

# Amazon

## Exam Questions AWS-Certified-Cloud-Practitioner

Amazon AWS Certified Cloud Practitioner



**NEW QUESTION 1**

- (Topic 2)

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

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

**Answer: C**

**Explanation:**

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

**NEW QUESTION 2**

- (Topic 2)

A company wants an in-memory data store that is compatible with open source in the cloud. Which AWS service should the company use?

- A. Amazon DynamoDB
- B. Amazon ElastiCache
- C. Amazon Elastic Block Store (Amazon EBS)
- D. Amazon Redshift

**Answer: B**

**Explanation:**

Amazon ElastiCache is a fully managed in-memory data store service that is compatible with open source engines such as Redis and Memcached<sup>1</sup>. It provides fast and scalable performance for applications that require high throughput and low latency<sup>1</sup>. Amazon DynamoDB is a fully managed NoSQL database service that provides consistent and single-digit millisecond latency at any scale<sup>2</sup>. Amazon EBS is a block storage service that provides persistent and durable storage volumes for Amazon EC2 instances<sup>3</sup>. Amazon Redshift is a fully managed data warehouse service that allows users to run complex analytic queries using SQL<sup>4</sup>.

**NEW QUESTION 3**

- (Topic 2)

A company wants to use Amazon EC2 instances for a stable production workload that will run for 1 year. Which instance purchasing option meets these requirements MOST cost-effectively?

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

**Answer: B**

**Explanation:**

B is correct because Reserved Instances are the instance purchasing option that offers the most cost-effective way to use Amazon EC2 instances for a stable production workload that will run for 1 year, as they provide significant discounts compared to On-Demand Instances in exchange for a commitment to use a specific amount of computing power for a period of time. A is incorrect because Dedicated Hosts are the instance purchasing option that allows customers to use physical servers that are fully dedicated to their use, which is more expensive and less flexible than Reserved Instances. C is incorrect because On-Demand Instances are the instance purchasing option that allows customers to pay for compute capacity by the hour or second with no long-term commitments, which is more suitable for short-term, variable, and unpredictable workloads. D is incorrect because Spot Instances are the instance purchasing option that allows customers to bid on spare Amazon EC2 computing capacity, which is more suitable for flexible, scalable, and fault-tolerant workloads that can tolerate interruptions.

**NEW QUESTION 4**

- (Topic 2)

Which AWS Cloud design principle does a company follow by using AWS CloudTrail?

- A. Recover automatically.
- B. Perform operations as code.
- C. Measure efficiency.
- D. Ensure traceability.

**Answer: D**

**Explanation:**

The company follows the AWS Cloud design principle of ensuring traceability by using AWS CloudTrail. AWS CloudTrail is a service that records the API calls and events made by or on behalf of the AWS account. The company can use AWS CloudTrail to monitor, audit, and analyze the activity and changes in their AWS resources and applications. AWS CloudTrail helps the company to achieve compliance, security, governance, and operational efficiency. Recovering automatically, performing operations as code, and measuring efficiency are other AWS Cloud design principles, but they are not directly related to using AWS CloudTrail. Recovering automatically means that the company can design their cloud workloads to handle failures gracefully and resume normal operations without manual intervention. Performing operations as code means that the company can automate the creation, configuration, and management of their cloud resources using scripts or templates. Measuring efficiency means that the company can monitor and optimize the performance and utilization of their cloud resources and

applications34

#### NEW QUESTION 5

- (Topic 1)

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

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

**Answer: C**

#### Explanation:

AWS Shield is a managed Distributed Denial of Service (DDoS) protection service that safeguards applications running on AWS. AWS Shield provides always-on detection and automatic inline mitigations that minimize application downtime and latency, so there is no need to engage AWS Support to benefit from DDoS protection.

#### NEW QUESTION 6

- (Topic 1)

Which AWS features will meet these requirements? (Select TWO.)

- A. Security groups
- B. Network ACLs
- C. S3 bucket policies
- D. IAM user policies
- E. S3 bucket versioning

**Answer: CD**

#### Explanation:

The correct answers are C and D because S3 bucket policies and IAM user policies are AWS features that will meet the requirements. S3 bucket policies are access policies that can be attached to Amazon S3 buckets to grant or deny permissions to the bucket and the objects it contains. S3 bucket policies can be used to control who has permission to read, write, or delete objects that the company stores in the S3 bucket. IAM user policies are access policies that can be attached to IAM users to grant or deny permissions to AWS resources and actions. IAM user policies can be used to control who has permission to read, write, or delete objects that the company stores in the S3 bucket. The other options are incorrect because they are not AWS features that will meet the requirements. Security groups and network ACLs are AWS features that act as firewalls to control inbound and outbound traffic to and from Amazon EC2 instances and subnets. Security groups and network ACLs do not control who has permission to read, write, or delete objects that the company stores in the S3 bucket. S3 bucket versioning is an AWS feature that enables users to keep multiple versions of the same object in the same bucket. S3 bucket versioning can be used to recover from accidental overwrites or deletions of objects, but it does not control who has permission to read, write, or delete objects that the company stores in the S3 bucket. Reference: Using Bucket Policies and User Policies, Security Groups for Your VPC, Network ACLs, [Using Versioning]

#### NEW QUESTION 7

- (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 8

- (Topic 1)

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

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

**Answer: D**

#### Explanation:

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

#### NEW QUESTION 9

- (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 10

- (Topic 1)

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

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

**Answer:** AD

#### Explanation:

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

#### NEW QUESTION 10

- (Topic 1)

A company needs to identify the last time that a specific user accessed the AWS Management Console.

Which AWS service will provide this information?

- A. Amazon Cognito
- B. AWS CloudTrail
- C. Amazon Inspector
- D. Amazon GuardDuty

**Answer:** B

#### Explanation:

AWS CloudTrail is the service that will provide the information about the last time that a specific user accessed the AWS Management Console. AWS CloudTrail is a service that records the API calls and events made by or on behalf of your AWS account. You can use AWS CloudTrail to view, search, and download the history of AWS console sign-in events, which include the user name, date, time, source IP address, and other details of the sign-in activity. Amazon Cognito, Amazon Inspector, and Amazon GuardDuty are not services that will provide this information. Amazon Cognito is a service that provides user authentication and authorization for web and mobile applications. Amazon Inspector is a service that assesses the security and compliance of your applications running on AWS. Amazon GuardDuty is a service that monitors your AWS account and workloads for malicious or unauthorized activity.

#### NEW QUESTION 11

- (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.

**NEW QUESTION 12**

- (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 13**

- (Topic 1)

A company has two AWS accounts in an organization in AWS Organizations for consolidated billing. All of the company's AWS resources are hosted in one AWS Region.

Account A has purchased five Amazon EC2 Standard Reserved Instances (RIs) and has four EC2 instances running. Account B has not purchased any RIs and also has four EC2 instances running. Which statement is true regarding pricing for these eight instances?

- A. The eight instances will be charged as regular instances.
- B. Four instances will be charged as RIs, and four will be charged as regular instances.
- C. Five instances will be charged as RIs, and three will be charged as regular instances.
- D. The eight instances will be charged as RIs.

**Answer: B**

**Explanation:**

The statement that is true regarding pricing for these eight instances is: four instances will be charged as RIs, and four will be charged as regular instances. Amazon EC2 Reserved Instances (RIs) are a pricing model that allows users to reserve EC2 instances for a specific term and benefit from discounted hourly rates and capacity reservation. RIs are purchased for a specific AWS Region, and can be shared across multiple accounts in an organization in AWS Organizations for consolidated billing. However, RIs are applied on a first-come, first-served basis, and there is no guarantee that all instances in the organization will be charged at the RI rate. In this case, Account A has purchased five RIs and has four instances running, so all four instances will be charged at the RI rate. Account B has not purchased any RIs and also has four instances running, so all four instances will be charged at the regular rate. The remaining RI in Account A will not be applied to any instance in Account B, and will be wasted.

**NEW QUESTION 14**

- (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 17**

- (Topic 1)

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

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

**Answer: C**

**Explanation:**

The correct answer is C because Amazon QuickSight is an AWS service that will provide the dashboards and charts for the insights from business data. Amazon QuickSight is a fully managed, scalable, and serverless business intelligence service that enables users to create and share interactive dashboards and charts. Amazon QuickSight can connect to various data sources, such as Amazon S3, Amazon RDS, Amazon Redshift, and more. Amazon QuickSight also provides users with machine learning insights, such as anomaly detection, forecasting, and natural language narratives. The other options are incorrect because they are



not AWS services that will provide the dashboards and charts for the insights from business data. Amazon Macie is an AWS service that helps users discover, classify, and protect sensitive data stored in Amazon S3. Amazon Aurora is an AWS service that provides a relational database that is compatible with MySQL and PostgreSQL. AWS CloudTrail is an AWS service that enables users to track user activity and API usage across their AWS account. Reference: Amazon QuickSight FAQs

**NEW QUESTION 21**

- (Topic 1)

An ecommerce company has migrated its IT infrastructure from an on-premises data center to the AWS Cloud. Which AWS service is used to track, record, and audit configuration changes made to AWS resources?

- A. AWS Shield
- B. AWS Config
- C. AWS IAM
- D. Amazon Inspector

**Answer: B**

**Explanation:**

AWS Config is a service that enables you to assess, audit, and evaluate the configurations of your AWS resources. AWS Config continuously monitors and records your AWS resource configurations and allows you to automate the evaluation of recorded configurations against desired configurations. With AWS Config, you can review changes in configurations and relationships between AWS resources, dive into detailed resource configuration histories, and determine your overall compliance against the configurations specified in your internal guidelines<sup>3</sup>.

**NEW QUESTION 23**

- (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 28**

- (Topic 1)

A company plans to migrate to AWS and wants to create cost estimates for its AWS use cases. Which AWS service or tool can the company use to meet these requirements?

- A. AWS Pricing Calculator
- B. Amazon CloudWatch
- C. AWS Cost Explorer
- D. AWS Budgets

**Answer: A**

**Explanation:**

AWS Pricing Calculator is a web-based planning tool that customers can use to create estimates for their AWS use cases. They can use it to model their solutions before building them, explore the AWS service price points, and review the calculations behind their estimates. Therefore, the correct answer is A. You can learn more about AWS Pricing Calculator and how it works from this page.

**NEW QUESTION 29**

- (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 34**

- (Topic 1)

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

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

**Answer: B**

**Explanation:**

Pay-as-you-go pricing is the feature of the AWS Cloud that gives users the ability to pay based on current needs rather than forecasted needs. Pay-as-you-go pricing means that users only pay for the AWS services and resources they use, without any upfront or long-term commitments. This allows users to scale up or down their usage depending on their changing business requirements, and avoid paying for idle or unused capacity. Pay-as-you-go pricing also enables users to benefit from the economies of scale and lower costs of AWS as they grow their business<sup>5</sup>

**NEW QUESTION 39**

- (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 42**

- (Topic 1)

Which AWS service or tool can be used to consolidate payments for a company with multiple AWS accounts?

- A. AWS Cost and Usage Report
- B. AWS Organizations
- C. Cost Explorer
- D. AWS Budgets

**Answer: B**

**Explanation:**

AWS Organizations is an account management service that enables you to consolidate multiple AWS accounts into an organization that you create and centrally manage. AWS Organizations includes consolidated billing and account management capabilities that enable you to better meet the budgetary, security, and compliance needs of your business<sup>1</sup>.

**NEW QUESTION 44**

- (Topic 1)

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

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

**Answer: D**

**Explanation:**

An Amazon Machine Image (AMI) is a deployable Amazon EC2 instance template that is prepackaged with software and security requirements. It provides the information required to launch an instance, which is a virtual server in the cloud. You can use an AMI to launch as many instances as you need. You can also create your own custom AMIs or use AMIs shared by other AWS users<sup>1</sup>.

**NEW QUESTION 48**

- (Topic 1)

Which AWS service or feature captures information about the network traffic to and from an Amazon EC2 instance?

- A. VPC Reachability Analyzer
- B. Amazon Athena
- C. VPC Flow Logs
- D. AWS X-Ray

**Answer: C**

**Explanation:**

The correct answer is C because VPC Flow Logs is an AWS service or feature that captures information about the network traffic to and from an Amazon EC2 instance. VPC Flow Logs is a feature that enables customers to capture information about the IP traffic going to and from network interfaces in their VPC. VPC Flow Logs can help customers to monitor and troubleshoot connectivity issues, such as traffic not reaching an instance or traffic being rejected by a security group. The other options are incorrect because they are not AWS services or features that capture information about the network traffic to and from an Amazon EC2 instance. VPC Reachability Analyzer is an AWS service or feature that enables customers to perform connectivity testing between resources in their VPC and identify configuration issues that prevent connectivity. Amazon Athena is an AWS service that enables customers to query data stored in Amazon S3 using standard SQL. AWS X-Ray is an AWS service that enables customers to analyze and debug distributed applications, such as those built using a microservices architecture.

Reference: VPC Flow Logs

#### NEW QUESTION 49

- (Topic 1)

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

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

**Answer:** AD

#### Explanation:

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

#### NEW QUESTION 54

- (Topic 1)

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

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

**Answer:** C

#### Explanation:

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

#### NEW QUESTION 58

- (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 63

- (Topic 1)

Which AWS network services or features allow CIDR block notation when providing an IP address range? (Select TWO.)

- A. Security groups
- B. Amazon Machine Image (AMI)
- C. Network access control list (network ACL)
- D. AWS Budgets
- E. Amazon Elastic Block Store (Amazon EBS)

**Answer:** AC

#### Explanation:



Security groups and network access control lists (network ACLs) are two AWS network services or features that allow CIDR block notation when providing an IP address range. Security groups act as a firewall for associated Amazon EC2 instances, controlling both inbound and outbound traffic at the instance level. Network ACLs act as a firewall for associated subnets, controlling both inbound and outbound traffic at the subnet level. Both security groups and network ACLs use CIDR block notation to specify the IP address ranges that are allowed or denied

**NEW QUESTION 67**

- (Topic 1)

Which statements represent the cost-effectiveness of the AWS Cloud? (Select TWO.)

- A. Users can trade fixed expenses for variable expenses.
- B. Users can deploy all over the world in minutes.
- C. AWS offers increased speed and agility.
- D. AWS is responsible for patching the infrastructure.
- E. Users benefit from economies of scale.

**Answer:** AE

**Explanation:**

The statements that represent the cost-effectiveness of the AWS Cloud are:

? Users can trade fixed expenses for variable expenses. By using the AWS Cloud, users can pay only for the resources they use, instead of investing in fixed and upfront costs for hardware and software. This can lower the total cost of ownership and increase the return on investment.

? Users benefit from economies of scale. By using the AWS Cloud, users can leverage the massive scale and efficiency of AWS to access lower prices and higher performance. AWS passes the cost savings to the users through price reductions and innovations. AWS Cloud Value Framework

**NEW QUESTION 69**

- (Topic 3)

A company needs a graph database service that is scalable and highly available.

Which AWS service meets these requirements?

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

**Answer:** D

**Explanation:**

The AWS service that meets the requirements of providing a graph database service that is scalable and highly available is Amazon Neptune. Amazon Neptune is a fast, reliable, and fully managed graph database service that supports property graph and RDF graph models. Amazon Neptune is designed to store billions of relationships and query the graph with milliseconds latency. Amazon Neptune also offers high availability and durability by replicating six copies of the data across three Availability Zones and continuously backing up the data to Amazon S3. Amazon Aurora, Amazon Redshift, and Amazon DynamoDB are other AWS services that provide relational or non- relational database solutions, but they do not support graph database models.

**NEW QUESTION 71**

- (Topic 3)

Which VPC component provides a layer of security at the subnet level?

- A. Security groups
- B. Network ACLs
- C. NAT gateways
- D. Route tables

**Answer:** B

**Explanation:**

Network ACLs are a feature that provide a layer of security at the subnet level by acting as a firewall to control traffic in and out of one or more subnets. Network ACLs can be configured with rules that allow or deny traffic based on the source and destination IP addresses, ports, and protocols. Security groups are a feature that provide a layer of security at the instance level by acting as a firewall to control traffic to and from one or more instances. Security groups can be configured with rules that allow or deny traffic based on the source and destination IP addresses, ports, protocols, and security groups. NAT gateways are a feature that enable instances in a private subnet to connect to the internet or other AWS services, but prevent the internet from initiating a connection with those instances. Route tables are a feature that determine where network traffic from a subnet or gateway is directed.

**NEW QUESTION 76**

- (Topic 3)

A company wants to migrate a database from an on-premises environment to Amazon RDS.

After the migration is complete, which management task will the company still be responsible for?

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

**Answer:** B

**Explanation:**

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

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

**NEW QUESTION 81**

- (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 83**

- (Topic 3)

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

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

**Answer:** B

**Explanation:**

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

**NEW QUESTION 87**

- (Topic 3)

Which AWS Cloud benefit describes the ability to acquire resources as they are needed and release resources when they are no longer needed?

- A. Economies of scale
- B. Elasticity
- C. Agility
- D. Security

**Answer:** B

**Explanation:**

The AWS Cloud benefit that describes the ability to acquire resources as they are needed and release resources when they are no longer needed is elasticity. Elasticity means that users can quickly add and remove resources to match the demand of their applications, and only pay for what they use. Elasticity enables users to handle unpredictable workloads, reduce costs, and improve performance<sup>1</sup>. Economies of scale, agility, and security are other benefits of the AWS Cloud, but they do not describe the specific ability of acquiring and releasing resources on demand.

**NEW QUESTION 88**

- (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 92**

- (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 95**

- (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 96**

- (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 98**

- (Topic 2)

A company is running an application on AWS. The company wants to identify and prevent the accidental

Which AWS service or feature will meet these requirements?

- A. Amazon GuardDuty
- B. Network ACL
- C. AWS WAF
- D. AWS Network Firewall

**Answer:** A

**Explanation:**

Amazon GuardDuty is a threat detection service that continuously monitors for malicious activity and unauthorized behavior to protect your AWS accounts,

workloads, and data stored in Amazon S3. With the cloud, the collection and aggregation of account and network activities is simplified, but it can be time consuming for security teams to continuously analyze event log data for potential threats. With GuardDuty, you can automate anomaly detection and get actionable findings to help you protect your AWS resources<sup>4</sup>.

**NEW QUESTION 101**

- (Topic 2)

What is an AWS responsibility under the AWS shared responsibility model?

- A. Configure the security group rules that determine which ports are open on an Amazon EC2 Linux instance.
  - B. Ensure the security of the internal network in the AWS data centers.
  - C. Patch the guest operating system with the latest security patches on Amazon EC2.
  - D. Turn on server-side encryption for Amazon S3 buckets.
- A company wants to deploy its critical application on AWS and maintain high availability.

**Answer:** B

**Explanation:**

Under the AWS shared responsibility model, AWS is responsible for ensuring the security of the internal network in the AWS data centers, as well as the physical security of the hardware and facilities that run AWS services. AWS customers are responsible for configuring the security group rules that determine which ports are open on an EC2 Linux instance, patching the guest operating system with the latest security patches on EC2, and turning on server-side encryption for S3 buckets. Source: AWS Shared Responsibility Model

**NEW QUESTION 105**

- (Topic 2)

A company wants its workload to perform consistently and correctly. Which benefit of AWS Cloud computing does this goal represent?

- A. Security
- B. Elasticity
- C. Pay-as-you-go pricing
- D. Reliability

**Answer:** D

**Explanation:**

Reliability is the benefit of AWS Cloud computing that ensures the workload performs consistently and correctly. According to the AWS Cloud Practitioner Essentials course, reliability means "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."<sup>1</sup> Elasticity, security, and pay-as-you-go pricing are also benefits of AWS Cloud computing, but they do not directly relate to the goal of consistent and correct performance.

**NEW QUESTION 108**

- (Topic 2)

A company wants to move its data warehouse application to the AWS Cloud. The company wants to run and scale its analytics services without needing to provision and manage data warehouse clusters. Which AWS service will meet these requirements?

- A. Amazon Redshift provisioned data warehouse
- B. Amazon Redshift Serverless
- C. Amazon Athena
- D. Amazon S3

**Answer:** B

**Explanation:**

Amazon Redshift Serverless is the AWS service that will meet the requirements of the company that wants to move its data warehouse application to the AWS Cloud and run and scale its analytics services without needing to provision and manage data warehouse clusters. Amazon Redshift Serverless is a new feature of Amazon Redshift, which is a fully managed data warehouse service that allows customers to run complex queries and analytics on large volumes of structured and semi-structured data. Amazon Redshift Serverless automatically scales the compute and storage resources based on the workload demand, and customers only pay for the resources they consume. Amazon Redshift Serverless also simplifies the management and maintenance of the data warehouse, as customers do not need to worry about choosing the right cluster size, resizing the cluster, or distributing the data across the nodes. Amazon Redshift provisioned data warehouse, Amazon Athena, and Amazon S3 are not the best services to meet the requirements of the company. Amazon Redshift provisioned data warehouse requires customers to choose the number and type of nodes for their cluster, and manually resize the cluster if their workload changes. Amazon Athena is a serverless query service that allows customers to analyze data stored in Amazon S3 using standard SQL, but it is not a data warehouse service that can store and organize the data. Amazon S3 is a scalable object storage service that can store any amount and type of data, but it is not a data warehouse service that can run complex queries and analytics on the data.

**NEW QUESTION 109**

- (Topic 2)

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

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

**Answer:** AB

**Explanation:**

The combination of AWS services that the application can use to migrate to a microservices-based application are Amazon Simple Queue Service (Amazon SQS)



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

**NEW QUESTION 114**

- (Topic 2)

Which AWS service is designed to help users orchestrate a workflow process for a set of AWS Lambda functions?

- A. Amazon DynamoDB
- B. AWS CodePipeline
- C. AWS Batch
- D. AWS Step Functions

**Answer:** D

**Explanation:**

The AWS service that is designed to help users orchestrate a workflow process for a set of AWS Lambda functions is AWS Step Functions. AWS Step Functions is a service that helps users coordinate multiple AWS services into serverless workflows that can be triggered by events, such as messages, API calls, or schedules. AWS Step Functions allows users to create and visualize complex workflows that can include branching, parallel execution, error handling, retries, and timeouts. AWS Step Functions can integrate with AWS Lambda to orchestrate a sequence of Lambda functions that perform different tasks or logic. Amazon DynamoDB, AWS CodePipeline, and AWS Batch are not the best services to use for orchestrating a workflow process for a set of AWS Lambda functions. Amazon DynamoDB is a fully managed NoSQL database service that provides fast and consistent performance, scalability, and flexibility. AWS CodePipeline is a fully managed continuous delivery service that helps users automate the release process of their applications. AWS Batch is a fully managed service that helps users run batch computing workloads on the AWS Cloud.

**NEW QUESTION 115**

- (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 120**

- (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. 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 124**

- (Topic 2)

Which perspective of the AWS Cloud Adoption Framework (AWS CAF) connects technology and business?

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

**Answer:** D

**Explanation:**

The perspective of the AWS Cloud Adoption Framework (AWS CAF) that connects technology and business is governance. The governance perspective focuses on the alignment of the IT strategy and processes with the business strategy and goals, as well as the management of the IT budget, risk, and compliance. The governance perspective capabilities are portfolio management, business performance management, and IT governance. The governance perspective helps organizations ensure that their cloud adoption delivers the expected business value and outcomes, and that their cloud solutions are secure, reliable, and compliant. Operations, people, and security are other perspectives of the AWS CAF, but they do not directly connect technology and business. The operations perspective focuses on the management and monitoring of the cloud resources and applications, as well as the automation and optimization of the operational processes. The people perspective focuses on the development and empowerment of the human resources, as well as the transformation of the organizational culture and structure. The security perspective focuses on the protection of the information assets and systems in the cloud, as well as the implementation of the security policies and controls.

**NEW QUESTION 126**

- (Topic 2)

A company has a single Amazon EC2 instance. The company wants to adopt a highly available architecture. What can the company do to meet this requirement?

- A. Scale vertically to a larger EC2 instance size.
- B. Scale horizontally across multiple Availability Zones.
- C. Purchase an EC2 Dedicated Instance.
- D. Change the EC2 instance family to a compute optimized instance.

**Answer:** B

**Explanation:**

Scaling horizontally across multiple Availability Zones is a way to adopt a highly available architecture, as it increases the fault tolerance and resilience of the application. Scaling vertically to a larger EC2 instance size is a way to improve the performance of the application, but it does not improve the availability. Purchasing an EC2 Dedicated Instance is a way to isolate the instance from other AWS customers, but it does not improve the availability. Changing the EC2 instance family to a compute optimized instance is a way to optimize the instance type for the workload, but it does not improve the availability. These concepts are explained in the AWS Well-Architected Framework<sup>2</sup>.

**NEW QUESTION 129**

- (Topic 2)

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

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

**Answer:** A

**Explanation:**

Convertible Reserved Instances (RIs) are a type of Reserved Instance that allow you to change the attributes of the RI as long as the exchange results in the creation of Reserved Instances of equal or greater value. You can exchange Convertible RIs for other Convertible RIs from a different instance family, size, platform, tenancy, or scope (Region or Availability Zone)<sup>3</sup>.

**NEW QUESTION 134**

- (Topic 2)

A company has an AWS-hosted website located behind an Application Load Balancer. The company wants to safeguard the website from SQL injection or cross-site scripting.

Which AWS service should the company use?

- A. Amazon GuardDuty
- B. AWS WAF
- C. AWS Trusted Advisor
- D. Amazon Inspector

**Answer:** B

**Explanation:**

The company should use AWS WAF to safeguard the website from SQL injection or cross-site scripting. 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, AWS Trusted Advisor, and Amazon Inspector are not the best services to use for this purpose. Amazon GuardDuty is a threat detection service that monitors for malicious activity and unauthorized behavior across the AWS accounts and resources. AWS Trusted Advisor is a service that provides best practice recommendations for cost optimization, performance, security, and fault tolerance. Amazon Inspector is a service that assesses the security and compliance of applications running on Amazon EC2 instances<sup>12</sup>

**NEW QUESTION 137**

- (Topic 2)

An ecommerce company wants to design a highly available application that will be hosted on multiple Amazon EC2 instances.

How should the company deploy the EC2 instances to meet these requirements?

- A. Across multiple edge locations
- B. Across multiple VPCs
- C. Across multiple Availability Zones
- D. Across multiple AWS accounts

**Answer:** C

**Explanation:**

The company should deploy the EC2 instances across multiple Availability Zones to design a highly available application. Availability Zones are isolated locations within an AWS Region that are engineered to be fault-tolerant and operate independently of each other. By deploying the EC2 instances across multiple Availability Zones, the company can ensure that their application can withstand the failure of an entire Availability Zone and continue to operate with minimal disruption. Deploying the EC2 instances across multiple edge locations, VPCs, or AWS accounts will not provide the same level of availability and fault tolerance as Availability Zones. Edge locations are part of the Amazon CloudFront service, which is a content delivery network (CDN) that caches and serves web content to users. VPCs are virtual networks that isolate the AWS resources within an AWS Region. AWS accounts are the primary units of ownership and access control for AWS resources<sup>12</sup>

**NEW QUESTION 140**

- (Topic 2)

A company is running an application that is hosted on Amazon EC2 instances. The usage of the EC2 instances is higher during daytime hours than nighttime hours. The company wants to optimize the number of EC2 instances based on this usage pattern.

Which AWS service or instance purchasing option should the company use to meet these requirements?

- A. Spot Instances
- B. Reserved Instances
- C. AWS CloudFormation
- D. AWS Auto Scaling

**Answer:** D

**Explanation:**

AWS Auto Scaling is the AWS service that allows users to optimize the number of EC2 instances based on the usage pattern, as it automatically adjusts the capacity to maintain steady and predictable performance at the lowest possible cost. Spot Instances are a way to reduce the cost of EC2 instances by bidding on unused EC2 capacity, but they are not suitable for applications that require steady and reliable performance. Reserved Instances are a way to reduce the cost of EC2 instances by committing to a certain amount of usage for a period of time, but they are not flexible to adjust to the usage pattern. AWS CloudFormation is a way to automate the creation and management of AWS resources, but it does not optimize the number of EC2 instances based on the usage pattern. These concepts are explained in the AWS Cloud Practitioner Essentials course<sup>3</sup>.

**NEW QUESTION 144**

- (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
- 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."<sup>5</sup> 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 147**

- (Topic 1)

Which options does AWS make available for customers who want to learn about security in the cloud in an instructor-led setting? (Select TWO.)

- A. AWS Trusted Advisor
- B. AWS Online Tech Talks
- C. AWS Blog
- D. AWS Forums
- E. AWS Classroom Training

**Answer:** BE

**Explanation:**

The correct answers are B and E because AWS Online Tech Talks and AWS Classroom Training are options that AWS makes available for customers who want to learn about security in the cloud in an instructor-led setting. AWS Online Tech Talks are live, online presentations that cover a broad range of topics at varying technical levels. AWS Online Tech Talks are delivered by AWS experts and feature live Q&A sessions with the audience. AWS Classroom Training are in-person or virtual courses that are led by accredited AWS instructors. AWS Classroom Training offer hands-on labs, exercises, and best practices to help customers gain confidence and skills on AWS. The other options are incorrect because they are not options that AWS makes available for customers who want to learn about



security in the cloud in an instructor-led setting. AWS Trusted Advisor is an AWS service that provides real-time guidance to help customers follow AWS best practices for security, performance, cost optimization, and fault tolerance. AWS Blog is an AWS resource that provides news, announcements, and insights from AWS experts and customers. AWS Forums are AWS resources that enable customers to interact with other AWS users and get feedback and support. Reference: AWS Online Tech Talks, AWS Classroom Training

**NEW QUESTION 151**

- (Topic 1)

Which of the following is a characteristic of the AWS account root user?

- A. The root user is the only user that can be configured with multi-factor authentication (MFA).
- B. The root user is the only user that can access the AWS Management Console.
- C. The root user is the first sign-in identity that is available when an AWS account is created.
- D. The root user has a password that cannot be changed.

**Answer:** C

**Explanation:**

The AWS account root user is the first sign-in identity that is available when an AWS account is created. It has complete access to all AWS services and resources in the account. The root user email address and password are the same credentials that are used to sign in to the AWS Management Console<sup>4</sup>. The root user should be used only to perform a few account and service management tasks. For day-to-day tasks, it is recommended to use AWS Identity and Access Management (IAM) users or roles instead.

**NEW QUESTION 152**

- (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 153**

- (Topic 1)

Which benefit is included with an AWS Enterprise Support plan?

- A. AWS Partner Network (APN) support at no cost
- B. Designated support from an AWS technical account manager (TAM)
- C. On-site support from AWS engineers
- D. AWS managed compliance as code with AWS Config

**Answer:** B

**Explanation:**

AWS offers different support plans to meet the needs of different customers. The AWS Enterprise Support plan is the highest level of support that provides customers with concierge-like service, where the main focus is helping them achieve their outcomes and find success in the cloud. One of the benefits of the AWS Enterprise Support plan is that customers get designated support from an AWS technical account manager (TAM), who provides consultative architectural and operational guidance based on their applications and use cases. Therefore, the correct answer is B. You can learn more about AWS support plans and their benefits from this page.

**NEW QUESTION 157**

- (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 162**

- (Topic 1)

A retail company is building a new mobile app. The company is evaluating whether to build the app at an on-premises data center or in the AWS Cloud. responsibility model?

- A. Amazon FSx for Windows File Server
- B. Amazon Workspaces virtual Windows desktop
- C. AWS Directory Service for Microsoft Active Directory
- D. Amazon RDS for Microsoft SQL Server

**Answer:** C

**Explanation:**

AWS Directory Service for Microsoft Active Directory is the AWS service that provides a managed Microsoft Active Directory in the AWS Cloud. It enables the user to use their existing Active Directory users, groups, and policies to access AWS resources, such as Amazon EC2 instances, Amazon S3 buckets, and AWS Single Sign-On. It also integrates with other Microsoft applications and services, such as Microsoft SQL Server, Microsoft Office 365, and Microsoft SharePoint

**NEW QUESTION 163**

- (Topic 1)

A company wants to host its relational databases on AWS. The databases have predefined schemas that the company needs to replicate on AWS. Which AWS services could the company use for the databases? (Select TWO.)

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

**Answer:** AB

**Explanation:**

The correct answers are A and B because Amazon Aurora and Amazon RDS are AWS services that the company could use for the relational databases. Amazon Aurora is a relational database that is compatible with MySQL and PostgreSQL. Amazon Aurora is a fully managed, scalable, and high-performance service that offers up to five times the throughput of standard MySQL and up to three times the throughput of standard PostgreSQL. Amazon RDS is a service that enables users to set up, operate, and scale relational databases in the cloud. Amazon RDS supports six popular database engines: MySQL, PostgreSQL, Oracle, SQL Server, MariaDB, and Amazon Aurora. The other options are incorrect because they are not AWS services that the company could use for the relational databases. 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. Amazon DynamoDB is a key- value and document database. Reference: Amazon Aurora, Amazon RDS

**NEW QUESTION 166**

- (Topic 1)

A company needs to continuously monitor its environment to analyze network and account activity and identify potential security threats. Which AWS service should the company use to meet these requirements?

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

**Answer:** D

**Explanation:**

Amazon GuardDuty is a service that provides intelligent threat detection and continuous monitoring for the AWS environment. It analyzes network and account activity using machine learning and threat intelligence to identify potential security threats, such as unauthorized access, compromised credentials, malicious hosts, and reconnaissance activities. It also generates detailed and actionable findings that can be viewed on the AWS Management Console or sent to other AWS services, such as Amazon CloudWatch Events and AWS Lambda, for further analysis or remediation. Amazon GuardDuty OverviewAWS Certified Cloud Practitioner - aws.amazon.com

**NEW QUESTION 168**

- (Topic 1)

Who enables encryption of data at rest for Amazon Elastic Block Store (Amazon EBS)?

- A. AWS Support
- B. AWS customers
- C. AWS Key Management Service (AWS KMS)
- D. AWS Trusted Advisor

**Answer:** B

**Explanation:**

AWS customers are responsible for enabling encryption of data at rest for Amazon Elastic Block Store (Amazon EBS). Amazon EBS encryption offers a simple

encryption solution for your EBS volumes that does not require you to build, maintain, and secure your own key management infrastructure. You can encrypt both the boot and data volumes of your EC2 instances. You can use AWS Key Management Service (AWS KMS) customer master keys (CMKs) or your own CMKs to encrypt your volumes2.

**NEW QUESTION 170**

- (Topic 1)

Which services can be used to deploy applications on AWS? (Select TWO.)

- A. AWS Elastic Beanstalk
- B. AWS Config
- C. AWS OpsWorksQ
- D. AWS Application Discovery Service
- E. Amazon Kinesis

**Answer:** AC

**Explanation:**

The services that can be used to deploy applications on AWS are:

? AWS Elastic Beanstalk. This is a service that simplifies the deployment and management of web applications on AWS. Users can upload their application code and Elastic Beanstalk automatically handles the provisioning, scaling, load balancing, monitoring, and health checking of the resources needed to run the application. Users can also retain full control and access to the underlying resources and customize their configuration settings. Elastic Beanstalk supports multiple platforms, such as Java, .NET, PHP, Node.js, Python, Ruby, Go, and Docker. [AWS Elastic Beanstalk Overview] AWS Certified Cloud Practitioner - [aws.amazon.com](https://aws.amazon.com)

? AWS OpsWorks. This is a service that provides configuration management and automation for AWS resources. Users can define the application architecture and the configuration of each resource using Chef or Puppet, which are popular open- source automation platforms. OpsWorks then automatically creates and configures the resources according to the user's specifications. OpsWorks also provides features such as auto scaling, monitoring, and integration with other AWS services. OpsWorks has two offerings: OpsWorks for Chef Automate and OpsWorks for Puppet Enterprise. [AWS OpsWorks Overview] AWS Certified Cloud Practitioner - [aws.amazon.com](https://aws.amazon.com)

**NEW QUESTION 173**

- (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 differentAWS 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 176**

- (Topic 1)

A newly created IAM user has no IAM policy attached.

What will happen when the user logs in and attempts to view the AWS resources in the account?

- A. All AWS services will be read-only access by default.
- B. Access to all AWS resources will be denied.
- C. Access to the AWS billing services will be allowed.
- D. Access to AWS resources will be allowed through the AWS CLL

**Answer:** B

**Explanation:**

Access to all AWS resources will be denied if a newly created IAM user has no IAM policy attached and logs in and attempts to view the AWS resources in the account.

IAM policies are the way to grant permissions to IAM users, groups, and roles to access and manage AWS resources. By default, IAM users have no permissions, unless they are explicitly granted by an IAM policy. Therefore, a newly created IAM user without any IAM policy attached will not be able to view or perform any actions on the AWS resources in the account. Access to the AWS billing services and AWS CLI will also be denied, unless the user has the necessary permissions.

**NEW QUESTION 179**

- (Topic 1)

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

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

**Answer: B**

**Explanation:**

The company should use Amazon RDS with a MySQL database to meet the requirements of moving its workload to AWS so that the tasks of patching the database and taking backup snapshots of the data in the clusters will be completed automatically. Amazon RDS is a managed service that simplifies the setup, operation, and scaling of relational databases in the AWS Cloud. Amazon RDS automates common database administration tasks such as patching, backup, and recovery. Amazon RDS also supports MySQL and other popular database engines.

**NEW QUESTION 182**

- (Topic 1)

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

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

**Answer: B**

**Explanation:**

AWS Compute Optimizer is the AWS service or tool that provides recommendations to help users get rightsized Amazon EC2 instances based on historical workload usage data. AWS Compute Optimizer analyzes the configuration and performance characteristics of the EC2 instances and delivers recommendations for optimal instance types, sizes, and configurations. AWS Compute Optimizer helps users improve performance, reduce costs, and eliminate underutilized resources.

**NEW QUESTION 183**

- (Topic 1)

An Availability Zone consists of:

- A. one or more data centers in a single location.
- B. two or more data centers in multiple locations.
- C. one or more physical hosts in a single data center.
- D. two or more physical hosts in multiple data centers.

**Answer: A**

**Explanation:**

The correct answer is A because an Availability Zone consists of one or more data centers in a single location. An Availability Zone is an isolated location within an AWS Region that has independent power, cooling, and networking. Each Availability Zone has one or more data centers that host the physical servers and storage devices that run the AWS services. The other options are incorrect because they are not accurate descriptions of an Availability Zone. Two or more data centers in multiple locations are not an Availability Zone, but rather multiple Availability Zones within an AWS Region. One or more physical hosts in a single data center are not an Availability Zone, but rather the components of a data center within an Availability Zone. Two or more physical hosts in multiple data centers are not an Availability Zone, but rather the components of multiple data centers within one or more Availability Zones. Reference: [Regions, Availability Zones, and Local Zones]

**NEW QUESTION 185**

- (Topic 1)

A company needs a content delivery network that provides secure delivery of data, videos, applications, and APIs to users globally with low latency and high transfer speeds.

Which AWS service meets these requirements?

- A. Amazon CloudFront
- B. Elastic Load Balancing
- C. Amazon S3
- D. Amazon Elastic Transcoder

**Answer: A**

**Explanation:**

The correct answer is A because Amazon CloudFront is an AWS service that provides secure delivery of data, videos, applications, and APIs to users globally with low latency and high transfer speeds. Amazon CloudFront is a fast content delivery network (CDN) that integrates with other AWS services, such as Amazon S3, Amazon EC2, AWS Lambda, and AWS Shield. Amazon CloudFront delivers content through a worldwide network of edge locations that are located close to the end users. The other options are incorrect because they are not AWS services that provide secure delivery of data, videos, applications, and APIs to users globally with low latency and high transfer speeds. Elastic Load Balancing is an AWS service that distributes incoming traffic across multiple targets, such as Amazon EC2 instances, containers, and IP addresses. Amazon S3 is an AWS service that provides object storage for data of any size and type. Amazon Elastic Transcoder is an AWS service that converts media files from their original source format into different formats that will play on various devices. Reference: Amazon CloudFront FAQs

**NEW QUESTION 186**

- (Topic 3)

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

What is the MOST operationally efficient AWS solution for this scenario?

- A. Create an Amazon S3 bucket for each use
- B. Mount each bucket by using an S3 file system mounting utility.
- C. Configure and deploy an AWS Storage Gateway file gateway
- D. Connect each user's workstation to the file gateway.
- E. Move each user's working environment to Amazon Workspace
- F. Set up an Amazon WorkDocs account for each user.
- G. Deploy an Amazon EC2 instance and attach an Amazon Elastic Block Store (Amazon EBS) Provisioned IOPS volume
- H. Share the EBS volume directly with the users.

**Answer:** B

**Explanation:**

AWS Storage Gateway is a hybrid cloud storage service that allows you to extend your on-premises file storage capabilities to the AWS Cloud. AWS Storage Gateway file gateway enables you to store and access your files in Amazon S3 using industry-standard file protocols such as NFS and SMB. File gateway caches frequently accessed files locally, providing low-latency access to your data. File gateway also optimizes the transfer of data between your on-premises environment and AWS, minimizing the amount of bandwidth consumed. By using file gateway, you can retain the performance benefit of sharing content locally while leveraging the scalability, durability, and cost-effectiveness of Amazon S3.

References: AWS Storage Gateway, File Gateway

**NEW QUESTION 191**

- (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 193**

- (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 196**

- (Topic 3)

Which AWS service supports a hybrid architecture that gives users the ability to extend AWS infrastructure, AWS services, APIs, and tools to data centers, co-location environments, or on-premises facilities?

- A. AWS Snowmobile
- B. AWS Local Zones
- C. AWS Outposts
- D. AWS Fargate

**Answer:** C



**Explanation:**

AWS Outposts is a service that delivers AWS infrastructure and services to virtually any on-premises or edge location for a truly consistent hybrid experience. AWS Outposts allows 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. Run applications and workloads on premises using familiar AWS services, tools, and APIs<sup>2</sup>. AWS Outposts is the only AWS service that supports a hybrid architecture that gives users the ability to extend AWS infrastructure, AWS services, APIs, and tools to data centers, co- location environments, or on-premises facilities. References: On-Premises Infrastructure - AWS Outposts Family

**NEW QUESTION 197**

- (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<sup>2</sup>.

**NEW QUESTION 199**

- (Topic 3)

Which AWS services or features can a company use to connect the network of its on- premises data center to AWS? (Select TWO.)

- A. AWS VPN
- B. AWS Directory Service
- C. AWS Data Pipeline
- D. AWS Direct Connect
- E. AWS CloudHSM

**Answer:** AD

**Explanation:**

AWS VPN and AWS Direct Connect are two services that enable customers to connect their on-premises data center network to the AWS Cloud. AWS VPN establishes a secure and encrypted connection over the public internet, while AWS Direct Connect establishes a dedicated and private connection through a partner network. You can learn more about AWS VPN from [this webpage] or [this digital course]. You can learn more about AWS Direct Connect from [this webpage] or [this digital course].

**NEW QUESTION 201**

- (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<sup>4</sup>.

**NEW QUESTION 204**

- (Topic 3)

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

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

**Answer:** C

**Explanation:**

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

**NEW QUESTION 209**

- (Topic 3)

Which AWS service converts text to lifelike voices?

- A. Amazon Transcribe
- B. Amazon Rekognition
- C. Amazon Polly
- D. Amazon Textract

**Answer:** C

**Explanation:**

Amazon Polly is a service that turns text into lifelike speech, allowing you to create applications that talk, and build entirely new categories of speech-enabled products. Polly's Text-to-Speech (TTS) service uses advanced deep learning technologies to synthesize natural sounding human speech<sup>1</sup>. Amazon Polly supports dozens of languages and a wide range of natural-sounding voices. You can customize and control the speech output by using lexicons and SSML tags. You can also store and redistribute the speech output in standard audio formats like MP3 and OGG<sup>2</sup>.

Amazon Transcribe is a service that converts speech to text, enabling you to create text transcripts from audio or video files. It can recognize multiple speakers, different languages, accents, dialects, and background noises. It can also add punctuation and formatting to the transcripts. Amazon Transcribe is useful for applications such as subtitling, captioning, transcription, and voice search.

Amazon Rekognition is a service that provides image and video analysis using computer vision and deep learning. It can detect objects, faces, text, scenes, activities, and emotions in images and videos. It can also perform face recognition, face comparison, face search, celebrity recognition, and facial analysis. Amazon Rekognition is useful for applications such as security, social media, e-commerce, and media and entertainment.

Amazon Textract is a service that extracts text and data from scanned documents using optical character recognition (OCR) and machine learning. It can identify the contents of fields in forms and tables, as well as the relationships between them. It can also preserve the layout and structure of the original document.

Amazon Textract is useful for applications such as data entry, document management, compliance, and analytics. References:

? Text to Speech Software – Amazon Polly – Amazon Web Services

? What is Text to Speech – Amazon Web Services (AWS)

? AWS Amazon Polly - Text to Speech Converter - CodeCanyon

? Amazon's Text-To-Speech AI Service Sounds More Natural And ... - Forbes

? Working with AWS Amazon Polly Text-to-Speech (TTS) Service

? [Automatic Speech Recognition - Amazon Transcribe - AWS]

? [Amazon Rekognition – Video and Image - AWS]

? [Extract Text & Data - OCR - Amazon Textract - AWS]

**NEW QUESTION 212**

- (Topic 3)

A user has a stateful workload that will run on Amazon EC2 for the next 3 years. What is the MOST cost-effective pricing model for this workload?

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

**Answer:** B

**Explanation:**

Reserved Instances are a pricing model that offers significant discounts on Amazon EC2 usage compared to On-Demand Instances. Reserved Instances are suitable for stateful workloads that have predictable and consistent usage patterns for a long-term period. By committing to a one-year or three-year term, customers can reduce their total cost of ownership and optimize their cloud spend. Reserved Instances also provide capacity reservation, ensuring that customers have access to the EC2 instances they need when they need them. References: AWS Pricing Calculator, Amazon EC2 Pricing, [AWS Cloud Practitioner Essentials: Module 3 - Compute in the Cloud]

**NEW QUESTION 217**

- (Topic 3)

A company needs a bridge between technology and business to help evolve to a culture of continuous growth and learning.

Which perspective in the AWS Cloud Adoption Framework (AWS CAF) serves as this bridge?

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

**Answer:** A

**Explanation:**

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

**NEW QUESTION 220**

- (Topic 3)

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

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

**Answer:** A

**Explanation:**

Amazon S3 Glacier Deep Archive is the lowest-cost storage class in the cloud. It is designed for long-term data archiving that is rarely accessed. It offers a

retrieval time of 12 hours and a durability of 99.999999999% (11 9's). It is ideal for data that must be retained for 7 years or longer to meet regulatory compliance requirements.

**NEW QUESTION 225**

- (Topic 3)

A company is using Amazon DynamoDB.

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

- A. Patch the operating system
- B. Provision hosts
- C. Manage database access permissions.
- D. Secure the operating system

**Answer:** C

**Explanation:**

According to the AWS shared responsibility model, AWS is responsible for the security of the cloud, while customers are responsible for the security in the cloud. This means that AWS is responsible for the physical servers, networking, and operating system that run DynamoDB, while customers are responsible for the security of their data and access to the database. Customers need to manage database access permissions, such as creating and managing AWS Identity and Access Management (IAM) policies and roles, and using encryption and key management options to protect their data<sup>123</sup>. References: 1: Shared Responsibility Model - Amazon Web Services (AWS), 2: Security in Amazon DynamoDB - Amazon DynamoDB, 3: AWS Shared Responsibility Model - Introduction to DevOps ...

**NEW QUESTION 227**

- (Topic 3)

Amazon Elastic File System (Amazon EFS) and Amazon FSx offer which type of storage?

- A. File storage
- B. Object storage
- C. Block storage
- D. Instance store

**Answer:** A

**Explanation:**

Amazon Elastic File System (Amazon EFS) and Amazon FSx are AWS services that offer file storage. File storage is a type of storage that organizes data into files and folders that can be accessed and shared over a network. File storage is suitable for applications that require shared access to data, such as content management, media processing, and web serving. Amazon EFS provides a simple, scalable, and fully managed elastic file system that can be used with AWS Cloud services and on-premises resources. Amazon FSx provides fully managed third-party file systems, such as Windows File Server and Lustre, with native compatibility and high performance<sup>12</sup>

**NEW QUESTION 230**

- (Topic 3)

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

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

**Answer:** C

**Explanation:**

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

**NEW QUESTION 234**

- (Topic 3)

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

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

**Answer:** BC

**Explanation:**

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

#### NEW QUESTION 238

- (Topic 3)

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

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

**Answer:** B

#### Explanation:

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

#### NEW QUESTION 243

- (Topic 3)

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

Which tasks should the company perform to meet these requirements, according to the AWS Cloud Adoption Framework (AWS CAF)? (Select TWO.)

- A. Realign teams to focus on products and value streams.
- B. Create new value propositions with new products and services.
- C. Use agile methods to rapidly iterate and evolve.
- D. Use a new data and analytics platform to create actionable insights.
- E. Migrate and modernize legacy infrastructure.

**Answer:** AC

#### Explanation:

Realigning teams to focus on products and value streams, and using agile methods to rapidly iterate and evolve are tasks that the company should perform to meet the requirements of becoming more responsive to customer inquiries and feedback, according to the AWS Cloud Adoption Framework (AWS CAF). AWS CAF organizes guidance into six areas of focus, called perspectives: business, people, governance, platform, security, and operations. Each perspective is divided into capabilities, which describe the skills and processes to execute the transition effectively. The people perspective helps you prepare your organization for cloud adoption, and includes capabilities such as organizational change management, staff skills and readiness, and organizational alignment. The business perspective helps you align IT strategy with business strategy, and includes capabilities such as business case development, value proposition, and product ownership. Creating new value propositions with new products and services is a task that belongs to the business perspective, but it is not directly related to the requirement of becoming more responsive to customer inquiries and feedback. Using a new data and analytics platform to create actionable insights is a task that belongs to the platform perspective, which helps you design, implement, and optimize the architecture of the AWS environment. However, it is also not directly related to the requirement of becoming more responsive to customer inquiries and feedback. Migrating and modernizing legacy infrastructure is a task that belongs to the operations perspective, which helps you enable, run, use, operate, and recover IT workloads to the level agreed upon with your business stakeholders. However, it is also not directly related to the requirement of becoming more responsive to customer inquiries and feedback.

#### 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 248

- (Topic 3)

AWS has the ability to achieve lower pay-as-you-go pricing by aggregating usage across hundreds of thousands of users.

This describes which advantage of the AWS Cloud?

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

**Answer:** C

#### Explanation:

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

#### NEW QUESTION 251

- (Topic 3)



A company wants a customized assessment of its current on-premises environment. The company wants to understand its projected running costs in the AWS Cloud.

Which AWS service or tool will meet these requirements?

- A. AWS Trusted Advisor
- B. Amazon Inspector
- C. AWS Control Tower
- D. Migration Evaluator

**Answer:** D

**Explanation:**

Migration Evaluator is an AWS service that provides a customized assessment of your current on-premises environment and helps you build a data-driven business case for migration to AWS. Migration Evaluator collects and analyzes data from your on-premises servers, such as CPU, memory, disk, network, and utilization metrics, and compares them with the most cost-effective AWS alternatives. Migration Evaluator also helps you understand your existing software licenses and running costs, and provides recommendations for Bring Your Own License (BYOL) and License Included (LI) options in AWS. Migration Evaluator generates a detailed report that shows your projected running costs in the AWS Cloud, along with potential savings and benefits. You can use this report to support your decision-making and planning for cloud migration. References: Cloud Business Case & Migration Plan - Amazon Migration Evaluator - AWS, Getting started with Migration Evaluator

**NEW QUESTION 256**

- (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)

A company wants a time-series database service that makes it easier to store and analyze trillions of events each day.

Which AWS service will meet this requirement?

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

**Answer:** B

**Explanation:**

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

References:

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

**NEW QUESTION 264**

- (Topic 3)

Which task does AWS perform automatically?

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

**Answer:** B

**Explanation:**

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

**NEW QUESTION 266**

- (Topic 3)

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

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

**Answer: B**

**Explanation:**

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

**NEW QUESTION 270**

- (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 272**

- (Topic 3)

A company wants to use the AWS Cloud to deploy an application globally. Which architecture deployment model should the company use to meet this requirement?

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

**Answer: A**

**Explanation:**

The architecture deployment model that the company should use to meet this requirement is A. Multi-Region. A multi-region deployment model is a cloud computing architecture that distributes an application and its data across multiple geographic regions. A multi-region deployment model enables a company to achieve global reach, high availability, disaster recovery, and performance optimization. By deploying an application in multiple regions, a company can serve customers from the nearest region, reduce latency, increase redundancy, and comply with data sovereignty regulations<sup>12</sup>. A single-region deployment model is a cloud computing architecture that runs an application and its data within a single geographic region. A single-region deployment model is simpler and cheaper than a multi-region deployment model, but it has limited scalability, availability, and performance. A single-region deployment model may not be suitable for a company that wants to deploy an application globally, as it may face challenges such as network latency, regional outages, or regulatory compliance<sup>12</sup>. A multi-AZ (Availability Zone) deployment model is a cloud computing architecture that distributes an application and its data across multiple isolated locations within a single region. An Availability Zone is a physically separate location within an AWS Region that has independent power, cooling, and networking. A multi-AZ deployment model enhances the availability and durability of an application by providing redundancy and fault tolerance within a region<sup>34</sup>. A single-AZ deployment model is a cloud computing architecture that runs an application and its data within a single Availability Zone. A single-AZ deployment model is the simplest and most cost-effective option, but it has no redundancy or fault tolerance. A single-AZ deployment model may not be suitable for a company that wants to deploy an application globally, as it may face challenges such as network latency, regional outages, or regulatory compliance<sup>34</sup>. References: 1: AWS Cloud Computing - W3Schools 2: Understand the Different Cloud Computing Deployment Models Unit - Trailhead 3: Regions and Availability Zones - Amazon Elastic Compute Cloud 4: AWS Reference Architecture Diagrams

**NEW QUESTION 276**

- (Topic 3)

A company is building an application that needs to deliver images and videos globally with minimal latency. Which approach can the company use to accomplish this in a cost effective manner?

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

**Answer:** A

**Explanation:**

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

**NEW QUESTION 279**

- (Topic 3)

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

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

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

**Answer:** C

**Explanation:**

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

**NEW QUESTION 281**

- (Topic 3)

What does the concept of agility mean in AWS Cloud computing? (Select TWO.)

- A. The speed at which AWS resources are implemented
- B. The speed at which AWS creates new AWS Regions
- C. The ability to experiment quickly
- D. The elimination of wasted capacity
- E. The low cost of entry into cloud computing

**Answer:** AC

**Explanation:**

Agility in AWS Cloud computing means the ability to rapidly provision and deprovision AWS resources as needed, and the ability to experiment quickly with new ideas and solutions. Agility helps businesses to respond to changing customer demands, market opportunities, and competitive threats, and to innovate faster and cheaper. Agility also reduces the risk of failure, as businesses can test and validate their assumptions before committing to large-scale deployments. Some of the benefits of agility in AWS Cloud computing are:

? The speed at which AWS resources are implemented: AWS provides a variety of services and tools that allow you to create, configure, and launch AWS resources in minutes, using the AWS Management Console, the AWS Command Line Interface (AWS CLI), the AWS Software Development Kits (AWS SDKs), or the AWS CloudFormation templates. You can also use the AWS Cloud Development Kit (AWS CDK) to define your AWS resources as code using familiar programming languages, and synthesize them into AWS CloudFormation templates. You can also use the AWS Service Catalog to create and manage standardized portfolios of AWS resources that meet your organizational policies and best practices. AWS also offers on-demand, pay-as-you-go pricing models, so you only pay for the resources you use, and you can scale them up or down as your needs change

? The ability to experiment quickly: AWS enables you to experiment quickly with new ideas and solutions, without having to invest in upfront capital or long-term commitments. You can use AWS to create and test multiple prototypes, hypotheses, and minimum viable products (MVPs) in parallel, and measure their performance and feedback. You can also use AWS to leverage existing services and solutions, such as AWS Marketplace, AWS Solutions, and AWS Quick Starts, that can help you accelerate your innovation process. AWS also supports a culture of experimentation and learning, by providing tools and resources for continuous integration and delivery (CI/CD), testing, monitoring, and analytics.

References: Six advantages of cloud computing - Overview of Amazon Web Services, AWS Cloud Development Kit (AWS CDK), AWS Service Catalog, AWS Pricing, AWS CloudFormation, [Experimentation and Testing - AWS Well-Architected Framework], [AWS Marketplace], [AWS Solutions], [AWS Quick Starts], [AWS Developer Tools]

**NEW QUESTION 282**

.....

## Thank You for Trying Our Product

### We offer two products:

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

2nd - Questions and Answers in PDF Format

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

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

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