

# Exam Questions AWS-Certified-Cloud-Practitioner

Amazon AWS Certified Cloud Practitioner

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



#### NEW QUESTION 1

- (Topic 2)

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<sup>1</sup>. 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<sup>1</sup>.

#### NEW QUESTION 2

- (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<sup>1</sup>

#### NEW QUESTION 3

- (Topic 2)

A company needs Amazon EC2 instances for a workload that can tolerate interruptions.

Which EC2 instance purchasing option meets this requirement with the LARGEST discount compared to On-Demand prices?

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

**Answer:** A

#### Explanation:

Spot Instances are spare Amazon EC2 instances that are available at up to 90% discount compared to On-Demand prices. They are suitable for workloads that can tolerate interruptions, such as batch processing, data analysis, and testing. Spot Instances are allocated based on the current supply and demand, and can be reclaimed by AWS with a two-minute notice when the demand exceeds the supply<sup>5</sup>. Convertible Reserved Instances are a type of Reserved Instances that provide a significant discount (up to 54%) compared to On-Demand prices and a capacity reservation for Amazon EC2 instances. They are available in 1-year or 3-year terms and allow users to change the instance family, size, operating system, or tenancy during the term. Standard Reserved Instances are another type of Reserved Instances that provide a larger discount (up to 75%) compared to On-Demand prices and a capacity reservation for Amazon EC2 instances. They are available in 1-year or 3-year terms and do not allow users to change the instance attributes during the term. Dedicated Hosts are physical servers with Amazon EC2 instance capacity fully dedicated to the user's use. They are suitable for users who have specific server- bound software licenses or compliance requirements.

#### NEW QUESTION 4

- (Topic 2)

Which AWS service provides a highly accurate and easy-to-use enterprise search service that is powered by machine learning (ML)?

- A. Amazon Kendra
- B. Amazon SageMaker
- C. Amazon Augmented AI (Amazon A2I)
- D. Amazon Polly

**Answer:** A

#### Explanation:

Amazon Kendra is a service that provides a highly accurate and easy-to-use enterprise search service that is powered by machine learning. Kendra delivers powerful natural language search capabilities to your websites and applications so your end users can more easily find the information they need within the vast amount of content spread across your company. Amazon SageMaker is a service that provides a fully managed platform for data scientists and developers to quickly and easily build, train, and deploy machine learning models at any scale. Amazon Augmented AI (Amazon A2I) is a service that makes it easy to build the workflows required for human review of ML predictions. Amazon A2I brings human review to all developers, removing the undifferentiated heavy lifting associated with building human review systems or managing large numbers of human reviewers. Amazon Polly is a service that turns text into lifelike speech, allowing you to create applications that talk, and build entirely new categories of speech-enabled products. None of these services provide an enterprise search service that is powered by machine learning.

#### 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)

A company wants to develop a shopping application that records customer orders. The application needs to use an AWS managed database service to store data.

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

- A. Amazon RDS
- B. Amazon Redshift
- C. Amazon ElastiCache
- D. Amazon Neptune

**Answer:** A

#### Explanation:

A is correct because Amazon RDS is the AWS service that provides a managed relational database service that supports various database engines, such as MySQL, PostgreSQL, Oracle, and SQL Server. B is incorrect because Amazon Redshift is the AWS service that provides a managed data warehouse service that is optimized for analytical queries. C is incorrect because Amazon ElastiCache is the AWS service that provides a managed in-memory data store service that supports Redis and Memcached. D is incorrect because Amazon Neptune is the AWS service that provides a managed graph database service that supports property graph and RDF models.

#### NEW QUESTION 9

- (Topic 2)

What does "security of the cloud" refer to in the AWS shared responsibility model?

- A. Availability of AWS services such as Amazon EC2
- B. Security of the cloud infrastructure that runs all the AWS services
- C. Implementation of password policies for IAM users
- D. Security of customer environments by using AWS Network Firewall partners

**Answer:** B

**Explanation:**

Security of the cloud refers to the security of the cloud infrastructure that runs all the AWS services. This includes the hardware, software, networking, and facilities that AWS operates and manages. AWS is responsible for protecting the security of the cloud as part of the AWS shared responsibility model. Availability of AWS services such as Amazon EC2 refers to the ability of the services to be up and running and to meet the expected performance. Availability is part of the reliability pillar of the AWS Well-Architected Framework and is a shared responsibility between AWS and the customer. Implementation of password policies for IAM users refers to the security of the customer data and applications in the cloud. This includes the configuration and management of IAM user permissions, encryption keys, security group rules, network ACLs, and other aspects of access management. The customer is responsible for protecting the security in the cloud as part of the AWS shared responsibility model. Security of customer environments by using AWS Network Firewall partners refers to the security of the customer data and applications in the cloud. AWS Network Firewall is a managed service that provides network protection for Amazon VPCs. It allows customers to use AWS Marketplace partners to implement firewall rules and policies. The customer is responsible for protecting the security in the cloud as part of the AWS shared responsibility model.

**NEW QUESTION 10**

- (Topic 2)

Which design principle is included in the operational excellence pillar of the AWS Well-Architected Framework?

- A. Create annotated documentation.
- B. Anticipate failure.
- C. Ensure performance efficiency.
- D. Optimize costs.

**Answer:** A

**Explanation:**

Create annotated documentation is the design principle that is included in the operational excellence pillar of the AWS Well-Architected Framework. According to the

AWS Well-Architected Framework whitepaper, creating annotated documentation means "documenting your workload so that the team understands the architecture, how to operate the workload, and how the workload delivers value to customers."3 Anticipate failure, ensure performance efficiency, and optimize costs are design principles that belong to other pillars of the AWS Well-Architected Framework, such as reliability, performance efficiency, and cost optimization.

**NEW QUESTION 10**

- (Topic 2)

A company is preparing to launch a redesigned website on AWS. Users from around the world will download digital handbooks from the website.

Which AWS solution should the company use to provide these static files securely?

- A. Amazon Kinesis Data Streams
- B. Amazon CloudFront with Amazon S3
- C. Amazon EC2 instances with an Application Load Balancer
- D. Amazon Elastic File System (Amazon EFS)

**Answer:** B

**Explanation:**

Amazon CloudFront with Amazon S3 is a solution that allows you to provide static files securely to users from around the world. Amazon CloudFront is a fast content delivery network (CDN) service that securely delivers data, videos, applications, and APIs to customers globally with low latency, high transfer speeds, all within a developer-friendly environment. Amazon S3 is an object storage service that offers industry-leading scalability, data availability, security, and performance. You can use Amazon S3 to store and retrieve any amount of data from anywhere. You can also configure Amazon S3 to work with Amazon CloudFront to distribute your content to edge locations near your users for faster delivery and lower latency. Amazon Kinesis Data Streams is a service that enables you to build custom applications that process or analyze streaming data for specialized needs. This option is not relevant for providing static files securely. Amazon EC2 instances with an Application Load Balancer is a solution that allows you to distribute incoming traffic across multiple targets, such as EC2 instances, in multiple Availability Zones. This option is suitable for dynamic web applications, but not necessary for static files. Amazon Elastic File System (Amazon EFS) is a service that provides a simple, scalable, fully managed elastic NFS file system for use with AWS Cloud services and on-premises resources. This option is not relevant for providing static files securely.

**NEW QUESTION 12**

- (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 Memcached1. It provides fast and scalable performance for applications that require high throughput and low latency1. Amazon DynamoDB is a fully managed NoSQL database service that provides consistent and single-digit millisecond latency at any scale2. Amazon EBS is a block storage service that provides persistent and durable storage volumes for Amazon EC2 instances3. Amazon Redshift is a fully managed data warehouse service that allows users to run complex analytic queries using SQL4.

### NEW QUESTION 13

- (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 15

- (Topic 2)

A company wants to access a report about the estimated environmental impact of the company's AWS usage. Which AWS service or feature should the company use to meet this requirement?

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

**Answer:** C

#### Explanation:

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

### NEW QUESTION 18

- (Topic 2)

Which AWS service offers a global content delivery network (CDN) that helps companies securely deliver websites, videos, applications, and APIs at high speeds with low latency?

- A. Amazon EC2
- B. Amazon CloudFront
- C. Amazon CloudWatch
- D. AWS CloudFormation

**Answer:** B

#### Explanation:

Amazon CloudFront is the AWS service that offers a global content delivery network (CDN) that helps companies securely deliver websites, videos, applications, and APIs at high speeds with low latency. Amazon CloudFront is a web service that speeds up distribution of static and dynamic web content, such as HTML, CSS, JavaScript, and image files, to users. Amazon CloudFront uses a global network of edge locations, located near users' geographic locations, to cache and serve content with high availability and performance. Amazon CloudFront also provides features such as AWS Shield for DDoS protection, AWS Certificate Manager for SSL/TLS encryption, AWS WAF for web application firewall, and AWS Lambda@Edge for customizing content delivery with serverless code. Amazon EC2, Amazon CloudWatch, and AWS CloudFormation are not services that offer a global CDN. Amazon EC2 is a service that provides scalable compute capacity in the cloud. Amazon CloudWatch is a service that provides monitoring and observability for AWS resources and applications. AWS CloudFormation is a service that provides a common language to model and provision AWS resources and their dependencies.

### NEW QUESTION 23

- (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 27**

- (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 28**

- (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 31**

- (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 35**

- (Topic 1)

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

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

**Answer:** A

**Explanation:**

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

**NEW QUESTION 39**

- (Topic 1)

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

- A. Availability
- B. Reliability
- C. Scalability
- D. Responsive design
- E. Operational excellence

**Answer:** BE

**Explanation:**

The correct answers to the questions are B and E because reliability and operational excellence are pillars of the AWS Well-Architected Framework. The AWS Well-Architected Framework is a set of best practices and guidelines for designing and operating reliable, secure, efficient, and cost-effective systems in the cloud. The AWS Well-Architected Framework consists of five pillars: operational excellence, security, reliability, performance efficiency, and cost optimization. Each pillar has a set of design principles that describe the characteristics of a well-architected system. Reliability is the pillar that focuses on the ability of a system to recover from failures and meet business and customer demand. Operational excellence is the pillar that focuses on the ability of a system to run and monitor processes that support business outcomes and continually improve. The other options are incorrect because they are not pillars of the AWS Well-Architected Framework. Availability, scalability, and responsive design are important aspects of cloud architecture, but they are not separate pillars in the framework. Availability and scalability are related to the reliability and performance efficiency pillars, while responsive design is related to the customer experience and user interface. Reference: AWS Well-Architected Framework

**NEW QUESTION 43**

- (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 48**

- (Topic 1)

Which of the following are design principles for reliability in the AWS Cloud? (Select TWO.)

- A. Build architectures with tightly coupled resources.
- B. Use AWS Trusted Advisor to meet security best practices.
- C. Use automation to recover immediately from failure.
- D. Rightsize Amazon EC2 instances to ensure optimal performance.
- E. Simulate failures to test recovery processes.

**Answer:** CE

**Explanation:**

The design principles for reliability in the AWS Cloud are:

? Test recovery procedures. The best way to ensure that systems can recover from failures is to regularly test them using simulated scenarios. This can help identify gaps and improve the recovery process.

? Automatically recover from failure. By using automation, systems can detect and correct failures without human intervention. This can reduce the impact and duration of failures and improve the availability of the system.

? Scale horizontally to increase aggregate system availability. By adding more redundant resources to the system, the impact of individual resource failures can be reduced. This can also improve the performance and scalability of the system.

? Stop guessing capacity. By using monitoring and automation, systems can adjust the capacity based on the demand and performance metrics. This can prevent failures due to insufficient or excessive capacity and optimize the cost and efficiency of the system.

? Manage change in automation. By using automation, changes to the system can be applied in a consistent and controlled manner. This can reduce the risk of human errors and configuration drifts that can cause failures. AWS Well-Architected Framework

**NEW QUESTION 52**

- (Topic 1)

Which of the following is an AWS value proposition that describes a user's ability to scale infrastructure based on demand?

- A. Speed of innovation
- B. Resource elasticity
- C. Decoupled architecture

D. Global deployment

**Answer:** B

**Explanation:**

Resource elasticity is an AWS value proposition that describes a user's ability to scale infrastructure based on demand. Resource elasticity means that the user can provision or deprovision resources quickly and easily, without any upfront commitment or long-term contract. Resource elasticity can help the user optimize the cost and performance of the application, as well as respond to changing business needs and customer expectations. Resource elasticity can be achieved by using services such as Amazon EC2, Amazon S3, Amazon RDS, Amazon DynamoDB, Amazon ECS, and AWS Lambda. [AWS Cloud Value Framework] AWS Certified Cloud Practitioner - aws.amazon.com

**NEW QUESTION 57**

- (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 59**

- (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 60**

- (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 63**

- (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 67**

- (Topic 1)

A company wants to centrally manage security policies and billing services within a multi- account AWS environment. Which AWS service should the company use to meet these requirements?

- A. AWS Identity and Access Management (IAM)
- B. AWS Organizations
- C. AWS Resource Access Manager (AWS RAM)
- D. AWS Config

**Answer:** B

**Explanation:**

AWS Organizations is a service that helps you centrally manage and govern your environment as you grow and scale your AWS resources. You can use AWS Organizations to create groups of accounts and apply policies to them. You can also use AWS Organizations to consolidate billing for multiple accounts. Therefore, the correct answer is B. You can learn more about AWS Organizations and its features from this page.

**NEW QUESTION 69**

- (Topic 1)

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

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

**Answer:** AD

**Explanation:**

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

**NEW QUESTION 70**

- (Topic 1)

Which database engine is compatible with Amazon RDS?

- A. Apache Cassandra
- B. MongoDB
- C. Neo4j
- D. PostgreSQL

**Answer:** D

**Explanation:**

Amazon RDS supports six database engines: Amazon Aurora, MySQL, MariaDB, PostgreSQL, Oracle, and SQL Server. Apache Cassandra, MongoDB, and Neo4j are not compatible with Amazon RDS. Therefore, the correct answer is D. You can learn more about Amazon RDS and its supported database engines from this page.

**NEW QUESTION 75**

- (Topic 1)

Which activity can companies complete by using AWS Organizations?

- A. Troubleshoot the performance of applications.
- B. Manage service control policies (SCPs).
- C. Migrate applications to microservices.
- D. Monitor the performance of applications.

**Answer:** B

**Explanation:**

Managing service control policies (SCPs) is an activity that companies can complete by using AWS Organizations. AWS Organizations is a service that enables the user to consolidate multiple AWS accounts into an organization that can be managed as a single unit. AWS Organizations allows the user to create groups of accounts and apply policies to them, such as service control policies (SCPs) that specify the services and actions that users and roles can access in the accounts. AWS Organizations also enables the user to use consolidated billing, which combines the usage and charges from all the accounts in the organization into a single bill<sup>3</sup>.

#### NEW QUESTION 79

- (Topic 1)

Which task is the responsibility of a company that is using Amazon RDS?

- A. Provision the underlying infrastructure.
- B. Create IAM policies to control administrative access to the service.
- C. Install the cables to connect the hardware for compute and storage.
- D. Install and patch the RDS operating system.

**Answer:** B

#### Explanation:

The correct answer is B because AWS IAM policies can be used to control administrative access to the Amazon RDS service. The other options are incorrect because they are the responsibilities of AWS, not the company that is using Amazon RDS. AWS manages the provisioning, cabling, installation, and patching of the underlying infrastructure for Amazon RDS. Reference: Amazon RDS FAQs

#### NEW QUESTION 81

- (Topic 1)

Which of the following is a cloud benefit that AWS offers to its users?

- A. The ability to configure AWS data center hypervisors
- B. The ability to purchase hardware in advance of increased traffic
- C. The ability to deploy to AWS on a global scale
- D. Compliance audits for user IT environments

**Answer:** C

#### Explanation:

The ability to deploy to AWS on a global scale is a cloud benefit that AWS offers to its users. AWS has a global infrastructure that consists of AWS Regions, Availability Zones, and edge locations. Users can choose from multiple AWS Regions around the world to deploy their applications and data closer to their end users, while also meeting their compliance and regulatory requirements. Users can also leverage AWS services, such as Amazon CloudFront, Amazon Route 53, and AWS Global Accelerator, to improve the performance and availability of their global applications. AWS also provides tools and guidance to help users optimize their global deployments, such as AWS Well- Architected Framework, AWS CloudFormation, and AWS Migration Hub. AWS Global Infrastructure [AWS Cloud Value Framework] AWS Certified Cloud Practitioner - aws.amazon.com

#### NEW QUESTION 82

- (Topic 1)

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

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

**Answer:** C

#### Explanation:

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

#### NEW QUESTION 87

- (Topic 1)

Which best practice for cost governance does this example show?

- A. Resource controls
- B. Cost allocation
- C. Architecture optimization
- D. Tagging enforcement

**Answer:** C

#### Explanation:

Architecture optimization is the best practice for cost governance that this example shows. Architecture optimization is the process of designing and implementing AWS solutions that are efficient, scalable, and cost-effective. By using specific AWS services to improve efficiency and reduce cost, the company is following the architecture optimization best practice. Some of the techniques for architecture optimization include using the right size and type of resources, leveraging elasticity and scalability, choosing the most suitable storage class, and using serverless and managed services<sup>2</sup>.

#### NEW QUESTION 89

- (Topic 1)

Which activity is a customer responsibility in the AWS Cloud according to the AWS shared responsibility model?

- A. Ensuring network connectivity from AWS to the internet
- B. Patching and fixing flaws within the AWS Cloud infrastructure
- C. Ensuring the physical security of cloud data centers
- D. Ensuring Amazon EBS volumes are backed up

**Answer:** D

**Explanation:**

The AWS shared responsibility model describes how AWS and the customer share responsibility for security and compliance of the AWS environment. AWS is responsible for the security of the cloud, which includes the physical security of AWS facilities, the infrastructure, hardware, software, and networking that run AWS services. The customer is responsible for security in the cloud, which includes the configuration of security groups, the encryption of customer data on AWS, the management of AWS Lambda infrastructure, and the management of network throughput of each AWS Region. One of the customer responsibilities is to ensure that Amazon EBS volumes are backed up.

**NEW QUESTION 92**

- (Topic 1)

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

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

**Answer:** C

**Explanation:**

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

**NEW QUESTION 93**

- (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 98**

- (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 101**

SIMULATION - (Topic 1)

A company runs thousands of simultaneous simulations using AWS Batch. Each simulation is stateless, is fault tolerant, and runs for up to 3 hours.

Which pricing model enables the company to optimize costs and meet these requirements?

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

**Answer:** B

**Explanation:**

The correct answer is B because Spot Instances enable the company to optimize costs and meet the requirements. Spot Instances are spare 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 applications that can run for any duration. The other options are incorrect because they do not enable the company to optimize costs and meet the requirements. Reserved Instances are 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. On-Demand Instances are EC2 instances that are launched and billed at a fixed hourly rate. On-Demand Instances are suitable for short-term, irregular, or unpredictable workloads that cannot be interrupted. Dedicated Instances are EC2 instances that run on hardware that is dedicated to a single customer. Dedicated Instances are suitable for workloads that require regulatory compliance or data isolation. Reference: [Amazon EC2 Instance Purchasing Options]

**NEW QUESTION 103**

- (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 105**

- (Topic 1)

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

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

**Answer:** B

**Explanation:**

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

**NEW QUESTION 106**

- (Topic 1)

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

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

**Answer:** C

**Explanation:**

The correct answer is C because Amazon QuickSight is an AWS service that will provide the dashboards and charts for the insights from business data. Amazon QuickSight is a fully managed, scalable, and serverless business intelligence service that enables users to create and share interactive dashboards and charts. Amazon QuickSight can connect to various data sources, such as Amazon S3, Amazon RDS, Amazon Redshift, and more. Amazon QuickSight also provides users with machine learning insights, such as anomaly detection, forecasting, and natural language narratives. The other options are incorrect because they are not AWS services that will provide the dashboards and charts for the insights from business data. Amazon Macie is an AWS service that helps users discover, classify, and protect sensitive data stored in Amazon S3. Amazon Aurora is an AWS service that provides a relational database that is compatible with MySQL and PostgreSQL. AWS CloudTrail is an AWS service that enables users to track user activity and API usage across their AWS account. Reference: Amazon QuickSight FAQs

**NEW QUESTION 110**

- (Topic 1)

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

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

**Answer:** D

**Explanation:**

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

Reference: AWS Security Hub FAQs

**NEW QUESTION 111**

- (Topic 1)

Which AWS service or feature can be used to estimate costs before deployment?

- A. AWS Free Tier
- B. AWS Pricing Calculator
- C. AWS Billing and Cost Management
- D. AWS Cost and Usage Report

**Answer:** B

**Explanation:**

AWS Pricing Calculator can be used to estimate costs before deployment. AWS Pricing Calculator is a tool that helps the user to compare the cost of AWS services for different use cases and configurations. The user can create estimates for various AWS services, such as Amazon EC2, Amazon S3, Amazon RDS, and more. The user can also adjust the parameters, such as region, instance type, storage size, and duration, to see how they affect the cost. AWS Pricing Calculator provides a detailed breakdown of the estimated cost, as well as a summary of the key drivers of the cost.

**NEW QUESTION 115**

- (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 116**

- (Topic 1)

Which AWS service or tool helps to centrally manage billing and allow controlled access to resources across AWS accounts?

- A. AWS Identity and Access Management (IAM)
- B. AWS Organizations
- C. AWS Cost Explorer
- D. AWS Budgets

**Answer:** B

**Explanation:**

AWS Organizations helps to centrally manage billing and allow controlled access to resources across AWS accounts. AWS Organizations is a service that enables the user to consolidate multiple AWS accounts into an organization that can be managed as a single unit. AWS Organizations allows the user to create groups of accounts and apply policies to them, such as service control policies (SCPs) that specify the services and actions that users and roles can access in the accounts. AWS Organizations also enables the user to use consolidated billing, which combines the usage and charges from all the accounts in the organization into a single bill.

**NEW QUESTION 118**

- (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 122**

- (Topic 1)

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

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

**Answer:** B

**Explanation:**

Pay-as-you-go pricing is the feature of the AWS Cloud that gives users the ability to pay based on current needs rather than forecasted needs. Pay-as-you-go pricing means that users only pay for the AWS services and resources they use, without any upfront or long-term commitments. This allows users to scale up or down their usage depending on their changing business requirements, and avoid paying for idle or unused capacity. Pay-as-you-go pricing also enables users to benefit from the economies of scale and lower costs of AWS as they grow their business5

**NEW QUESTION 124**

- (Topic 1)

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

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

**Answer:** A

**Explanation:**

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

**NEW QUESTION 125**

- (Topic 1)

A company has been storing monthly reports in an Amazon S3 bucket. The company exports the report data into comma-separated values (.csv) files. A developer wants to write a simple query that can read all of these files and generate a summary report.

Which AWS service or feature should the developer use to meet these requirements with the LEAST amount of operational overhead?

- A. Amazon S3 Select
- B. Amazon Athena
- C. Amazon Redshift
- D. Amazon EC2

**Answer:** B

**Explanation:**

Amazon Athena is the AWS service that the developer should use to write a simple query that can read all of the .csv files stored in an Amazon S3 bucket and generate a summary report. Amazon Athena is an interactive query service that allows users to analyze data in Amazon S3 using standard SQL. Amazon Athena does not require any server setup or management, and users only pay for the queries they run. Amazon Athena can handle various data formats, including .csv, and can integrate with other AWS services such as Amazon QuickSight for data visualization

#### NEW QUESTION 129

- (Topic 1)

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

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

**Answer:** C

#### Explanation:

Designing for automatic failover to healthy resources is an AWS best practice when designing AWS workloads to be operational even when there are component failures. This means that you should architect your system to handle the loss of one or more components without impacting the availability or performance of your application. You can use various AWS services and features to achieve this, such as Auto Scaling, Elastic Load Balancing, Amazon Route 53, Amazon CloudFormation, and AWS CloudFormation4.

#### NEW QUESTION 133

- (Topic 1)

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

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

**Answer:** D

#### Explanation:

An Amazon Machine Image (AMI) is a deployable Amazon EC2 instance template that is prepackaged with software and security requirements. It provides the information required to launch an instance, which is a virtual server in the cloud. You can use an AMI to launch as many instances as you need. You can also create your own custom AMIs or use AMIs shared by other AWS users1.

#### NEW QUESTION 138

- (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 142

- (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 144

- (Topic 1)

A large company has a workload that requires hardware to remain on premises. The company wants to use the same management and control plane services that

it currently uses on AWS.

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

- A. AWS Device Farm
- B. AWS Fargate
- C. AWS Outposts
- D. AWS Ground Station

**Answer:** C

**Explanation:**

The correct answer is C because AWS Outposts is an AWS service that enables the company to meet the requirements. AWS Outposts is a fully managed service that extends AWS infrastructure, services, APIs, and tools to virtually any datacenter, co- location space, or on-premises facility. AWS Outposts allows customers to run their workloads on the same hardware and software that AWS uses in its cloud, while maintaining local access and control. The other options are incorrect because they are not AWS services that enable the company to meet the requirements. AWS Device Farm is an AWS service that enables customers to test their mobile and web applications on real devices in the AWS Cloud. AWS Fargate is an AWS service that enables customers to run containers without having to manage servers or clusters. AWS Ground Station is an AWS service that enables customers to communicate with satellites and downlink data from orbit.

Reference: AWS Outposts FAQs

**NEW QUESTION 148**

- (Topic 1)

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

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

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

**Answer:** B

**Explanation:**

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

**NEW QUESTION 152**

- (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.

**NEW QUESTION 156**

- (Topic 1)

A company deploys its application to multiple AWS Regions and configures automatic failover between those Regions.

Which cloud concept does this architecture represent?

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

**Answer:** B

**Explanation:**

Reliability is the cloud concept that this architecture represents. Reliability is the ability of a system to recover from infrastructure or service disruptions, dynamically acquire computing resources to meet demand, and mitigate disruptions such as misconfigurations or transient network issues. Deploying an application to multiple AWS Regions and configuring automatic failover between those Regions enhances the reliability of the application by reducing the impact of regional failures and increasing the availability of the application.

**NEW QUESTION 157**

- (Topic 1)

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

- A. Users can trade fixed expenses for variable expenses.
- B. Users can deploy all over the world in minutes.

- C. AWS offers increased speed and agility.
- D. AWS is responsible for patching the infrastructure.
- E. Users benefit from economies of scale.

**Answer:** AE

**Explanation:**

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

- ? Users can trade fixed expenses for variable expenses. By using the AWS Cloud, users can pay only for the resources they use, instead of investing in fixed and upfront costs for hardware and software. This can lower the total cost of ownership and increase the return on investment.
- ? Users benefit from economies of scale. By using the AWS Cloud, users can leverage the massive scale and efficiency of AWS to access lower prices and higher performance. AWS passes the cost savings to the users through price reductions and innovations. AWS Cloud Value Framework

**NEW QUESTION 160**

- (Topic 1)

Which AWS service uses a combination of publishers and subscribers?

- A. AWS Lambda
- B. Amazon Simple Notification Service (Amazon SNS)
- C. Amazon CloudWatch
- D. AWS CloudFormation

**Answer:** B

**Explanation:**

Amazon Simple Notification Service (Amazon SNS) is a service that provides fully managed pub/sub messaging. Pub/sub messaging is a pattern that uses a combination of publishers and subscribers. Publishers are entities that produce messages and send them to topics. Subscribers are entities that receive messages from topics. Topics are logical access points that act as communication channels between publishers and subscribers. Amazon SNS enables applications to decouple, scale, and coordinate the delivery of messages to multiple endpoints, such as email, SMS, mobile push notifications, Lambda functions, SQS queues, and HTTP/S endpoints. Amazon SNS OverviewAWS Certified Cloud Practitioner - aws.amazon.com

**NEW QUESTION 162**

- (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 164**

- (Topic 3)

A company has designed its AWS Cloud infrastructure to run its workloads effectively. The company also has protocols in place to continuously improve supporting processes.

Which pillar of the AWS Well-Architected Framework does this scenario represent?

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

**Answer:** D

**Explanation:**

The scenario represents the operational excellence pillar of the AWS Well-Architected Framework, which focuses on running and monitoring systems to deliver business value and continually improve supporting processes and procedures<sup>1</sup>. Security, performance efficiency, cost optimization, and reliability are the other four pillars of the framework<sup>1</sup>.

**NEW QUESTION 169**

- (Topic 3)

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

Which AWS service meets these requirements?

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

**Answer:** D

**Explanation:**

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

#### NEW QUESTION 172

- (Topic 3)

A company wants to grant users in one AWS account access to resources in another AWS account. The users do not currently have permission to access the resources.

Which AWS service will meet this requirement?

- A. IAM group
- B. IAM role
- C. IAM tag
- D. IAM Access Analyzer

**Answer: B**

#### Explanation:

IAM roles are a way to delegate access to resources in different AWS accounts. IAM roles allow users to assume a set of permissions for a limited time without having to create or share long-term credentials. IAM roles can be used to grant cross-account access by creating a trust relationship between the accounts and specifying the permissions that the role can perform. Users can then switch to the role and access the resources in the other account using temporary security credentials provided by the role. References: Cross account resource access in IAM, IAM tutorial: Delegate access across AWS accounts using IAM roles, How to Enable Cross-Account Access to the AWS Management Console

#### NEW QUESTION 175

- (Topic 3)

A company wants to migrate its workloads to AWS, but it lacks expertise in AWS Cloud computing.

Which AWS service or feature will help the company with its migration?

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

**Answer: D**

#### Explanation:

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

#### NEW QUESTION 178

- (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 181

- (Topic 3)

Which AWS service will allow a user to set custom cost and usage limits, and will alert when the thresholds are exceeded?

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

**Answer: B**

#### Explanation:

AWS Budgets allows you to set custom budgets that alert you when your costs or usage exceed (or are forecasted to exceed) your budgeted amount. You can also use AWS Budgets to set reservation utilization or coverage targets and receive alerts when your utilization drops below the threshold you define. AWS

Budgets provides you with a comprehensive view of your cost and usage, as well as your reservation utilization and coverage<sup>1</sup>.

#### NEW QUESTION 182

- (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 185

- (Topic 3)

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

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

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

**Answer:** C

#### Explanation:

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

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

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

Reference:

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

#### NEW QUESTION 189

- (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 193

- (Topic 3)

A company is running and managing its own Docker environment on Amazon EC2 instances. The company wants an alternative to help manage cluster size, scheduling, and environment maintenance.

Which AWS service meets these requirements?

- A. AWS Lambda
- B. Amazon RDS
- C. AWS Fargate
- D. Amazon Athena

**Answer:** C

**Explanation:**

AWS Fargate is a serverless compute engine for containers that works with both Amazon Elastic Container Service (Amazon ECS) and Amazon Elastic Kubernetes Service (Amazon EKS). AWS Fargate allows you to run containers without having to manage servers or clusters of Amazon EC2 instances. With AWS Fargate, you only pay for the compute resources you use to run your containers, and you don't need to worry about scaling, patching, securing, or maintaining the underlying infrastructure. AWS Fargate simplifies the deployment and management of containerized applications, and enables you to focus on building and running your applications instead of managing the infrastructure. References: AWS Fargate, What is AWS Fargate?

**NEW QUESTION 196**

- (Topic 3)

A company deployed an Amazon EC2 instance last week. A developer realizes that the EC2 instance is no longer running. The developer reviews a list of provisioned EC2 instances, and the EC2 instance is no longer on the list. What can the developer do to generate a recent history of the EC2 instance?

- A. Run Cost Explorer to identify the start time and end time of the EC2 instance.
- B. Use Amazon Inspector to find out when the EC2 instance was stopped.
- C. Perform a search in AWS CloudTrail to find all EC2 instance-related events.
- D. Use AWS Secrets Manager to display hidden termination logs of the EC2 instance.

**Answer:** C

**Explanation:**

AWS CloudTrail is a service that enables governance, compliance, operational auditing, and risk auditing of a customer's AWS account. AWS CloudTrail allows customers to track user activity and API usage across their AWS infrastructure. AWS CloudTrail can also provide a history of EC2 instance events, such as launch, stop, terminate, and reboot. Cost Explorer is a tool that enables customers to visualize, understand, and manage their AWS costs and usage over time. Amazon Inspector is an automated security assessment service that helps improve the security and compliance of applications deployed on AWS. AWS Secrets Manager helps customers protect secrets needed to access their applications, services, and IT resources.

**NEW QUESTION 198**

- (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 202**

- (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 206**

- (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 211**

- (Topic 2)

A company has an application that runs periodically in an on-premises environment. The application runs for a few hours most days, but runs for 8 hours a day for a week at the end of each month.

Which AWS service or feature should be used to host the application in the AWS Cloud?

- A. Amazon EC2 Standard Reserved Instances
- B. Amazon EC2 On-Demand Instances
- C. AWS Wavelength
- D. Application Load Balancer

**Answer:** B

**Explanation:**

Amazon EC2 On-Demand Instances are instances that you pay for by the second, with no long-term commitments or upfront payments<sup>4</sup>. This option is suitable for applications that have unpredictable or intermittent workloads, such as the one described in the question. Amazon EC2 Standard 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. 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. This option is not relevant for the application described in the question. Application Load Balancer is a type of load balancer that operates at the application layer and distributes traffic based on the content of the request. This option is not a service or feature to host the application, but rather to balance the traffic among multiple instances.

**NEW QUESTION 216**

- (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 217**

- (Topic 2)

A company wants to move its iOS application development and build activities to AWS. Which AWS service or resource should the company use for these activities?

- A. AWS CodeCommit
- B. Amazon EC2 M1 Mac instances
- C. AWS Amplify
- D. AWS App Runner

**Answer:** B

**Explanation:**

Amazon EC2 M1 Mac instances are the AWS service or resource that the company should use for its iOS application development and build activities, as they enable users to run macOS on AWS and access a broad and growing set of AWS services. AWS CodeCommit is a service that provides a fully managed source control service that hosts secure Git-based repositories. AWS Amplify is a set of tools and services that enable developers to build full-stack web and mobile applications using AWS. AWS App Runner is a service that makes it easy for developers to quickly deploy containerized web applications and APIs. These concepts are explained in the AWS Developer Tools page<sup>4</sup>.

**NEW QUESTION 221**

- (Topic 2)

Which AWS solution provides the ability for a company to run AWS services in the company's on-premises data center?

- A. AWS Direct Connect
- B. AWS Outposts
- C. AWS Systems Manager hybrid activations
- D. AWS Storage Gateway

**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. AWS Outposts enables you to run AWS services in your on-premises data center<sup>1</sup>.

**NEW QUESTION 224**

- (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 229**

- (Topic 2)

A company has a compliance requirement to record and evaluate configuration changes, as well as perform remediation actions on AWS resources. Which AWS service should the company use?

- A. AWS Config
- B. AWS Secrets Manager
- C. AWS CloudTrail
- D. AWS Trusted Advisor

**Answer:** A

**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. This can help you simplify compliance auditing, security analysis, change management, and operational troubleshooting<sup>1</sup>.

**NEW QUESTION 233**

- (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 235**

- (Topic 2)

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

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

**Answer:** C

**Explanation:**

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

**NEW QUESTION 236**

- (Topic 2)

An application runs on multiple Amazon EC2 instances that access a shared file system simultaneously. Which AWS storage service should be used?

- A. Amazon EBS
- B. Amazon EFS
- C. Amazon S3
- D. AWS Artifact

**Answer:** B

**Explanation:**

Amazon Elastic File System (Amazon EFS) is the AWS storage service that should be used for an application that runs on multiple Amazon EC2 instances that access a shared file system simultaneously. Amazon EFS is a fully managed service that provides a scalable, elastic, and highly available file system for Linux-based workloads. Amazon EFS supports the Network File System version 4 (NFSv4) protocol and allows multiple EC2 instances to read and write data to the same file system concurrently. Amazon EFS also integrates with other AWS services, such as AWS Backup, AWS CloudFormation, and AWS CloudTrail. For more information, see What is Amazon Elastic File System? and [Amazon EFS Use Cases].

**NEW QUESTION 238**

- (Topic 2)

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

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

**Answer:** B

**Explanation:**

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

**NEW QUESTION 243**

- (Topic 2)

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

- A. Security awareness and training
- B. Development of an IAM password policy
- C. Patching of the guest operating system
- D. Physical and environmental controls

**Answer:** D

**Explanation:**

Physical and environmental controls are entirely the responsibility of AWS, according to the AWS shared responsibility model. 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 and environmental controls of the AWS global infrastructure, such as power, cooling, fire suppression, and physical access. The customer is responsible for the security in the cloud, which includes the configuration and management of the AWS resources and applications. For more information, see [AWS Shared Responsibility Model] and [AWS Cloud Security].

**NEW QUESTION 247**

- (Topic 2)

A company needs to host a highly available application in the AWS Cloud. The application runs infrequently for short periods of time. Which AWS service will meet these requirements with the LEAST amount of operational overhead?

- A. Amazon EC2
- B. AWS Fargate
- C. AWS Lambda
- D. Amazon Aurora

**Answer:** C

**Explanation:**

The AWS service that will meet the requirements of the company that needs to host a highly available application in the AWS Cloud that runs infrequently for short

periods of time with the least amount of operational overhead is AWS Lambda. AWS Lambda is a serverless compute service that allows customers to run code without provisioning or managing servers. The company can use AWS Lambda to create and deploy their application as functions that are triggered by events, such as API calls, messages, or schedules. AWS Lambda automatically scales the compute resources based on the demand, and customers only pay for the compute time they consume. AWS Lambda also simplifies the management and maintenance of the application, as customers do not need to worry about the underlying infrastructure, security, or availability. Amazon EC2, AWS Fargate, and Amazon Aurora are not the best services to use for this purpose. Amazon EC2 is a service that provides scalable compute capacity in the cloud, and allows customers to launch and run virtual servers, called instances, with a variety of operating systems, configurations, and specifications. Amazon EC2 requires customers to provision and manage the instances, and pay for the instance hours they use, regardless of the application usage. AWS Fargate is a serverless compute engine for containers that allows customers to run containerized applications without managing servers or clusters. AWS Fargate requires customers to specify the amount of CPU and memory resources for each container, and pay for the resources they allocate, regardless of the application usage.

Amazon Aurora is a fully managed relational database service that provides high performance, availability, and compatibility. Amazon Aurora is not a compute service, and it is not suitable for hosting an application that runs infrequently for short periods of time<sup>12</sup>

#### NEW QUESTION 248

- (Topic 2)

A company is reviewing the design of an application that will be migrated from on premises to a single Amazon EC2 instance. What should the company do to make the application highly available?

- A. Provision additional EC2 instances in other Availability Zones.
- B. Configure an Application Load Balancer (ALB). Assign the EC2 instance as the ALB's target.
- C. Use an Amazon Machine Image (AMI) to create the EC2 instance.
- D. Provision the application by using an EC2 Spot Instance.

**Answer:** A

#### Explanation:

Provisioning additional EC2 instances in other Availability Zones is a way to make the application highly available, as it reduces the impact of failures and increases fault tolerance. Configuring an Application Load Balancer and assigning the EC2 instance as the ALB's target is a way to distribute traffic among multiple instances, but it does not make the application highly available if there is only one instance. Using an Amazon Machine Image to create the EC2 instance is a way to launch a virtual server with a preconfigured operating system and software, but it does not make the application highly available by itself. Provisioning the application by using an EC2 Spot Instance is a way to use spare EC2 capacity at up to 90% off the On-Demand price, but it does not make the application highly available, as Spot Instances can be interrupted by EC2 with a two-minute notification.

#### NEW QUESTION 250

- (Topic 2)

A company needs to design a solution for the efficient use of compute resources for an enterprise workload. The company needs to make informed decisions as its technology needs evolve.

Which pillar of the AWS Well-Architected Framework do these requirements represent?

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

**Answer:** B

#### Explanation:

Performance efficiency is the pillar of the AWS Well-Architected Framework that represents the requirements of designing a solution for the efficient use of compute resources for an enterprise workload and making informed decisions as the technology needs evolve. It focuses on using the right resources and services for the workload, monitoring performance, and continuously improving the efficiency of the solution. Operational excellence is the pillar of the AWS Well-Architected Framework that represents the ability to run and monitor systems to deliver business value and to continually improve supporting processes and procedures. Cost optimization is the pillar of the AWS Well-Architected Framework that represents the ability to run systems to deliver business value at the lowest price point. Reliability is the pillar of the AWS Well-Architected Framework that represents 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.

#### NEW QUESTION 251

- (Topic 2)

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

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

**Answer:** B

#### Explanation:

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

#### NEW QUESTION 252

- (Topic 2)

A user is moving a workload from a local data center to an architecture that is distributed between the local data center and the AWS Cloud.

Which type of migration is this?

- A. On-premises to cloud native
- B. Hybrid to cloud native

- C. On-premises to hybrid
- D. Cloud native to hybrid

**Answer:** C

**Explanation:**

C is correct because moving a workload from a local data center to an architecture that is distributed between the local data center and the AWS Cloud is an example of an on-premises to hybrid migration. A 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. A is incorrect because on-premises to cloud native migration is the process of moving a workload from a local data center to an architecture that is fully hosted and managed on the AWS Cloud. B is incorrect because hybrid to cloud native migration is the process of moving a workload from an architecture that is distributed between the local data center and the AWS Cloud to an architecture that is fully hosted and managed on the AWS Cloud. D is incorrect because cloud native to hybrid migration is the process of moving a workload from an architecture that is fully hosted and managed on the AWS Cloud to an architecture that is distributed between the local data center and the AWS Cloud.

**NEW QUESTION 254**

- (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 255**

- (Topic 2)

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

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

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

**Answer:** D

**Explanation:**

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

**NEW QUESTION 260**

- (Topic 2)

A company is planning a migration to the AWS Cloud and wants to examine the costs that are associated with different workloads.

Which AWS tool will meet these requirements?

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

**Answer:** C

**Explanation:**

The AWS tool that will meet the requirements of the company that is planning a migration to the AWS Cloud and wants to examine the costs that are associated with different workloads is AWS Pricing Calculator. AWS Pricing Calculator is a tool that helps customers estimate the cost of using AWS services based on their requirements and preferences. The company can use AWS Pricing Calculator to compare the costs of different AWS services and configurations, such as Amazon EC2, Amazon S3, Amazon RDS, and more. AWS Pricing Calculator also provides detailed breakdowns of the cost components, such as compute, storage, network, and data transfer. AWS Pricing Calculator helps customers plan and optimize their cloud budget and migration strategy. AWS Budgets, AWS Cost Explorer, and AWS Cost and Usage Report are not the best tools to use for this purpose. AWS Budgets is a tool that helps customers monitor and manage their AWS spending and usage against predefined budget limits and thresholds. AWS Cost Explorer is a tool that helps customers analyze and visualize their AWS spending and usage trends over time. AWS Cost and Usage Report is a tool that helps customers access comprehensive and granular information about their AWS costs and usage in a CSV or Parquet file. These tools are more useful for tracking and optimizing the existing AWS costs and usage, rather than estimating the costs of different workloads34

**NEW QUESTION 265**

- (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 266**

- (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 271**

- (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 274**

- (Topic 2)

A new AWS user who has little cloud experience wants to build an application by using AWS services. The user wants to learn how to implement specific AWS services from other customer examples. The user also wants to ask questions to AWS experts. Which AWS service or resource will meet these requirements?

- A. AWS Online Tech Talks
- B. AWS documentation
- C. AWS Marketplace
- D. AWS Health Dashboard

**Answer:** A

**Explanation:**

AWS Online Tech Talks are online presentations that cover a broad range of topics at varying technical levels and provide a live Q&A session with AWS experts. They are a great resource for new AWS users who want to learn how to implement specific AWS services from other customer examples and ask questions to AWS experts. AWS documentation, AWS Marketplace, and AWS Health Dashboard do not offer the same level of interactivity and guidance as AWS Online Tech Talks. Source: AWS Online Tech Talks

**NEW QUESTION 277**

- (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 279**

- (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 283**

- (Topic 2)

A company has set up a VPC in its AWS account and has created a subnet in the VPC. The company wants to make the subnet public.

Which AWS features should the company use to meet this requirement? (Select TWO.)

- A. Amazon VPC internet gateway
- B. Amazon VPC NAT gateway
- C. Amazon VPC route tables
- D. Amazon VPC network ACL
- E. Amazon EC2 security groups

**Answer:** AC

**Explanation:**

To make a subnet public, the company should use an Amazon VPC internet gateway and an Amazon VPC route table. An internet gateway is a horizontally scaled, redundant, and highly available VPC component that allows communication between your VPC and the internet. A route table contains a set of rules, called routes, that are used to determine where network traffic from your subnet or gateway is directed. To enable internet access for a subnet, you need to attach an internet gateway to your VPC and add a route to the internet gateway in the route table associated with the subnet.

**NEW QUESTION 284**

.....

## THANKS FOR TRYING THE DEMO OF OUR PRODUCT

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

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

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

## Money Back Guarantee

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

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