



## Amazon

### Exam Questions AWS-Certified-Cloud-Practitioner

Amazon AWS Certified Cloud Practitioner

## About ExamBible

### *Your Partner of IT Exam*

## Found in 1998

ExamBible is a company specialized on providing high quality IT exam practice study materials, especially Cisco CCNA, CCDA, CCNP, CCIE, Checkpoint CCSE, CompTIA A+, Network+ certification practice exams and so on. We guarantee that the candidates will not only pass any IT exam at the first attempt but also get profound understanding about the certificates they have got. There are so many alike companies in this industry, however, ExamBible has its unique advantages that other companies could not achieve.

## Our Advances

### \* 99.9% Uptime

All examinations will be up to date.

### \* 24/7 Quality Support

We will provide service round the clock.

### \* 100% Pass Rate

Our guarantee that you will pass the exam.

### \* Unique Gurantee

If you do not pass the exam at the first time, we will not only arrange FULL REFUND for you, but also provide you another exam of your claim, ABSOLUTELY FREE!

### NEW QUESTION 1

- (Topic 2)

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

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

**Answer:** D

#### **Explanation:**

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

### NEW QUESTION 2

- (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. 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 3

- (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 4

- (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.

### NEW QUESTION 5

- (Topic 2)

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

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

**Answer: C**

#### Explanation:

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

### NEW QUESTION 6

- (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 7

- (Topic 2)

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

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

**Answer: B**

#### Explanation:

Operations is an option that is a perspective that includes foundational capabilities of the AWS Cloud Adoption Framework (AWS CAF). Operations is one of the six perspectives of the AWS CAF, along with business, people, governance, platform, and security. Operations focuses on the processes and procedures to support the ongoing management and maintenance of the cloud-based IT assets. It covers topics such as monitoring, backup and recovery, change management, incident management, and automation. Sustainability is not a perspective of the AWS CAF, but a concept that refers to the ability of a system to operate in an environmentally friendly and socially responsible manner. Performance efficiency is not a perspective of the AWS CAF, but a pillar of the AWS Well-Architected Framework. It focuses on using the right resources and services for the workload, monitoring performance, and continuously improving the efficiency of the solution. Reliability is not a perspective of the AWS CAF, but a pillar of the AWS Well-Architected Framework. It focuses on 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 8

- (Topic 2)

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

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

**Answer: C**

#### Explanation:

The AWS service that requires the customer to patch the guest operating system is Amazon EC2. Amazon EC2 is a service that provides scalable compute

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

#### NEW QUESTION 9

- (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 10

- (Topic 1)

A company needs to run its existing custom, nonproduction workloads in the AWS Cloud quickly and cost-effectively. The workloads can recover from interruptions easily. Which pricing model should the company use?

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

**Answer:** C

#### Explanation:

The correct answer is C because Spot Instances are the pricing model that enables the company to run its existing custom, nonproduction workloads in the AWS Cloud quickly and cost-effectively. Spot Instances are spare Amazon EC2 instances that are available at up to 90% discount compared to On-Demand prices. Spot Instances are suitable for stateless, fault-tolerant, and flexible workloads that can recover from interruptions easily. The other options are incorrect because they are not the pricing model that enables the company to run its existing custom, nonproduction workloads in the AWS Cloud quickly and cost-effectively. Reserved Instances are Amazon EC2 instances that are reserved for a specific period of time (one or three years) in exchange for a lower hourly rate. Reserved Instances are suitable for steady-state or predictable workloads that run for a long duration. On-Demand Instances are Amazon 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 Hosts are physical servers that are dedicated to a single customer. Dedicated Hosts are suitable for workloads that require regulatory compliance or data isolation. Reference: Amazon EC2 Instance Purchasing Options

#### NEW QUESTION 10

- (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 13

- (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 16**

- (Topic 1)

What is an Availability Zone?

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

**Answer:** B

**Explanation:**

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

**NEW QUESTION 18**

- (Topic 1)

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

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

**Answer:** D

**Explanation:**

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

**NEW QUESTION 19**

- (Topic 1)

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

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

**Answer:** B

**Explanation:**

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

**NEW QUESTION 22**

- (Topic 1)

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

- A. Vertical scaling
- B. Manual failure recovery
- C. Testing recovery procedures

D. Changing infrastructure manually

**Answer:** C

**Explanation:**

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

**NEW QUESTION 25**

- (Topic 1)

An auditor needs to find out whether a specific AWS service is compliant with specific compliance frameworks. Which AWS service will provide this information?

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

**Answer:** A

**Explanation:**

AWS Artifact is the service that will provide the information about whether a specific AWS service is compliant with specific compliance frameworks. AWS Artifact is a self-service portal that allows you to access, review, and download AWS security and compliance reports and agreements. You can use AWS Artifact to verify the compliance status of AWS services across various regions and compliance programs, such as ISO, PCI, SOC, FedRAMP, HIPAA, and more<sup>12</sup>

**NEW QUESTION 26**

- (Topic 1)

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

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

**Answer:** B

**Explanation:**

The AWS Well-Architected Framework is a set of best practices and guidelines for designing and operating systems in the cloud. The framework consists of five pillars: operational excellence, security, reliability, performance efficiency, and cost optimization. The concept of elasticity represents a system's ability to adapt to changes in demand by scaling resources up or down automatically. Therefore, the correct answer is B. You can learn more about the AWS Well-Architected Framework and its pillars from this page.

**NEW QUESTION 29**

- (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 30**

- (Topic 1)

Which AWS service can report how AWS resource configurations have changed over time?

- A. AWS CloudTrail
- B. Amazon CloudWatch
- C. AWS Config
- D. Amazon Inspector

**Answer:** C

**Explanation:**

AWS Config is a service that enables users to assess, audit, and evaluate the configurations of AWS resources. It continuously monitors and records the configuration changes of the resources and evaluates them against desired configurations and best practices. It also provides a detailed view of the resource configuration history and relationships, as well as compliance reports and notifications. AWS Config can help users maintain consistent and secure configurations,

troubleshoot issues, and simplify compliance auditing. AWS Config Overview AWS Certified Cloud Practitioner - aws.amazon.com

### NEW QUESTION 35

- (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 36

- (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 41

- (Topic 1)

Which AWS database service provides in-memory data storage?

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

**Answer:** B

#### Explanation:

The correct answer is B because Amazon ElastiCache is a service that provides in-memory data storage. Amazon ElastiCache is a fully managed, scalable, and high-performance service that supports two popular open-source in-memory engines: Redis and Memcached. Amazon ElastiCache allows users to store and retrieve data from fast, low-latency, and high-throughput in-memory systems. Users can use Amazon ElastiCache to improve the performance of their applications by caching frequently accessed data, reducing database load, and enabling real-time data processing. The other options are incorrect because they are not services that provide in-memory data storage. Amazon DynamoDB is a service that provides key-value and document data storage. Amazon RDS is a service that provides relational data storage. Amazon Timestream is a service that provides time series data storage. Reference: Amazon ElastiCache FAQs

### NEW QUESTION 44

- (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 48

- (Topic 1)

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

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

**Answer: B**

#### Explanation:

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

#### NEW QUESTION 49

- (Topic 1)

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

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

**Answer: B**

#### Explanation:

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

#### NEW QUESTION 52

- (Topic 1)

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

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

**Answer: D**

#### Explanation:

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

#### NEW QUESTION 55

- (Topic 1)

What is the total amount of storage offered by Amazon S3?

- A. WOMB
- B. 5 GB
- C. 5 TB
- D. Unlimited

**Answer: D**

#### Explanation:

Amazon S3 offers unlimited storage for any amount of data. You can store as many objects as you want, and each object can be as large as 5 terabytes. You pay only for the storage space that you actually use, and there are no minimum commitments or upfront fees. Amazon S3 also provides high durability, availability, scalability, and security for your data.

#### NEW QUESTION 60

- (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 63**

- (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 65**

- (Topic 3)

Which AWS service or tool helps users visualize, understand, and manage spending and usage over time?

- A. AWS Organizations
- B. AWS Pricing Calculator
- C. AWS Cost Explorer
- D. AWS Service Catalog

**Answer:** C

**Explanation:**

AWS Cost Explorer is the AWS service or tool that helps users visualize, understand, and manage spending and usage over time. AWS Cost Explorer is a web-based interface that allows users to access interactive graphs and tables that display their AWS costs and usage data. Users can create custom reports that analyze cost and usage data by various dimensions, such as service, region, account, tag, and more. Users can also view historical data for up to the last 12 months, forecast future costs for up to the next 12 months, and get recommendations for cost optimization. AWS Cost Explorer also provides preconfigured views that show common cost and usage scenarios, such as monthly spend by service, daily spend by linked account, and Reserved Instance utilization. Users can use AWS Cost Explorer to monitor their AWS spending and usage trends, identify cost drivers and anomalies, and optimize their resource allocation and budget planning. References: Cloud Cost Analysis - AWS Cost Explorer - AWS, Analyzing your costs with AWS Cost Explorer

**NEW QUESTION 70**

- (Topic 3)

A company needs to securely store important credentials that an application uses to connect users to a database.

Which AWS service can meet this requirement with the MINIMAL amount of operational overhead?

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

**Answer:** C

**Explanation:**

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

**NEW QUESTION 75**

- (Topic 3)

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

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

**Answer:** B

**Explanation:**

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

#### NEW QUESTION 76

- (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 80

- (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 84

- (Topic 2)

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

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

**Answer:** AE

#### Explanation:

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

#### NEW QUESTION 87

- (Topic 2)

Which AWS service is always available free of charge to users?

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

D. Amazon ElastiCacheA company has only basic knowledge of AWS technologies.

**Answer:** B

**Explanation:**

AWS Identity and Access Management (IAM) is a web service that helps you securely control access to AWS resources for your users. You use IAM to control who can use your AWS resources (authentication) and what resources they can use and in what ways (authorization). IAM is always available free of charge to users<sup>4</sup>.

**NEW QUESTION 91**

- (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 95**

- (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 98**

- (Topic 2)

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

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

**Answer:** AB

**Explanation:**

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

**NEW QUESTION 103**

- (Topic 2)

Which AWS service can defend against DDoS attacks?

- A. AWS Firewall Manager
- B. AWS Shield Standard
- C. AWS WAF
- D. Amazon Inspector

**Answer:** B

**Explanation:**

AWS Shield Standard is a service that provides protection against Distributed Denial of Service (DDoS) attacks for all AWS customers at no additional charge. It automatically detects and mitigates the most common and frequently occurring network and transport layer DDoS attacks that target AWS resources, such as Amazon EC2 instances, Elastic Load Balancers, Amazon CloudFront distributions, and Amazon Route 53 hosted zones. AWS Firewall Manager is a service that allows users to centrally configure and manage firewall rules across their AWS accounts and resources, such as AWS WAF web ACLs, AWS Shield Advanced protections, and Amazon VPC security groups. AWS WAF is a web application firewall that helps protect web applications from common web exploits, such as SQL injection, cross-site scripting, and bot attacks. Amazon Inspector is an automated security assessment service that helps improve the security and compliance of applications deployed on AWS. It analyzes the behavior of the applications and checks for vulnerabilities, exposures, and deviations from best practices.

**NEW QUESTION 106**

- (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.

**NEW QUESTION 109**

- (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 110**

- (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 111**

- (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 115**

- (Topic 2)

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

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

**Answer: D**

**Explanation:**

Reliability is the benefit of AWS Cloud computing that ensures the workload performs consistently and correctly. According to the AWS Cloud Practitioner Essentials course, reliability means "the ability of a system to recover from infrastructure or service disruptions, dynamically acquire computing resources to meet demand, and mitigate disruptions such as misconfigurations or transient network issues."<sup>1</sup> Elasticity, security, and pay-as-you-go pricing are also benefits of AWS Cloud computing, but they do not directly relate to the goal of consistent and correct performance.

**NEW QUESTION 119**

- (Topic 2)

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

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

**Answer: D**

**Explanation:**

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

**NEW QUESTION 122**

- (Topic 2)

A company wants to move its data warehouse application to the AWS Cloud. The company wants to run and scale its analytics services without needing to provision and manage data warehouse clusters.

Which AWS service will meet these requirements?

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

**Answer: B**

**Explanation:**

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

#### NEW QUESTION 127

- (Topic 2)

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

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

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

**Answer: D**

#### Explanation:

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

#### NEW QUESTION 129

- (Topic 2)

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 132

- (Topic 2)

Which encryption types can be used to protect objects at rest in Amazon S3? (Select TWO.)

- A. Server-side encryption with AmazonS3 managed encryption keys (SSE-S3)
- B. Server-side encryption with AWS KMSmanaged keys (SSE-KMS)
- C. TLS
- D. SSL
- E. Transparent Data Encryption (TDE)

**Answer: AB**

#### Explanation:

Server-side encryption with Amazon S3 managed encryption keys (SSE-S3) and server-side encryption with AWS KMS managed keys (SSE-KMS) are the encryption types that can be used to protect objects at rest in Amazon S3. Server-side encryption means that Amazon S3 encrypts the objects before saving them on disks and decrypts them when they are downloaded. SSE-S3 uses one master key per bucket that is managed by Amazon S3. SSE-KMS uses a customer master key (CMK) that is stored in AWS Key Management Service (AWS KMS) and provides additional benefits, such as audit trails and key rotation. For more information, see Protecting Data Using Server-Side Encryption and Protecting Data Using Encryption.

#### NEW QUESTION 134

- (Topic 2)

Which of the following is the customer's responsibility, according to the AWS shared responsibility model?

- A. Identity and access management
- B. Hard drive initialization
- C. Protection of data center hardware
- D. Security of Availability Zones

**Answer: A**

#### Explanation:

Identity and access management is the customer's responsibility, according to the AWS shared responsibility model. This means that the customer is responsible for managing user access to the AWS resources, using tools such as AWS Identity and Access Management (IAM), AWS Single Sign-On (SSO), and AWS Organizations. The customer is also responsible for securing their data in transit and at rest, using encryption, key management, and other methods. Hard drive initialization, protection of data center hardware, and security of Availability Zones are AWS's responsibility, as they are part of the infrastructure, physical security, and network security that AWS provides to the customer12

### NEW QUESTION 135

- (Topic 2)

A company needs to host a web server on Amazon EC2 instances for at least 1 year. The web server cannot tolerate interruption. Which EC2 instance purchasing option will meet these requirements MOST cost-effectively?

- A. On-Demand Instances
- B. Partial Upfront Reserved Instances
- C. Spot Instances
- D. No Upfront Reserved Instances

**Answer: B**

#### Explanation:

The most cost-effective EC2 instance purchasing option for the company that needs to host a web server on Amazon EC2 instances for at least 1 year and cannot tolerate interruption is Partial Upfront Reserved Instances. Reserved Instances are a pricing model that offer significant discounts compared to On-Demand Instances in exchange for a commitment to use a specific amount of compute capacity for a fixed period of time (1 or 3 years). Partial Upfront Reserved Instances require customers to pay a portion of the total cost upfront, and the remaining cost in monthly installments over the term. This option offers a lower effective hourly rate than No Upfront Reserved Instances, which require no upfront payment but have higher monthly payments. On-Demand Instances and Spot Instances are not the best options for the company. On-Demand Instances are a pricing model that offer the most flexibility and no long-term commitment, but have the highest hourly rate. Spot Instances are a pricing model that offer the lowest cost, but are subject to interruption based on supply and demand<sup>34</sup>

### NEW QUESTION 139

- (Topic 2)

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

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

**Answer: D**

#### Explanation:

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

### NEW QUESTION 142

- (Topic 2)

Which aspect of security is the customer's responsibility, according to the AWS shared responsibility model?

- A. Patch and configuration management
- B. Service and communications protection or zone security
- C. Physical and environmental controls
- D. Awareness and training

**Answer: A**

#### Explanation:

According to the AWS shared responsibility model, AWS is responsible for the security of the cloud, while the customer is responsible for the security in the cloud. This means that AWS provides the physical and environmental controls, the service and communications protection, and the awareness and training for its employees, while the customer provides the patch and configuration management, the identity and access management, the data encryption, and the firewall configuration for its resources<sup>3</sup>.

### NEW QUESTION 146

- (Topic 2)

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

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

**Answer: C**

#### Explanation:

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

#### NEW QUESTION 147

- (Topic 2)

Which AWS services can a company use to host and run a MySQL database? (Select TWO.)

- A. Amazon RDS
- B. Amazon DynamoDB
- C. Amazon S3
- D. Amazon EC2
- E. Amazon MQ

**Answer:** AD

#### Explanation:

Amazon RDS and Amazon EC2 are two AWS services that you can use to host and run a MySQL database. Amazon RDS is a service that makes it easy to set up, operate, and scale a relational database in the cloud. You can use Amazon RDS to launch a MySQL database instance and let Amazon RDS manage common database tasks such as backups, patching, scaling, and replication<sup>6</sup>. Amazon EC2 is a service that provides secure, resizable compute capacity in the cloud. You can use Amazon EC2 to launch a virtual server and install MySQL software on it. You have complete control over your database configuration, but you are responsible for managing and maintaining the database software and the underlying infrastructure<sup>7</sup>. Amazon DynamoDB is a key-value and document database that delivers single-digit millisecond performance at any scale. Amazon S3 is an object storage service that offers industry-leading scalability, data availability, security, and performance. Amazon MQ is a managed message broker service for Apache ActiveMQ. None of these services can help you host and run a MySQL database.

#### NEW QUESTION 148

- (Topic 2)

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

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

**Answer:** B

#### Explanation:

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

#### NEW QUESTION 150

- (Topic 2)

A company that is planning to migrate to the AWS Cloud is based in an isolated area that has limited internet connectivity. The company needs to perform local data processing on premises. The company needs a solution that can operate without a stable internet connection. Which AWS service will meet these requirements?

- A. Amazon S3
- B. AWS Snowball Edge
- C. AWS StorageGateway
- D. AWS Backup

**Answer:** B

#### Explanation:

AWS Snowball Edge is a service that provides a physical device that can store up to 100 TB of data and perform local data processing on premises. It enables users to transfer data to and from the AWS Cloud in areas with limited or no internet connectivity. It also supports AWS Greengrass, which allows users to run AWS Lambda functions and other AWS services locally without a stable internet connection. Amazon S3 is a storage service that provides scalable, durable, and secure object storage. It requires a stable internet connection to transfer data to and from the AWS Cloud. AWS Storage Gateway is a service that provides a hybrid storage solution that connects on-premises applications to AWS Cloud storage services, such as Amazon S3, Amazon S3 Glacier, and Amazon EBS. It requires a stable internet connection to synchronize data between the on-premises and cloud storage. AWS Backup is a service that provides a centralized and automated solution to back up data across AWS services and on-premises resources. It requires a stable internet connection to transfer data to and from the AWS Cloud.

#### NEW QUESTION 151

- (Topic 2)

Which AWS solution should the company use to meet this requirement?

- A. AWS Config
- B. AWS software development kits (SDKs)
- C. AWS Service Catalog
- D. AWS AppSync

**Answer:** C

**Explanation:**

AWS Service Catalog is a service that allows you to create and manage catalogs of IT services that are approved for use on AWS. You can use AWS Service Catalog to centrally manage commonly deployed IT services and help your organization achieve consistent governance and meet your compliance requirements, while enabling users to quickly deploy only the approved IT services they need<sup>1</sup>. AWS Config is a service that enables you to assess, audit, and evaluate the configurations of your AWS resources. AWS software development kits (SDKs) are tools that enable you to easily integrate your applications with AWS services using your preferred programming language. AWS AppSync is a service that simplifies application development by letting you create a flexible API to securely access, manipulate, and combine data from one or more data sources. None of these services can help you limit your employees' AWS access to a portfolio of predefined AWS resources.

**NEW QUESTION 155**

- (Topic 2)

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

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

**Answer:** BC

**Explanation:**

Basing the selection of Amazon EC2 instance types on past utilization patterns is a way to right size the AWS resources and optimize the performance and cost. Using Amazon S3 Lifecycle policies to move objects that users access infrequently to lower-cost storage tiers is another way to reduce the storage costs and align them with the business value of the data. These two actions are recommended by the AWS Cost Optimization Pillar<sup>1</sup>. Switching from Amazon RDS to Amazon DynamoDB is not necessarily a cost-saving action, as it depends on the use case and the data model. Using Multi-AZ deployments for Amazon RDS is a way to improve the availability and durability of the database, but it also increases the cost. Replacing existing Amazon EC2 instances with AWS Elastic Beanstalk is a way to simplify the deployment and management of the application, but it does not affect the cost of the underlying EC2 instances.

**NEW QUESTION 156**

- (Topic 2)

A company has an Amazon S3 bucket containing images of scanned financial invoices. The company is building an artificial intelligence (AI)-based application on AWS. The company wants the application to identify and read total balance amounts on the invoices.

Which AWS service will meet these requirements?

- A. Amazon Forecast
- B. Amazon Textract
- C. Amazon Rekognition
- D. Amazon Lex

**Answer:** B

**Explanation:**

Amazon Textract is a service that automatically extracts text and data from scanned documents. Amazon Textract goes beyond simple optical character recognition (OCR) to also identify the contents of fields in forms and information stored in tables. Amazon Textract can analyze images of scanned financial invoices and extract the total balance amounts, as well as other relevant information, such as invoice number, date, vendor name, etc<sup>5</sup>.

**NEW QUESTION 158**

- (Topic 2)

A company provides a web-based ecommerce service that runs in two Availability Zones within a single AWS Region. The web service distributes content that is stored in the Amazon S3 Standard storage class. The company wants to improve the web service's performance globally.

What should the company do to meet this requirement?

- A. Change the S3 storage class to S3 Intelligent-Tiering.
- B. Deploy an Amazon CloudFront distribution to cache web server content in edge locations.
- C. Use Amazon API Gateway for the web service.
- D. Migrate the website ecommerce servers to Amazon EC2 with enhanced networking.

**Answer:** B

**Explanation:**

Amazon CloudFront is a fast content delivery network (CDN) service that securely delivers data, videos, applications, and APIs to customers globally with low latency, high transfer speeds, all within a developer-friendly environment. CloudFront can cache web server content in edge locations, which are located closer to the end users, to improve the web service's performance globally<sup>2</sup>.

**NEW QUESTION 160**

- (Topic 1)

A company is designing an identity access management solution for an application. The company wants users to be able to use their social media, email, or online shopping accounts to access the application.

Which AWS service provides this functionality?

- A. AWS IAM Identity Center (AWS Single Sign-On)
- B. AWS Config
- C. Amazon Cognito
- D. AWS Identity and Access Management (IAM)

**Answer:** C

**Explanation:**

The correct answer is C because Amazon Cognito provides identity federation and user authentication for web and mobile applications. Amazon Cognito allows users to sign in with their social media, email, or online shopping accounts. The other options are incorrect because they do not provide identity federation or user authentication. AWS IAM Identity Center (AWS Single Sign-On) is a service that enables users to access multiple AWS accounts and applications with a single sign-on experience. AWS Config is a service that enables users to assess, audit, and evaluate the configurations of their AWS resources. AWS Identity and Access Management (IAM) is a service that enables users to manage access to AWS resources using users, groups, roles, and policies.

Reference: Amazon Cognito FAQs

**NEW QUESTION 161**

- (Topic 1)

A company is configuring its AWS Cloud environment. The company's administrators need to group users together and apply permissions to the group. Which AWS service or feature can the company use to meet these requirements?

- A. AWS Organizations
- B. Resource groups
- C. Resource tagging
- D. AWS Identity and Access Management (IAM)

**Answer: D**

**Explanation:**

The AWS service or feature that the company can use to group users together and apply permissions to the group is AWS Identity and Access Management (IAM). AWS IAM is a service that enables users to create and manage users, groups, roles, and permissions for AWS services and resources. Users can use IAM groups to organize multiple users that have similar access requirements, and attach policies to the groups that define the permissions for the users in the group. This simplifies the management and administration of user access

**NEW QUESTION 163**

- (Topic 1)

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

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

**Answer: B**

**Explanation:**

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

**NEW QUESTION 164**

- (Topic 1)

Which factors affect costs in the AWS Cloud? (Select TWO.)

- A. The number of unused AWS Lambda functions
- B. The number of configured Amazon S3 buckets
- C. Inbound data transfers without acceleration
- D. Outbound data transfers without acceleration
- E. Compute resources that are currently in use

**Answer: DE**

**Explanation:**

Outbound data transfers without acceleration and compute resources that are currently in use are the factors that affect costs in the AWS Cloud. Outbound data transfers without acceleration refer to the amount of data that is transferred from AWS to the internet, without using any service that can optimize the speed and cost of the data transfer, such as AWS Global Accelerator or Amazon CloudFront. Outbound data transfers are charged at different rates depending on the source and destination AWS Regions, and the volume of data transferred. Compute resources that are currently in use refer to the AWS services and resources that provide computing capacity, such as Amazon EC2 instances, AWS Lambda functions, or Amazon ECS tasks. Compute resources are charged based on the type, size, and configuration of the resources, and the duration and frequency of their usage.

**NEW QUESTION 168**

- (Topic 1)

How can an AWS user conduct security assessments of Amazon EC2 instances, NAT gateways, and Elastic Load Balancers in a way that is approved by AWS?

- A. Flood a target with requests.
- B. Use Amazon Inspector.
- C. Perform penetration testing.
- D. Use the AWS Service Health Dashboard.

**Answer:** B

**Explanation:**

Amazon Inspector is an automated security assessment service that helps improve the security and compliance of applications deployed on AWS. Amazon Inspector automatically assesses applications for exposure, vulnerabilities, and deviations from best practices. After performing an assessment, Amazon Inspector produces a detailed list of security findings prioritized by level of severity.

**NEW QUESTION 171**

- (Topic 1)

A company needs to configure rules to identify threats and protect applications from malicious network access. Which AWS service should the company use to meet these requirements?

- A. AWS Identity and Access Management (IAM)
- B. Amazon QuickSight
- C. AWS WAF
- D. Amazon Detective

**Answer:** C

**Explanation:**

AWS WAF is the AWS service that the company should use to configure rules to identify threats and protect applications from malicious network access. AWS WAF is a web application firewall that helps to filter, monitor, and block malicious web requests based on customizable rules. AWS WAF can be integrated with other AWS services, such as Amazon CloudFront, Amazon API Gateway, and Application Load Balancer. For more information, see [What is AWS WAF?](#) and [How AWS WAF Works](#).

**NEW QUESTION 172**

- (Topic 1)

A company wants its Amazon EC2 instances to operate in a highly available environment, even if there is a natural disaster in a particular geographic area. Which solution achieves this goal?

- A. Use EC2 instances in a single Availability Zone.
- B. Use EC2 instances in multiple AWS Regions.
- C. Use EC2 instances in multiple edge locations.
- D. Use Amazon CloudFront with the EC2 instances configured as the source.

**Answer:** B

**Explanation:**

To achieve high availability in the event of a natural disaster, the company should use EC2 instances in multiple AWS Regions. AWS Regions are geographically isolated areas that consist of multiple Availability Zones. Availability Zones are physically separate locations within an AWS Region that are engineered to be isolated from failures. By using EC2 instances in multiple AWS Regions, the company can ensure that its applications can continue to run even if one Region is affected by a disaster. [AWS Global Infrastructure](#)  
[AWS Well-Architected Framework](#)

**NEW QUESTION 175**

- (Topic 1)

Which AWS service gives users the ability to provision a dedicated and private network connection from their internal network to AWS?

- A. AWS CloudHSM
- B. AWS Direct Connect
- C. AWS VPN
- D. Amazon Connect

**Answer:** B

**Explanation:**

AWS Direct Connect gives users the ability to provision a dedicated and private network connection from their internal network to AWS. AWS Direct Connect links the user's internal network to an AWS Direct Connect location over a standard Ethernet fiber-optic cable. One end of the cable is connected to the user's router, the other to an AWS Direct Connect router. With this connection in place, the user can create virtual interfaces directly to the AWS cloud and Amazon Virtual Private Cloud (Amazon VPC), bypassing internet service providers in the network path.

**NEW QUESTION 176**

- (Topic 1)

Which AWS service or feature is used to Troubleshoot network connectivity issues between Amazon EC2 instances?

- A. AWS Certificate Manager (ACM)
- B. Internet gateway
- C. VPC Flow Logs
- D. AWS CloudHSM

**Answer:** C

**Explanation:**

VPC Flow Logs is the AWS service or feature that is used to troubleshoot network connectivity issues between Amazon EC2 instances. VPC Flow Logs is a

feature that enables users to capture information about the IP traffic going to and from network interfaces in their VPC. VPC Flow Logs can help users monitor and diagnose network-related issues, such as traffic not reaching an instance, or an instance not responding to requests. VPC Flow Logs can be published to Amazon CloudWatch Logs, Amazon S3, or Amazon Kinesis Data Firehose for analysis and storage.

#### NEW QUESTION 181

- (Topic 1)

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

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

**Answer:** D

#### Explanation:

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

#### NEW QUESTION 185

- (Topic 1)

Which AWS service can a company use to perform complex analytical queries?

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

**Answer:** C

#### Explanation:

Amazon Redshift is a fully managed, petabyte-scale data warehouse service in the cloud. You can start with just a few hundred gigabytes of data and scale to a petabyte or more. This enables you to use your data to acquire new insights for your business and customers. Amazon Redshift is designed for complex analytical queries that often involve aggregations and joins across very large tables. Amazon Redshift supports standard SQL and integrates with many existing business intelligence tools<sup>1</sup>.

#### NEW QUESTION 190

- (Topic 1)

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

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

**Answer:** B

#### Explanation:

AWS customers are responsible for enabling encryption of data at rest for Amazon Elastic Block Store (Amazon EBS). Amazon EBS encryption offers a simple encryption solution for your EBS volumes that does not require you to build, maintain, and secure your own key management infrastructure. You can encrypt both the boot and data volumes of your EC2 instances. You can use AWS Key Management Service (AWS KMS) customer master keys (CMKs) or your own CMKs to encrypt your volumes<sup>2</sup>.

#### NEW QUESTION 191

- (Topic 1)

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

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

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

**Answer:** BC

#### Explanation:

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

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

? Write the code for the Lambda functions to define the application logic. Lambda functions are units of code that can be written in any supported programming

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

#### NEW QUESTION 192

- (Topic 1)

A company's information security manager is supervising a move to AWS and wants to ensure that AWS best practices are followed. The manager has concerns about the potential misuse of AWS account root user credentials.

Which of the following is an AWS best practice for using the AWS account root user credentials?

- A. Allow only the manager to use the account root user credentials for normal activities.
- B. Use the account root user credentials only for Amazon EC2 instances from the AWS Free Tier.
- C. Use the account root user credentials only when they alone must be used to perform a required function.
- D. Use the account root user credentials only for the creation of private VPC subnets.

**Answer: C**

#### Explanation:

The AWS best practice for using the AWS account root user credentials is to use them only when they alone must be used to perform a required function. The AWS account root user credentials have full access to all the resources in the account, and therefore pose a security risk if compromised or misused. You should create individual IAM users with the minimum necessary permissions for everyday tasks, and use AWS Organizations to manage multiple accounts. You should also enable multi-factor authentication (MFA) and rotate the password for the root user regularly. Some of the functions that require the root user credentials are changing the account name, closing the account, changing the support plan, and restoring an IAM user's access.

#### NEW QUESTION 194

- (Topic 1)

Which of the following are advantages of moving to the AWS Cloud? (Select TWO.)

- A. The ability to turn over the responsibility for all security to AWS.
- B. The ability to use the pay-as-you-go model.
- C. The ability to have full control over the physical infrastructure.
- D. No longer having to guess what capacity will be required.
- E. No longer worrying about users access controls.

**Answer: BD**

#### Explanation:

The advantages of moving to the AWS Cloud are the ability to use the pay- as-you-go model and no longer having to guess what capacity will be required. The pay-as- you-go model allows the user to pay only for the resources they use, without any upfront or long-term commitments. This reduces the cost and risk of over-provisioning or under- provisioning resources. No longer having to guess what capacity will be required means that the user can scale their resources up or down according to the demand, without wasting money on idle resources or losing customers due to insufficient capacity.

#### NEW QUESTION 199

- (Topic 1)

What does the Amazon S3 Intelligent-Tiering storage class offer?

- A. Payment flexibility by reserving storage capacity
- B. Long-term retention of data by copying the data to an encrypted Amazon Elastic Block Store (AmazonEBS) volume
- C. Automatic cost savings by moving objects between tiers based on access pattern changes
- D. Secure, durable, and lowest cost storage for data archival

**Answer: C**

#### Explanation:

The Amazon S3 Intelligent-Tiering storage class offers automatic cost savings by moving objects between tiers based on access pattern changes. This storage class is designed for data with unknown or changing access patterns. It has two access tiers: frequent access and infrequent access. Objects are stored in the frequent access tier by default, and are moved to the infrequent access tier after 30 consecutive days of no access. If an object in the infrequent access tier is accessed, it is moved back to the frequent access tier. There are no retrieval fees in S3 Intelligent-Tiering, and no additional tiering fees when objects are moved between access tiers within the S3 Intelligent-Tiering storage class.

#### NEW QUESTION 203

- (Topic 1)

Using Amazon Elastic Container Service (Amazon ECS) to break down a monolithic architecture into microservices is an example of:

- A. a loosely coupled architecture.
- B. a tightly coupled architecture.
- C. a stateless architecture.
- D. a stateful architecture.

**Answer: A**

#### Explanation:

Using Amazon Elastic Container Service (Amazon ECS) to break down a monolithic architecture into microservices is an example of a loosely coupled architecture. A loosely coupled architecture is one where the components are independent and can communicate with each other through well-defined interfaces. This allows for greater scalability, flexibility, and resilience. A tightly coupled architecture is one where the components are interdependent and rely on each other for functionality. This can lead to increased complexity, fragility, and difficulty in changing or scaling the system. Amazon ECS Overview AWS Well-Architected Framework

### NEW QUESTION 208

- (Topic 1)

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

- A. Patch AWS network devices.
- B. Set user password rules.
- C. Provide physical security for compute resources.
- D. Configure security groups.
- E. Patch the operating system of an Amazon EC2 instance.

**Answer:** AC

#### Explanation:

The correct answers are A and C because patching AWS network devices and providing physical security for compute resources are tasks that are the responsibility of AWS, according to the AWS shared responsibility model. The AWS shared responsibility model is a framework that defines the division of responsibilities between AWS and the customer for security and compliance. AWS is responsible for the security of the cloud, which includes the global infrastructure, such as the regions, availability zones, and edge locations; the hardware, software, networking, and facilities that run the AWS services; and the virtualization layer that separates the customer instances and storage. The customer is responsible for the security in the cloud, which includes the customer data, the guest operating systems, the applications, the identity and access management, the firewall configuration, and the encryption. The other options are incorrect because they are tasks that are the responsibility of the customer, according to the AWS shared responsibility model. Setting user password rules, configuring security groups, and patching the operating system of an Amazon EC2 instance are all tasks that the customer has to perform to secure their AWS environment.

Reference: AWS Shared Responsibility Model

### NEW QUESTION 213

- (Topic 1)

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

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

**Answer:** B

#### Explanation:

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

### NEW QUESTION 216

- (Topic 1)

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

- A. Pay for compute resources in advance.
- B. Make data-driven decisions to determine cloud architectural design.
- C. Emphasize manual processes to allow for changes.
- D. Test systems at production scale.
- E. Refine operational procedures infrequently.

**Answer:** BD

#### Explanation:

The correct answers are B and D because making data-driven decisions to determine cloud architectural design and testing systems at production scale are AWS Cloud design principles. Making data-driven decisions to determine cloud architectural design means that users should collect and analyze data from their AWS resources and applications to optimize their performance, availability, security, and cost. Testing systems at production scale means that users should simulate real-world scenarios and load conditions to validate the functionality, reliability, and scalability of their systems. The other options are incorrect because they are not AWS Cloud design principles. Paying for compute resources in advance means that users have to invest heavily in data centers and servers before they know how they will use them. This is not a cloud design principle, but rather a traditional IT model. Emphasizing manual processes to allow for changes means that users have to rely on human intervention and coordination to perform operational tasks and updates. This is not a cloud design principle, but rather a source of inefficiency and error. Refining operational procedures infrequently means that users have to stick to the same methods and practices without adapting to the changing needs and feedback. This is not a cloud design principle, but rather a hindrance to innovation and improvement. Reference: AWS Well-Architected Framework

### NEW QUESTION 217

- (Topic 1)

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

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

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

**Answer:** B

#### Explanation:

Amazon Personalize is a fully managed machine learning service that customers can use to generate personalized recommendations for their users. It can also generate user segments based on the users' affinity for certain items or item metadata. Amazon Personalize uses the customers' data to train and deploy custom

recommendation models that can be integrated into their applications. Therefore, the correct answer is B. You can learn more about Amazon Personalize and its use cases from this page.

#### NEW QUESTION 222

- (Topic 3)

A company wants to build a new web application by using AWS services. The application must meet the on-demand load for periods of heavy activity. Which AWS services or resources provide the necessary workload adjustments to meet these requirements? (Select TWO.)

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

**Answer:** BD

#### Explanation:

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

#### NEW QUESTION 223

- (Topic 3)

A company needs to set a maximum spending limit on AWS services each month. The company also needs to set up alerts for when the company reaches its spending limit.

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

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

**Answer:** D

#### Explanation:

AWS Budgets is a service that helps you plan your service usage, service costs, and instance reservations, and track how close your plan is to your budgeted amount. You can set custom budgets that alert you when you exceed (or are forecasted to exceed) your budgeted thresholds. You can also use AWS Budgets to set a maximum spending limit on AWS services each month and set up alerts for when you reach your spending limit. Cost Explorer is a service that enables you to visualize, understand, and manage your AWS costs and usage over time. You can use Cost Explorer to view charts and graphs that show how your costs are trending, identify areas that need further inquiry, and see the impact of your cost management actions. However, Cost Explorer does not allow you to set a maximum spending limit or alerts for your AWS services. AWS Trusted Advisor is a service that provides you real time guidance to help you provision your resources following AWS best practices, including security and performance. It can help you monitor for cost optimization opportunities, such as unused or underutilized resources, but it does not allow you to set a maximum spending limit or alerts for your AWS services. Service Quotas is a service that enables you to view and manage your quotas, also referred to as limits, from a central location. Quotas, also referred to as limits, are the maximum number of resources that you can create in your AWS account. However, Service Quotas does not allow you to set a maximum spending limit or alerts for your AWS services.

#### NEW QUESTION 224

- (Topic 3)

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

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

**Answer:** C

#### Explanation:

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

#### NEW QUESTION 227

- (Topic 3)

A company wants a key-value NoSQL database that is fully managed and serverless. Which AWS service will meet these requirements?

- A. Amazon DynamoDB
- B. Amazon RDS
- C. Amazon Aurora
- D. Amazon Memory DB for Redis

**Answer:** A

#### Explanation:

Amazon DynamoDB is a key-value and document database that delivers single-digit millisecond performance at any scale. It is a fully managed, serverless

database that does not require provisioning, patching, or backup. It offers built-in security, backup and restore, and in-memory caching<sup>3</sup>. Amazon RDS is a relational database service that makes it easy to set up, operate, and scale a relational database in the cloud. It provides cost-efficient and resizable capacity while automating time-consuming administration tasks such as hardware provisioning, database setup, patching, and backups. However, it is not a key-value NoSQL database, and it is not serverless, as it requires you to choose an instance type and size<sup>4</sup>. Amazon Aurora is a MySQL and PostgreSQL-compatible relational database built for the cloud, that combines the performance and availability of traditional enterprise databases with the simplicity and cost-effectiveness of open source databases. However, it is also not a key-value NoSQL database, and it is not serverless, as it requires you to choose an instance type and size. Amazon MemoryDB for Redis is a Redis-compatible, durable, in-memory database service that delivers ultra-fast performance and multi-AZ reliability for the most demanding applications. However, it is also not a key-value NoSQL database, and it is not serverless, as it requires you to choose a node type and size.

#### NEW QUESTION 231

- (Topic 3)

Which characteristic of the AWS Cloud helps users eliminate underutilized CPU capacity'?

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

**Answer: B**

#### Explanation:

Elasticity is a characteristic of the AWS Cloud that helps users eliminate underutilized CPU capacity. Elasticity refers to the ability to dynamically provision and de-provision computing resources as per demand, ensuring that the application or service always has the required resources to operate efficiently. Elasticity helps users optimize performance and costs, as they only pay for the resources they use and avoid wasting resources when the demand is low<sup>3,4,5</sup>. References: 3: Which characteristic of the aws cloud helps users eliminate ..., 4: AWS Elastic Load Balancing and Application Load Balancer, 5: Which characteristic of the AWS Cloud helps users eliminate ...

#### NEW QUESTION 232

- (Topic 3)

Which AWS service converts text to lifelike voices?

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

**Answer: C**

#### Explanation:

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

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

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

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

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

- ? Text to Speech Software – Amazon Polly – Amazon Web Services
- ? What is Text to Speech – Amazon Web Services (AWS)
- ? AWS Amazon Polly - Text to Speech Converter - CodeCanyon
- ? Amazon's Text-To-Speech AI Service Sounds More Natural And ... - Forbes
- ? Working with AWS Amazon Polly Text-to-Speech (TTS) Service
- ? [Automatic Speech Recognition - Amazon Transcribe - AWS]
- ? [Amazon Rekognition – Video and Image - AWS]
- ? [Extract Text & Data - OCR - Amazon Textract - AWS]

#### NEW QUESTION 234

- (Topic 3)

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

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

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

**Answer: A**

#### Explanation:

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

#### NEW QUESTION 237

- (Topic 3)

An ecommerce company wants to provide relevant product recommendations to its customers. The recommendations will include products that are frequently purchased with other products that the customer already purchased. The recommendations also will include products of a specific color and products from the customer's favorite brand.

Which AWS service or feature should the company use to meet these requirements with the LEAST development effort?

- A. Amazon Comprehend
- B. Amazon Forecast
- C. Amazon Personalize
- D. Amazon SageMaker Studio

**Answer: C**

**Explanation:**

Amazon Personalize is a service that provides real-time personalized recommendations based on the user's behavior, preferences, and context. It can also incorporate metadata such as product color and brand to generate more relevant recommendations. Amazon Comprehend is a natural language processing (NLP) service that can analyze text for entities, sentiments, topics, and more. Amazon Forecast is a service that provides accurate time-series forecasting based on machine learning. Amazon SageMaker Studio is a web-based integrated development environment (IDE) for machine learning.

**NEW QUESTION 241**

- (Topic 3)

What can a cloud practitioner use to retrieve AWS security and compliance documents and submit them as evidence to an auditor or regulator?

- A. AWS Certificate Manager
- B. AWS Systems Manager
- C. AWS Artifact
- D. Amazon Inspector

**Answer: C**

**Explanation:**

AWS Artifact is a service that provides on-demand access to AWS security and compliance documents, such as AWS ISO certifications, Payment Card Industry (PCI) reports, and Service Organization Control (SOC) reports. You can download these documents and submit them as evidence to your auditors or regulators to demonstrate the security and compliance of the AWS infrastructure and services that you use. AWS Artifact also allows you to review, accept, and manage AWS agreements, such as the Business Associate Addendum (BAA) for customers who are subject to the Health Insurance Portability and Accountability Act (HIPAA).  
References: AWS Artifact, What is AWS Artifact?

**NEW QUESTION 242**

- (Topic 3)

According to security best practices, how should an Amazon EC2 instance be given access to an Amazon S3 bucket?

- A. Hard code an IAM user's secret key and access key directly in the application, and upload the file.
- B. Store the IAM user's secret key and access key in a text file on the EC2 instance, read the keys, then upload the file.
- C. Have the EC2 instance assume a role to obtain the privileges to upload the file.
- D. Modify the S3 bucket policy so that any service can upload to it at any time.

**Answer: C**

**Explanation:**

According to security best practices, the best way to give an Amazon EC2 instance access to an Amazon S3 bucket is to have the EC2 instance assume a role to obtain the privileges to upload the file. A role is an AWS Identity and Access Management (IAM) entity that defines a set of permissions for making AWS service requests. You can use roles to delegate access to users, applications, or services that don't normally have access to your AWS resources. For example, you can create a role that allows EC2 instances to access S3 buckets, and then attach the role to the EC2 instance. This way, the EC2 instance can assume the role and obtain temporary security credentials to access the S3 bucket. This method is more secure and scalable than storing or hardcoding IAM user credentials on the EC2 instance, as it avoids the risk of exposing or compromising the credentials. It also allows you to manage the permissions centrally and dynamically, and to audit the access using AWS CloudTrail. For more information on how to create and use roles for EC2 instances, see [Using an IAM role to grant permissions to applications running on Amazon EC2 instances](#)

The other options are not recommended for security reasons. Hardcoding or storing IAM user credentials on the EC2 instance is a bad practice, as it exposes the credentials to potential attackers or unauthorized users who can access the instance or the application code. It also makes it difficult to rotate or revoke the credentials, and to track the usage of the credentials. Modifying the S3 bucket policy to allow any service to upload to it at any time is also a bad practice, as it opens the bucket to potential data breaches, data loss, or data corruption. It also violates the principle of least privilege, which states that you should grant only the minimum permissions necessary for a task.

References: [Using an IAM role to grant permissions to applications running on Amazon EC2 instances](#)

**NEW QUESTION 244**

- (Topic 3)

Which scenarios represent the concept of elasticity on AWS? (Select TWO.)

- A. Scaling the number of Amazon EC2 instances based on traffic
- B. Resizing Amazon RDS instances as business needs change
- C. Automatically directing traffic to less-utilized Amazon EC2 instances
- D. Using AWS compliance documents to accelerate the compliance process
- E. Having the ability to create and govern environments using code

**Answer: AB**

**Explanation:**

These are two scenarios that represent the concept of elasticity on AWS. Elasticity means the ability to adjust the resources and capacity of the system in response to changes in demand or environment. Scaling the number of Amazon EC2 instances based on traffic means using services such as AWS Auto Scaling or Elastic Load Balancing to add or remove instances as the traffic increases or decreases. Resizing Amazon RDS instances as business needs change means

using the Amazon RDS console or API to modify the instance type, storage type, or storage size of the database as the workload grows or shrinks. You can learn more about the concept of elasticity on AWS from [this webpage] or [this digital course].

#### NEW QUESTION 248

- (Topic 3)

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

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

**Answer: C**

#### Explanation:

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

#### NEW QUESTION 253

- (Topic 3)

Which of the following is a fully managed graph database service on AWS?

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

**Answer: D**

#### Explanation:

Amazon Neptune is a fully managed graph database service on AWS. A graph database is a type of database that stores and queries data as a network of nodes and edges, representing entities and relationships. Graph databases are useful for applications that deal with highly connected data, such as social networks, recommendation engines, fraud detection, and knowledge graphs<sup>45</sup>. Amazon Neptune is a fast, reliable, and scalable graph database service that supports two popular graph models: property graphs and RDF. Amazon Neptune also supports two open standards for querying graphs: Apache TinkerPop Gremlin and SPARQL. Amazon Neptune handles the heavy lifting of managing the database, such as provisioning, patching, backup, recovery, encryption, and replication<sup>456</sup>. References: 4: Managed Graph Database - Amazon Neptune - AWS, 5: Amazon Neptune – A Fully Managed Graph Database Service, 6: Working with AWS Neptune. Neptune is a fully-managed graph ... - Medium

#### NEW QUESTION 256

- (Topic 3)

A company wants to establish a private network connection between AWS and its corporate network. Which AWS service or feature will meet this requirement?

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

**Answer: C**

#### Explanation:

AWS Direct Connect is a cloud service solution that makes it easy to establish a dedicated network connection from your premises to AWS. Using AWS Direct Connect, you can establish private connectivity between AWS and your datacenter, office, or colocation environment, which in many cases can reduce your network costs, increase bandwidth throughput, and provide a more consistent network experience than internet-based connections<sup>12</sup>. References: 1: Dedicated Network Connection - AWS Direct Connect - AWS, 2: What is AWS Direct Connect? - AWS Direct Connect

#### NEW QUESTION 258

- (Topic 3)

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

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

**Answer: BC**

#### Explanation:

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

**NEW QUESTION 262**

- (Topic 3)

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

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

**Answer:** A**Explanation:**

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

**NEW QUESTION 267**

.....

## Relate Links

**100% Pass Your AWS-Certified-Cloud-Practitioner Exam with ExamBible Prep Materials**

<https://www.exambible.com/AWS-Certified-Cloud-Practitioner-exam/>

## Contact us

We are proud of our high-quality customer service, which serves you around the clock 24/7.

Viste - <https://www.exambible.com/>