrally. You can use AWS Organizations to automate account creation, apply policies to control access and permissions, and consolidate billing across your accounts. You can also use AWS Organizations to prevent users from creating Amazon EC2 instances in certain regions or with certain configurations<sup>2</sup>

**NEW QUESTION 226**

- (Topic 3)

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

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

**Answer:** C

**Explanation:**

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

**NEW QUESTION 229**

- (Topic 3)

A company wants a key-value NoSQL database that is fully managed and serverless. Which AWS service will meet these requirements?

- A. Amazon DynamoDB
- B. Amazon RDS
- C. Amazon Aurora
- D. Amazon Memory DB for Redis

**Answer:** A

**Explanation:**

Amazon DynamoDB is a key-value and document database that delivers single-digit millisecond performance at any scale. It is a fully managed, serverless database that does not require provisioning, patching, or backup. It offers built-in security, backup and restore, and in-memory caching<sup>3</sup>. Amazon RDS is a relational database service that makes it easy to set up, operate, and scale a relational database in the cloud. It provides cost-efficient and resizable capacity while automating time-consuming administration tasks such as hardware provisioning, database setup, patching, and backups. However, it is not a key-value NoSQL database, and it is not serverless, as it requires you to choose an instance type and size<sup>4</sup>. Amazon Aurora is a MySQL and PostgreSQL-compatible relational database built for the cloud, that combines the performance and availability of traditional enterprise databases with the simplicity and cost-effectiveness of open source databases. However, it is also not a key-value NoSQL database, and it is not serverless, as it requires you to choose an instance type and size. Amazon MemoryDB for Redis is a Redis-compatible, durable, in-memory database service that delivers ultra-fast performance and multi-AZ reliability for the most demanding applications. However, it is also not a key-value NoSQL database, and it is not serverless, as it requires you to choose a node type and size.

**NEW QUESTION 230**

- (Topic 3)

A developer who has no AWS Cloud experience wants to use AWS technology to build a web application.

Which AWS service should the developer use to start building the application?

- A. Amazon SageMaker
- B. AWS Lambda
- C. Amazon Lightsail
- D. Amazon Elastic Container Service (Amazon ECS)

**Answer:** C

**Explanation:**

Amazon Lightsail is an easy-to-use cloud platform that offers everything you need to build an application or website, plus a cost-effective, monthly plan<sup>1</sup>. It is

designed for developers who have little or no prior cloud experience and want to launch and manage applications on AWS with minimal complexity<sup>2</sup>. Amazon SageMaker is a service for building, training, and deploying machine learning models<sup>3</sup>. AWS Lambda is a service that lets you run code without provisioning or managing servers<sup>4</sup>. Amazon Elastic Container Service (Amazon ECS) is a fully managed container orchestration service.

#### NEW QUESTION 233

- (Topic 3)

A company needs to set a maximum spending limit on AWS services each month. The company also needs to set up alerts for when the company reaches its spending limit.

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

- A. Cost Explorer
- B. AWS Trusted Advisor
- C. Service Quotas
- D. AWS Budgets

**Answer:** D

#### **Explanation:**

AWS Budgets is a service that helps you plan your service usage, service costs, and instance reservations, and track how close your plan is to your budgeted amount. You can set custom budgets that alert you when you exceed (or are forecasted to exceed) your budgeted thresholds. You can also use AWS Budgets to set a maximum spending limit on AWS services each month and set up alerts for when you reach your spending limit. Cost Explorer is a service that enables you to visualize, understand, and manage your AWS costs and usage over time. You can use Cost Explorer to view charts and graphs that show how your costs are trending, identify areas that need further inquiry, and see the impact of your cost management actions. However, Cost Explorer does not allow you to set a maximum spending limit or alerts for your AWS services. AWS Trusted Advisor is a service that provides you real time guidance to help you provision your resources following AWS best practices, including security and performance. It can help you monitor for cost optimization opportunities, such as unused or underutilized resources, but it does not allow you to set a maximum spending limit or alerts for your AWS services. Service Quotas is a service that enables you to view and manage your quotas, also referred to as limits, from a central location. Quotas, also referred to as limits, are the maximum number of resources that you can create in your AWS account. However, Service Quotas does not allow you to set a maximum spending limit or alerts for your AWS services.

#### NEW QUESTION 235

.....

## Thank You for Trying Our Product

### We offer two products:

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

2nd - Questions and Answers in PDF Format

### CLF-C02 Practice Exam Features:

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

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