

# Exam Questions CLF-C01

AWS Certified Cloud Practitioner

<https://www.2passeasy.com/dumps/CLF-C01/>



### NEW QUESTION 1

- (Topic 3)

A company runs business applications in an on-premises data center and in the AWS Cloud. The company needs a shared file system that can be available to both environments.

Which AWS service meets these requirements?

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

**Answer:** D

#### Explanation:

Amazon Elastic File System (Amazon EFS) is a service that provides a simple, scalable, fully managed elastic NFS file system for use with AWS Cloud services and on-premises resources. 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. You can use Amazon EFS to create a shared file system that can be available to both your on-premises data center and your AWS Cloud environment. Amazon Elastic Block Store (Amazon EBS) is a service that provides persistent block storage volumes for use with Amazon EC2 instances in the AWS Cloud. Each Amazon EBS volume is automatically replicated within its Availability Zone to protect you from component failure, offering high availability and durability. However, Amazon EBS volumes are not shared file systems, and they cannot be available to both your on-premises data center and your AWS Cloud environment. Amazon S3 is a service that provides object storage through a web services interface. You can use Amazon S3 to store and protect any amount of data for a range of use cases, such as data lakes, websites, mobile applications, backup and restore, archive, enterprise applications, IoT devices, and big data analytics. However, Amazon S3 is not a shared file system, and it cannot be available to both your on-premises data center and your AWS Cloud environment without additional configuration. Amazon ElastiCache is a service that enables you to seamlessly set up, run, and scale popular open-source compatible in-memory data stores in the cloud. You can use Amazon ElastiCache to improve the performance of your applications by allowing you to retrieve information from fast, managed, in-memory data stores, instead of relying entirely on slower disk-based databases. However, Amazon ElastiCache is not a shared file system, and it cannot be available to both your on-premises data center and your AWS Cloud environment.

### NEW QUESTION 2

- (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 3

- (Topic 3)

A company wants to monitor for misconfigured security groups that are allowing unrestricted access to specific ports.

Which AWS service will meet this requirement?

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

**Answer:** A

#### Explanation:

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

#### NEW QUESTION 4

- (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 5

- (Topic 3)

Which AWS service provides this functionality?

- A. AWS IAM Identity Center (AWS Single Sign-On)
- B. AWS Systems Manager
- C. AWS Config
- D. AWS Control Tower

**Answer:** D

#### Explanation:

AWS Control Tower is a service that provides an easy way to set up and govern a secure, multi-account AWS environment. It automates the creation of accounts, organizational units, policies, and best practices based on the AWS Well-Architected Framework. AWS IAM Identity Center (AWS Single Sign-On) is a service that enables users to centrally manage access to multiple AWS accounts and business applications using a single sign-on experience. AWS Systems Manager is a service that provides operational management for AWS resources and applications. AWS Config is a service that enables users to assess, audit, and evaluate the configurations of AWS resources.

#### NEW QUESTION 6

- (Topic 3)

A company is migrating to the AWS Cloud to meet storage needs. The company wants to optimize costs based on the amount of storage that the company uses. Which AWS offering or benefit will meet these requirements MOST cost-effectively?

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

**Answer:** D

#### Explanation:

Volume-based discounts are an AWS offering or benefit that can help the company optimize costs based on the amount of storage that the company uses.

Volume-based discounts are discounts that AWS provides for some storage services, such as Amazon S3 and Amazon EBS, when the company stores a large amount of data. The more data the company stores, the lower the price per GB. For example, Amazon S3 offers six storage classes, each with a different price per GB. The price per GB decreases as the amount of data stored in each storage class increases

#### NEW QUESTION 7

- (Topic 3)

A company wants to make an upfront commitment for continued use of its production Amazon EC2 instances in exchange for a reduced overall cost. Which pricing options meet these requirements with the LOWEST cost? (Select TWO.)

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

**Answer:** CD

#### Explanation:

Reserved Instances (RIs) are a pricing model that allows you to reserve EC2 instances for a specified period of time (one or three years) and receive a significant discount compared to On-Demand pricing. RIs are suitable for workloads that have predictable usage patterns and require a long-term commitment. You can choose between three payment options: All Upfront, Partial Upfront, or No Upfront. The more you pay upfront, the greater the discount<sup>1</sup>.

Savings Plans are a flexible pricing model that can help you reduce your EC2 costs by up to 72% compared to On-Demand pricing, in exchange for a commitment to a consistent amount of usage (measured in \$/hour) for a one or three year term. Savings Plans apply to usage across EC2, AWS Lambda, and AWS Fargate. You can choose between two types of Savings Plans: Compute Savings Plans and EC2 Instance Savings Plans. Compute Savings Plans offer the most flexibility and apply to any instance family, size, OS, tenancy, or region. EC2 Instance Savings Plans offer the highest discount and apply to a specific instance family within a region<sup>2</sup>.

Spot Instances are a pricing model that allows you to bid for unused EC2 capacity in the AWS cloud and are available at a discount of up to 90% compared to On-Demand pricing. Spot Instances are suitable for fault-tolerant or stateless workloads that can run on heterogeneous hardware and have flexible start and end times. However, Spot Instances are not guaranteed and can be interrupted by AWS at any time if the demand for capacity increases or your bid price is lower than

the current Spot price<sup>3</sup>.

On-Demand Instances are a pricing model that allows you to pay for compute capacity by the hour or second with no long-term commitments. On-Demand Instances are suitable for short-term, spiky, or unpredictable workloads that cannot be interrupted, or for applications that are being developed or tested on EC2 for the first time. However, On-Demand Instances are the most expensive option among the four pricing models<sup>4</sup>.

Dedicated Hosts are physical EC2 servers fully dedicated for your use. Dedicated Hosts can help you reduce costs by allowing you to use your existing server-bound software licenses, such as Windows Server, SQL Server, and SUSE Linux Enterprise Server. Dedicated Hosts can be purchased On-Demand or as part of Savings Plans. Dedicated Hosts are suitable for workloads that need to run on dedicated physical servers or have strict licensing requirements. However, Dedicated Hosts are not the lowest cost option among the four pricing models.

#### NEW QUESTION 8

- (Topic 3)

Which option is a customer responsibility under the AWS shared responsibility model?

- A. Maintenance of underlying hardware of Amazon EC2 instances
- B. Application data security
- C. Physical security of data centers
- D. Maintenance of VPC components

**Answer:** B

#### Explanation:

The option that is a customer responsibility under the AWS shared responsibility model is B. Application data security.

According to the AWS shared responsibility model, AWS is responsible for the security of the cloud, while the customer is responsible for the security in the cloud. This means that AWS manages the security of the underlying infrastructure, such as the hardware, software, networking, and facilities that run the AWS services, while the customer manages the security of their applications, data, and resources that they use on top of AWS<sup>12</sup>. Application data security is one of the customer responsibilities under the AWS shared responsibility model. This means that the customer is responsible for protecting their application data from unauthorized access, modification, deletion, or leakage. The customer can use various AWS services and features to help with application data security, such as encryption, key management, access control, logging, and auditing<sup>12</sup>. Maintenance of underlying hardware of Amazon EC2 instances is not a customer responsibility under the AWS shared responsibility model. This is part of the AWS responsibility to secure the cloud. AWS manages the physical servers that host the Amazon EC2 instances and ensures that they are updated, patched, and replaced as needed<sup>13</sup>.

Physical security of data centers is not a customer responsibility under the AWS shared responsibility model. This is also part of the AWS responsibility to secure the cloud. AWS operates and controls the facilities where the AWS services are hosted and ensures that they are protected from unauthorized access, environmental hazards, fire, and theft<sup>14</sup>. Maintenance of VPC components is not a customer responsibility under the AWS shared responsibility model. This is a shared responsibility between AWS and the customer. AWS provides the VPC service and ensures that it is secure and reliable, while the customer configures and manages their own VPCs and related components, such as subnets, route tables, security groups, network ACLs, gateways, and endpoints<sup>15</sup>.

References:

1: Shared Responsibility Model - Amazon Web Services (AWS) 2: AWS Cloud Computing - W3Schools 3: [Amazon EC2 FAQs - Amazon Web Services] 4: [AWS Security - Amazon Web Services] 5: [Amazon Virtual Private Cloud (VPC) - Amazon Web Services]

#### NEW QUESTION 9

- (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 10

- (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 10

- (Topic 3)

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

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

**Answer:** C

**Explanation:**

AWS Shield Advanced is a service that provides high levels of detection and near-real-time (NRT) mitigation against large and sophisticated distributed denial of service (DDoS) attacks on applications running on AWS. AWS Shield Advanced also provides you with 24x7 access to the AWS DDoS Response Team (DRT) and protection against DDoS attacks of any size or duration<sup>1</sup>. Amazon GuardDuty is a service that provides threat detection for your AWS accounts and workloads, but it does not offer DDoS protection<sup>3</sup>. Amazon Inspector is a service that helps you improve the security and compliance of your applications deployed on AWS by automatically assessing them for vulnerabilities and deviations from best practices. Amazon Macie is a service that uses machine learning and pattern matching to discover and protect your sensitive data in AWS.

**NEW QUESTION 11**

- (Topic 3)

A company encourages its teams to test failure scenarios regularly and to validate their understanding of the impact of potential failures.  
Which pillar of the AWS Well-Architected Framework does this philosophy represent?

- A. Operational excellence
- B. Cost optimization
- C. Performance efficiency
- D. Security

**Answer:** A

**Explanation:**

This is the pillar of the AWS Well-Architected Framework that represents the philosophy of testing failure scenarios regularly and validating the understanding of the impact of potential failures. The operational excellence pillar covers the best practices for designing, running, monitoring, and improving systems in the AWS Cloud. Testing failure scenarios is one of the ways to improve the system's resilience, reliability, and recovery. You can learn more about the operational excellence pillar from this whitepaper or this digital course.

**NEW QUESTION 12**

- (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 14**

- (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 15**

- (Topic 3)

What is a benefit of using AWS serverless computing?

- A. Application deployment and management are not required

- B. Application security will be fully managed by AWS
- C. Monitoring and logging are not needed
- D. Management of infrastructure is offloaded to AWS

**Answer:** D

**Explanation:**

AWS serverless computing is a way of building and running applications without thinking about servers. AWS manages the infrastructure for you, so you don't have to provision, scale, patch, or monitor servers. You only pay for the compute time you consume, and you can focus on your application logic instead of managing servers<sup>12</sup>. References: Serverless Computing – Amazon Web Services, AWS Serverless Computing, Benefits, Architecture and Use-cases - XenonStack

**NEW QUESTION 18**

- (Topic 3)

A company is running its application in the AWS Cloud and wants to protect against a DDoS attack. The company's security team wants near real-time visibility into DDoS attacks.

Which AWS service or traffic filter will meet these requirements with the MOST features for DDoS protection?

- A. AWS Shield Advanced
- B. AWS Shield
- C. Amazon GuardDuty
- D. Network ACLs

**Answer:** A

**Explanation:**

AWS Shield Advanced is a managed Distributed Denial of Service (DDoS) protection service that safeguards applications running on AWS. AWS Shield Advanced provides you with 24x7 access to the AWS DDoS Response Team (DRT) and protection against DDoS attacks of any size or duration. AWS Shield Advanced also provides near real-time visibility into attacks, advanced attack mitigation capabilities, and integration with AWS WAF and AWS Firewall Manager<sup>1</sup>. AWS Shield is a standard service that provides always-on detection and automatic inline mitigations to minimize application downtime and latency, but it does not offer the same level of features and support as AWS Shield Advanced<sup>2</sup>. Amazon GuardDuty is a threat detection service that continuously monitors for malicious activity and unauthorized behavior, but it does not provide DDoS protection<sup>3</sup>. Network ACLs are stateless filters that can be associated with a subnet to control the traffic to and from the subnet, but they are not designed to protect against DDoS attacks

**NEW QUESTION 23**

- (Topic 3)

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

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

Which S3 feature will meet these requirements?

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

**Answer:** D

**Explanation:**

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

**NEW QUESTION 27**

- (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 28**

- (Topic 3)

Which AWS service or feature allows a user to establish a dedicated network connection between a company's on-premises data center and the AWS Cloud?

- A. AWS Direct Connect

- B. VPC peering
- C. AWS VPN
- D. Amazon Route 53

**Answer:** A

**Explanation:**

AWS Direct Connect is an AWS service that allows users to establish a dedicated network connection between their on-premises data center and the AWS Cloud. This connection bypasses the public internet and provides more predictable network performance, reduced bandwidth costs, and increased security. Users can choose from different port speeds and connection types, and use AWS Direct Connect to access AWS services in any AWS Region globally. Users can also use AWS Direct Connect in conjunction with AWS VPN to create a hybrid network architecture that combines the benefits of both private and public connectivity. References: AWS Direct Connect, [AWS Cloud Practitioner Essentials: Module 3 - Compute in the Cloud]

**NEW QUESTION 32**

- (Topic 3)

Which benefits can customers gain by using AWS Marketplace? (Select TWO.)

- A. Speed of business
- B. Fewer legal objections
- C. Ability to pay with credit cards
- D. No requirement for product licenses for any products
- E. Free use of all services for the first hour

**Answer:** AB

**Explanation:**

AWS Marketplace is a digital catalog that offers thousands of software products and solutions from independent software vendors (ISVs) and AWS partners. Customers can use AWS Marketplace to find, buy, and deploy software on AWS. Some of the benefits of using AWS Marketplace are:  
? Speed of business: You can quickly and easily discover and deploy software that meets your business needs, without having to go through lengthy procurement processes. You can also use AWS Marketplace to test and compare different solutions before making a purchase decision.  
? Fewer legal objections: You can benefit from standardized contract terms and conditions that are pre-negotiated between AWS and the ISVs. This reduces the time and effort required to review and approve legal agreements.

**NEW QUESTION 33**

- (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 36**

- (Topic 3)

In the AWS shared responsibility model, which tasks are the responsibility of AWS? (Select TWO.)

- A. Patch an Amazon EC2 instance operating system.
- B. Configure a security group.
- C. Monitor the health of an Availability Zone.
- D. Protect the infrastructure that runs Amazon EC2 instances.
- E. Manage access to the data in an Amazon S3 bucket

**Answer:** CD

**Explanation:**

According to the AWS shared responsibility model, AWS is responsible for the security of the cloud, which includes the tasks of monitoring the health of an Availability Zone and protecting the infrastructure that runs Amazon EC2 instances. An Availability Zone is a physically isolated location within an AWS Region that has its own power, cooling, and network connectivity. AWS monitors the health and performance of each Availability Zone and notifies customers of any issues or disruptions. AWS also protects the infrastructure that runs AWS services, such as Amazon EC2, by implementing physical, environmental, and operational security measures. AWS is not responsible for patching an Amazon EC2 instance operating system, configuring a security group, or managing access to the data in an Amazon S3 bucket. These are the customer's responsibilities for security in the cloud. The customer must ensure that the operating system and applications on their EC2 instances are up to date and secure. The customer must also configure the security group rules that control the inbound and outbound traffic for their EC2 instances. The customer must also manage the access permissions and encryption settings for their S3 buckets and objects<sup>2</sup>

**NEW QUESTION 41**

- (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 43**

- (Topic 3)

A company needs to deploy applications in the AWS Cloud as quickly as possible. The company also needs to minimize the complexity that is related to the management of AWS resources.

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

- A. AWS config
- B. AWS Elastic Beanstalk
- C. Amazon EC2
- D. Amazon Personalize

**Answer:** B

**Explanation:**

AWS Elastic Beanstalk is the AWS service that allows customers to deploy applications in the AWS Cloud as quickly as possible. AWS Elastic Beanstalk automatically handles the deployment, from capacity provisioning, load balancing, and auto-scaling to application health monitoring. Customers can upload their code and Elastic Beanstalk will take care of the rest<sup>1</sup>. AWS Elastic Beanstalk also minimizes the complexity that is related to the management of AWS resources. Customers can retain full control of the underlying AWS resources powering their applications and adjust the settings to suit their needs<sup>1</sup>. Customers can also use the AWS Management Console, the AWS Command Line Interface (AWS CLI), or APIs to manage their applications<sup>1</sup>.

AWS Config is the AWS service that enables customers to assess, audit, and evaluate the configurations of their AWS resources. AWS Config continuously monitors and records the configuration changes of the resources and evaluates them against desired configurations or best practices<sup>2</sup>. AWS Config does not help customers deploy applications in the AWS Cloud as quickly as possible or minimize the complexity that is related to the management of AWS resources.

Amazon EC2 is the AWS service that provides secure, resizable compute capacity in the cloud. Customers can launch virtual servers called instances and choose from various configurations of CPU, memory, storage, and networking resources<sup>3</sup>. Amazon EC2 does not automatically handle the deployment or management of AWS resources for customers. Customers have to manually provision, configure, monitor, and scale their instances and other related resources.

Amazon Personalize is the AWS service that enables customers to create personalized recommendations for their users based on their behavior and preferences. Amazon Personalize uses machine learning to analyze data and deliver real-time recommendations<sup>4</sup>. Amazon Personalize does not help customers deploy applications in the AWS Cloud as quickly as possible or minimize the complexity that is related to the management of AWS resources.

**NEW QUESTION 44**

- (Topic 3)

A software engineer wants to launch a virtual machine (VM) and MySQL database on AWS.

Which AWS service will meet these requirements with the LEAST operational effort?

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

**Answer:** B

**Explanation:**

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

**NEW QUESTION 49**

- (Topic 3)

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

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

**Answer:** D

**Explanation:**

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

**NEW QUESTION 50**

- (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 55**

- (Topic 3)

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

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

**Answer:** C

**Explanation:**

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

**NEW QUESTION 60**

- (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 62**

- (Topic 3)

A company needs a fully managed file server that natively supports Microsoft workloads and file systems. The file server must also support the SMB protocol.

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

- A. Amazon Elastic File System (Amazon EFS)
- B. Amazon FSx for Lustre
- C. Amazon FSx for Windows File Server
- D. Amazon Elastic Block Store (Amazon EBS)

**Answer:** C

**Explanation:**

Amazon FSx for Windows File Server is a fully managed file server that supports Microsoft workloads and file systems, including the SMB protocol. It provides features such as user quotas, end-user file restore, and Microsoft Active Directory integration. Amazon EFS is a fully managed file system that supports the NFS protocol, not SMB. Amazon FSx for Lustre is a fully managed file system that supports high-performance computing workloads, not Microsoft workloads. Amazon EBS is a block storage service that does not provide a file system or SMB support. References: Amazon FSx for Windows File Server, Amazon FSx for Lustre, Amazon EFS, Amazon EBS

#### NEW QUESTION 65

- (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 70

- (Topic 3)

Which AWS services can be used to store files? (Select TWO.)

- A. Amazon S3
- B. AWS Lambda
- C. Amazon Elastic Block Store (Amazon EBS)
- D. Amazon SageMaker
- E. AWS Storage Gateway

**Answer:** AC

#### Explanation:

Amazon S3 and Amazon EBS are two AWS services that can be used to store files . Amazon S3 is an object storage service that offers high scalability, durability, availability, and performance. Amazon EBS is a block storage service that provides persistent and low-latency storage volumes for Amazon EC2 instances. AWS Lambda, Amazon SageMaker, and AWS Storage Gateway are other AWS services that have different purposes, such as serverless computing, machine learning, and hybrid cloud storage .

#### NEW QUESTION 71

- (Topic 3)

Which AWS service provides storage that can be mounted across multiple Amazon EC2 instances?

- A. Amazon Workspaces
- B. Amazon Elastic File System (Amazon EFS)
- C. AWS Database Migration Service (AWS DMS)
- D. AWS Snowball Edge

**Answer:** B

#### Explanation:

Amazon EFS is a fully managed service that provides scalable and elastic file storage for multiple Amazon EC2 instances. Amazon EFS supports the Network File System (NFS) protocol, which allows multiple EC2 instances to access the same file system concurrently. You can learn more about Amazon EFS from this webpage or this digital course.

#### NEW QUESTION 72

- (Topic 3)

A company has created an AWS Cost and Usage Report and wants to visualize the report. Which AWS service should the company use to ingest and display this information?

- A. Amazon QuickSight
- B. Amazon Pinpoint
- C. Amazon Neptune
- D. Amazon Kinesis

**Answer:** A

#### Explanation:

Amazon QuickSight is an AWS service that provides business intelligence and data visualization capabilities. Amazon QuickSight enables you to ingest, analyze, and display data from various sources, such as AWS Cost and Usage Reports, Amazon S3, Amazon Athena, Amazon Redshift, and Amazon RDS. You can use Amazon QuickSight to create interactive dashboards and charts that show insights and trends from your data. You can also share your dashboards and charts with other users or embed them into your applications.

#### NEW QUESTION 74

- (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 78**

- (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 termination<sup>12</sup>. 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 logs<sup>13</sup>.

**NEW QUESTION 82**

- (Topic 3)

A company is running its application in the AWS Cloud. The company wants to periodically review its AWS account for cost optimization opportunities. Which AWS service or tool can the company use to meet these requirements?

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

**Answer:** A

**Explanation:**

AWS Cost Explorer is an AWS service or tool that the company can use to periodically review its AWS account for cost optimization opportunities. AWS Cost Explorer is a tool that enables the company to visualize, understand, and manage their AWS costs and usage over time. The company can use AWS Cost Explorer to access interactive graphs and tables that show the breakdown of their costs and usage by service, region, account, tag, and more. The company can also use AWS Cost Explorer to forecast their future costs, identify trends and anomalies, and discover potential savings by using Reserved Instances or Savings Plans.

**NEW QUESTION 83**

- (Topic 3)

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

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

**Answer:** B

**Explanation:**

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

**NEW QUESTION 88**

- (Topic 3)

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

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

**Answer:** CE

**Explanation:**

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

#### NEW QUESTION 91

- (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 93

- (Topic 3)

A company wants to migrate its high-performance computing (HPC) application to Amazon EC2 instances. The application has multiple components. The application must have fault tolerance and must have the ability to fail over automatically.

Which AWS infrastructure solution will meet these requirements with the LEAST latency between components?

- A. Multiple AWS Regions
- B. Multiple edge locations
- C. Multiple Availability Zones
- D. Regional edge caches

**Answer:** C

#### Explanation:

Using EC2 instances in multiple Availability Zones is an AWS infrastructure solution that meets the requirements of migrating a high performance computing (HPC) application to AWS with fault tolerance and failover capabilities, and with the least latency between components. An Availability Zone is a physically isolated location within an AWS Region that has its own power, cooling, and network connectivity. EC2 instances within the same Region can communicate with each other using low-latency private IP addresses. By using EC2 instances in multiple Availability Zones, the company can achieve fault tolerance and failover for their HPC application, because they can distribute the workload and data across different locations that are independent of each other. If one Availability Zone becomes unavailable or impaired, the company can redirect the traffic and data to another Availability Zone without affecting the performance and availability of the application<sup>5</sup>

#### NEW QUESTION 95

- (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 97

- (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 102**

- (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 104**

- (Topic 3)

A company wants to migrate its database to a managed AWS service that is compatible with PostgreSQL.

Which AWS services will meet these requirements? (Select TWO)

- A. Amazon Athena
- B. Amazon RDS
- C. Amazon EC2
- D. Amazon DynamoDB
- E. Amazon Aurora

**Answer:** BE

**Explanation:**

Amazon RDS and Amazon Aurora are both managed AWS services that support the PostgreSQL database engine. Amazon RDS makes it easier to set up, operate, and scale PostgreSQL deployments on the cloud, while Amazon Aurora is a cloud-native database engine that is compatible with PostgreSQL and offers higher performance and availability. Amazon Athena is a serverless query service that does not support PostgreSQL, but can analyze data in Amazon S3 using standard SQL. Amazon EC2 is a compute service that allows users to launch virtual machines, but does not provide any database management features. Amazon DynamoDB is a NoSQL database service that is not compatible with PostgreSQL, but offers fast and consistent performance at any scale. References: Hosted PostgreSQL - Amazon RDS for PostgreSQL - AWS, Amazon RDS for PostgreSQL - Amazon Relational Database Service, AWS PostgreSQL: Managed or Self-Managed? - NetApp, AWS Announces Amazon Aurora Supports PostgreSQL 12 - InfoQ, Amazon Aurora vs PostgreSQL | What are the differences? - StackShare

**NEW QUESTION 107**

- (Topic 3)

A company has 5 TB of data stored in Amazon S3. The company plans to occasionally run queries on the data for analysis.

Which AWS service should the company use to run these queries in the MOST cost- effective manner?

- A. Amazon Redshift
- B. Amazon Athena
- C. Amazon Kinesis
- D. Amazon RDS

**Answer:** B

**Explanation:**

Amazon Athena is a serverless, interactive analytics service that allows users to run SQL queries on data stored in Amazon S3. It is ideal for occasional queries on large datasets, as it does not require any server provisioning, configuration, or management. Users only pay for the queries they run, based on the amount of data scanned. Amazon Athena supports various data formats, such as CSV, JSON, Parquet, ORC, and Avro, and integrates with AWS Glue Data Catalog to create and manage schemas. Amazon Athena also supports querying data from other sources, such as on- premises or other cloud systems, using data connectors<sup>1</sup>.

Amazon Redshift is a fully managed data warehouse service that allows users to run complex analytical queries on petabyte-scale data. However, it requires users to provision and maintain clusters of nodes, and pay for the storage and compute capacity they use. Amazon Redshift is more suitable for frequent and consistent

queries on structured or semi-structured data<sup>2</sup>.

Amazon Kinesis is a platform for streaming data on AWS, enabling users to collect, process, and analyze real-time data. It is not designed for querying data stored in Amazon S3. Amazon Kinesis consists of four services: Kinesis Data Streams, Kinesis Data Firehose, Kinesis Data Analytics, and Kinesis Video Streams<sup>3</sup>.

Amazon RDS is a relational database service that provides six database engines: Amazon Aurora, PostgreSQL, MySQL, MariaDB, Oracle Database, and SQL Server. It simplifies database administration tasks such as backup, patching, scaling, and replication. However, it is not optimized for querying data stored in Amazon S3. Amazon RDS is more suitable for transactional workloads that require high performance and availability<sup>4</sup>.

References:

? Interactive SQL - Serverless Query Service - Amazon Athena - AWS

? [Amazon Redshift – Data Warehouse Solution - AWS]

? [Amazon Kinesis - Streaming Data Platform - AWS]

? [Amazon Relational Database Service (RDS) – AWS]

#### NEW QUESTION 108

- (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 110

- (Topic 3)

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

- A. Amazon EC2
- B. AWS Elastic Beanstalk
- C. AWS CodeBuild
- D. Amazon Personalize

**Answer:** B

#### Explanation:

AWS Elastic Beanstalk is a service that allows you to deploy and manage applications on AWS without manually creating and configuring the required resources, such as EC2 instances, load balancers, security groups, databases, and more. AWS Elastic Beanstalk automatically handles the provisioning, scaling, load balancing, health monitoring, and updating of your application, while giving you full control over the underlying AWS resources if needed. AWS Elastic Beanstalk supports a variety of platforms and languages, such as Java, .NET, PHP, Node.js, Python, Ruby, Go, and Docker. You can use the AWS Management Console, the AWS CLI, the AWS SDKs, or the AWS Elastic Beanstalk API to create and manage your applications. You can also use AWS CodeStar, AWS CodeCommit, AWS CodeBuild, AWS CodeDeploy, and AWS CodePipeline to integrate AWS Elastic Beanstalk with your development and deployment workflows<sup>12</sup>

#### NEW QUESTION 115

- (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 120

- (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 121**

- (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 124**

- (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 129**

- (Topic 3)

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

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

**Answer:** A

**Explanation:**

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

### NEW QUESTION 130

- (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 131

- (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 134

- (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 136

- (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 141

- (Topic 3)

A company wants a list of all users in its AWS account, the status of all of the users' access keys, and if multi-factor authentication (MFA) has been configured. Which AWS service or feature will meet these requirements?

- A. AWS Key Management Service (AWS KMS)
- B. IAM Access Analyzer
- C. IAM credential report
- D. Amazon CloudWatch

**Answer: C**

#### Explanation:

IAM credential report is a feature that allows you to generate and download a report that lists all IAM users in your AWS account and the status of their various credentials, including access keys and MFA devices. You can use this report to audit the security status of your IAM users and ensure that they follow the best practices for using AWS1.

AWS Key Management Service (AWS KMS) is a service that allows you to create and manage encryption keys to protect your data. It does not provide information about IAM users or their credentials2.

IAM Access Analyzer is a feature that helps you identify the resources in your AWS account, such as S3 buckets or IAM roles, that are shared with an external entity. It does not provide information about IAM users or their credentials3.

Amazon CloudWatch is a service that monitors and collects metrics, logs, and events from your AWS resources and applications. It does not provide information about IAM users or their credentials4.

References:

- ? Getting credential reports for your AWS account - AWS Identity and Access Management
- ? AWS Key Management Service - Amazon Web Services
- ? IAM Access Analyzer - AWS Identity and Access Management
- ? Amazon CloudWatch - Amazon Web Services

#### NEW QUESTION 143

- (Topic 3)

A company is migrating its data center to AWS. The company needs an AWS Support plan that provides chat access to a cloud support engineer 24 hours a day, 7 days a week. The company does not require access to infrastructure event management.

What is the MOST cost-effective AWS Support plan that meets these requirements?

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

**Answer: B**

#### Explanation:

AWS Business Support is the most cost-effective AWS Support plan that provides chat access to a cloud support engineer 24/7. AWS Business Support also offers phone and email support, as well as a response time of less than one hour for urgent issues. AWS Business Support does not include access to infrastructure event management, which is a feature of AWS Enterprise Support. AWS Enterprise Support is more expensive and provides additional benefits, such as a technical account manager, a support concierge, and a response time of less than 15 minutes for critical issues. AWS Developer Support and AWS Basic Support do not provide chat access to a cloud support engineer. AWS Developer Support provides email support and a response time of less than 12 hours for general guidance issues. AWS Basic Support provides customer service and account support, as well as access to forums and documentation1

#### NEW QUESTION 146

- (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 151

- (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 152**

- (Topic 3)

A company must be able to develop, test, and launch an application in the AWS Cloud quickly. Which advantage of cloud computing will meet these requirements?

- A. Stop guessing capacity
- B. Trade fixed expense for variable expense
- C. Achieve economies of scale
- D. Increase speed and agility

**Answer:** D

**Explanation:**

One of the benefits of cloud computing is that it enables customers to increase speed and agility in developing, testing, and launching applications. Cloud computing provides on-demand access to a variety of IT resources, such as compute, storage, networking, databases, and analytics, without requiring upfront investments or long-term commitments. Customers can provision and release resources in minutes, scale up and down as needed, and experiment with new technologies and features. This allows customers to accelerate their innovation cycles, deliver faster time-to-market, and respond to changing customer needs and demands

**NEW QUESTION 156**

- (Topic 3)

A company is running an Amazon EC2 instance in a VPC.

An ecommerce company is using Amazon EC2 Auto Scaling groups to manage a fleet of web servers running on Amazon EC2.

This architecture follows which AWS Well-Architected Framework best practice?

- A. Secure the workload
- B. Decouple infrastructure components
- C. Design for failure
- D. Think parallel

**Answer:** C

**Explanation:**

Design for failure is one of the best practices of the AWS Well-Architected Framework. It means that the architecture should be resilient and fault-tolerant, and able to handle failures without impacting the availability and performance of the applications. By using Amazon EC2 Auto Scaling groups, the ecommerce company can design for failure by automatically scaling the number of EC2 instances up or down based on demand or health status. Amazon EC2 Auto Scaling groups can also distribute the EC2 instances across multiple Availability Zones, which are isolated locations within an AWS Region that have independent power, cooling, and network connectivity. This way, the company can ensure that their web servers can handle traffic spikes, recover from failures, and provide a consistent user experience

**NEW QUESTION 161**

- (Topic 3)

Which of the following are general AWS Cloud design principles described in the AWS Well-Architected Framework?

- A. Consolidate key components into monolithic architectures.
- B. Test systems at production scale.
- C. Provision more capacity than a workload is expected to need.
- D. Drive architecture design based on data collected about the workload behavior and requirements.
- E. Make AWS Cloud architectural decisions static, one-time events.

**Answer:** BD

**Explanation:**

These are two of the general AWS Cloud design principles described in the AWS Well-Architected Framework. Testing systems at production scale means using tools such as AWS CloudFormation, AWS CodeDeploy, and AWS X-Ray to simulate real-world scenarios and measure the performance, scalability, and availability of the system. Driving architecture design based on data means using tools such as Amazon CloudWatch, AWS CloudTrail, and AWS Config to collect and analyze metrics, logs, and events about the system and use the insights to optimize the system's design and operation. You can learn more about the AWS Well-Architected Framework from this whitepaper or [this digital course].

**NEW QUESTION 166**

- (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 168**

- (Topic 3)

A company needs to set up user authentication for a new application. Users must be able to sign in directly with a user name and password, or through a third-party provider.

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

- A. AWS IAM Identity Center (AWS Single Sign-On)
- B. AWS Signer
- C. Amazon Cognito
- D. AWS Directory Service

**Answer:** C

**Explanation:**

Amazon Cognito is a service that provides user authentication and authorization for web and mobile applications. You can use Amazon Cognito to enable users to sign in directly with a user name and password, or through a third-party provider, such as Facebook, Google, or Amazon. You can also use Amazon Cognito to manage user profiles, preferences, and security settings.

**NEW QUESTION 169**

- (Topic 3)

A company wants to use guidelines from the AWS Well-Architected Framework to limit human error and facilitate consistent responses to events.

Which of the following is a Well-Architected design principle that will meet these requirements?

- A. Use AWS CodeDeploy.
- B. Perform operations as code.
- C. Migrate workloads to a Dedicated Host.
- D. Use AWS Compute Optimizer.

**Answer:** B

**Explanation:**

This is a design principle of the operational excellence pillar of the AWS Well-Architected Framework. Performing operations as code means using scripts, templates, or automation tools to perform routine tasks, such as provisioning, configuration, deployment, and monitoring. This reduces human error, increases consistency, and enables faster recovery from failures. You can learn more about the operational excellence pillar from this whitepaper or this digital course.

**NEW QUESTION 173**

- (Topic 3)

A company is building an application that needs to deliver images and videos globally with minimal latency.

Which approach can the company use to accomplish this in a cost effective manner?

- A. Deliver the content through Amazon CloudFront.
- B. Store the content on Amazon S3 and enable S3 cross-region replication.
- C. Implement a VPN across multiple AWS Regions.
- D. Deliver the content through AWS PrivateLink.

**Answer:** A

**Explanation:**

Amazon CloudFront is a fast content delivery network (CDN) service that securely delivers data, videos, applications, and APIs to customers globally with low latency, high transfer speeds, all within a developer-friendly environment. It works seamlessly with services including AWS Shield for DDoS mitigation, Amazon S3, Elastic Load Balancing or Amazon EC2 as origins for your applications, and Lambda@Edge to run custom code closer to customers' users and to customize the user experience. By using CloudFront, you can cache your content at the edge locations that are closest to your end users, reducing the network latency and improving the performance of your application. CloudFront also offers a pay-as-you-go pricing model, so you only pay for the data transfer and requests that you use.

**NEW QUESTION 178**

- (Topic 3)

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

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

**Answer:** D

**Explanation:**

Amazon ElastiCache is a fully managed in-memory data store and cache service that delivers sub-millisecond response times to applications. You can use ElastiCache as a primary data store for your applications, or as a cache to improve the performance of your existing databases. ElastiCache supports two popular open-source in-memory engines: Redis and Memcached5.

**NEW QUESTION 183**

- (Topic 3)

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

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

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

**Answer:** A

**Explanation:**

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

A multi-region deployment model is a cloud computing architecture that distributes an application and its data across multiple geographic regions. A multi-region deployment model enables a company to achieve global reach, high availability, disaster recovery, and performance optimization. By deploying an application in multiple regions, a company can serve customers from the nearest region, reduce latency, increase redundancy, and comply with data sovereignty regulations12.

A single-region deployment model is a cloud computing architecture that runs an application and its data within a single geographic region. A single-region deployment model is simpler and cheaper than a multi-region deployment model, but it has limited scalability, availability, and performance. A single-region deployment model may not be suitable for a company that wants to deploy an application globally, as it may face challenges such as network latency, regional outages, or regulatory compliance12.

A multi-AZ (Availability Zone) deployment model is a cloud computing architecture that distributes an application and its data across multiple isolated locations within a single region. An Availability Zone is a physically separate location within an AWS Region that has independent power, cooling, and networking. A multi-AZ deployment model enhances the availability and durability of an application by providing redundancy and fault tolerance within a region34.

A single-AZ deployment model is a cloud computing architecture that runs an application and its data within a single Availability Zone. A single-AZ deployment model is the simplest and most cost-effective option, but it has no redundancy or fault tolerance. A single-AZ deployment model may not be suitable for a company that wants to deploy an application globally, as it may face challenges such as network latency, regional outages, or regulatory compliance34.

References:

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

**NEW QUESTION 186**

- (Topic 3)

An ecommerce company wants to distribute traffic between the Amazon EC2 instances that host its website.

Which AWS service or resource will meet these requirements?

- A. Application Load Balancer
- B. AWS WAF
- C. AWS CloudHSM
- D. AWS Direct Connect

**Answer:** A

**Explanation:**

This is the AWS service or resource that will meet the requirements of distributing traffic between the Amazon EC2 instances that host the website. Application Load Balancer is a type of Elastic Load Balancing that distributes incoming application traffic across multiple targets, such as Amazon EC2 instances, containers, IP addresses, and Lambda functions. Application Load Balancer operates at the application layer (layer 7) of the OSI model and supports advanced features such as path-based routing, host-based routing, health checks, and SSL termination. You can learn more about Application Load Balancer from [this webpage] or [this digital course].

**NEW QUESTION 190**

- (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 192**

- (Topic 3)

A company wants to migrate to AWS and use the same security software it uses on premises. The security software vendor offers its security software as a service on AWS.

Where can the company purchase the security solution?

- A. AWS Partner Solutions Finder
- B. AWS Support Center
- C. AWS Management Console
- D. AWS Marketplace

**Answer:** D

**Explanation:**

AWS Marketplace is an online store that helps customers find, buy, and immediately start using the software and services that run on AWS. Customers can choose from a wide range of software products in popular categories such as security, networking, storage, machine learning, business intelligence, database, and DevOps. Customers can also use AWS Marketplace to purchase software as a service (SaaS) solutions that are integrated with AWS. Customers can benefit from simplified procurement, billing, and deployment processes, as well as flexible pricing options and free trials. Customers can also leverage AWS Marketplace to discover and subscribe to solutions offered by AWS Partners, such as the security software vendor mentioned in the question. References: AWS Marketplace, [AWS Marketplace: Software as a Service (SaaS)], [AWS Cloud Practitioner Essentials: Module 6 - AWS Pricing, Billing, and Support]

**NEW QUESTION 195**

- (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.

**NEW QUESTION 199**

- (Topic 3)

A company is looking for a managed machine learning (ML) service that can recommend products based on a customer's previous behaviors. Which AWS service meets this requirement?

- A. Amazon Personalize
- B. Amazon SageMaker
- C. Amazon Pinpoint
- D. Amazon Comprehend

**Answer:** A

**Explanation:**

The AWS service that meets the requirement of providing a managed machine learning (ML) service that can recommend products based on a customer's previous behaviors is Amazon Personalize. Amazon Personalize is a fully managed service that enables developers to create personalized recommendations for customers using their own data. Amazon Personalize can automatically process and examine the data, identify what is meaningful, select the right algorithms, and train and optimize a personalized recommendation model. Amazon SageMaker, Amazon Pinpoint, and Amazon Comprehend are other AWS services related to machine learning, but they do not provide the specific functionality of product recommendation.

**NEW QUESTION 201**

- (Topic 3)

Which task does AWS perform automatically?

- A. Encrypt data that is stored in Amazon DynamoDB.
- B. Patch Amazon EC2 instances.
- C. Encrypt user network traffic.
- D. Create TLS certificates for users' websites.

**Answer:** B

**Explanation:**

AWS performs some tasks automatically to help you manage and secure your AWS resources. One of these tasks is patching Amazon EC2 instances. AWS provides two options for patching your EC2 instances: managed instances and patch baselines. Managed instances are a group of EC2 instances or on-premises servers that you can manage using AWS Systems Manager. Patch baselines define the patches that AWS Systems Manager applies to your instances. You can use AWS Systems Manager to automate the process of patching your instances based on a schedule or a maintenance window.

**NEW QUESTION 205**

- (Topic 3)

A user needs a relational database but does not have the resources to manage the hardware, resiliency, and replication. Which AWS service option meets the user's requirements'?

- A. Run MySQL on Amazon Elastic Container Service (Amazon ECS)

- B. Run MySQL on Amazon EC2
- C. Choose Amazon RDS for MySQL
- D. Choose Amazon ElastiCache for Redis

**Answer:** C

**Explanation:**

Amazon RDS for MySQL is a fully managed, open-source cloud database service that allows you to easily operate and scale your relational database of choice, including MySQL. With Amazon RDS for MySQL, you don't have to worry about the hardware, resiliency, and replication of your database, as Amazon RDS handles these tasks for you. Amazon RDS for MySQL also provides features such as automated backups, multi-AZ deployments, read replicas, encryption, monitoring, and more. Amazon RDS for MySQL is compatible with the MySQL Community Edition versions 5.7 and 8.0, which means that you can use the same code, applications, and tools that you already use with MySQL4567. References: 4: Hosted MySQL - Amazon RDS for MySQL - AWS, 5: Amazon RDS for MySQL - Amazon Relational Database Service, 6: Amazon RDS for MySQL —, 7: Managed SQL Database - Amazon Relational Database Service (RDS) - AWS

**NEW QUESTION 210**

- (Topic 3)

A company wants a time-series database service that makes it easier to store and analyze trillions of events each day. Which AWS service will meet this requirement?

- A. Amazon Neptune
- B. Amazon Timestream
- C. Amazon Forecast
- D. Amazon DocumentDB (with MongoDB compatibility)

**Answer:** B

**Explanation:**

Amazon Timestream is a fast, scalable, and serverless time-series database service for IoT and other operational applications that makes it easy to store and analyze trillions of events per day up to 1,000 times faster and at as little as 1/10th the cost of relational databases1. Amazon Timestream saves you time and cost in managing the lifecycle of time series data, and its purpose-built query engine lets you access and analyze recent and historical data together with a single query1. Amazon Timestream has built-in time series analytics functions, helping you identify trends and patterns in near real time1. The other options are not suitable for storing and analyzing trillions of events per day. Amazon Neptune is a graph database service that supports highly connected data sets. Amazon Forecast is a machine learning service that generates accurate forecasts based on historical data. Amazon DocumentDB (with MongoDB compatibility) is a document database service that supports MongoDB workloads.

References:

? 1: Time Series Database – Amazon Timestream – Amazon Web Services

**NEW QUESTION 215**

- (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 define67.

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

breached89. 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 219**

- (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 behalf3

? 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 AWS4

#### NEW QUESTION 224

- (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](https://aws.amazon.com), Regions, Availability Zones, and Local Zones - Amazon Relational Database Service

#### NEW QUESTION 227

- (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 228

- (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 232

- (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 234

- (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 239

- (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. 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 240

- (Topic 3)

A company is building an application on AWS. The application needs to comply with credit card regulatory requirements. The company needs proof that the AWS services and deployment are in compliance. Which actions should the company take to meet these requirements? (Select TWO.)

- A. Use Amazon Inspector to submit the application for certification.
- B. Ensure that the application's underlying hardware components comply with requirements.
- C. Use AWS Artifact to access AWS documents about the compliance of the services.
- D. Get the compliance of the application certified by a company assessor.
- E. Use AWS Security Hub to certify the compliance of the application.

**Answer: CD**

#### Explanation:

Using AWS Artifact to access AWS documents about the compliance of the services, and getting the compliance of the application certified by a company assessor are actions that the company should take to meet the requirements of complying with credit card regulatory requirements. AWS Artifact is a service that 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. AWS Artifact can help you demonstrate compliance with credit card regulatory requirements by providing you with proof that the AWS services and deployment are in compliance. Getting the compliance of the application certified by a company assessor is an action that the company should take to ensure that the application meets the specific requirements of the credit card industry. A company assessor is an independent third-party entity that is qualified to assess the compliance of the application with the relevant standards and regulations. Using Amazon Inspector to submit the application for certification is not an action that the company should take, because Amazon Inspector is a service that helps you improve the security and compliance of your applications deployed on AWS by automatically assessing them for vulnerabilities and deviations from best practices, but it does not provide certification for the applications. Ensuring that the application's underlying hardware components comply with requirements is not an action that the company should take, because the application is deployed on AWS, and AWS is responsible for the security and compliance of the underlying hardware components. This is part of the shared responsibility model, where AWS is responsible for security of the cloud, and customers are responsible for security in the cloud. Using AWS Security Hub to certify the compliance of the application is not an action that the company should take, because AWS Security Hub is a service that gives you a comprehensive view of your security posture across your AWS accounts and helps you check your environment against security industry standards and best practices, but it does not provide certification for the applications.

#### NEW QUESTION 242

- (Topic 3)

Which AWS services are supported by Savings Plans? (Select TWO.)

- A. Amazon EC2
- B. Amazon RDS
- C. Amazon SageMaker
- D. Amazon Redshift
- E. Amazon DynamoDB

**Answer:** AC

**Explanation:**

The AWS services that are supported by Savings Plans are:

? Amazon EC2: Amazon EC2 is a service that provides scalable computing capacity in the AWS cloud. You can use Amazon EC2 to launch virtual servers, configure security and networking, and manage storage. Amazon EC2 is eligible for both Compute Savings Plans and EC2 Instance Savings Plans<sup>12</sup>.

? Amazon SageMaker: Amazon SageMaker is a service that helps you build and deploy machine learning models. You can use Amazon SageMaker to access Jupyter notebooks, use common machine learning algorithms, train and tune models, and deploy them to a hosted environment. Amazon SageMaker is eligible for SageMaker Savings Plans<sup>13</sup>.

The other options are not supported by Savings Plans. Amazon RDS, Amazon Redshift, and Amazon DynamoDB are database services that are eligible for Reserved Instances, but not Savings Plans<sup>4</sup>.

**NEW QUESTION 247**

- (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). Network ACLs can be configured with rules that allow or deny traffic based on the source and destination IP addresses, ports, and protocols<sup>1</sup>. AWS Security Hub is a service that provides a comprehensive view of the security posture of AWS accounts and resources<sup>2</sup>. Security groups are features that act as firewalls for controlling traffic at the instance level<sup>3</sup>. AWS WAF is a web application firewall that helps protect web applications from common web exploits<sup>4</sup>.

**NEW QUESTION 252**

- (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>345</sup>.

**NEW QUESTION 255**

- (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 257

- (Topic 3)

A company has deployed an Amazon EC2 instance.

Which option is an AWS responsibility under the AWS shared responsibility model?

- A. Managing and encrypting application data
- B. Installing updates and security patches of guest operating system
- C. Configuration of infrastructure devices
- D. Configuration of security groups on each instance

**Answer: C**

#### Explanation:

According to the AWS shared responsibility model, AWS is responsible for protecting the infrastructure that runs all of the services offered in the AWS Cloud, such as data centers, hardware, software, networking, and facilities<sup>1</sup>. This includes the configuration of infrastructure devices, such as routers, switches, firewalls, and load balancers<sup>2</sup>. Customers are responsible for managing their data, applications, operating systems, security groups, and other aspects of their AWS environment<sup>1</sup>. Therefore, options A, B, and D are customer responsibilities, not AWS responsibilities. References: 1: AWS Well-Architected Framework - Elasticity; 2: Reactive Systems on AWS - Elastic

#### NEW QUESTION 258

- (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 263

- (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 264

- (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 269

- (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 272**

- (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 274**

.....

## THANKS FOR TRYING THE DEMO OF OUR PRODUCT

Visit Our Site to Purchase the Full Set of Actual CLF-C01 Exam Questions With Answers.

We Also Provide Practice Exam Software That Simulates Real Exam Environment And Has Many Self-Assessment Features. Order the CLF-C01 Product From:

<https://www.2passeasy.com/dumps/CLF-C01/>

### Money Back Guarantee

#### **CLF-C01 Practice Exam Features:**

- \* CLF-C01 Questions and Answers Updated Frequently
- \* CLF-C01 Practice Questions Verified by Expert Senior Certified Staff
- \* CLF-C01 Most Realistic Questions that Guarantee you a Pass on Your FirstTry
- \* CLF-C01 Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year