

# Exam Questions AWS-Certified-Cloud-Practitioner

Amazon AWS Certified Cloud Practitioner

<https://www.2passeasy.com/dumps/AWS-Certified-Cloud-Practitioner/>



### NEW QUESTION 1

- (Topic 2)

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

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

**Answer:** D

#### Explanation:

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

### NEW QUESTION 2

- (Topic 2)

A company manages factory machines in real time. The company wants to use AWS technology to deploy its monitoring applications as close to the factory machines as possible.

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

- A. AWS Outposts
- B. Amazon EC2
- C. AWS App Runner
- D. AWS Batch

**Answer:** A

#### Explanation:

AWS Outposts is a fully managed service that extends AWS infrastructure, AWS services, APIs, and tools to virtually any datacenter, co-location space, or on-premises facility for a truly consistent hybrid experience. AWS Outposts enables you to run AWS services in your on-premises data center<sup>1</sup>.

### NEW QUESTION 3

- (Topic 2)

Which AWS service or tool provides on-demand access to AWS security and compliance reports and AWS online agreements?

- A. AWS Artifact
- B. AWS Trusted Advisor
- C. Amazon Inspector
- D. AWS Billing console

**Answer:** A

#### Explanation:

AWS Artifact is the AWS service or tool that provides on-demand access to AWS security and compliance reports and AWS online agreements. AWS Trusted Advisor is a tool that provides real-time guidance to help users provision their resources following AWS best practices. Amazon Inspector is a service that helps users improve the security and compliance of their applications. AWS Billing console is a tool that helps users manage their AWS costs and usage. These concepts are explained in the AWS Cloud Practitioner Essentials course<sup>3</sup>.

### NEW QUESTION 4

- (Topic 2)

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

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

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

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

**Answer:** B

#### Explanation:

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

### NEW QUESTION 5

- (Topic 2)

A company wants to access a report about the estimated environmental impact of the company's AWS usage.

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

- A. AWS Organizations
- B. IAM policy
- C. AWS Billing console
- D. Amazon Simple Notification Service (Amazon SNS)

**Answer:** C

**Explanation:**

The company should use the AWS Billing console to access a report about the estimated environmental impact of the company's AWS usage. The AWS Billing console provides customers with various tools and reports to manage and monitor their AWS costs and usage. One of the reports available in the AWS Billing console is the AWS Sustainability Dashboard, which shows the estimated carbon footprint and energy mix of the customer's AWS usage. The company can use this dashboard to measure and improve the sustainability of their cloud workloads. AWS Organizations, IAM policy, and Amazon Simple Notification Service (Amazon SNS) are not services or features that can provide a report about the estimated environmental impact of the company's AWS usage. AWS Organizations is a service that enables customers to centrally manage and govern their AWS accounts. IAM policy is a document that defines the permissions for an IAM identity (user, group, or role) or an AWS resource. Amazon SNS is a fully managed pub/sub messaging service that enables customers to send messages to subscribers or other AWS services.

**NEW QUESTION 6**

- (Topic 2)

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

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

**Answer:** C

**Explanation:**

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

**NEW QUESTION 7**

- (Topic 1)

A company uses Amazon Aurora as its database service. The company wants to encrypt its databases and database backups. Which party manages the encryption of the database clusters and database snapshots, according to the AWS shared responsibility model?

- A. AWS
- B. The company
- C. AWS Marketplace partners
- D. Third-party partners

**Answer:** A

**Explanation:**

AWS manages the encryption of the database clusters and database snapshots for Amazon Aurora, as well as the encryption keys. This is part of the AWS shared responsibility model, where AWS is responsible for the security of the cloud, and the customer is responsible for the security in the cloud. Encryption is one of the security features that AWS provides to protect the data at rest and in transit. For more information, see Amazon Aurora FAQs and AWS Shared Responsibility Model.

**NEW QUESTION 8**

- (Topic 1)

What is an Availability Zone?

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

**Answer:** B

**Explanation:**

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

**NEW QUESTION 9**

- (Topic 1)

Which AWS Support plan assigns an AWS concierge agent to a company's account?

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

**Answer:** D

**Explanation:**

AWS Enterprise Support is the AWS Support plan that assigns an AWS concierge agent to a company's account. AWS Enterprise Support is the highest level of support that AWS offers, and it provides the most comprehensive and personalized assistance. An AWS concierge agent is a dedicated technical account manager who acts as a single point of contact for the company and helps to optimize the AWS environment, resolve issues, and access AWS experts. For more information, see [AWS Support Plans] and [AWS Concierge Support].

**NEW QUESTION 10**

- (Topic 1)

A company is migrating a relational database server to the AWS Cloud. The company wants to minimize administrative overhead of database maintenance tasks. Which AWS service will meet these requirements?

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

**Answer:** D

**Explanation:**

Amazon RDS is the AWS service that will meet the requirements of migrating a relational database server to the AWS Cloud and minimizing administrative overhead of database maintenance tasks. Amazon RDS is a fully managed relational database service that handles routine database tasks, such as provisioning, patching, backup, recovery, failure detection, and repair. Amazon RDS supports several database engines, such as MySQL, PostgreSQL, Oracle, SQL Server, and Amazon Aurora5.

**NEW QUESTION 10**

- (Topic 1)

A company wants to establish a security layer in its VPC that will act as a firewall to control subnet traffic. Which AWS service or feature will meet this requirement?

- A. Routing tables
- B. Network access control lists (network ACLs)
- C. Security groups
- D. Amazon GuardDuty

**Answer:** C

**Explanation:**

Security groups are the service or feature that meets the requirement of establishing a security layer in a VPC that will act as a firewall to control subnet traffic. Security groups are stateful firewalls that control the inbound and outbound traffic at the instance level. You can assign one or more security groups to each instance in a VPC, and specify the rules that allow or deny traffic based on the protocol, port, and source or destination. Security groups are associated with network interfaces, and therefore apply to all the instances in the subnets that use those network interfaces. Routing tables are used to direct traffic between subnets and gateways, not to filter traffic. Network ACLs are stateless firewalls that control the inbound and outbound traffic at the subnet level, but they are less granular and more cumbersome to manage than security groups. Amazon GuardDuty is a threat detection service that monitors your AWS account and workloads for malicious or unauthorized activity, not a firewall service.

**NEW QUESTION 13**

- (Topic 1)

A company needs to store data across multiple Availability Zones in an AWS Region. The data will not be accessed regularly but must be immediately retrievable. Which Amazon Elastic File System (Amazon EFS) storage class meets these requirements MOST cost effectively?

- A. EFS Standard
- B. EFS Standard-Infrequent Access(EFS Standard-IA)
- C. EFS One Zone
- D. EFS One Zone-Infrequent Access (EFS One Zone-IA)

**Answer:** B

**Explanation:**

EFS Standard-Infrequent Access (EFS Standard-IA) is the storage class that meets the requirements of storing data across multiple Availability Zones in an AWS Region, that will not be accessed regularly but must be immediately retrievable, most cost- effectively. EFS Standard-IA is designed for files that are accessed less frequently, but still require the same high performance, low latency, and high availability as EFS Standard. EFS Standard-IA has a lower storage cost than EFS Standard, but charges a small additional fee for each access. EFS One Zone and EFS One Zone-IA store data in a single Availability Zone, which reduces the availability and durability compared to EFS Standard and EFS Standard-IA.

**NEW QUESTION 14**

- (Topic 1)

In which of the following AWS services should database credentials be stored for maximum security?

- A. AWS Identity and Access Management (IAM)
- B. AWS Secrets Manager
- C. Amazon S3

D. AWS Key Management Service (AWS KMS)

**Answer:** B

**Explanation:**

AWS Secrets Manager is the AWS service where database credentials should be stored for maximum security. AWS Secrets Manager helps to protect the secrets, such as database credentials, passwords, API keys, and tokens, that are used to access applications, services, and resources. AWS Secrets Manager enables secure storage, encryption, rotation, and retrieval of the secrets. AWS Secrets Manager also integrates with other AWS services, such as AWS Identity and Access Management (IAM), AWS Key Management Service (AWS KMS), and AWS Lambda. For more information, see [What is AWS Secrets Manager?] and [Getting Started with AWS Secrets Manager].

**NEW QUESTION 15**

- (Topic 1)

Which design principle is achieved by following the reliability pillar of the AWS Well- Architected Framework?

- A. Vertical scaling
- B. Manual failure recovery
- C. Testing recovery procedures
- D. Changing infrastructure manually

**Answer:** C

**Explanation:**

Testing recovery procedures is the design principle that is achieved by following the reliability pillar of the AWS Well-Architected Framework. The reliability pillar focuses on the ability of a system to recover from failures and prevent disruptions. Testing recovery procedures helps to ensure that the system can handle different failure scenarios and restore normal operations as quickly as possible. Testing recovery procedures also helps to identify and mitigate any risks or gaps in the system design and implementation. For more information, see [Reliability Pillar] and [Testing for Reliability].

**NEW QUESTION 20**

- (Topic 1)

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

What is the MOST operationally efficient solution to delegate permissions?

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

**Answer:** A

**Explanation:**

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

**NEW QUESTION 23**

- (Topic 1)

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

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

**Answer:** A

**Explanation:**

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

**NEW QUESTION 24**

- (Topic 1)

Which design principle should be considered when architecting in the AWS Cloud?

- A. Think of servers as non-disposable resources.
- B. Use synchronous integration of services.
- C. Design loosely coupled components.
- D. Implement the least permissive rules for security groups.

**Answer:** C

**Explanation:**

Designing loosely coupled components is a design principle that should be considered when architecting in the AWS Cloud. Loose coupling is a way of designing systems to reduce interdependencies and minimize the impact of changes. Loose coupling allows components to interact with each other through well-defined interfaces, rather than direct references. This reduces the risk of failures and errors propagating across the system, and enables greater scalability, availability, and maintainability<sup>5</sup>.

#### NEW QUESTION 29

- (Topic 1)

Which task requires the use of AWS account root user credentials?

- A. The deletion of IAM users
- B. The change to a different AWS Support plan
- C. The creation of an organization in AWS Organizations
- D. The deletion of Amazon EC2 instances

**Answer:** C

#### Explanation:

The creation of an organization in AWS Organizations requires the use of AWS account root user credentials. The AWS account root user is the email address that was used to create the AWS account. The root user has complete access to all AWS services and resources in the account, and can perform sensitive tasks such as changing the account settings, closing the account, or creating an organization. The root user credentials should be used sparingly and securely, and only for tasks that cannot be performed by IAM users or roles<sup>4</sup>

#### NEW QUESTION 33

- (Topic 1)

Which AWS Support plan provides customers with access to an AWS technical account manager (TAM)?

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

**Answer:** D

#### Explanation:

The correct answer is D because AWS Enterprise Support is the support plan that provides customers with access to an AWS technical account manager (TAM). AWS Enterprise Support is the highest level of support plan offered by AWS, and it provides customers with the most comprehensive and personalized support experience. An AWS TAM is a dedicated technical resource who works closely with customers to understand their business and technical needs, provide proactive guidance, and coordinate support across AWS teams. The other options are incorrect because they are not support plans that provide customers with access to an AWS TAM. AWS Basic Support is the default and free support plan that provides customers with access to online documentation, forums, and account information. AWS Developer Support is the lowest level of paid support plan that provides customers with access to technical support during business hours, general guidance, and best practice recommendations. AWS Business Support is the intermediate level of paid support plan that provides customers with access to technical support 24/7, system health checks, architectural guidance, and case management. Reference: AWS Support Plans

#### NEW QUESTION 34

- (Topic 1)

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

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

**Answer:** C

#### Explanation:

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

#### NEW QUESTION 36

- (Topic 1)

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

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

**Answer:** B

#### Explanation:

Economies of scale is the advantage of AWS Cloud computing that minimizes variable costs. Economies of scale refers to the reduction in the cost per unit as the output increases. AWS Cloud computing leverages economies of scale by providing a large pool of shared resources that can be accessed on demand and paid for as needed. AWS Cloud computing also passes the cost savings to the customers by offering lower prices and discounts. For more information, see Economies of Scale and AWS Pricing.

#### NEW QUESTION 38

- (Topic 1)

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

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

**Answer:** D

**Explanation:**

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

**NEW QUESTION 39**

- (Topic 1)

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

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

**Answer:** A

**Explanation:**

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

**NEW QUESTION 41**

- (Topic 1)

Which AWS service is a key-value database that provides sub-millisecond latency on a large scale?

- A. Amazon DynamoDB
- B. Amazon Aurora
- C. Amazon DocumentDB (with MongoDB compatibility)
- D. Amazon Neptune

**Answer:** A

**Explanation:**

The correct answer is A because Amazon DynamoDB is a key-value database that provides sub-millisecond latency on a large scale. Amazon DynamoDB is a fully managed, serverless, and scalable NoSQL database service that supports both key-value and document data models. The other options are incorrect because they are not key-value databases. Amazon Aurora is a relational database that is compatible with MySQL and PostgreSQL. Amazon DocumentDB (with MongoDB compatibility) is a document database that is compatible with MongoDB. Amazon Neptune is a graph database that supports property graph and RDF models. Reference: Amazon DynamoDB FAQs

**NEW QUESTION 46**

- (Topic 1)

company wants to protect its AWS Cloud information, systems, and assets while performing risk assessment and mitigation tasks. Which pillar of the AWS Well-Architected Framework is supported by these goals?

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

**Answer:** B

**Explanation:**

The pillar of the AWS Well-Architected Framework that is supported by the goals of protecting AWS Cloud information, systems, and assets while performing risk assessment and mitigation tasks is security. Security is the ability to protect information, systems, and assets while delivering business value through risk assessments and mitigation strategies. The security pillar covers topics such as identity and access management, data protection, infrastructure protection, detective controls, incident response, and compliance

**NEW QUESTION 50**

- (Topic 1)

Which pillar of the AWS Well-Architected Framework includes a design principle about measuring the overall efficiency of workloads in terms of business value?

- A. Operational excellence
- B. Security
- C. Reliability

D. Cost optimization

**Answer:** A

**Explanation:**

The operational excellence pillar of the AWS Well-Architected Framework includes a design principle about measuring the overall efficiency of workloads in terms of business value. This principle states that you should monitor and measure key performance indicators (KPIs) and set targets and thresholds that align with your business goals. You should also use feedback loops to continuously improve your processes and procedures<sup>1</sup>.

**NEW QUESTION 51**

- (Topic 1)

When designing AWS workloads to be operational even when there are component failures, what is an AWS best practice?

- A. Perform quarterly disaster recovery tests.
- B. Place the main component on the us-east-1 Region.
- C. Design for automatic failover to healthy resources.
- D. Design workloads to fit on a single Amazon EC2 instance.

**Answer:** C

**Explanation:**

Designing for automatic failover to healthy resources is an AWS best practice when designing AWS workloads to be operational even when there are component failures. This means that you should architect your system to handle the loss of one or more components without impacting the availability or performance of your application. You can use various AWS services and features to achieve this, such as Auto Scaling, Elastic Load Balancing, Amazon Route 53, Amazon CloudFormation, and AWS CloudFormation<sup>4</sup>.

**NEW QUESTION 54**

- (Topic 1)

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

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

**Answer:** B

**Explanation:**

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

**NEW QUESTION 56**

- (Topic 1)

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

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

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

**Answer:** B

**Explanation:**

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

**NEW QUESTION 58**

- (Topic 1)

A company deploys its application to multiple AWS Regions and configures automatic failover between those Regions. Which cloud concept does this architecture represent?

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

**Answer:** B

**Explanation:**

Reliability is the cloud concept that this architecture represents. Reliability is the ability of a system to recover from infrastructure or service disruptions,

dynamically acquire computing resources to meet demand, and mitigate disruptions such as misconfigurations or transient network issues. Deploying an application to multiple AWS Regions and configuring automatic failover between those Regions enhances the reliability of the application by reducing the impact of regional failures and increasing the availability of the application<sup>4</sup>

#### NEW QUESTION 60

- (Topic 3)

Which database engines does Amazon Aurora support? (Select TWO.)

- A. Oracle
- B. Microsoft SQL Server
- C. MySQL
- D. PostgreSQL
- E. MongoDB

**Answer:** CD

#### Explanation:

Amazon Aurora is a relational database service that is compatible with MySQL and PostgreSQL engines. It delivers up to five times the performance of MySQL and up to three times the performance of PostgreSQL. It also provides high availability, scalability, security, and durability<sup>1</sup>

#### NEW QUESTION 64

- (Topic 3)

A company wants to migrate its workloads to AWS, but it lacks expertise in AWS Cloud computing. Which AWS service or feature will help the company with its migration?

- A. AWS Trusted Advisor
- B. AWS Consulting Partners
- C. AWS Artifacts
- D. AWS Managed Services

**Answer:** D

#### Explanation:

AWS Managed Services is a service that provides operational management for AWS infrastructure and applications. It helps users migrate their workloads to AWS and provides ongoing support, security, compliance, and automation. AWS Trusted Advisor is a service that provides best practices and recommendations for cost optimization, performance, security, and fault tolerance. AWS Consulting Partners are professional services firms that help customers design, architect, build, migrate, and manage their workloads and applications on AWS. AWS Artifacts is a service that provides on-demand access to AWS compliance reports and select online agreements.

#### NEW QUESTION 67

- (Topic 3)

Which AWS service or feature gives users the ability to capture information about network traffic in a VPC?

- A. VPC Flow Logs
- B. Amazon Inspector
- C. VPC route tables
- D. AWS CloudTrail

**Answer:** A

#### Explanation:

VPC Flow Logs is a feature that enables you to capture information about the IP traffic going to and from network interfaces in your VPC. Flow log data can be published to Amazon CloudWatch Logs, Amazon S3, or Amazon Kinesis Data Firehose. You can use VPC Flow Logs to diagnose network issues, monitor traffic patterns, detect security anomalies, and comply with auditing requirements<sup>34</sup>. References: Logging IP traffic using VPC Flow Logs - Amazon Virtual Private Cloud, New – VPC Traffic Mirroring – Capture & Inspect Network Traffic | AWS News Blog

#### NEW QUESTION 72

- (Topic 3)

A customer runs an On-Demand Amazon Linux EC2 instance for 3 hours, 5 minutes, and 6 seconds. For how much time will the customer be billed?

- A. 3 hours, 5 minutes
- B. 3 hours, 5 minutes, and 6 seconds
- C. 3 hours, 6 minutes
- D. 4 hours

**Answer:** C

#### Explanation:

Amazon EC2 usage is calculated by either the hour or the second based on the size of the instance, operating system, and the AWS Region where the instances are launched. Pricing is per instance-hour consumed for each instance, from the time an instance is launched until it's terminated or stopped. Each partial instance-hour consumed is billed per-second for Linux instances and as a full hour for all other instance types<sup>1</sup>. Therefore, the customer will be billed for 3 hours and 6 minutes for running an On-Demand Amazon Linux EC2 instance for 3 hours, 5 minutes, and 6 seconds. References: Understand Amazon EC2 instance-hours billing

#### NEW QUESTION 74

- (Topic 3)

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

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

**Answer:** D

**Explanation:**

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

**NEW QUESTION 79**

- (Topic 3)

A company is assessing its AWS Business Support plan to determine if the plan still meets the company's needs. The company is considering switching to AWS Enterprise Support.

Which additional benefit will the company receive with AWS Enterprise Support?

- A. A full set of AWS Trusted Advisor checks
- B. Phone, email, and chat access to cloud support engineers 24 hours a day, 7 days a week
- C. A designated technical account manager (TAM) to assist in monitoring and optimization
- D. A consultative review and architecture guidance for the company's applications

**Answer:** C

**Explanation:**

The additional benefit that the company will receive with AWS Enterprise Support is C. A designated technical account manager (TAM) to assist in monitoring and optimization.

A TAM is a dedicated point of contact who works with the customer to understand their use cases, applications, and goals, and provides proactive guidance and best practices to help them optimize their AWS environment. A TAM also helps the customer with case management, escalations, service updates, and feature requests<sup>12</sup>.

A full set of AWS Trusted Advisor checks is available for customers with Business, Enterprise On-Ramp, or Enterprise Support plans<sup>1</sup>. Phone, email, and chat access to cloud support engineers 24/7 is available for customers with Business, Enterprise On-Ramp, or Enterprise Support plans<sup>1</sup>. A consultative review and architecture guidance for the company's applications is available for customers with Enterprise On-Ramp or Enterprise Support plans<sup>1</sup>. Therefore, these benefits are not exclusive to AWS Enterprise Support.

Reference:

1: AWS Support Plan Comparison | Developer, Business, Enterprise ...

**NEW QUESTION 82**

- (Topic 3)

A company needs to run some of its workloads on premises to comply with regulatory guidelines. The company wants to use the AWS Cloud to run workloads that are not required to be on premises. The company also wants to be able to use the same API calls for the on-premises workloads and the cloud workloads.

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

- A. Dedicated Hosts
- B. AWS Outposts
- C. Availability Zones
- D. AWS Wavelength

**Answer:** B

**Explanation:**

AWS Outposts is a fully managed service that extends AWS infrastructure, AWS services, APIs, and tools to virtually any datacenter, co-location space, or on-premises facility for a truly consistent hybrid experience<sup>1</sup>. AWS Outposts enables customers to run workloads on premises using the same AWS APIs, tools, and services that they use in the cloud<sup>2</sup>. Dedicated Hosts are physical servers with EC2 instance capacity fully dedicated to a customer's use<sup>3</sup>. Availability Zones are one or more discrete data centers, each with redundant power, networking, and connectivity, housed in separate facilities within an AWS Region<sup>4</sup>. AWS Wavelength is an AWS Infrastructure offering optimized for mobile edge computing applications.

**NEW QUESTION 86**

- (Topic 2)

A company wants to migrate to the AWS Cloud. The company needs the ability to acquire resources when the resources are necessary.

The company also needs the ability to release those resources when the resources are no longer necessary.

Which architecture concept of the AWS Cloud meets these requirements?

- A. Elasticity
- B. Availability
- C. Reliability
- D. Durability

**Answer:** A

**Explanation:**

The architecture concept of the AWS Cloud that meets the requirements of the company that wants to migrate to the AWS Cloud and needs the ability to acquire and release resources as needed is elasticity. Elasticity means that AWS customers can quickly and easily provision and scale up or down AWS resources as their demand changes, without any upfront costs or long-term commitments. AWS provides various tools and services that enable customers to achieve elasticity, such as Amazon EC2 Auto Scaling, Amazon CloudWatch, and AWS CloudFormation. Elasticity helps customers optimize their performance, availability, and cost efficiency. Availability, reliability, and durability are other architecture concepts of the AWS Cloud, but they are not directly related to the ability to acquire and release resources as needed. Availability means that AWS customers can access their AWS resources and applications whenever and wherever they need them. Reliability means that AWS customers can depend on their AWS resources and applications to function correctly and consistently. Durability means that AWS customers can preserve their data and objects for long periods of time without loss or corruption<sup>12</sup>

**NEW QUESTION 90**

- (Topic 2)

A company wants to securely store Amazon RDS database credentials and automatically rotate user passwords periodically. Which AWS service or capability will meet these requirements?

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

**Answer:** C

**Explanation:**

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

**NEW QUESTION 92**

- (Topic 2)

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

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

**Answer:** B

**Explanation:**

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

**NEW QUESTION 94**

- (Topic 2)

A company runs a database on Amazon Aurora in the us-east-1 Region. The company has a disaster recovery requirement that the database be available in another Region.

Which solution meets this requirement with minimal disruption to the database operations?

- A. Perform an Aurora Multi-AZ deployment.
- B. Deploy Aurora cross-Region read replicas.
- C. Create Amazon Elastic Block Store (Amazon EBS) volume snapshots for Aurora and copy them to another Region.
- D. Deploy Aurora Replicas.

**Answer:** B

**Explanation:**

The solution that meets the requirement of the company that runs a database on Amazon Aurora in the us-east-1 Region and has a disaster recovery requirement that the database be available in another Region with minimal disruption to the database operations is to deploy Aurora cross-Region read replicas. Aurora cross-Region read replicas are secondary Aurora clusters that are created in a different AWS Region from the primary Aurora cluster, and are kept in sync with the primary cluster using physical replication. The company can use Aurora cross-Region read replicas to improve the availability and durability of the database, as well as to reduce the recovery time objective (RTO) and recovery point objective (RPO) in case of a regional disaster. Performing an Aurora Multi-AZ deployment, creating Amazon EBS volume snapshots for Aurora and copying them to another Region, and deploying Aurora Replicas are not the best solutions for this requirement. An Aurora Multi-AZ deployment is a configuration that creates one or more Aurora Replicas within the same AWS Region as the primary Aurora cluster, and provides automatic failover in case of an Availability Zone outage. However, this does not provide cross-Region disaster recovery. Creating Amazon EBS volume snapshots for Aurora and copying them to another Region is a manual process that requires stopping the database, creating the snapshots, copying them to the target Region, and restoring them to a new Aurora cluster. This process can cause significant downtime and data loss. Deploying Aurora Replicas is a configuration that creates one or more secondary Aurora clusters within the same AWS Region as the primary Aurora cluster, and provides read scaling and high availability. However, this does not provide cross-Region disaster recovery.

**NEW QUESTION 95**

- (Topic 2)

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

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

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

**Answer: B**

**Explanation:**

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

**NEW QUESTION 98**

- (Topic 2)

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

Which AWS feature or purchasing option will meet these requirements?

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

**Answer: B**

**Explanation:**

Consolidated billing is an AWS feature that allows users to combine the usage and costs of multiple AWS accounts into a single bill. This enables users to obtain billing discounts that are based on the company's use of AWS services, such as volume pricing tiers, Reserved Instance discounts, and Savings Plans discounts<sup>5</sup>. Resource tagging is an AWS feature that allows users to assign metadata to AWS resources, such as EC2 instances, S3 buckets, and Lambda functions. This enables users to organize, track, and manage their AWS resources, such as filtering, grouping, and reporting. Pay-as-you-go pricing is an AWS pricing model that allows users to pay only for the resources and services they use, without any upfront or long-term commitments. This enables users to lower their costs by scaling up or down as needed, and avoiding over-provisioning or under-utilization. Spot Instances are spare EC2 instances that are available at up to 90% discount compared to On-Demand prices. They are suitable for workloads that can tolerate interruptions, such as batch processing, data analysis, and testing. Spot Instances are allocated based on the current supply and demand, and can be reclaimed by AWS with a two-minute notice when the demand exceeds the supply.

**NEW QUESTION 99**

- (Topic 2)

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

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

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

**Answer: B**

**Explanation:**

AWS WAF is a web application firewall that helps protect your web applications or APIs against common web exploits that may affect availability, compromise security, or consume excessive resources. AWS WAF gives you control over how traffic reaches your applications by enabling you to create security rules that block common attack patterns, such as SQL injection or cross-site scripting, and rules that filter out specific traffic patterns you define<sup>2</sup>. You can use AWS WAF to create a custom rule that blocks SQL injection attacks on your website.

**NEW QUESTION 101**

- (Topic 2)

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

(Select TWO.)

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

**Answer: AB**

**Explanation:**

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

**NEW QUESTION 102**

- (Topic 2)

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

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

**Answer: C**

**Explanation:**

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

**NEW QUESTION 107**

- (Topic 2)

A company is setting up AWS Identity and Access Management (IAM) on an AWS account. Which recommendation complies with IAM security best practices?

- A. Use the account root user access keys for administrative tasks.
- B. Grant broad permissions so that all company employees can access the resources they need.
- C. Turn on multi-factor authentication (MFA) for added security during the login process.
- D. Avoid rotating credentials to prevent issues in production applications.

**Answer: C**

**Explanation:**

C is correct because turning on multi-factor authentication (MFA) for added security during the login process is one of the IAM security best practices recommended by AWS. MFA adds an extra layer of protection on top of the user name and password, making it harder for attackers to access the AWS account. A is incorrect because using the account root user access keys for administrative tasks is not a good practice, as the root user has full access to all the resources in the AWS account and can cause irreparable damage if compromised. AWS recommends creating individual IAM users with the least privilege principle and using roles for applications that run on Amazon EC2 instances. B is incorrect because granting broad permissions so that all company employees can access the resources they need is not a good practice, as it increases the risk of unauthorized or accidental actions on the AWS resources. AWS recommends granting only the permissions that are required to perform a task and using groups to assign permissions to IAM users. D is incorrect because avoiding rotating credentials to prevent issues in production applications is not a good practice, as it increases the risk of credential leakage or compromise. AWS recommends rotating credentials regularly and using temporary security credentials from AWS STS when possible.

**NEW QUESTION 112**

- (Topic 2)

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

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

**Answer: D**

**Explanation:**

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

**NEW QUESTION 113**

- (Topic 2)

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

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

**Answer: B**

**Explanation:**

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

**NEW QUESTION 115**

- (Topic 2)

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

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

**Answer:** B

**Explanation:**

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

**NEW QUESTION 117**

- (Topic 2)

A company needs to launch an Amazon EC2 instance.

Which of the following can the company use during the launch process to configure the root volume of the EC2 instance?

- A. Amazon EC2 Auto Scaling
- B. Amazon Data Lifecycle Manager (Amazon DLM)
- C. Amazon Machine Image (AMI)
- D. Amazon Elastic Block Store (Amazon EBS) volume

**Answer:** C

**Explanation:**

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

**NEW QUESTION 122**

- (Topic 2)

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

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

**Answer:** D

**Explanation:**

Amazon ElastiCache is a service that offers fully managed in-memory data store and cache services that deliver sub-millisecond response times to applications. You can use Amazon ElastiCache to improve the performance of your applications by retrieving data from fast, managed, in-memory data stores, instead of relying entirely on slower disk-based databases. Amazon Aurora is a relational database service that combines the performance and availability of high-end commercial databases with the simplicity and cost-effectiveness of open source databases. Amazon RDS is a service that makes it easy to set up, operate, and scale a relational database in the cloud. Amazon DynamoDB is a key-value and document database that delivers single-digit millisecond performance at any scale. None of these services are in-memory data store services.

**NEW QUESTION 124**

- (Topic 2)

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

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

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

**Answer:** D

**Explanation:**

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

**NEW QUESTION 129**

- (Topic 2)

A company has an application workload that is stateless by design and can sustain occasional downtime. The application performs massively parallel computations.

Which Amazon EC2 pricing model should the company choose for its application to reduce cost?

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

**Answer:** B

**Explanation:**

Amazon EC2 Spot Instances let you take advantage of unused EC2 capacity in the AWS cloud. Spot Instances are available at up to a 90% discount compared to On-Demand prices. You can use Spot Instances for various stateless, fault-tolerant, or flexible applications such as big data, containerized workloads, CI/CD, web servers, high-performance computing (HPC), and other test & development workloads. Spot Instances are well-suited for massively parallel computations, as they can provide large amounts of compute capacity at a low cost, and can be interrupted with a two-minute notice<sup>3</sup>

**NEW QUESTION 134**

- (Topic 2)

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

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

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

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

**Answer:** C

**Explanation:**

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

**NEW QUESTION 137**

- (Topic 2)

A company is running workloads for multiple departments within a single VPC. The company needs to be able to bill each department for its resource usage.

Which action should the company take to accomplish this goal with the LEAST operational overhead?

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

**Answer:** A

**Explanation:**

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

**NEW QUESTION 141**

- (Topic 2)

A company moves a workload to AWS to run on Amazon EC2 instances. The company needs to run the workload in the most cost-effective way.

What can the company do to meet this requirement?

- A. Use AWS Key Management Service (AWS KMS).
- B. Use multiple AWS accounts and consolidated billing.
- C. Use AWS CloudFormation to deploy the infrastructure.
- D. Rightsized all the EC2 instances that are used in the deployment.

**Answer:** D

**Explanation:**

Rightsizing all the EC2 instances that are used in the deployment is the best way to run the workload in the most cost-effective way. Rightsizing means choosing the optimal instance type and size for the workload based on the performance and capacity requirements. Rightsizing helps to avoid over-provisioning or under-provisioning of the EC2 instances, which can result in wasted resources or poor performance. Rightsizing also helps to take advantage of the different pricing models and features that AWS offers, such as On-Demand, Reserved, and Spot Instances, and Auto Scaling. For more information, see Rightsizing Your Instances and [Cost Optimization with AWS].

**NEW QUESTION 145**

- (Topic 2)

A company plans to migrate its on-premises workload to AWS. Before the migration, the company needs to estimate its future AWS service costs.

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

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

**Answer: C**

**Explanation:**

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

**NEW QUESTION 150**

- (Topic 2)

A company wants to create multiple isolated networks in the same AWS account. Which AWS service or component will provide this functionality?

- A. AWS Transit Gateway
- B. Internet gateway
- C. Amazon VPC
- D. Amazon EC2

**Answer: C**

**Explanation:**

Amazon Virtual Private Cloud (Amazon VPC) is the AWS service that allows customers to create multiple isolated networks in the same AWS account. A VPC is a logically isolated section of the AWS Cloud where customers can launch AWS resources in a virtual network that they define. Customers can create multiple VPCs within an AWS account, each with its own IP address range, subnets, route tables, security groups, network access control lists, gateways, and other components. AWS Transit Gateway, Internet gateway, and Amazon EC2 are not services or components that provide the functionality of creating multiple isolated networks in the same AWS account. AWS Transit Gateway is a service that enables customers to connect their Amazon VPCs and their on-premises networks to a single gateway. An Internet gateway is a component that enables communication between instances in a VPC and the Internet. Amazon EC2 is a service that provides scalable compute capacity in the cloud<sup>34</sup>

**NEW QUESTION 153**

- (Topic 2)

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

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

**Answer: B**

**Explanation:**

The AWS service that will meet the requirements of the company that is hosting a web application on Amazon EC2 instances and wants to implement custom conditions to filter and control inbound web traffic is AWS WAF. AWS WAF is a web application firewall that helps protect web applications from common web exploits that could affect availability, compromise security, or consume excessive resources. The company can use AWS WAF to create custom rules that block malicious requests that match certain patterns, such as SQL injection or cross-site scripting. AWS WAF can be applied to web applications that are behind an Application Load Balancer, Amazon CloudFront, or Amazon API Gateway. Amazon GuardDuty, Amazon Macie, and AWS Shield are not the best services to use for this purpose. Amazon GuardDuty is a threat detection service that monitors for malicious activity and unauthorized behavior across the AWS accounts and resources. Amazon Macie is a data security and data privacy service that uses machine learning and pattern matching to discover, classify, and protect sensitive data stored in Amazon S3. AWS Shield is a managed distributed denial of service (DDoS) protection service that safeguards web applications running on AWS. These services are more useful for detecting and preventing different types of threats and attacks, rather than filtering and controlling inbound web traffic based on custom conditions.

**NEW QUESTION 156**

- (Topic 2)

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

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

**Answer: C**

**Explanation:**

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

and access AWS resources, and accelerate their innovation and time to market.

#### NEW QUESTION 161

- (Topic 2)

Which actions are examples of a company's effort to right size its AWS resources to control cloud costs? (Select TWO.)

- A. Switch from Amazon RDS to Amazon DynamoDB to accommodate NoSQL dataset
- B. Q
- C. Base the selection of Amazon EC2 instance types on past utilization patterns.
- D. Use Amazon S3 Lifecycle policies to move objects that users access infrequently to lower-cost storage tiers.
- E. Use Multi-AZ deployments for Amazon RDS.
- F. Replace existing Amazon EC2 instances with AWS Elastic Beanstalk.

**Answer:** BC

#### Explanation:

Basing the selection of Amazon EC2 instance types on past utilization patterns is a way to right size the AWS resources and optimize the performance and cost. Using Amazon S3 Lifecycle policies to move objects that users access infrequently to lower-cost storage tiers is another way to reduce the storage costs and align them with the business value of the data. These two actions are recommended by the AWS Cost Optimization Pillar1. Switching from Amazon RDS to Amazon DynamoDB is not necessarily a cost-saving action, as it depends on the use case and the data model. Using Multi-AZ deployments for Amazon RDS is a way to improve the availability and durability of the database, but it also increases the cost. Replacing existing Amazon EC2 instances with AWS Elastic Beanstalk is a way to simplify the deployment and management of the application, but it does not affect the cost of the underlying EC2 instances.

#### NEW QUESTION 165

- (Topic 2)

Which AWS services or tools are designed to protect a workload from SQL injections, cross-site scripting, and DDoS attacks? (Select TWO.)

- A. VPC endpoint
- B. Virtual private gateway Q
- C. AWS Shield Standard
- D. AWS Config
- E. AWS WAF

**Answer:** C

#### Explanation:

AWS Shield Standard and AWS WAF are the AWS services or tools that are designed to protect a workload from SQL injections, cross-site scripting, and DDoS attacks.

According to the AWS Shield Developer Guide, "AWS Shield is a managed Distributed Denial of Service (DDoS) protection service that safeguards applications running on AWS. AWS Shield provides always-on detection and automatic inline mitigations that minimize application downtime and latency, so there is no need to engage AWS Support to benefit from DDoS protection."5 According to the AWS WAF Developer Guide, "AWS WAF is a web application firewall that helps protect your web applications or APIs against common web exploits that may affect availability, compromise security, or consume excessive resources. AWS WAF gives you control over how traffic reaches your applications by enabling you to create security rules that block common attack patterns, such as SQL injection or cross-site scripting, and rules that filter out specific traffic patterns you define." VPC endpoint, virtual private gateway, and AWS Config are not designed to protect a workload from these types of attacks.

#### NEW QUESTION 169

- (Topic 2)

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

Which AWS service will meet this requirement?

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

**Answer:** A

#### Explanation:

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

#### NEW QUESTION 173

- (Topic 2)

Which AWS service is always free of charge for users?

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

**Answer:** D

#### Explanation:

AWS Identity and Access Management (IAM) is a service that allows users to manage access to AWS resources and services. It enables users to create and

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

#### NEW QUESTION 174

- (Topic 1)

Which pillar of the AWS Well-Architected Framework focuses on the ability to run workloads effectively, gain insight into operations, and continuously improve supporting processes and procedures?

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

**Answer: C**

#### Explanation:

The AWS Well-Architected Framework is a set of best practices and guidelines for designing and operating systems in the cloud. The framework consists of five pillars: operational excellence, security, reliability, performance efficiency, and cost optimization. The operational excellence pillar focuses on the ability to run workloads effectively, gain insight into operations, and continuously improve supporting processes and procedures. Therefore, the correct answer is C. You can learn more about the AWS Well-Architected Framework and its pillars from this page.

#### NEW QUESTION 177

- (Topic 1)

A company wants to ensure that two Amazon EC2 instances are in separate data centers with minimal communication latency between the data centers. How can the company meet this requirement?

- A. Place the EC2 instances in two separate AWS Regions connected with a VPC peering connection.
- B. Place the EC2 instances in two separate Availability Zones within the same AWS Region.
- C. Place one EC2 instance on premises and the other in an AWS Region.
- D. Then connect them by using an AWS VPN connection.
- E. Place both EC2 instances in a placement group for dedicated bandwidth.

**Answer: B**

#### Explanation:

The correct answer is B because placing the EC2 instances in two separate Availability Zones within the same AWS Region is the best way to meet the requirement. Availability Zones are isolated locations within an AWS Region that have independent power, cooling, and networking. Users can launch their resources, such as Amazon EC2 instances, in multiple Availability Zones to increase the fault tolerance and resilience of their applications. Availability Zones within the same AWS Region are connected with low-latency, high-throughput, and highly redundant networking. The other options are incorrect because they are not the best ways to meet the requirement. Placing the EC2 instances in two separate AWS Regions connected with a VPC peering connection is not the best way to meet the requirement because AWS Regions are geographically dispersed and may have higher communication latency between them than Availability Zones within the same AWS Region. VPC peering connection is a networking connection between two VPCs that enables users to route traffic between them using private IP addresses. Placing one EC2 instance on premises and the other in an AWS Region, and then connecting them by using an AWS VPN connection is not the best way to meet the requirement because on-premises and AWS Region are geographically dispersed and may have higher communication latency between them than Availability Zones within the same AWS Region. AWS VPN connection is a secure and encrypted connection between a user's network and their VPC. Placing both EC2 instances in a placement group for dedicated bandwidth is not the best way to meet the requirement because a placement group is a logical grouping of instances within a single Availability Zone that enables users to launch instances with specific performance characteristics. A placement group does not ensure that the instances are in separate data centers, and it does not provide low-latency communication between instances in different Availability Zones. Reference: [Regions, Availability Zones, and Local Zones], [VPC Peering], [AWS VPN], [Placement Groups]

#### NEW QUESTION 182

- (Topic 1)

Which design principles support the reliability pillar of the AWS Well-Architected Framework? (Select TWO.)

- A. Perform operations as code.
- B. Enable traceability.
- C. Automatically scale to meet demand.
- D. Deploy resources globally to improve response time.
- E. Automatically recover from failure.

**Answer: CE**

#### Explanation:

The design principles that support the reliability pillar of the AWS Well-Architected Framework are: automatically scale to meet demand, and automatically recover from failure. These principles help users design systems that can handle changes in load, avoid disruptions, and resume normal operations quickly. Automatically scaling to meet demand means adjusting the capacity of the system based on the current and anticipated workload, using services such as AWS Auto Scaling, Amazon EC2, and AWS Lambda. Automatically recovering from failure means detecting and resolving issues, using services such as Amazon CloudWatch, AWS CloudFormation, and AWS CloudTrail.

#### NEW QUESTION 185

- (Topic 1)

A company is using AWS Lambda functions to build an application.

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

- A. Patch the servers where the Lambda functions are deployed.
- B. Establish the IAM permissions that define who can run the Lambda functions.
- C. Write the code for the Lambda functions to define the application logic.
- D. Deploy Amazon EC2 instances to support the Lambda functions.
- E. Scale out the Lambda functions when the load increases.

**Answer:** BC

**Explanation:**

According to the AWS shared responsibility model, AWS is responsible for the security of the cloud, while the user is responsible for the security in the cloud. This means that AWS manages the security and maintenance of the underlying infrastructure, such as the servers, networks, and operating systems, while the user manages the security and configuration of the resources and applications that run on AWS. For AWS Lambda functions, the tasks that are the user's responsibility are:

? Establish the IAM permissions that define who can run the Lambda functions. IAM is a service that enables users to manage access and permissions for AWS resources and users. Users can create IAM policies, roles, and users to grant or deny permissions to run Lambda functions, invoke other AWS services, or access AWS resources from Lambda functions. [AWS Lambda Permissions] AWS Certified Cloud Practitioner - aws.amazon.com

? Write the code for the Lambda functions to define the application logic. Lambda functions are units of code that can be written in any supported programming language, such as Python, Node.js, Java, or Go. Users can write the code for the Lambda functions using the AWS Management Console, the AWS Command Line Interface (AWS CLI), the AWS SDKs, or any code editor of their choice. Users can also use AWS Lambda Layers to share and manage common code and dependencies across multiple functions. [AWS Lambda Overview] AWS Certified Cloud Practitioner - aws.amazon.com

**NEW QUESTION 187**

- (Topic 1)

A company needs to test a new application that was written in Python. The code will activate when new images are stored in an Amazon S3 bucket. The application will put a watermark on each image and then will store the images in a different S3 bucket. Which AWS service should the company use to conduct the test with the LEAST amount of operational overhead?

- A. Amazon EC2
- B. AWS CodeDeploy
- C. AWS Lambda
- D. Amazon Lightsail

**Answer:** C

**Explanation:**

AWS Lambda is a compute service that lets you run code without provisioning or managing servers. AWS Lambda executes your code only when needed and scales automatically, from a few requests per day to thousands per second. You pay only for the compute time you consume - there is no charge when your code is not running. With AWS Lambda, you can run code for virtually any type of application or backend service - all with zero administration. AWS Lambda runs your code on a high-availability compute infrastructure and performs all of the administration of the compute resources, including server and operating system maintenance, capacity provisioning and automatic scaling, code monitoring and logging

**NEW QUESTION 190**

- (Topic 1)

Which task is the responsibility of AWS, according to the AWS shared responsibility model?

- A. Set up multi-factor authentication (MFA) for each Workspaces user account.
- B. Ensure the environmental safety and security of the AWS infrastructure that hosts Workspaces.
- C. Provide security for Workspaces user accounts through AWS Identity and Access Management (IAM).
- D. Configure AWS CloudTrail to log API calls and user activity. A company stores data in an Amazon S3 bucket.
- E. The company must control who has permission to read, write, or delete objects that the company stores in the S3 bucket.

**Answer:** B

**Explanation:**

The correct answer is B because ensuring the environmental safety and security of the AWS infrastructure that hosts Workspaces is the responsibility of AWS, according to the AWS shared responsibility model. The AWS shared responsibility model is a framework that defines the division of responsibilities between AWS and the customer for security and compliance. AWS is responsible for the security of the cloud, which includes the global infrastructure, such as the regions, availability zones, and edge locations; the hardware, software, networking, and facilities that run the AWS services; and the virtualization layer that separates the customer instances and storage. The customer is responsible for the security in the cloud, which includes the customer data, the guest operating systems, the applications, the identity and access management, the firewall configuration, and the encryption. The other options are incorrect because they are the responsibility of the customer, according to the AWS shared responsibility model. Setting up multi-factor authentication (MFA) for each Workspaces user account, providing security for Workspaces user accounts through AWS Identity and Access Management (IAM), configuring AWS CloudTrail to log API calls and user activity, and encrypting data at rest and in transit are all tasks that the customer has to perform to secure their Workspaces environment. Reference: AWS Shared Responsibility Model, Amazon WorkSpaces Security

**NEW QUESTION 195**

- (Topic 1)

Which of the following are user authentication services managed by AWS? (Select TWO.)

- A. Amazon Cognito
- B. AWS Lambda
- C. AWS License Manager
- D. AWS Identity and Access Management (IAM)
- E. AWS CodeStar

**Answer:** AD

**Explanation:**

The user authentication services managed by AWS are: Amazon Cognito and AWS Identity and Access Management (IAM). These services help users securely

manage and control access to their AWS resources and applications. Amazon Cognito is a service that provides user sign-up, sign-in, and access control for web and mobile applications. Amazon Cognito supports various identity providers, such as Facebook, Google, and Amazon, as well as custom user pools. AWS IAM is a service that enables users to create and manage users, groups, roles, and permissions for AWS services and resources. AWS IAM supports various authentication methods, such as passwords, access keys, and multi-factor authentication (MFA)

#### NEW QUESTION 199

- (Topic 1)

A company is designing a web application that will run on Amazon EC2 instances.

Which AWS services and features will improve availability and reduce the impact of failures for this application?

(Select TWO.)

- A. Amazon EC2 Auto Scaling for the EC2 instances
- B. VPC subnet ACLs to check the health of a service
- C. Resources that are distributed across multiple Availability Zones
- D. Configuration of AWS Server Migration Service (AWS SMS) to move the EC2 instances to a different AWS Region
- E. Resources that are distributed across multiple AWS points of presence

**Answer:** AC

#### Explanation:

The correct answers are A and C because Amazon EC2 Auto Scaling and resources that are distributed across multiple Availability Zones are AWS services and features that will improve availability and reduce the impact of failures for the web application. Amazon EC2 Auto Scaling is a service that enables users to automatically adjust the number of Amazon EC2 instances in response to changes in demand or performance. Amazon EC2 Auto Scaling helps users to maintain optimal availability and performance of their applications by adding or removing instances as needed. Resources that are distributed across multiple Availability Zones are AWS features that enable users to increase the fault tolerance and resilience of their applications. Availability Zones are isolated locations within an AWS Region that have independent power, cooling, and networking. Users can launch their resources, such as Amazon EC2 instances, in multiple Availability Zones to protect their applications from the failure of a single location. The other options are incorrect because they are not AWS services and features that will improve availability and reduce the impact of failures for the web application. VPC subnet ACLs are AWS features that enable users to control the inbound and outbound traffic to and from their subnets within a VPC. VPC subnet ACLs do not check the health of a service, but rather filter the network traffic based on rules. Configuration of AWS Server Migration Service (AWS SMS) is an AWS service that enables users to migrate their on-premises servers to AWS. Configuration of AWS SMS does not help to move the Amazon EC2 instances to a different AWS Region, but rather to migrate the servers from the source environment to AWS. Resources that are distributed across multiple AWS points of presence are AWS features that enable users to deliver content to their end users with low latency and high performance. AWS points of presence are edge locations that are part of the AWS Global Infrastructure. Users can use services such as Amazon CloudFront and AWS Global Accelerator to distribute their content across multiple AWS points of presence. Reference: Amazon EC2 Auto Scaling, [Regions, Availability Zones, and Local Zones]

#### NEW QUESTION 202

- (Topic 1)

A developer needs to build an application for a retail company. The application must provide real-time product recommendations that are based on machine learning.

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

- A. AWS Health Dashboard
- B. Amazon Personalize
- C. Amazon Forecast
- D. Amazon Transcribe

**Answer:** B

#### Explanation:

Amazon Personalize is a fully managed machine learning service that customers can use to generate personalized recommendations for their users. It can also generate user segments based on the users' affinity for certain items or item metadata. Amazon Personalize uses the customers' data to train and deploy custom recommendation models that can be integrated into their applications. Therefore, the correct answer is B. You can learn more about Amazon Personalize and its use cases from this page.

#### NEW QUESTION 206

- (Topic 3)

A company wants to build a new web application by using AWS services. The application must meet the on-demand load for periods of heavy activity.

Which AWS services or resources provide the necessary workload adjustments to meet these requirements? (Select TWO.)

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

**Answer:** BD

#### Explanation:

Amazon EC2 Auto Scaling helps you ensure that you have the correct number of Amazon EC2 instances available to handle the load for your application. You create collections of EC2 instances, called Auto Scaling groups. You can specify the minimum number of instances in each Auto Scaling group, and Amazon EC2 Auto Scaling ensures that your group never goes below this size. You can specify the maximum number of instances in each Auto Scaling group, and Amazon EC2 Auto Scaling ensures that your group never goes above this size. AWS Lambda lets you run code without provisioning or managing servers. You pay only for the compute time you consume. With Lambda, you can run code for virtually any type of application or backend service - all with zero administration. Just upload your code and Lambda takes care of everything required to run and scale your code with high availability. You can set up your code to automatically trigger from other AWS services or call it directly from any web or mobile app.

#### NEW QUESTION 211

- (Topic 3)

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

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

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

**Answer: B**

**Explanation:**

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

References: What is Amazon EC2?, AWS Systems Manager Patch Manager

**NEW QUESTION 215**

- (Topic 3)

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

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

**Answer: B**

**Explanation:**

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

**NEW QUESTION 217**

- (Topic 3)

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

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

**Answer: B**

**Explanation:**

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

**NEW QUESTION 220**

- (Topic 3)

Which AWS Support plan is the minimum recommended tier for users who have production workloads on AWS?

- A. AWS Developer Support
- B. AWS Enterprise Support
- C. AWS Business Support
- D. AWS Enterprise On-Ramp Support

**Answer: C**

**Explanation:**

AWS Business Support is the minimum recommended tier for users who have production workloads on AWS. AWS Business Support provides 24x7 access to cloud support engineers via phone, chat, or email, as well as a guaranteed response time of less than one hour for urgent issues. AWS Business Support also includes access to AWS Trusted Advisor, a tool that provides real-time guidance to help you provision your resources following AWS best practices.

**NEW QUESTION 221**

- (Topic 3)

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

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

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

**Answer:** C

**Explanation:**

Amazon Lightsail is an easy-to-use cloud platform that offers everything you need to build an application or website, plus a cost-effective, monthly plan<sup>1</sup>. It is designed for developers who have little or no prior cloud experience and want to launch and manage applications on AWS with minimal complexity<sup>2</sup>. Amazon SageMaker is a service for building, training, and deploying machine learning models<sup>3</sup>. AWS Lambda is a service that lets you run code without provisioning or managing servers<sup>4</sup>. Amazon Elastic Container Service (Amazon ECS) is a fully managed container orchestration service.

**NEW QUESTION 226**

- (Topic 3)

A company needs to block SQL injection attacks.  
Which AWS service or feature can meet this requirement?

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

**Answer:** A

**Explanation:**

AWS WAF is a web application firewall that helps protect web applications from common web exploits, such as SQL injection attacks. It allows customers to create custom rules that block malicious requests. AWS Shield is a managed service that protects against distributed denial of service (DDoS) attacks, not SQL injection attacks. Network ACLs and security groups are network-level security features that filter traffic based on IP addresses and ports, not web requests or SQL queries. References: [AWS WAF], [AWS Shield], [Network ACLs], [Security groups]

**NEW QUESTION 231**

- (Topic 3)

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

- A. Physical connectivity among Availability Zones
- B. Network switch maintenance
- C. Hardware updates and firmware patches
- D. Amazon EC2 updates and security patches

**Answer:** D

**Explanation:**

According to the AWS shared responsibility model, customers are responsible for managing their data, applications, operating systems, security groups, and other aspects of their AWS environment. This includes installing updates and security patches of the guest operating system and any application software or utilities installed by the customer on the instances. 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. This includes the physical connectivity among Availability Zones, the network switch maintenance, and the hardware updates and firmware patches. Therefore, option D is the correct answer, and options A, B, and C are AWS responsibilities, not customer responsibilities. References: : AWS Well-Architected Framework - Elasticity; : Reactive Systems on AWS - Elastic

**NEW QUESTION 235**

- (Topic 3)

Which AWS service is a continuous delivery and deployment solution?

- A. AWSAppSync
- B. AWS CodePipeline
- C. AWS Cloud9
- D. AWS CodeCommit

**Answer:** B

**Explanation:**

AWS CodePipeline is a continuous delivery and deployment service that automates the release process of software applications across different stages, such as source code, build, test, and deploy<sup>2</sup>. AWSAppSync, AWS Cloud9, and AWS CodeCommit are other AWS services related to application development, but they do not provide continuous delivery and deployment solutions<sup>34</sup>.

**NEW QUESTION 239**

- (Topic 3)

Which AWS services make use of global edge locations'? (Select TWO.)

- A. AWS Fargate
- B. Amazon CloudFront
- C. AWS Global Accelerator
- D. AWS Wavelength
- E. Amazon VPC

**Answer:** BC

**Explanation:**

Amazon CloudFront and AWS Global Accelerator are two AWS services that make use of global edge locations. Edge locations are AWS sites that are deployed worldwide in major cities and places with a high population. Edge locations are used to cache data and reduce latency for end-user access<sup>1</sup>. Amazon CloudFront is a content delivery network (CDN) service that securely delivers data, videos, applications, and APIs to customers globally with low latency and high transfer speeds. Amazon CloudFront uses a global network of over 200 edge locations and 13 regional edge caches to cache your content closer to your viewers, improving performance and reducing costs<sup>23</sup>. AWS Global Accelerator is a networking service that improves the availability and performance of your applications with local or global users. AWS Global Accelerator uses the AWS global network to route user traffic to the optimal endpoint based on health, performance, and policies. AWS Global Accelerator uses over 100 edge locations to bring your application endpoints closer to your users, reducing network hops and improving user experience<sup>45</sup>. References: 1: AWS for the Edge - Amazon Web Services (AWS), 2: Content Delivery Network (CDN) - Amazon CloudFront - AWS, 3: Amazon CloudFront Documentation, 4: AWS Global Accelerator - Amazon Web Services, 5: AWS Global Accelerator Documentation

#### NEW QUESTION 241

- (Topic 3)

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

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

**Answer:** CD

#### Explanation:

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

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

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

#### NEW QUESTION 243

- (Topic 3)

Which service enables customers to audit API calls in their AWS accounts'?

- A. AWS CloudTrail
- B. AWS Trusted Advisor
- C. Amazon Inspector
- D. AWS X-Ray

**Answer:** A

#### Explanation:

AWS CloudTrail is a service that provides a record of actions taken by a user, role, or an AWS service in your AWS account. CloudTrail captures all API calls for AWS services as events, including calls from the AWS Management Console, AWS SDKs, command line tools, and higher-level AWS services. You can use CloudTrail to monitor, audit, and troubleshoot your AWS account activity<sup>34</sup>. AWS Trusted Advisor is a service that provides best practices recommendations for cost optimization, performance, security, and fault tolerance in your AWS account<sup>5</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<sup>6</sup>. AWS X-Ray is a service that helps you analyze and debug your applications by collecting data about the requests that your application serves, and providing tools to view, filter, and gain insights into that data<sup>7</sup>. References: Logging AWS Audit Manager API calls with CloudTrail, Logging AWS Account Management API calls using AWS CloudTrail, Review API calls in your AWS account using CloudTrail, Monitor the usage of AWS API calls using Amazon CloudWatch, Which service enables customers to audit API calls in their AWS ...

#### NEW QUESTION 244

- (Topic 3)

A company wants to define a central data protection policy that works across AWS services for compute, storage, and database resources. Which AWS service will meet this requirement?

- A. AWS Batch
- B. AWS Elastic Disaster Recovery
- C. AWS Backup
- D. Amazon FSx

**Answer:** C

#### Explanation:

The AWS service that will meet this requirement is C. AWS Backup.

AWS Backup is a service that allows you to define a central data protection policy that works across AWS services for compute, storage, and database resources. You can use AWS Backup to create backup plans that specify the frequency, retention, and lifecycle of your backups, and apply them to your AWS resources using tags or resource IDs. AWS Backup supports various AWS services, such as Amazon EC2, Amazon EBS, Amazon RDS, Amazon DynamoDB, Amazon EFS, Amazon FSx, and AWS Storage Gateway<sup>12</sup>. AWS Batch is a service that allows you to run batch computing workloads on AWS. AWS Batch does not provide a central data protection policy, but rather enables you to optimize the allocation and utilization of your compute resources<sup>3</sup>.

AWS Elastic Disaster Recovery is a service that allows you to prepare for and recover from disasters using AWS. AWS Elastic Disaster Recovery does not provide a central data protection policy, but rather helps you minimize downtime and data loss by replicating your applications and data to AWS<sup>4</sup>.

Amazon FSx is a service that provides fully managed file storage for Windows and Linux applications. Amazon FSx does not provide a central data protection policy, but rather offers features such as encryption, snapshots, backups, and replication to protect your file systems<sup>5</sup>.

References:

1: AWS Backup – Centralized backup across AWS services 3: AWS Batch – Run Batch Computing Jobs on AWS 2: Data Protection Reference Architectures with AWS Backup 4: AWS Elastic Disaster Recovery – Prepare for and recover from disasters using AWS 5: Amazon FSx – Fully managed file storage for Windows and Linux applications

#### NEW QUESTION 245

- (Topic 3)

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

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

**Answer:** A

#### Explanation:

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

#### NEW QUESTION 247

- (Topic 3)

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

Which AWS service will meet these requirements?

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

**Answer:** A

#### Explanation:

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

#### NEW QUESTION 249

- (Topic 3)

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

Which AWS Cloud benefit does this scenario describe?

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

**Answer:** B

#### Explanation:

The trade of infrastructure expenses for operating expenses is one of the benefits of the AWS Cloud. By moving to the AWS Cloud, the company can avoid the upfront costs of purchasing and maintaining on-premises infrastructure, such as servers, storage, network, and software. Instead, the company can pay only for the AWS resources and services that they use, as they use them. This reduces the risk and complexity of planning and managing IT infrastructure, and allows the company to focus on innovation and growth. Increased speed to market, massive economies of scale, and the ability to go global in minutes are also benefits of the AWS Cloud, but they are not the best ones to describe this scenario. Increased speed to market means that the company can launch new products and services faster by using AWS services and tools. Massive economies of scale means that the company can benefit from the lower costs and higher performance that AWS achieves by operating at a large scale. The ability to go global in minutes means that the company can deploy their applications and data in multiple regions and availability zones around the world to reach their customers faster and improve performance and reliability<sup>5</sup>

#### NEW QUESTION 253

- (Topic 3)

A company is building an application in the AWS Cloud. The company wants to use temporary credentials for the application to access other AWS resources.

Which AWS service will meet these requirements?

- A. AWS Key Management Service (Aws KMS)
- B. AWS CloudHSM
- C. Amazon Cognito
- D. AWS Security Token Service (Aws STS)

**Answer:** D

#### Explanation:

AWS Security Token Service (AWS STS) is a service that provides temporary security credentials to users or applications that need to access AWS resources. The temporary credentials have a limited lifetime and can be configured to last from a few minutes to several hours. The credentials are not stored with the user or

application, but are generated dynamically and provided on request. The credentials work almost identically to long-term access key credentials, but have the advantage of not requiring distribution, rotation, or revocation<sup>1</sup>.

AWS Key Management Service (AWS KMS) is a service that provides encryption and decryption services for data and keys. It does not provide temporary security credentials<sup>2</sup>. AWS CloudHSM is a service that provides hardware security modules (HSMs) for cryptographic operations and key management. It does not provide temporary security credentials<sup>3</sup>.

Amazon Cognito is a service that provides user authentication and authorization for web and mobile applications. It can also provide temporary security credentials for authenticated users, but not for applications<sup>4</sup>.

#### NEW QUESTION 256

- (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 257

- (Topic 3)

A company wants its Amazon EC2 instances to share the same geographic area but use multiple independent underlying power sources. Which solution achieves this goal?

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

**Answer:** C

#### Explanation:

The solution that achieves the goal of having Amazon EC2 instances share the same geographic area but use multiple independent underlying power sources is to use EC2 instances in multiple Availability Zones in the same AWS Region. An Availability Zone is a physically isolated location within an AWS Region that has its own power, cooling, and network connectivity. An AWS Region is a geographical area that consists of two or more Availability Zones. By using multiple Availability Zones, users can increase the fault tolerance and resilience of their applications, as well as reduce latency for end users<sup>3</sup>. Using EC2 instances in a single Availability Zone, multiple AWS Regions, or the same edge location and the same AWS Region would not meet the requirement of having multiple independent power sources.

#### NEW QUESTION 261

- (Topic 3)

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

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

**Answer:** AC

#### Explanation:

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

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

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

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

References:

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

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

#### NEW QUESTION 263

- (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 268**

- (Topic 3)

Which AWS service can a company use to visually design and build serverless applications?

- A. AWS Lambda
- B. AWS Batch
- C. AWS Application Composer
- D. AWS App Runner

**Answer:** C

**Explanation:**

AWS Application Composer is a service that allows users to visually design and build serverless applications. Users can drag and drop components, such as AWS Lambda functions, Amazon API Gateway endpoints, Amazon DynamoDB tables, and Amazon S3 buckets, to create a serverless application architecture. Users can also configure the properties, permissions, and dependencies of each component, and deploy the application to their AWS account with a few clicks. AWS Application Composer simplifies the design and configuration of serverless applications, and reduces the need to write code or use AWS CloudFormation templates. References: AWS Application Composer, AWS releases Application Composer to make serverless 'easier' but initial scope is limited

**NEW QUESTION 269**

- (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 270**

- (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 273**

- (Topic 3)

A company is moving to the AWS Cloud to reduce operational overhead for its application infrastructure.

Which IT operation will the company still be responsible for after the migration to AWS?

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

**Answer:** D

**Explanation:**

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

**NEW QUESTION 274**

- (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 276**

- (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 279**

.....

## THANKS FOR TRYING THE DEMO OF OUR PRODUCT

Visit Our Site to Purchase the Full Set of Actual AWS-Certified-Cloud-Practitioner Exam Questions With Answers.

We Also Provide Practice Exam Software That Simulates Real Exam Environment And Has Many Self-Assessment Features. Order the AWS-Certified-Cloud-Practitioner Product From:

<https://www.2passeasy.com/dumps/AWS-Certified-Cloud-Practitioner/>

## Money Back Guarantee

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

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