

## AWS-Certified-Cloud-Practitioner Dumps

## Amazon AWS Certified Cloud Practitioner

<https://www.certleader.com/AWS-Certified-Cloud-Practitioner-dumps.html>



**NEW QUESTION 1**

- (Topic 2)

A company wants to create a chatbot and integrate the chatbot with its current web application. Which AWS service will meet these requirements?

- A. AmazonKendra
- B. Amazon Lex
- C. AmazonTextract
- D. AmazonPolly

**Answer:** B

**Explanation:**

The AWS service that will meet the requirements of the company that wants to create a chatbot and integrate the chatbot with its current web application is Amazon Lex. Amazon Lex is a service that helps customers build conversational interfaces using voice and text. The company can use Amazon Lex to create a chatbot that can understand natural language and respond to user requests, using the same deep learning technologies that power Amazon Alexa. Amazon Lex also provides easy integration with other AWS services, such as Amazon Comprehend, Amazon Polly, and AWS Lambda, as well as popular platforms, such as Facebook Messenger, Slack, and Twilio. Amazon Lex helps customers create engaging and interactive chatbots for their web applications. Amazon Kendra, Amazon Textract, and Amazon Polly are not the best services to use for this purpose. Amazon Kendra is a service that helps customers provide accurate and natural answers to natural language queries using machine learning. Amazon Textract is a service that helps customers extract text and data from scanned documents using optical character recognition (OCR) and machine learning. Amazon Polly is a service that helps customers convert text into lifelike speech using deep learning. These services are more useful for different types of natural language processing and generation tasks, rather than creating and integrating chatbots.

**NEW QUESTION 2**

- (Topic 2)

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

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

**Answer:** D

**Explanation:**

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

**NEW QUESTION 3**

- (Topic 2)

Which task is the responsibility of AWS when using AWS services?

- A. Management of IAM user permissions
- B. Creation of security group rules for outbound access
- C. Maintenance of physical and environmental controls
- D. Application of Amazon EC2 operating system patches

**Answer:** C

**Explanation:**

AWS is responsible for maintaining the physical and environmental controls of the AWS Cloud, such as power, cooling, fire suppression, and physical security. The customer is responsible for managing the IAM user permissions, creating security group rules for outbound access, applying Amazon EC2 operating system patches, and other aspects of security in the cloud.

**NEW QUESTION 4**

- (Topic 2)

A company migrated its core application onto multiple workloads in the AWS Cloud. The company wants to improve the application's reliability. Which cloud design principle should the company implement to achieve this goal?

- A. Maximize utilization.
- B. Decouple the components.
- C. Rightsize the resources.
- D. Adopt a consumption model.

**Answer:** B

**Explanation:**

Decoupling the components of an application means reducing the dependencies and interactions between them, which can improve the application's reliability, scalability, and performance. Decoupling can be achieved by using services such as Amazon Simple Queue Service (Amazon SQS), Amazon Simple Notification Service (Amazon SNS), and AWS Lambda.

**NEW QUESTION 5**

- (Topic 2)

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

- A. Configure AWS Identity and Access Management (IAM).
- B. Configure security groups on Amazon EC2 instances.
- C. Secure the access of physical AWS facilities.
- D. Patch applications that run on Amazon EC2 instances.
- E. Perform infrastructure patching and maintenance.

**Answer:** CE

**Explanation:**

The tasks that are the responsibility of AWS according to the AWS shared responsibility model are securing the access of physical AWS facilities and performing infrastructure patching and maintenance. The AWS shared responsibility model 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 physical security of the hardware, software, networking, and facilities that run the AWS services. AWS is also responsible for the maintenance and patching of the infrastructure that supports the AWS services. The customer is responsible for the security in the cloud, which includes the configuration and management of the AWS resources and applications that they use. Configuring AWS Identity and Access Management (IAM), configuring security groups on Amazon EC2 instances, and patching applications that run on Amazon EC2 instances are tasks that are the responsibility of the customer, not AWS.

**NEW QUESTION 6**

- (Topic 2)

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

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

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

**Answer:** A

**Explanation:**

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

**NEW QUESTION 7**

- (Topic 2)

A company is collecting user behavior patterns to identify how to meet goals for sustainability impact.

Which guidelines are best practices for the company to implement to meet these goals? (Select TWO.)

- A. Scale infrastructure with user load.
- B. Maximize the geographic distance between workloads and user locations.
- C. Eliminate creation and maintenance of unused assets.
- D. Scale resources with excess capacity and remove auto scaling.
- E. Scale infrastructure based on the number of users.

**Answer:** AC

**Explanation:**

To meet the goals for sustainability impact, the company should follow the best practices of scaling infrastructure with user load and eliminating creation and maintenance of unused assets. Scaling infrastructure with user load means adjusting the capacity of the infrastructure to match the demand of the users, which can reduce the energy consumption and carbon footprint of the system. Eliminating creation and maintenance of unused assets means avoiding the waste of resources and money on assets that are not needed or used, which can also improve the environmental and economic efficiency of the system<sup>3</sup>.

**NEW QUESTION 8**

- (Topic 2)

Which AWS service or tool helps companies measure the environmental impact of their AWS usage?

- A. AWS customer carbon footprint tool
- B. AWS Compute Optimizer
- C. Sustainability pillar
- D. OS-Climate (Open Source Climate Data Commons)

**Answer:** A

**Explanation:**

AWS customer carbon footprint tool is an AWS service or tool that helps companies measure the environmental impact of their AWS usage. It allows users to estimate the carbon emissions associated with their AWS resources and services, such as EC2, S3, and Lambda. It also provides recommendations and best practices to reduce the carbon footprint and improve the sustainability of their AWS workloads<sup>4</sup>. AWS Compute Optimizer is an AWS service that helps users optimize the performance and cost of their EC2 instances and Auto Scaling groups. It provides recommendations for optimal instance types, sizes, and configurations based on the workload characteristics and utilization metrics. It does not help users measure the environmental impact of their AWS usage. Sustainability pillar is a concept that refers to the ability of a system to operate in an environmentally friendly and socially responsible manner. It is not an AWS service or tool that helps users measure the environmental impact of their AWS usage. OS-Climate (Open Source Climate Data Commons) is an initiative that aims to provide open source data, tools, and platforms to accelerate climate action and innovation. It is not an AWS service or tool that helps users measure the environmental impact of their AWS usage.

**NEW QUESTION 9**

- (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 10**

- (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 10**

- (Topic 2)

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

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

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

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

**Answer:** B

**Explanation:**

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

**NEW QUESTION 14**

- (Topic 2)

Which options are perspectives that include foundational capabilities of the AWS Cloud Adoption Framework (AWS CAF)? (Select TWO.)

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

**Answer:** CD

**Explanation:**

The options that are perspectives that include foundational capabilities of the AWS Cloud Adoption Framework (AWS CAF) are operations and performance efficiency. The AWS CAF is a guidance that helps organizations design and travel an accelerated path to successful cloud adoption. The AWS CAF organizes the cloud adoption process into six areas of focus, called perspectives, which are business, people, governance, platform, security, and operations. Each perspective is divided into capabilities, which are further divided into skills and responsibilities. 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 operations perspective capabilities are operations support, operations integration, and service management. The performance efficiency perspective focuses on the selection and configuration of the right cloud resources and services to meet the performance requirements of the applications, as well as the continuous improvement and innovation of the cloud solutions. The performance efficiency perspective capabilities are selection, review, and monitoring. Sustainability, security, and reliability are not perspectives of the AWS CAF, but they are aspects of the AWS Well-Architected Framework. The AWS Well-Architected Framework is a guidance that helps users build and operate secure, reliable, efficient, and cost-effective systems in the cloud. The AWS Well-Architected Framework consists of five pillars, which are operational excellence, security, reliability, performance efficiency, and cost optimization. Sustainability is a cross-cutting theme that applies to all the pillars, and refers to the environmental and social impact of the cloud solutions.

**NEW QUESTION 17**

- (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 22**

- (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 applications<sup>34</sup>

**NEW QUESTION 24**

- (Topic 1)

Who is responsible for decommissioning end-of-life underlying storage devices that are used to host data on AWS?

- A. Customer
- B. AWS
- C. Account creator
- D. Auditing team

**Answer: B**

**Explanation:**

AWS is responsible for decommissioning end-of-life underlying storage devices that are used to host data on AWS. AWS follows strict and audited data destruction processes to ensure that customer data is not exposed to unauthorized individuals or devices when an AWS storage device reaches the end of its useful life. AWS uses techniques detailed in DoD 5220.22-M ("National Industrial Security Program Operating Manual") or NIST 800-88 ("Guidelines for Media Sanitization") to destroy data as part of the decommissioning process<sup>3</sup>.

**NEW QUESTION 27**

- (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<sup>3</sup>.

**NEW QUESTION 30**

- (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 31**

- (Topic 1)

Which statement describes a characteristic of the AWS global infrastructure?

- A. Edge locations contain multiple AWS Regions.
- B. AWS Regions contain multiple Regional edge caches.
- C. Availability Zones contain multiple data centers.
- D. Each data center contains multiple edge locations.

**Answer:** C

**Explanation:**

Availability Zones contain multiple data centers. This is a characteristic of the AWS global infrastructure, which consists of AWS Regions, Availability Zones, and edge locations. AWS Regions are geographically isolated areas that contain multiple Availability Zones. 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. Edge locations are sites that are located closer to the end users and provide caching and content delivery services. AWS Global Infrastructure AWS Certified Cloud Practitioner - aws.amazon.com

**NEW QUESTION 32**

- (Topic 1)

A company has an application with robust hardware requirements. The application must be accessed by students who are using lightweight, low-cost laptops. Which AWS service will help the company deploy the application without investing in backend infrastructure or high end client hardware?

- A. Amazon AppStream 2.0
- B. AWS AppSync
- C. Amazon WorkLink
- D. AWS Elastic Beanstalk

**Answer:** A

**Explanation:**

The correct answer is A because Amazon AppStream 2.0 is a service that will help the company deploy the application without investing in backend infrastructure or high end client hardware. Amazon AppStream 2.0 is a fully managed, secure application streaming service that allows customers to stream desktop applications from AWS to any device running a web browser. Amazon AppStream 2.0 handles the provisioning, scaling, patching, and maintenance of the backend infrastructure, and delivers high performance and responsive user experience. The other options are incorrect because they are not services that will help the company deploy the application without investing in backend infrastructure or high end client hardware. AWS AppSync is a service that enables customers to create flexible APIs for synchronizing data across multiple data sources. Amazon WorkLink is a service that enables customers to provide secure, one-click access to internal websites and web apps from mobile devices. AWS Elastic Beanstalk is a service that enables customers to deploy and manage web applications using popular platforms such as Java, .NET, PHP, and Node.js. Reference: [Amazon AppStream 2.0 FAQs]

**NEW QUESTION 35**

- (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 39**

- (Topic 1)

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

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

**Answer:** D

**Explanation:**

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

**NEW QUESTION 43**

- (Topic 1)

A company recently migrated to the AWS Cloud. The company needs to determine whether its newly imported Amazon EC2 instances are the appropriate size and type.

Which AWS services can provide this information to the company? {Select TWO.}

- A. AWS Auto Scaling
- B. AWS Control Tower
- C. AWS Trusted Advisor
- D. AWS Compute Optimizer
- E. Amazon Forecast

**Answer:** CD

**Explanation:**

AWS Trusted Advisor and AWS Compute Optimizer are the AWS services that can provide information to the company about whether its newly imported Amazon EC2 instances are the appropriate size and type. AWS Trusted Advisor is an online tool that provides best practices recommendations in five categories: cost optimization, performance, security, fault tolerance, and service limits. AWS Trusted Advisor can help users identify underutilized or idle EC2 instances, and suggest ways to reduce costs and improve performance. AWS Compute Optimizer is a service that analyzes the configuration and utilization metrics of 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 46**

- (Topic 1)

A cloud practitioner is analyzing Amazon EC2 instance performance and usage to provide recommendations for potential cost savings.

Which cloud concept does this analysis demonstrate?

- A. Auto scaling
- B. Rightsizing
- C. Load balancing
- D. High availability

**Answer:** B

**Explanation:**

Rightsizing is the cloud concept that this analysis demonstrates. Rightsizing is the process of optimizing the performance and cost of your AWS resources by selecting the most appropriate type, size, and configuration based on your workload requirements and usage patterns. Rightsizing can help you achieve potential cost savings by reducing the over-provisioning or under-utilization of your resources. You can use various AWS tools and services, such as AWS Cost Explorer, AWS Compute Optimizer, and AWS Trusted Advisor, to analyze your resource utilization and performance metrics, and receive recommendations for rightsizing.

**NEW QUESTION 49**

- (Topic 1)

A company wants to establish a security layer in its VPC that will act as a firewall to control subnet traffic.

Which AWS service or feature will meet this requirement?

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

**Answer:** C

**Explanation:**

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

**NEW QUESTION 53**

- (Topic 1)

A company's application stores data in an Amazon S3 bucket. The company has an AWS Lambda function that processes data in the S3 bucket. The company needs to invoke the function once a day at a specific time. Which AWS service should the company use to meet this requirement?

- A. AWS Managed Services (AMS)
- B. AWS CodeStar
- C. Amazon EventBridge
- D. AWS Step Functions

**Answer: C**

**Explanation:**

Amazon EventBridge is the service that the company should use to meet the requirement of invoking the Lambda function once a day at a specific time. Amazon EventBridge is a serverless event bus service that allows you to easily connect your applications with data from AWS services, SaaS applications, and your own applications. You can use Amazon EventBridge to create rules that match events and route them to targets such as AWS Lambda functions, Amazon SNS topics, Amazon SQS queues, or other AWS services. You can also use Amazon EventBridge to create scheduled rules that trigger your targets at a specific time or interval, such as once a day. AWS Managed Services (AMS), AWS CodeStar, and AWS Step Functions are not services that the company should use to meet this requirement. AMS is a service that provides operational management for your AWS infrastructure and applications. AWS CodeStar is a service that provides a unified user interface for managing software development projects on AWS. AWS Step Functions is a service that coordinates multiple AWS services into serverless workflows.

**NEW QUESTION 57**

- (Topic 1)

A company wants to deploy and manage a Docker-based application on AWS. Which solution meets these requirements with the LEAST amount of operational overhead?

- A. An open-source Docker orchestrator on Amazon EC2 instances
- B. AWS AppSync
- C. Amazon Elastic Container Registry (Amazon ECR)
- D. Amazon Elastic Container Service (Amazon ECS)

**Answer: D**

**Explanation:**

Amazon Elastic Container Service (Amazon ECS) is a solution that meets the requirements of deploying and managing a Docker-based application on AWS with the least amount of operational overhead. Amazon ECS is a fully managed container orchestration service that makes it easy to run, scale, and secure Docker container applications on AWS. Amazon ECS eliminates the need for you to install, operate, and scale your own cluster management infrastructure. With simple API calls, you can launch and stop container-enabled applications, query the complete state of your cluster, and access many familiar features like security groups, Elastic Load Balancing, EBS volumes, and IAM roles.

**NEW QUESTION 61**

- (Topic 1)

Which AWS service or feature offers HTTP attack protection to users running public-facing web applications?

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

**Answer: D**

**Explanation:**

AWS WAF is the AWS service or feature that offers HTTP attack protection to users running public-facing web applications. AWS WAF is a web application firewall that helps users protect their web applications from common web exploits, such as SQL injection, cross-site scripting, and bot attacks. Users can create custom rules to define the web traffic that they want to allow, block, or count. Users can also use AWS Managed Rules, which are pre-configured rules that are curated and maintained by AWS or AWS Marketplace Sellers. AWS WAF can be integrated with other AWS services, such as Amazon CloudFront, Amazon API Gateway, and Application Load Balancer, to provide comprehensive security for web applications. [AWS WAF Overview] AWS Certified Cloud Practitioner - aws.amazon.com

**NEW QUESTION 64**

- (Topic 1)

When a user wants to utilize their existing per-socket, per-core, or per-virtual machine software licenses for a Microsoft Windows server running on AWS, which Amazon EC2 instance type is required?

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

**Answer: C**

**Explanation:**

The correct answer is C because Dedicated Hosts are Amazon EC2 instances that are required when a user wants to utilize their existing per-socket, per-core, or per-virtual machine software licenses for a Microsoft Windows server running on AWS. Dedicated Hosts are physical servers that are dedicated to a single customer. Dedicated Hosts allow customers to use their existing server-bound software licenses, such as Windows Server, SQL Server, and SUSE Linux Enterprise Server, subject to their license terms. The other options are incorrect because they are not Amazon EC2 instances that are required when a user wants to utilize their existing per-socket, per-core, or per-virtual machine software licenses for a Microsoft Windows server running on AWS. Spot Instances are spare Amazon EC2 instances that are available at up to 90% discount compared to On-Demand prices. Spot Instances are suitable for stateless, fault-tolerant, and flexible workloads that can recover from interruptions easily. Dedicated Instances are Amazon EC2 instances that run on hardware that is dedicated to a single customer, but not to a specific physical server. Dedicated Instances do not allow customers to use their existing server-bound software licenses. Reserved Instances are Amazon EC2 instances that are reserved for a specific period of time (one or three years) in exchange for a lower hourly rate. Reserved Instances

are suitable for steady-state or predictable workloads that run for a long duration. Reserved Instances do not allow customers to use their existing server-bound software licenses. Reference: Dedicated Hosts, Amazon EC2 Instance Purchasing Options

**NEW QUESTION 67**

- (Topic 1)

A company is hosting a web application in a Docker container on Amazon EC2. AWS is responsible for which of the following tasks?

- A. Scaling the web application and services developed with Docker
- B. Provisioning or scheduling containers to run on clusters and maintain their availability
- C. Performing hardware maintenance in the AWS facilities that run the AWS Cloud
- D. Managing the guest operating system, including updates and security patches

**Answer: C**

**Explanation:**

AWS is responsible for performing hardware maintenance in the AWS facilities that run the AWS Cloud. This is part of the shared responsibility model, where AWS is responsible for the security of the cloud, and the customer is responsible for security in the cloud. AWS is also responsible for the global infrastructure that runs all of the services offered in the AWS Cloud, including the hardware, software, networking, and facilities that run AWS Cloud services<sup>3</sup>. The customer is responsible for the guest operating system, including updates and security patches, as well as the web application and services developed with Docker<sup>4</sup>.

**NEW QUESTION 68**

- (Topic 1)

A company needs to store data across multiple Availability Zones in an AWS Region. The data will not be accessed regularly but must be immediately retrievable.

Which Amazon Elastic File System (Amazon EFS) storage class meets these requirements MOST cost effectively?

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

**Answer: B**

**Explanation:**

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

**NEW QUESTION 69**

- (Topic 1)

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

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

**Answer: B**

**Explanation:**

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

**NEW QUESTION 74**

- (Topic 1)

Which AWS service or tool provides users with the ability to monitor AWS service quotas?

- A. AWS CloudTrail
- B. AWS Cost and Usage Reports
- C. AWS Trusted Advisor
- D. AWS Budgets

**Answer: C**

**Explanation:**

The correct answer is C because AWS Trusted Advisor is an AWS service or tool that provides users with the ability to monitor AWS service quotas. AWS Trusted Advisor is an online tool that provides users with real-time guidance to help them provision their resources following AWS best practices. One of the categories of checks that AWS Trusted Advisor performs is service limits, which monitors the usage of each AWS service and alerts users when they are close to reaching the default limit. The other options are incorrect because they are not AWS services or tools that provide users with the ability to monitor AWS service quotas. AWS CloudTrail is a service that enables users to track user activity and API usage across their AWS account. AWS Cost and Usage Reports is a tool that enables users to access comprehensive information about their AWS costs and usage. AWS Budgets is a tool that enables users to plan their service usage, costs, and reservations. Reference: [AWS Trusted Advisor FAQs]

**NEW QUESTION 77**

- (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 81**

- (Topic 1)

A company has a social media platform in which users upload and share photos with other users. The company wants to identify and remove inappropriate photos. The company has no machine learning (ML) scientists and must build this detection capability with no ML expertise.

Which AWS service should the company use to build this capability?

- A. Amazon SageMaker
- B. Amazon Textract
- C. Amazon Rekognition
- D. Amazon Comprehend

**Answer:** C

**Explanation:**

Amazon Rekognition is the AWS service that the company should use to build the capability of identifying and removing inappropriate photos. Amazon Rekognition is a service that uses deep learning technology to analyze images and videos for various purposes, such as face detection, object recognition, text extraction, and content moderation. Amazon Rekognition can help users detect unsafe or inappropriate content in images and videos, such as nudity, violence, or drugs, and provide confidence scores for each label. Amazon Rekognition does not require any machine learning expertise, and users can easily integrate it with other AWS services

**NEW QUESTION 85**

- (Topic 1)

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

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

**Answer:** A

**Explanation:**

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

**NEW QUESTION 87**

- (Topic 1)

An auditor needs to find out whether a specific AWS service is compliant with specific compliance frameworks.

Which AWS service will provide this information?

- A. AWS Artifact
- B. AWS Trusted Advisor
- C. Amazon GuardDuty
- D. AWS Certificate Manager (ACM)

**Answer:** A

**Explanation:**

AWS Artifact is the service that will provide the information about whether a specific AWS service is compliant with specific compliance frameworks. AWS Artifact is a self-service portal that allows you to access, review, and download AWS security and compliance reports and agreements. You can use AWS Artifact to verify the compliance status of AWS services across various regions and compliance programs, such as ISO, PCI, SOC, FedRAMP, HIPAA, and more.

**NEW QUESTION 88**

- (Topic 1)

Which AWS Well-Architected Framework concept represents a system's ability to remain functional when the system encounters operational problems?

- A. Consistency
- B. Elasticity
- C. Durability
- D. Latency

**Answer:** B

**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 concept of elasticity represents a system's ability to adapt to changes in demand by scaling resources up or down automatically. Therefore, the correct answer is B. You can learn more about the AWS Well-Architected Framework and its pillars from this page.

**NEW QUESTION 89**

- (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 94**

- (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 98**

- (Topic 1)

What can a user accomplish using AWS CloudTrail?

- A. Generate an IAM user credentials report.
- B. Record API calls made to AWS services.
- C. Assess the compliance of AWS resource configurations with policies and guidelines.
- D. Ensure that Amazon EC2 instances are patched with the latest security update
- E. A company uses Amazon Workspaces.

**Answer:** B

**Explanation:**

AWS CloudTrail is an AWS service that enables users to accomplish the task of recording API calls made to AWS services. AWS CloudTrail is a service that tracks user activity and API usage across the AWS account. AWS CloudTrail records the details of every API call made to AWS services, such as the identity of the caller, the time of the call, the source IP address of the caller, the parameters and responses of the call, and more. Users can use AWS CloudTrail to audit, monitor, and troubleshoot their AWS resources and actions. The other options are incorrect because they are not tasks that users can accomplish using AWS CloudTrail. Generating an IAM user credentials report is a task that users can accomplish using IAM, which is an AWS service that enables users to manage access and permissions to AWS resources and services. Assessing the compliance of AWS resource configurations with policies and guidelines is a task that users can accomplish using AWS Config, which is an AWS service that enables users to assess, audit, and evaluate the configurations of their AWS resources. Ensuring that Amazon EC2 instances are patched with the latest security updates is a task that users can accomplish using AWS Systems Manager, which is an AWS service that enables users to automate operational tasks, manage configuration and compliance, and monitor system health and performance. Reference: AWS CloudTrail FAQs

**NEW QUESTION 103**

- (Topic 1)

What are the characteristics of Availability Zones? (Select TWO.)

- A. All Availability Zones in an AWS Region are interconnected with high-bandwidth, low- latency networking
- B. Availability Zones are physically separated by a minimum of distance of 150 km (100 miles).
- C. All traffic between Availability Zones is encrypted.
- D. Availability Zones within an AWS Region share redundant power, networking, and connectivity.
- E. Every Availability Zone contains a single data center.

**Answer:** AD

**Explanation:**

Availability Zones are physically separate locations within an AWS Region that are engineered to be isolated from failures. Each Availability Zone has independent power, cooling, and physical security, and is connected to other Availability Zones in the same Region by a low-latency network. Therefore, the correct answers are A and D. You can learn more about Availability Zones and their characteristics from this page.

**NEW QUESTION 107**

- (Topic 1)

Which AWS services or features can control VPC traffic? (Select TWO.)

- A. Security groups
- B. AWS Direct Connect
- C. Amazon GuardDuty
- D. Network ACLs
- E. Amazon Connect

**Answer:** AD

**Explanation:**

The AWS services or features that can control VPC traffic are security groups and network ACLs. Security groups are stateful firewalls that control the inbound and outbound traffic at the instance level. You can assign one or more security groups to each instance in a VPC, and specify the rules that allow or deny traffic based on the protocol, port, and source or destination. Network ACLs are stateless firewalls that control the inbound and outbound traffic at the subnet level. You can associate one network ACL with each subnet in a VPC, and specify the rules that allow or deny traffic based on the protocol, port, and source or destination. AWS Direct Connect, Amazon GuardDuty, and Amazon Connect are not services or features that can control VPC traffic. AWS Direct Connect is a service that establishes a dedicated network connection between your premises and AWS. Amazon GuardDuty is a service that monitors your AWS account and workloads for malicious or unauthorized activity. Amazon Connect is a service that provides a cloud-based contact center solution.

**NEW QUESTION 110**

- (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 114**

- (Topic 1)

Which of the following is an advantage of AWS Cloud computing?

- A. Trade security for elasticity.
- B. Trade operational excellence for agility.
- C. Trade fixed expenses for variable expenses.
- D. Trade elasticity for performance.

**Answer:** C

**Explanation:**

The correct answer is C because AWS Cloud computing allows customers to trade fixed expenses for variable expenses. This means that customers only pay for the resources they use, and can scale up or down as needed. The other options are incorrect because they are not advantages of AWS Cloud computing. Trade security for elasticity means that customers have to compromise on the protection of their data and applications in order to adjust their capacity quickly. Trade operational excellence for agility means that customers have to sacrifice the quality and reliability of their operations in order to respond to changing needs faster. Trade elasticity for performance means that customers have to limit their ability to scale up or down in order to achieve higher speed and efficiency. Reference: What is Cloud Computing?

**NEW QUESTION 118**

- (Topic 1)

A company's user base needs to remotely access virtual desktop computers from the internet. Which AWS service provides this functionality?

- A. Amazon Connect
- B. Amazon Cognito
- C. Amazon Workspaces
- D. Amazon Upstream 2.0

**Answer:** C

**Explanation:**

Amazon Workspaces is the AWS service that provides the functionality of remotely accessing virtual desktop computers from the internet. Amazon Workspaces is a fully managed, secure desktop-as-a-service (DaaS) solution that allows users to provision cloud-based virtual desktops and access them from anywhere, using any supported device. Amazon Workspaces helps users reduce the complexity and cost of managing and maintaining physical desktops, and provides a consistent and secure user experience.

**NEW QUESTION 121**

- (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.

**NEW QUESTION 122**

- (Topic 1)

Which AWS service meets this requirement?

- A. AWS CloudFormation
- B. AWS Elastic Beanstalk
- C. AWS Cloud9
- D. AWS CloudShell

**Answer:** A

**Explanation:**

AWS CloudFormation is a service that gives developers and businesses an easy way to create a collection of related AWS and third-party resources, and provision and manage them in an orderly and predictable fashion. You can use AWS CloudFormation's sample templates or create your own templates to describe the AWS and third-party resources, and any associated dependencies or runtime parameters, required to run your application.

**NEW QUESTION 124**

- (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 127**

- (Topic 1)

A company needs to migrate all of its development teams to a cloud-based integrated development environment (IDE). Which AWS service should the company use?

- A. AWS CodeBuild
- B. AWS Cloud9
- C. AWS OpsWorks
- D. AWS Cloud Development Kit (AWS CDK)

**Answer:** B

**Explanation:**

The correct answer is B because AWS Cloud9 is an AWS service that enables users to run their existing custom, nonproduction workloads in the AWS Cloud quickly and cost-effectively. AWS Cloud9 is a cloud-based integrated development environment (IDE) that allows users to write, run, and debug code from a web browser. AWS Cloud9 supports multiple programming languages, such as Python, Java, Node.js, and more. AWS Cloud9 also provides users with a terminal that can access AWS services and resources, such as Amazon EC2 instances, AWS Lambda functions, and AWS CloudFormation stacks. The other options are incorrect because they are not AWS services that enable users to run their existing custom, nonproduction workloads in the AWS Cloud quickly and cost-effectively. AWS CodeBuild is an AWS service that enables users to compile, test, and package their code for deployment. AWS OpsWorks is an AWS service that enables users to configure and manage their applications using Chef or Puppet. AWS Cloud Development Kit (AWS CDK) is an AWS service that enables users to define and provision their cloud infrastructure using familiar programming languages, such as TypeScript, Python, Java, and C#. Reference: AWS Cloud9 FAQs

**NEW QUESTION 132**

- (Topic 1)

Which of the following are components of an AWS Site-to-Site VPN connection? (Select TWO.)

- A. AWS Storage Gateway
- B. Virtual private gateway
- C. NAT gateway
- D. Customer gateway
- E. Internet gateway

**Answer:** BD

**Explanation:**

The correct answers are B and D because a virtual private gateway and a customer gateway are components of an AWS Site-to-Site VPN connection. A virtual private gateway is the AWS side of the VPN connection that attaches to the customer's VPC. A customer gateway is the customer side of the VPN connection that resides in the customer's network. The other options are incorrect because they are not components of an AWS Site-to-Site VPN connection. AWS Storage Gateway is a service that connects on-premises software applications with cloud-based storage. NAT gateway is a service that enables instances in a private subnet to connect to the internet or other AWS services, but prevents the internet from initiating a connection with those instances. Internet gateway is a service that enables communication between instances in a VPC and the internet. Reference: [What is AWS Site-to-Site VPN?]

**NEW QUESTION 133**

- (Topic 1)

What are some advantages of using Amazon EC2 instances to host applications in the AWS Cloud instead of on premises? (Select TWO.)

- A. EC2 includes operating system patch management
- B. EC2 integrates with Amazon VPC
- C. AWS CloudTrail, and AWS Identity and Access Management (IAM)
- D. EC2 has a 100% service level agreement (SLA).
- E. EC2 has a flexible, pay-as-you-go pricing model.
- F. EC2 has automatic storage cost optimization.

**Answer:** BD

**Explanation:**

Some of the advantages of using Amazon EC2 instances to host applications in the AWS Cloud instead of on premises are:

? EC2 integrates with Amazon VPC, AWS CloudTrail, and AWS Identity and Access Management (IAM). Amazon VPC lets you provision a logically isolated section of the AWS Cloud where you can launch AWS resources in a virtual network that you define. AWS CloudTrail enables governance, compliance, operational auditing, and risk auditing of your AWS account. AWS IAM enables you to manage access to AWS services and resources securely. Therefore, the correct answer is B. You can learn more about Amazon EC2 and its integration with other AWS services from this page.

? EC2 has a flexible, pay-as-you-go pricing model. You only pay for the compute capacity you use, and you can scale up and down as needed. You can also choose from different pricing options, such as On-Demand, Savings Plans, Reserved Instances, and Spot Instances, to optimize your costs. Therefore, the correct answer is D. You can learn more about Amazon EC2 pricing from this page.

The other options are incorrect because:

? EC2 does not include operating system patch management. You are responsible for managing and maintaining your own operating systems on EC2 instances. You can use AWS Systems Manager to automate common maintenance tasks, such as applying patches, or use Amazon EC2 Image Builder to create and maintain secure images. Therefore, the incorrect answer is A.

? EC2 does not have a 100% service level agreement (SLA). The EC2 SLA guarantees 99.99% availability for each EC2 Region, not for each individual instance. Therefore, the incorrect answer is C.

? EC2 does not have automatic storage cost optimization. You are responsible for choosing the right storage option for your EC2 instances, such as Amazon Elastic Block Store (EBS) or Amazon Elastic File System (EFS), and monitoring and optimizing your storage costs. You can use AWS Cost Explorer or AWS Trusted Advisor to analyze and reduce your storage spending. Therefore, the incorrect answer is E.

**NEW QUESTION 134**

- (Topic 1)

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

- A. Operational excellence
- B. Security
- C. Reliability
- D. Cost optimization

**Answer:** A

**Explanation:**

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

**NEW QUESTION 136**

- (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 140**

- (Topic 1)

A company is running applications on Amazon EC2 instances in the same AWS account for several different projects. The company wants to track the infrastructure costs for each of the projects separately. The company must conduct this tracking with the least possible impact to the existing infrastructure and with no additional cost.

What should the company do to meet these requirements?

- A. Use a different EC2 instance type for each project.
- B. Publish project-specific custom Amazon CloudWatch metrics for each application.
- C. Deploy EC2 instances for each project in a separate AWS account.
- D. Use cost allocation tags with values that are specific to each project.

**Answer: D**

**Explanation:**

The correct answer is D because cost allocation tags are a way to track the infrastructure costs for each of the projects separately. Cost allocation tags are key-value pairs that can be attached to AWS resources, such as EC2 instances, and used to categorize and group them for billing purposes. The other options are incorrect because they do not meet the requirements of the question. Use a different EC2 instance type for each project does not help to track the costs for each project, and may impact the performance and compatibility of the applications. Publish project-specific custom Amazon CloudWatch metrics for each application does not help to track the costs for each project, and may incur additional charges for using CloudWatch. Deploy EC2 instances for each project in a separate AWS account does help to track the costs for each project, but it impacts the existing infrastructure and incurs additional charges for using multiple accounts.

Reference: Using Cost Allocation Tags

**NEW QUESTION 144**

- (Topic 1)

Which AWS service should a cloud practitioner use to receive real-time guidance for provisioning resources, based on AWS best practices related to security, cost optimization, and service limits?

- A. AWS Trusted Advisor
- B. AWS Config
- C. AWS Security Hub
- D. AWS Systems Manager

**Answer: A**

**Explanation:**

AWS Trusted Advisor is the AWS service that provides real-time guidance for provisioning resources, based on AWS best practices related to security, cost optimization, and service limits. AWS Trusted Advisor inspects the user's AWS environment and provides recommendations for improving performance, security, and reliability, reducing costs, and following best practices. AWS Trusted Advisor also alerts the user when they are approaching or exceeding their service limits, and helps them request limit increases.

**NEW QUESTION 149**

- (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 154**

- (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 155**

- (Topic 1)

A user wants to identify any security group that is allowing unrestricted incoming SSH traffic.

Which AWS service can be used to accomplish this goal?

- A. Amazon Cognito
- B. AWS Shield
- C. Amazon Macie
- D. AWS Trusted Advisor

**Answer: D**

**Explanation:**

The correct answer to the question is D because AWS Trusted Advisor is an AWS service that can be used to accomplish the goal of identifying any security group that is allowing unrestricted incoming SSH traffic. AWS Trusted Advisor is a service that provides customers with recommendations that help them follow AWS best practices. Trusted Advisor evaluates the customer's AWS environment and identifies ways to optimize their AWS infrastructure, improve security and performance, reduce costs, and monitor service quotas. One of the checks that Trusted Advisor performs is the Security Groups - Specific Ports Unrestricted check, which flags security groups that allow unrestricted access to specific ports, such as port 22 for SSH. Customers can use this check to review and modify their security group rules to restrict SSH access to only authorized sources. Reference: Security Groups - Specific Ports Unrestricted

**NEW QUESTION 159**

- (Topic 1)

Which of the following is an advantage that users experience when they move on-premises workloads to the AWS Cloud?

- A. Elimination of expenses for running and maintaining data centers
- B. Price discounts that are identical to discounts from hardware providers
- C. Distribution of all operational controls to AWS
- D. Elimination of operational expenses

**Answer: A**

**Explanation:**

The advantage that users experience when they move on-premises workloads to the AWS Cloud is: elimination of expenses for running and maintaining data centers. By moving on-premises workloads to the AWS Cloud, users can reduce or eliminate the costs associated with owning and operating physical servers, storage, network equipment, and facilities. These costs include hardware purchase, maintenance, repair, power, cooling, security, and staff. Users can also benefit from the pay-as-you-go pricing model of AWS, which allows them to pay only for the resources they use, and scale up or down as needed.

**NEW QUESTION 162**

- (Topic 1)

Which of the following is a benefit of decoupling an AWS Cloud architecture?

- A. Reduced latency
- B. Ability to upgrade components independently
- C. Decreased costs
- D. Fewer components to manage

**Answer: B**

**Explanation:**

A benefit of decoupling an AWS Cloud architecture is the ability to upgrade components independently. Decoupling is a way of designing systems to reduce interdependencies and minimize the impact of changes. Decoupling allows components to interact with each other through well-defined interfaces, rather than direct references. This reduces the risk of failures and errors propagating across the system, and enables greater scalability, availability, and maintainability. By decoupling an AWS Cloud architecture, the user can upgrade or modify one component without affecting the other components<sup>5</sup>.

**NEW QUESTION 165**

- (Topic 3)

Which task can only an AWS account root user perform?

- A. Changing the AWS Support plan
- B. Deleting AWS resources

- C. Creating an Amazon EC2 instance key pair
- D. Configuring AWS WAF

**Answer:** A

**Explanation:**

The AWS account root user is the email address that you use to sign up for AWS. The root user has complete access to all AWS services and resources in the account. The root user can perform tasks that only the root user can do, such as changing the AWS Support plan, closing the account, and restoring IAM user permissions<sup>34</sup>

**NEW QUESTION 168**

- (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<sup>5</sup>. 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 169**

- (Topic 3)

A company needs to securely store important credentials that an application uses to connect users to a database. Which AWS service can meet this requirement with the MINIMAL amount of operational overhead?

- A. AWS Key Management Service (AWS KMS)
- B. AWS Config
- C. AWS Secrets Manager
- D. Amazon GuardDuty

**Answer:** C

**Explanation:**

AWS Secrets Manager is a service that helps you protect secrets needed to access your applications, services, and IT resources. You can use AWS Secrets Manager to store, rotate, and retrieve database credentials, API keys, and other secrets throughout their lifecycle. AWS Secrets Manager eliminates the need to hardcode sensitive information in plain text, and reduces the risk of unauthorized access or leakage. AWS Secrets Manager also integrates with other AWS services, such as AWS Lambda, Amazon RDS, and AWS CloudFormation, to simplify the management of secrets across your environment<sup>5</sup>

**NEW QUESTION 174**

- (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 175**

- (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 180**

- (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 184**

- (Topic 3)

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

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

**Answer: B**

**Explanation:**

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

**NEW QUESTION 189**

- (Topic 2)

A company wants to migrate to the AWS Cloud. The company needs the ability to acquire resources when the resources are necessary. The company also needs the ability to release those resources when the resources are no longer necessary. Which architecture concept of the AWS Cloud meets these requirements?

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

**Answer: A**

**Explanation:**

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

**NEW QUESTION 190**

- (Topic 2)

A company wants to migrate its application to AWS. The company wants to replace upfront expenses with variable payment that is based on usage. What should the company do to meet these requirements?

- A. Use pay-as-you-go pricing.
- B. Purchase Reserved Instances.
- C. Pay less by using more.
- D. Rightsize instances.

**Answer:** A

**Explanation:**

Pay-as-you-go pricing is one of the main benefits of AWS. With pay-as-you-go pricing, you pay only for what you use, when you use it. There are no long-term contracts, termination fees, or complex licensing. You replace upfront expenses with lower variable costs and pay only for the resources you consume.

**NEW QUESTION 192**

- (Topic 2)

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

- A. Patch management
- B. Cost optimization
- C. Business technology strategy
- D. Physical and environmental controls

**Answer:** B

**Explanation:**

The AWS Well-Architected Framework helps you understand the pros and cons of decisions you make while building systems on AWS. By using the Framework, you will learn architectural best practices for designing and operating reliable, secure, efficient, and cost-effective systems in the cloud. The Framework consists of five pillars: operational excellence, security, reliability, performance efficiency, and cost optimization<sup>2</sup>.

**NEW QUESTION 195**

- (Topic 2)

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

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

**Answer:** C

**Explanation:**

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

**NEW QUESTION 196**

- (Topic 2)

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

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

**Answer:** B

**Explanation:**

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

**NEW QUESTION 197**

- (Topic 2)

Which design principles should a company apply to AWS Cloud workloads to maximize sustainability and minimize environmental impact? (Select TWO.)

- A. Maximize utilization of Amazon EC2 instances.
- B. Minimize utilization of Amazon EC2 instances.
- C. Minimize usage of managed services.
- D. Force frequent application reinstallations by users.
- E. Reduce the need for users to reinstall applications.

**Answer:** AE

**Explanation:**

To maximize sustainability and minimize environmental impact, a company should apply the following design principles to AWS Cloud workloads: maximize utilization of Amazon EC2 instances and reduce the need for users to reinstall applications. Maximizing utilization of Amazon EC2 instances means that the company can optimize the performance and efficiency of their compute resources, and avoid wasting energy and money on idle or underutilized instances. The company can use features such as Amazon EC2 Auto Scaling, Amazon EC2 Spot Instances, and AWS Compute Optimizer to automatically adjust the number and type of instances based on demand, cost, and performance. Reducing the need for users to reinstall applications means that the company can minimize the amount of data and bandwidth required to deliver their applications to users, and avoid unnecessary downloads and updates that consume energy and resources. The company can use services such as Amazon CloudFront, AWS AppStream 2.0, and AWS Amplify to deliver their applications faster, more securely, and more efficiently to users across the globe. Minimizing utilization of Amazon EC2 instances, minimizing usage of managed services, and forcing frequent application reinstallations by users are not design principles that would maximize sustainability and minimize environmental impact. Minimizing utilization of Amazon EC2 instances would reduce the performance and efficiency of the compute resources, and potentially increase the costs and complexity of the cloud workloads. Minimizing usage of managed services would increase the operational overhead and responsibility of the company, and potentially expose them to more security and reliability risks. Forcing frequent application reinstallations by users would increase the amount of data and bandwidth required to deliver the applications to users, and potentially degrade the user experience and satisfaction.

**NEW QUESTION 201**

- (Topic 2)

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

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

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

**Answer: B**

**Explanation:**

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

**NEW QUESTION 203**

- (Topic 2)

Which AWS service is used to temporarily provide federated security credentials to a

- A. Amazon GuardDuty
- B. AWS Simple Token Service (AWS STS)
- C. AWS Secrets Manager
- D. AWS Certificate Manager

**Answer: B**

**Explanation:**

The AWS service that is used to temporarily provide federated security credentials to a user is AWS Security Token Service (AWS STS). AWS STS is a service that enables customers to request temporary, limited-privilege credentials for AWS Identity and Access Management (IAM) users or for users that they authenticate (federated users). The company can use AWS STS to grant federated users access to AWS resources without creating permanent IAM users or sharing long-term credentials. AWS STS helps customers manage and secure access to their AWS resources for federated users. Amazon GuardDuty, AWS Secrets Manager, and AWS Certificate Manager 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 Secrets Manager is a service that helps customers manage and rotate secrets, such as database credentials, API keys, and passwords. AWS Certificate Manager is a service that helps customers provision, manage, and deploy public and private Secure Sockets Layer/Transport Layer Security (SSL/TLS) certificates for use with AWS services and internal connected resources. These services are more useful for different types of security and compliance tasks, rather than providing temporary federated security credentials to a user.

**NEW QUESTION 204**

- (Topic 2)

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

Which solution will meet these requirements?

- A. Use EC2 instances across multiple Availability Zones in the same AWS Region.
- B. Use Amazon CloudFront as the database for the EC2 instances.
- C. Use EC2 instances in the same edge location and the same Availability Zone.
- D. Use EC2 instances in AWS OpsWorks stacks in different AWS Regions.

**Answer: A**

**Explanation:**

Using EC2 instances across multiple Availability Zones in the same AWS Region is a solution that meets the requirements of sharing the same geographic area but using redundant underlying power sources. Availability Zones are isolated locations within an AWS Region that have independent power, cooling, and physical security. They are connected through low-latency, high-throughput, and highly redundant networking. By launching EC2 instances in different Availability Zones, users can increase the fault tolerance and availability of their applications. Amazon CloudFront is a content delivery network (CDN) service that speeds up the delivery of web content and media to end users by caching it at the edge locations closer to them. It is not a database service and cannot be used to store operational data for EC2 instances. Edge locations are sites that are part of the Amazon CloudFront network and are located in many cities around the world. They are not the same as Availability Zones and do not provide redundancy for EC2 instances. AWS OpsWorks is a configuration management service that allows users to automate the deployment and management of applications using Chef or Puppet. It can be used to create stacks that span multiple AWS Regions, but this would not meet the requirement of sharing the same geographic area.

**NEW QUESTION 207**

- (Topic 2)

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

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

**Answer: C**

**Explanation:**

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

**NEW QUESTION 212**

- (Topic 2)

A company has an environment that includes Amazon EC2 instances, Amazon Lightsail, and on-premises servers. The company wants to automate the security updates for its operating systems and applications.

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

- A. Use AWS Shield to identify and manage security events.
- B. Connect to each server by using a remote desktop connectio
- C. Run an update script.
- D. Use the AWS Systems Manager Patch Manager capability.
- E. Schedule Amazon GuardDuty to run on a nightly basis.

**Answer: C**

**Explanation:**

AWS Systems Manager Patch Manager is a capability that allows users to automate the security updates for their operating systems and applications. It enables users to scan their instances for missing patches, define patch baselines, schedule patching windows, and monitor patch compliance. It supports Amazon EC2 instances, Amazon Lightsail instances, and on-premises servers. AWS Shield is a service that provides protection against Distributed Denial of Service (DDoS) attacks for AWS resources and services. It does not automate the security updates for operating systems and applications. Connecting to each server by using a remote desktop connection and running an update script is a manual and time-consuming solution that requires a lot of operational effort. It is not a recommended best practice for automating the security updates for operating systems and applications. Amazon GuardDuty is a service that provides intelligent threat detection and continuous monitoring for AWS accounts and resources. It does not automate the security updates for operating systems and applications.

**NEW QUESTION 213**

- (Topic 2)

A company is using AWS Organizations to configure AWS accounts.

A company is planning its migration to the AWS Cloud. The company is identifying its capability gaps by using the AWS Cloud Adoption Framework (AWS CAF) perspectives.

Which phase of the cloud transformation journey includes these identification activities?

- A. Envision
- B. Align
- C. Scale
- D. Launch

**Answer: A**

**Explanation:**

The Envision phase of the cloud transformation journey is where the company defines its vision, business drivers, and desired outcomes for the cloud adoption. The company also identifies its capability gaps by using the AWS Cloud Adoption Framework (AWS CAF) perspectives, which are business, people, governance, platform, security, and operations2.

**NEW QUESTION 216**

- (Topic 2)

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:**

The AWS service or tool that provides recommendations to help users get rightsized Amazon EC2 instances based on historical workload usage data is AWS Compute Optimizer. AWS Compute Optimizer is a service that analyzes the configuration and performance of the AWS resources, such as Amazon EC2 instances, and provides recommendations for optimal resource types and sizes based on the workload patterns and metrics. AWS Compute Optimizer helps users improve the performance, availability, and cost efficiency of their AWS resources. AWS Pricing Calculator, AWS App Runner, and AWS Systems Manager are not the best services or tools to use for this purpose. AWS Pricing Calculator is a tool that helps users estimate the cost of using AWS services based on their

requirements and preferences. AWS App Runner is a service that helps users easily and quickly deploy web applications and APIs without managing any infrastructure. AWS Systems Manager is a service that helps users automate and manage the configuration and operation of their AWS resources and applications<sup>34</sup>

**NEW QUESTION 219**

- (Topic 2)

A company needs to centralize its operational data. The company also needs to automate tasks across all of its Amazon EC2 instances. Which AWS service can the company use to meet these requirements?

- A. AWS Trusted Advisor
- B. AWS Systems Manager
- C. AWS CodeDeploy
- D. AWS Elastic Beanstalk

**Answer: B**

**Explanation:**

AWS Systems Manager is a service that enables users to centralize and automate the management of their AWS resources. It provides a unified user interface to view operational data, such as inventory, patch compliance, and performance metrics. It also allows users to automate common and repetitive tasks, such as patching, backup, and configuration management, across all of their Amazon EC2 instances<sup>1</sup>. AWS Trusted Advisor is a service that provides best practices and recommendations to optimize the performance, security, and cost of AWS resources<sup>2</sup>. AWS CodeDeploy is a service that automates the deployment of code and applications to Amazon EC2 instances or other compute services<sup>3</sup>. AWS Elastic Beanstalk is a service that simplifies the deployment and management of web applications using popular platforms, such as Java, PHP, and Node.js<sup>4</sup>.

**NEW QUESTION 220**

- (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 221**

- (Topic 2)

A company needs to launch an Amazon EC2 instance.

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

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

**Answer: C**

**Explanation:**

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

**NEW QUESTION 226**

- (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 231**

- (Topic 2)

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

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

**Answer: D**

**Explanation:**

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

**NEW QUESTION 232**

- (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 233**

- (Topic 2)

A company has developed a distributed application that recovers gracefully from interruptions. The application periodically processes large volumes of data by using multiple Amazon EC2 instances. The application is sometimes idle for months.

Which EC2 instance purchasing option is MOST cost-effective for this use case?

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

**Answer: B**

**Explanation:**

Spot Instances are instances that use spare EC2 capacity that is available for up to 90% off the On-Demand price. Because Spot Instances can be interrupted by EC2 with two minutes of notification when EC2 needs the capacity back, you can use them for applications that have flexible start and end times, or that can withstand interruptions. This option is most cost-effective for the use case described in the question. Reserved Instances are instances that you purchase for a one-year or three-year term, and pay a lower hourly rate compared to On-Demand Instances. This option is suitable for applications that have steady state or predictable usage. Dedicated Instances are instances that run on hardware that's dedicated to a single customer within an Amazon VPC. This option is suitable for applications that have stringent regulatory or compliance requirements. On-Demand Instances are instances that you pay for by the second, with no long-term commitments or upfront payments. This option is suitable for applications that have unpredictable or intermittent workloads.

**NEW QUESTION 235**

- (Topic 2)

Which benefit of AWS Cloud computing provides lower latency between users and applications?

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

**Answer: C**

**Explanation:**

Global reach is the benefit of AWS Cloud computing that provides lower latency between users and applications. Global reach means that AWS customers can deploy their applications and data in multiple regions around the world, and deliver them to users with high performance and availability. AWS has the largest global infrastructure of any cloud provider, with 25 geographic regions and 81 Availability Zones, as well as 216 Points of Presence in 84 cities across 42 countries. Customers can choose the optimal locations for their applications and data based on their business requirements, such as compliance, data sovereignty, and customer proximity. Agility, economies of scale, and pay-as-you-go pricing are other benefits of AWS Cloud computing, but they do not directly provide lower latency between users and applications. Agility means that AWS customers can quickly and easily provision and scale up or down AWS resources as needed, without upfront costs or long-term commitments. Economies of scale means that AWS customers can benefit from the lower costs and higher efficiency that AWS achieves by operating at a massive scale and passing the savings to the customers. Pay-as-you-go pricing means that AWS customers only pay for the AWS resources they use, without any upfront costs or long-term contracts.

**NEW QUESTION 238**

- (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 243**

- (Topic 2)

Which options are common stakeholders for the AWS Cloud Adoption Framework (AWS CAF) platform perspective? (Select TWO.)

- A. Chief financial officers (CFOs)
- B. IT architects
- C. Chief information officers (CIOs)
- D. Chief data officers (CDOs)
- E. Engineers

**Answer:** BE

**Explanation:**

The common stakeholders for the AWS Cloud Adoption Framework (AWS CAF) platform perspective are IT architects and engineers. The AWS CAF is a guidance that helps organizations design and travel an accelerated path to successful cloud adoption. The AWS CAF organizes the cloud adoption process into six areas of focus, called perspectives, which are business, people, governance, platform, security, and operations. Each perspective is divided into capabilities, which are further divided into skills and responsibilities. The platform perspective focuses on the provisioning and management of the cloud infrastructure and services that support the business applications. The platform perspective capabilities are design, implementation, and optimization. The stakeholders for the platform perspective are the IT architects and engineers who are responsible for designing, implementing, and optimizing the cloud platform. Chief financial officers (CFOs), chief information officers (CIOs), and chief data officers (CDOs) are not the common stakeholders for the AWS CAF platform perspective. CFOs are the common stakeholders for the AWS CAF business perspective, which focuses on the value realization of the cloud adoption. CIOs are the common stakeholders for the AWS CAF governance perspective, which focuses on the alignment of the IT strategy and processes with the business strategy and goals. CDOs are the common stakeholders for the AWS CAF security perspective, which focuses on the protection of the information assets and systems in the cloud.

**NEW QUESTION 244**

- (Topic 2)

Which task can a company perform by using security groups in the AWS Cloud?

- A. Allow access to an Amazon EC2 instance through only a specific port.
- B. Deny access to malicious IP addresses at a subnet level.
- C. Protect data that is cached by Amazon CloudFront.
- D. Apply a stateless firewall to an Amazon EC2 instance.

**Answer:** A

**Explanation:**

Security groups are virtual firewalls that control the inbound and outbound traffic for Amazon EC2 instances. They can be used to allow access to an Amazon EC2 instance through only a specific port, such as port 22 for SSH or port 80 for HTTP. Security groups cannot deny access to malicious IP addresses at a subnet level, as they only allow or deny traffic based on the rules defined by the customer. To block malicious IP addresses, customers can use network ACLs, which are stateless firewalls that can be applied to subnets. Security groups cannot protect data that is cached by Amazon CloudFront, as they only apply to EC2 instances. To protect data that is cached by Amazon CloudFront, customers can use encryption, signed URLs, or signed cookies. Security groups are not stateless firewalls, as they track the state of the traffic and automatically allow the response traffic to flow back to the source. Stateless firewalls do not track the state of the traffic and require rules for both inbound and outbound traffic.

**NEW QUESTION 248**

- (Topic 2)

A company wants guidance to optimize the cost and performance of its current AWS environment. Which AWS service or tool should the company use to identify areas for optimization?

- A. Amazon QuickSight
- B. AWS Trusted Advisor
- C. AWS Organizations
- D. AWS Budgets

**Answer: B**

**Explanation:**

AWS Trusted Advisor is the AWS service or tool that the company should use to identify areas for optimization. According to the AWS Trusted Advisor User Guide, "AWS Trusted Advisor is an online tool that provides you real time guidance to help you provision your resources following AWS best practices. AWS Trusted Advisor checks help optimize your AWS infrastructure, increase security and performance, reduce your overall costs, and monitor service limits." Amazon QuickSight, AWS Organizations, and AWS Budgets are not designed to provide optimization recommendations for the current AWS environment.

**NEW QUESTION 250**

- (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 254**

- (Topic 2)

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

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

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

**Answer: B**

**Explanation:**

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

**NEW QUESTION 258**

- (Topic 2)

A manufacturing company has a critical application that runs at a remote site that has a slow internet connection. The company wants to migrate the workload to AWS. The application is sensitive to latency and interruptions in connectivity. The company wants a solution that can host this application with minimum latency. Which AWS service or feature should the company use to meet these requirements?

- A. Availability Zones
- B. AWS Local Zones
- C. AWS Wavelength
- D. AWS Outposts

**Answer: D**

**Explanation:**

AWS Outposts is a service that offers fully managed and configurable compute and storage racks built with AWS-designed hardware that allow you to run your workloads on premises and seamlessly connect to AWS services in the cloud. AWS Outposts is ideal for workloads that require low latency, local data processing, or local data storage. With AWS Outposts, you can use the same AWS APIs, tools, and infrastructure across on premises and the cloud to deliver a truly consistent hybrid experience. Availability Zones are isolated locations within each AWS Region that are engineered to be fault-tolerant and provide high availability. AWS Local Zones are extensions of AWS Regions that are placed closer to large population, industry, and IT centers where no AWS Region exists today. AWS Wavelength is a service that enables developers to build applications that deliver ultra-low latency to mobile devices and users by deploying AWS compute and storage at the edge of the 5G network. None of these services or features can help you host a critical application with minimum latency at a remote site that has a slow internet connection.

**NEW QUESTION 262**

- (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 266**

- (Topic 2)

How should the company deploy the application to meet these requirements?

- A. In a single Availability Zone
- B. On AWS Direct Connect
- C. On Reserved Instances
- D. In multiple Availability Zones

**Answer: D**

**Explanation:**

Deploying the application in multiple Availability Zones is the best way to ensure high availability for the application. Availability Zones are isolated locations within an AWS Region that are engineered to be fault-tolerant from failures in other Availability Zones. By deploying the application in multiple Availability Zones, the company can reduce the impact of outages and increase the resilience of the application. Deploying the application in a single Availability Zone, on AWS Direct Connect, or on Reserved Instances does not provide the same level of high availability as deploying the application in multiple Availability Zones. Source: Availability Zones

**NEW QUESTION 270**

- (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 275**

- (Topic 2)

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

What can the company do to meet this requirement?

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

**Answer: D**

**Explanation:**

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

Instances and [Cost Optimization with AWS].

**NEW QUESTION 279**

- (Topic 2)

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

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

**Answer: C**

**Explanation:**

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

**NEW QUESTION 282**

.....

## Thank You for Trying Our Product

\* 100% Pass or Money Back

All our products come with a 90-day Money Back Guarantee.

\* One year free update

You can enjoy free update one year. 24x7 online support.

\* Trusted by Millions

We currently serve more than 30,000,000 customers.

\* Shop Securely

All transactions are protected by VeriSign!

**100% Pass Your AWS-Certified-Cloud-Practitioner Exam with Our Prep Materials Via below:**

<https://www.certleader.com/AWS-Certified-Cloud-Practitioner-dumps.html>