



Amazon-Web-Services

Exam Questions CLF-C02

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 3)

A company plans to migrate to the AWS Cloud. The company is gathering information about its on-premises infrastructure and requires information such as the hostname, IP address, and MAC address.

Which AWS service will meet these requirements?

- A. AWS DataSync
- B. AWS Application Migration Service
- C. AWS Application Discovery Service
- D. AWS Database Migration Service (AWS DMS)

Answer: C

Explanation:

AWS Application Discovery Service is a service that helps you plan your migration to the AWS Cloud by collecting usage and configuration data about your on-premises servers and databases. This data includes information such as the hostname, IP address, and MAC address of each server, as well as the performance metrics, network connections, and processes running on them. You can use AWS Application Discovery Service to discover your on-premises inventory, map the dependencies between servers and applications, and estimate the cost and effort of migrating to AWS. You can also export the data to other AWS services, such as AWS Migration Hub and AWS Database Migration Service, to support your migration tasks. AWS Application Discovery Service offers two ways of performing discovery: agentless discovery and agent-based discovery. Agentless discovery uses a virtual appliance that you deploy on your VMware vCenter to collect data from your virtual machines and hosts. Agent-based discovery uses an agent that you install on each of your physical or virtual servers to collect data. You can choose the method that best suits your environment and needs. AWS DataSync is a service that helps you transfer data between your on-premises storage and AWS storage services, such as Amazon S3, Amazon EFS, and Amazon FSx for Windows File Server. AWS DataSync does not collect information about your on-premises infrastructure, but rather focuses on optimizing the data transfer speed, security, and reliability. AWS Application Migration Service is a service that helps you migrate your applications from your on-premises or cloud environment to AWS without making any changes to the applications, their architecture, or the migrated servers. AWS Application Migration Service does not collect information about your on-premises infrastructure, but rather uses a lightweight agent to replicate your servers as Amazon Machine Images (AMIs) and launch them as EC2 instances on AWS. AWS Database Migration Service is a service that helps you migrate your databases from your on-premises or cloud environment to AWS, either as a one-time migration or as a continuous replication. AWS Database Migration Service does not collect information about your on-premises infrastructure, but rather uses a source and a target endpoint to connect to your databases and transfer the data. References: AWS Application Discovery Service, AWS DataSync, AWS Application Migration Service, [AWS Database Migration Service]

NEW QUESTION 2

- (Topic 3)

A developer has been hired by a large company and needs AWS credentials. Which are security best practices that should be followed? (Select TWO.)

- A. Grant the developer access to only the AWS resources needed to perform the job.
- B. Share the AWS account root user credentials with the developer.
- C. Add the developer to the administrator's group in AWS IAM.
- D. Configure a password policy that ensures the developer's password cannot be changed.
- E. Ensure the account password policy requires a minimum length.

Answer: AE

Explanation:

The security best practices that should be followed are A and E.

* A. Grant the developer access to only the AWS resources needed to perform the job. This is an example of the principle of least privilege, which means giving the minimum permissions necessary to achieve a task. This reduces the risk of unauthorized access, data leakage, or accidental damage to AWS resources. You can use AWS Identity and Access Management (IAM) to create users, groups, roles, and policies that grant fine-grained access to AWS resources¹².

* E. Ensure the account password policy requires a minimum length. This is a basic security measure that helps prevent brute-force attacks or guessing of passwords. A longer password is harder to crack than a shorter one. You can use IAM to configure a password policy that enforces a minimum password length, as well as other requirements such as complexity, expiration, and history³⁴.

* B. Share the AWS account root user credentials with the developer. This is a bad practice that should be avoided. The root user has full access to all AWS resources and services, and can perform sensitive actions such as changing billing information, closing the account, or deleting all resources. Sharing the root user credentials exposes your account to potential compromise or misuse. You should never share your root user credentials with anyone, and use them only for account administration tasks⁵.

* C. Add the developer to the administrator's group in IAM. This is also a bad practice that should be avoided. The administrator's group has full access to all AWS resources and services, which is more than what a developer needs to perform their job. Adding the developer to the administrator's group violates the principle of least privilege and increases the risk of unauthorized access, data leakage, or accidental damage to AWS resources. You should create a custom group for the developer that grants only the necessary permissions for their role¹².

* D. Configure a password policy that ensures the developer's password cannot be changed. This is another bad practice that should be avoided. Preventing the developer from changing their password reduces their ability to protect their credentials and comply with security policies. For example, if the developer's password is compromised, they cannot change it to prevent further unauthorized access. Or if the company requires periodic password rotation, they cannot update their password to meet this requirement. You should allow the developer to change their password as needed, and enforce a password policy that sets reasonable rules for password management³⁴.

NEW QUESTION 3

- (Topic 3)

A company wants to migrate its on-premises workloads to the AWS Cloud. The company wants to separate workloads for chargeback to different departments. Which AWS services or features will meet these requirements? (Select TWO.)

- A. Placement groups
- B. Consolidated billing
- C. Edge locations
- D. AWS Config
- E. Multiple AWS accounts

Answer: BE

Explanation:

Consolidated billing is a feature of AWS Organizations that enables customers to consolidate billing and payment for multiple AWS accounts. With consolidated billing, customers can group multiple AWS accounts under one payer account, making it easier to manage billing and track costs across multiple accounts. Consolidated billing also offers benefits such as volume discounts, Reserved Instance discounts, and Savings Plans discounts. Consolidated billing is offered at no additional cost.

Multiple AWS accounts is a feature of AWS Organizations that enables customers to create and manage multiple AWS accounts from a central location. With multiple AWS accounts, customers can isolate workloads for different departments, projects, or environments, and apply granular access controls and policies to each account. Multiple AWS accounts also helps customers improve security, compliance, and governance of their AWS resources⁵⁶. References: 5:

Consolidated billing for AWS Organizations - AWS

Billing, 6: Understanding Consolidated Bills - AWS Billing, 7: AWS Consolidated Billing: Tutorial & Best Practices, 8: Simplifying Your Bills With Consolidated Billing on AWS - Aimably, 9: AWS Consolidated Billing - W3Schools

NEW QUESTION 4

- (Topic 3)

An ecommerce company has migrated its IT infrastructure from an on-premises data center to the AWS Cloud. Which cost is the company's direct responsibility?

- A. Cost of application software licenses
- B. Cost of the hardware infrastructure on AWS
- C. Cost of power for the AWS servers
- D. Cost of physical security for the AWS data center

Answer: A

Explanation:

The cost of application software licenses is the company's direct responsibility when it migrates its IT infrastructure from an on-premises data center to the AWS Cloud. Application software licenses are the agreements that grant users the right to use specific software products, such as operating systems, databases, or applications. Depending on the type and terms of the license, users may need to pay a fee to the software vendor or provider to use the software legally and access its features and updates. When users migrate their IT infrastructure to the AWS Cloud, they can choose to buy new licenses from AWS, bring their own licenses (BYOL), or use a combination of both. However, regardless of the option they choose, they are still responsible for complying with the license terms and paying the license fees to the software vendor or provider. AWS does not charge users for the application software licenses they bring or buy, but only for the AWS resources they use to run their applications. Therefore, the cost of application software licenses is the only cost among the options that is the company's direct responsibility. The other costs are either included in the AWS service fees or covered by AWS.

References: AWS License Manager Pricing, Software licensing: The blind spot in public cloud costs, Cost Optimization tips for SQL Server Licenses on AWS, Microsoft Licensing on AWS

NEW QUESTION 5

- (Topic 3)

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

A company wants to give its customers the ability to view specific data that is hosted in Amazon S3 buckets. The company wants to keep control over the full datasets that the company shares with the customers.

Which S3 feature will meet these requirements?

- A. S3 Storage Lens
- B. S3 Cross-Region Replication (CRR)
- C. S3 Versioning
- D. S3 Access Points

Answer: D

Explanation:

S3 Access Points are a feature of Amazon S3 that allows you to easily manage access to specific data that is hosted in S3 buckets. S3 Access Points are unique hostnames that customers can use to access data in S3 buckets. You can create multiple access points for a single bucket, each with its own name and permissions. You can use S3 Access Points to provide different levels of access to different groups of customers, such as read-only or write-only access. You can also use S3 Access Points to enforce encryption or logging requirements for specific data. S3 Access Points help you keep control over the full datasets that you share with your customers, while simplifying the access management and improving the performance and scalability of your applications.

NEW QUESTION 6

- (Topic 3)

A cloud practitioner needs to obtain AWS compliance reports before migrating an environment to the AWS Cloud How can these reports be generated?

- A. Contact the AWS Compliance team
- B. Download the reports from AWS Artifact
- C. Open a case with AWS Support
- D. Generate the reports with Amazon Made

Answer: B

Explanation:

AWS Artifact is a service that provides on-demand access to security and compliance reports from AWS and Independent Software Vendors (ISVs) who sell their products on AWS Marketplace. You can use AWS Artifact to download auditor-issued reports, certifications, accreditations, and other third-party attestations of AWS compliance with various standards and regulations, such as PCI-DSS, HIPAA, FedRAMP, GDPR, and more¹²³⁴. You can also use AWS Artifact to review, accept, and manage your agreements with AWS and apply them to current and future accounts within your organization². References: 1: Cloud Compliance - Amazon Web Services

(AWS), 2: Security Compliance Management - AWS Artifact - AWS, 3: AWS Compliance Contact Us - Amazon Web Services, 4: AWS SECURITY AND COMPLIANCE QUICK REFERENCE GUIDE

NEW QUESTION 7

- (Topic 3)

Which AWS service or feature offers security for a VPC by acting as a firewall to control traffic in and out of subnets?

- A. AWS Security Hub
- B. Security groups
- C. Network ACL
- D. AWSWAF

Answer: C

Explanation:

A network access control list (network ACL) is a feature that acts as a firewall for controlling traffic in and out of one or more subnets in a virtual private cloud (VPC). AWS Security Hub is a service that provides a comprehensive view of the security posture of AWS accounts and resources. Security groups are features that act as firewalls for controlling traffic at the instance level. AWS WAF is a web application firewall that helps protect web applications from common web exploits.

NEW QUESTION 8

- (Topic 3)

A company is storing sensitive customer data in an Amazon S3 bucket. The company wants to protect the data from accidental deletion or overwriting. Which S3 feature should the company use to meet these requirements?

- A. S3 Lifecycle rules
- B. S3 Versioning
- C. S3 bucket policies
- D. S3 server-side encryption

Answer: B

Explanation:

S3 Versioning is a feature that allows you to keep multiple versions of an object in the same bucket. You can use S3 Versioning to protect your data from accidental deletion or overwriting by enabling it on a bucket or a specific object. S3 Versioning also allows you to restore previous versions of an object if needed. S3 Lifecycle rules are used to automate the transition of objects between storage classes or to expire objects after a certain period of time. S3 bucket policies are used to control access to the objects in a bucket. S3 server-side encryption is used to encrypt the data at rest in S3. References: S3 Versioning, S3 Lifecycle rules, S3 bucket policies, S3 server-side encryption

NEW QUESTION 9

- (Topic 3)

Which cloud computing advantage is a company applying when it uses AWS Regions to increase application availability to users in different countries?

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

Answer: D

Explanation:

Global reach is a cloud computing advantage that a company can apply when it uses AWS Regions to increase application availability to users in different countries. Global reach refers to the ability to deploy applications and services in multiple geographic locations around the world, and to serve customers with low latency and high performance. AWS has the largest and most reliable global infrastructure of any cloud provider, with 25 Regions and 81 Availability Zones across the Americas, Europe, Asia Pacific, Africa, and the Middle East¹²³. By using AWS Regions, a company can choose the best location for its application based on customer proximity, compliance requirements, and disaster recovery strategies²³. References: 1: AWS Global Infrastructure - Amazon Web Services (AWS), 2: Regions and Availability Zones - Amazon Elastic Compute Cloud, 3: AWS Infrastructure: Regions and Availability Zones Explained

NEW QUESTION 10

- (Topic 3)

Which Amazon EC2 instance pricing model can provide discounts of up to 90%?

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

Answer: D

Explanation:

Spot Instances are Amazon EC2 instances that are available at a discounted price compared to On-Demand pricing. Spot Instances use spare EC2 capacity that is not being used by other customers, and the price fluctuates based on supply and demand. Customers can request Spot Instances for their applications and specify the maximum price they are willing to pay per hour. If the Spot price is lower than the customer's bid, the Spot Instance is launched and the customer pays the current Spot price. However, if the Spot price rises above the customer's bid, the Spot Instance is terminated by AWS and the customer is charged for the partial hour of usage. Therefore, Spot Instances can provide discounts of up to 90% or more, but they are not suitable for applications that require continuous or predictable availability. Spot Instances are recommended for applications that are flexible, fault-tolerant, or have low priority, such as batch processing, data analysis, or testing and development.

NEW QUESTION 10

- (Topic 3)

A company has a large number of Linux Amazon EC2 instances across several Availability Zones in an AWS Region. Applications that run on the EC2 instances need access to a common set of files.

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

- A. AWS Backup

- B. Amazon Elastic File System (Amazon EFS)
- C. Amazon Elastic Block Store (Amazon EBS)
- D. AWS Snowball Edge Storage Optimized

Answer: B

Explanation:

Amazon Elastic File System (Amazon EFS) is a service that provides a scalable and elastic file system for Linux-based workloads. It can be mounted on multiple Amazon EC2 instances across different Availability Zones within a region, allowing applications to access a common set of files¹. AWS Backup is a service that provides a centralized and automated way to back up data across AWS services. Amazon Elastic Block Store (Amazon EBS) is a service that provides persistent block storage volumes for Amazon EC2 instances. AWS Snowball Edge Storage Optimized is a device that provides a petabyte-scale data transport and edge computing solution.

NEW QUESTION 15

- (Topic 3)

A company is operating several factories where it builds products. The company needs the ability to process data, store data, and run applications with local system interdependencies that require low latency.

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

- A. AWS IoT Greengrass
- B. AWS Lambda
- C. AWS Outposts
- D. AWS Snowball Edge

Answer: C

Explanation:

AWS Outposts is a service that provides fully managed AWS infrastructure and services on premises. It allows users to run applications that require low latency and local data processing, while seamlessly connecting to the AWS Cloud for a consistent hybrid experience. AWS IoT Greengrass is a service that provides local compute, messaging, data caching, sync, and ML inference capabilities for connected devices. AWS Lambda is a service that allows users to run code without provisioning or managing servers. AWS Snowball Edge is a device that provides a petabyte-scale data transport and edge computing solution.

NEW QUESTION 18

- (Topic 3)

Elasticity in the AWS Cloud refers to which of the following? (Select TWO.)

- A. How quickly an Amazon EC2 instance can be restarted
- B. The ability to rightsized resources as demand shifts
- C. The maximum amount of RAM an Amazon EC2 instance can use
- D. The pay-as-you-go billing model
- E. How easily resources can be procured when they are needed

Answer: BE

Explanation:

Elasticity in the AWS Cloud refers to the ability to acquire resources as you need them and release resources when you no longer need them. In the cloud, you want to do this automatically¹. This means that you can rightsized resources as demand shifts, and you can easily procure resources when they are needed. Elasticity is not related to how quickly an Amazon EC2 instance can be restarted, the maximum amount of RAM an Amazon EC2 instance can use, or the pay-as-you-go billing model. These are aspects of scalability, performance, and cost, respectively².

For more information on elasticity, you can refer to the following sources:

- ? Elasticity - AWS Well-Architected Framework
- ? Elastic - Reactive Systems on AWS
- ? What is the difference between scalability and elasticity?

NEW QUESTION 21

- (Topic 3)

A company is running a monolithic on-premises application that does not scale and is difficult to maintain. The company has a plan to migrate the application to AWS and divide the application into microservices.

Which best practice of the AWS Well-Architected Framework is the company following with this plan?

- A. Integrate functional testing as part of AWS deployment.
- B. Use automation to deploy changes.
- C. Deploy the application to multiple locations.
- D. Implement loosely coupled dependencies.

Answer: D

Explanation:

The company is following the best practice of implementing loosely coupled dependencies by migrating the application to AWS and dividing the application into microservices. Loosely coupled dependencies are a design principle of the AWS Well-Architected Framework that helps to reduce the interdependencies between components and improve the scalability, reliability, and performance of the system. By breaking down the monolithic application into smaller, independent, and modular services, the company can reduce the complexity and maintenance costs, increase the agility and flexibility, and enable faster and more frequent deployments. AWS CloudFormation is an AWS service that provides the ability to manage infrastructure as code. Infrastructure as code is a process of defining and provisioning AWS resources using code or templates, rather than manual actions or scripts. AWS CloudFormation allows users to create and update stacks of AWS resources based on predefined templates that describe the desired state and configuration of the resources. AWS CloudFormation automates and simplifies the deployment and management of AWS resources, and ensures consistency and repeatability across different environments and regions. AWS CloudFormation also supports rollback, change sets, drift detection, and nested stacks features that help users to monitor and control the changes to their infrastructure.

References: Implementing Loosely Coupled Dependencies, What is AWS CloudFormation?

NEW QUESTION 24

- (Topic 3)

A company wants to migrate its database to a managed AWS service that is compatible with PostgreSQL. Which AWS services will meet these requirements? (Select TWO)

- A. Amazon Athena
- B. Amazon RDS
- C. Amazon EC2
- D. Amazon DynamoDB
- E. Amazon Aurora

Answer: BE

Explanation:

Amazon RDS and Amazon Aurora are both managed AWS services that support the PostgreSQL database engine. Amazon RDS makes it easier to set up, operate, and scale PostgreSQL deployments on the cloud, while Amazon Aurora is a cloud-native database engine that is compatible with PostgreSQL and offers higher performance and availability. Amazon Athena is a serverless query service that does not support PostgreSQL, but can analyze data in Amazon S3 using standard SQL. Amazon EC2 is a compute service that allows users to launch virtual machines, but does not provide any database management features. Amazon DynamoDB is a NoSQL database service that is not compatible with PostgreSQL, but offers fast and consistent performance at any scale. References: Hosted PostgreSQL - Amazon RDS for PostgreSQL - AWS, Amazon RDS for PostgreSQL - Amazon Relational Database Service, AWS PostgreSQL: Managed or Self-Managed? - NetApp, AWS Announces Amazon Aurora Supports PostgreSQL 12 - InfoQ, Amazon Aurora vs PostgreSQL | What are the differences? - StackShare

NEW QUESTION 27

- (Topic 3)

What is a customer responsibility when using AWS Lambda according to the AWS shared responsibility model?

- A. Managing the code within the Lambda function
- B. Confirming that the hardware is working in the data center
- C. Patching the operating system
- D. Shutting down Lambda functions when they are no longer in use

Answer: A

Explanation:

According to the AWS shared responsibility model, AWS is responsible for the security of the cloud, while customers are responsible for the security in the cloud. This means that AWS is responsible for the physical servers, networking, and operating system that run Lambda functions, while customers are responsible for the security of their code and AWS IAM to the Lambda service and within their function. Customers need to manage the code within the Lambda function, such as writing, testing, debugging, deploying, and updating the code, as well as ensuring that the code does not contain any vulnerabilities or malicious code that could compromise the security or performance of the function. References: 2: AWS Lambda - Amazon Web Services (AWS), 3: AWS Lambda Documentation, 1: Amazon CLF-C02: What is customer responsibility under AWS ... - PUPUWEB

NEW QUESTION 30

- (Topic 3)

A company wants to run its workload on Amazon EC2 instances for more than 1 year. This workload will run continuously. Which option offers a discounted hourly rate compared to the hourly rate of On-Demand Instances?

- A. AWS Graviton processor
- B. Dedicated Hosts
- C. EC2 Instance Savings Plans
- D. Amazon EC2 Auto Scaling instances

Answer: C

Explanation:

EC2 Instance Savings Plans are a flexible pricing model that offer discounted hourly rates on Amazon EC2 instance usage for a 1 or 3 year term. EC2 Instance Savings Plans provide savings up to 72% off On-Demand rates, in exchange for a commitment to a specific instance family in a chosen AWS Region (for example, M5 in Virginia). These plans automatically apply to usage regardless of size (for example, m5.xlarge, m5.2xlarge, etc.), OS (for example, Windows, Linux, etc.), and tenancy (Host, Dedicated, Default) within the specified family in a Region. With an EC2 Instance Savings Plan, you can change your instance size within the instance family (for example, from c5.xlarge to c5.2xlarge) or the operating system (for example, from Windows to Linux), or move from Dedicated tenancy to Default and continue to receive the discounted rate provided by your EC2 Instance Savings Plan. References: 4: Compute Savings Plans – Amazon Web Services, 5: What are Savings Plans? - Savings Plans, 6: How To Cut Your AWS Bill With Savings Plans (and avoid some common ...), 7: AWS Savings Plans vs Reserved Instances - GorillaStack

NEW QUESTION 33

- (Topic 3)

A company wants to query its server logs to gain insights about its customers' experiences. Which AWS service will store this data MOST cost-effectively?

- A. Amazon Aurora
- B. Amazon Elastic File System (Amazon EFS)
- C. Amazon Elastic Block Store (Amazon EBS)
- D. Amazon S3

Answer: D

Explanation:

Amazon S3 is an AWS service that provides scalable, durable, and cost-effective object storage in the cloud. Amazon S3 can store any amount and type of data, such as server logs, and offers various storage classes with different performance and pricing characteristics. Amazon S3 is the most cost-effective option for storing server logs, as it offers low-cost storage classes, such as S3 Standard-Infrequent Access (S3 Standard-IA) and S3 Intelligent-Tiering, that are suitable for

infrequently accessed or changing access patterns data. Amazon S3 also integrates with other AWS services, such as Amazon Athena and Amazon OpenSearch Service, that can query the server logs directly from S3 without requiring any additional data loading or transformation. References: Amazon S3, Amazon S3 Storage Classes, Querying Data in Amazon S3

NEW QUESTION 38

- (Topic 3)

A company wants to generate a list of IAM users. The company also wants to view the status of various credentials that are associated with the users, such as password, access keys, and multi-factor authentication (MFA) devices. Which AWS service or feature will meet these requirements?

- A. IAM credential report
- B. AWS IAM Identity Center (AWS Single Sign-On)
- C. AWS Identity and Access Management Access Analyzer
- D. AWS Cost and Usage Report

Answer: A

Explanation:

An IAM credential report is a feature of AWS Identity and Access Management (IAM) that allows you to view and download a report that lists all IAM users in your account and the status of their various credentials, such as passwords, access keys, and MFA devices. You can use this report to audit the security status of your IAM users and ensure that they follow the best practices for credential management¹. References: 1: AWS Documentation - IAM User Guide - Getting credential reports for your AWS account

NEW QUESTION 43

- (Topic 3)

Which task must a user perform by using the AWS account root user credentials?

- A. Make changes to AWS production resources.
- B. Change AWS Support plans.
- C. Access AWS Cost and Usage Reports.
- D. Grant auditors' access to an AWS account for a compliance audit.

Answer: B

Explanation:

Changing AWS Support plans is a task that must be performed by using the AWS account root user credentials. The root user is the email address that you used to sign up for AWS. It has complete access to all AWS services and resources in the account. You should use the root user only to perform a few account and service management tasks, such as changing AWS Support plans, closing the account, or changing the account name or email address. Making changes to AWS production resources, accessing AWS Cost and Usage Reports, and granting auditors access to an AWS account for a compliance audit are tasks that can be performed by using IAM users or roles, which are entities that you create in AWS to delegate permissions to access AWS services and resources.

NEW QUESTION 47

- (Topic 3)

Which AWS service can a company use to find security and compliance reports, including International Organization for Standardization (ISO) reports?

- A. AWS Artifact
- B. Amazon CloudWatch
- C. AWS Config
- D. AWS Audit Manager

Answer: A

Explanation:

AWS Artifact is a self-service portal that provides on-demand access to AWS security and compliance reports and select online agreements. You can use AWS Artifact to download AWS service audit reports, such as ISO, PCI, and SOC, and to accept and manage agreements with AWS, such as the Business Associate Addendum (BAA).

NEW QUESTION 52

- (Topic 3)

A company deployed an application on an Amazon EC2 instance. The application ran as expected for 6 months. In the past week, users have reported latency issues. A system administrator found that the CPU utilization was at 100% during business hours. The company wants a scalable solution to meet demand. Which AWS service or feature should the company use to handle the load for its application during periods of high demand?

- A. Auto Scaling groups
- B. AWS Global Accelerator
- C. Amazon Route 53
- D. An Elastic IP address

Answer: A

Explanation:

Auto Scaling groups are a feature that allows users to automatically scale the number of Amazon EC2 instances up or down based on demand or a predefined schedule. Auto Scaling groups can help improve the performance and availability of applications by adjusting the capacity in response to traffic fluctuations¹. AWS Global Accelerator is a service that improves the availability and performance of applications by routing traffic through AWS edge locations². Amazon Route 53 is a service that provides scalable and reliable domain name system (DNS) service³. An Elastic IP address is a static IPv4 address that can be associated with an Amazon EC2 instance⁴.

NEW QUESTION 55

- (Topic 3)

A company is migrating its data center to AWS. The company needs an AWS Support plan that provides chat access to a cloud support engineer 24 hours a day, 7 days a week. The company does not require access to infrastructure event management.

What is the MOST cost-effective AWS Support plan that meets these requirements?

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

Answer: B

Explanation:

AWS Business Support is the most cost-effective AWS Support plan that provides chat access to a cloud support engineer 24/7. AWS Business Support also offers phone and email support, as well as a response time of less than one hour for urgent issues. AWS Business Support does not include access to infrastructure event management, which is a feature of AWS Enterprise Support. AWS Enterprise Support is more expensive and provides additional benefits, such as a technical account manager, a support concierge, and a response time of less than 15 minutes for critical issues. AWS Developer Support and AWS Basic Support do not provide chat access to a cloud support engineer. AWS Developer Support provides email support and a response time of less than 12 hours for general guidance issues. AWS Basic Support provides customer service and account support, as well as access to forums and documentation¹

NEW QUESTION 59

- (Topic 3)

An auditor is preparing for an annual security audit. The auditor requests certification details for a company's AWS hosted resources across multiple Availability Zones in the us-east-1 Region.

How should the company respond to the auditor's request?

- A. Open an AWS Support ticket to request that the AWS technical account manager (TAM) respond and help the auditor.
- B. Open an AWS Support ticket to request that the auditor receive approval to conduct an onsite assessment of the AWS data centers in which the company operates.
- C. Explain to the auditor that AWS does not need to be audited because the company's application is hosted in multiple Availability Zones.
- D. Use AWS Artifact to download the applicable report for AWS security control
- E. Provide the report to the auditor.

Answer: D

Explanation:

AWS Artifact is your go-to, central resource for compliance-related information that matters to you. It provides on-demand access to AWS' security and compliance reports and select online agreements. Reports available in AWS Artifact include our Service Organization Control (SOC) reports, Payment Card Industry (PCI) reports, and certifications from accreditation bodies across geographies and compliance verticals that validate the implementation and operating effectiveness of AWS security controls. Agreements available in AWS Artifact include the Business Associate Addendum (BAA) and the Nondisclosure Agreement (NDA). You can use AWS Artifact to download the applicable report for AWS security controls and provide it to the auditor.

NEW QUESTION 60

- (Topic 3)

Which of the following are general AWS Cloud design principles described in the AWS Well-Architected Framework?

- A. Consolidate key components into monolithic architectures.
- B. Test systems at production scale.
- C. Provision more capacity than a workload is expected to need.
- D. Drive architecture design based on data collected about the workload behavior and requirements.
- E. Make AWS Cloud architectural decisions static, one-time events.

Answer: BD

Explanation:

These are two of the general AWS Cloud design principles described in the AWS Well-Architected Framework. Testing systems at production scale means using tools such as AWS CloudFormation, AWS CodeDeploy, and AWS X-Ray to simulate real-world scenarios and measure the performance, scalability, and availability of the system. Driving architecture design based on data means using tools such as Amazon CloudWatch, AWS CloudTrail, and AWS Config to collect and analyze metrics, logs, and events about the system and use the insights to optimize the system's design and operation. You can learn more about the AWS Well-Architected Framework from this whitepaper or [this digital course].

NEW QUESTION 61

- (Topic 3)

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

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

Answer: D

Explanation:

Amazon ElastiCache is a fully managed in-memory data store and cache service that delivers sub-millisecond response times to applications. You can use ElastiCache as a primary data store for your applications, or as a cache to improve the performance of your existing databases. ElastiCache supports two popular open-source in-memory engines: Redis and Memcached⁵.

NEW QUESTION 64

- (Topic 3)

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

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

Answer: A

Explanation:

The architecture deployment model that the company should use to meet this requirement is A. Multi-Region.

A multi-region deployment model is a cloud computing architecture that distributes an application and its data across multiple geographic regions. A multi-region deployment model enables a company to achieve global reach, high availability, disaster recovery, and performance optimization. By deploying an application in multiple regions, a company can serve customers from the nearest region, reduce latency, increase redundancy, and comply with data sovereignty regulations¹². A single-region deployment model is a cloud computing architecture that runs an application and its data within a single geographic region. A single-region deployment model is simpler and cheaper than a multi-region deployment model, but it has limited scalability, availability, and performance. A single-region deployment model may not be suitable for a company that wants to deploy an application globally, as it may face challenges such as network latency, regional outages, or regulatory compliance¹².

A multi-AZ (Availability Zone) deployment model is a cloud computing architecture that distributes an application and its data across multiple isolated locations within a single region. An Availability Zone is a physically separate location within an AWS Region that has independent power, cooling, and networking. A multi-AZ deployment model enhances the availability and durability of an application by providing redundancy and fault tolerance within a region³⁴.

A single-AZ deployment model is a cloud computing architecture that runs an application and its data within a single Availability Zone. A single-AZ deployment model is the simplest and most cost-effective option, but it has no redundancy or fault tolerance. A single-AZ deployment model may not be suitable for a company that wants to deploy an application globally, as it may face challenges such as network latency, regional outages, or regulatory compliance³⁴.

References:

1: AWS Cloud Computing - W3Schools 2: Understand the Different Cloud Computing Deployment Models Unit - Trailhead 3: Regions and Availability Zones - Amazon Elastic Compute Cloud 4: AWS Reference Architecture Diagrams

NEW QUESTION 68

- (Topic 3)

A company is looking for a managed machine learning (ML) service that can recommend products based on a customer's previous behaviors.
Which AWS service meets this requirement?

- A. Amazon Personalize
- B. Amazon SageMaker
- C. Amazon Pinpoint
- D. Amazon Comprehend

Answer: A

Explanation:

The AWS service that meets the requirement of providing a managed machine learning (ML) service that can recommend products based on a customer's previous behaviors is Amazon Personalize. Amazon Personalize is a fully managed service that enables developers to create personalized recommendations for customers using their own data. Amazon Personalize can automatically process and examine the data, identify what is meaningful, select the right algorithms, and train and optimize a personalized recommendation model². Amazon SageMaker, Amazon Pinpoint, and Amazon Comprehend are other AWS services related to machine learning, but they do not provide the specific functionality of product recommendation.

NEW QUESTION 73

- (Topic 3)

A user needs a relational database but does not have the resources to manage the hardware, resiliency, and replication.
Which AWS service option meets the user's requirements'?

- A. Run MySQL on Amazon Elastic Container Service (Amazon ECS)
- B. Run MySQL on Amazon EC2
- C. Choose Amazon RDS for MySQL
- D. Choose Amazon ElastiCache for Redis

Answer: C

Explanation:

Amazon RDS for MySQL is a fully managed, open-source cloud database service that allows you to easily operate and scale your relational database of choice, including MySQL. With Amazon RDS for MySQL, you don't have to worry about the hardware, resiliency, and replication of your database, as Amazon RDS handles these tasks for you. Amazon RDS for MySQL also provides features such as automated backups, multi-AZ deployments, read replicas, encryption, monitoring, and more. Amazon RDS for MySQL is compatible with the MySQL Community Edition versions 5.7 and 8.0, which means that you can use the same code, applications, and tools that you already use with MySQL⁴⁵⁶⁷. References: 4: Hosted MySQL - Amazon RDS for MySQL - AWS, 5: Amazon RDS for MySQL - Amazon Relational Database Service, 6: Amazon RDS for MySQL —, 7: Managed SQL Database - Amazon Relational Database Service (RDS) - AWS

NEW QUESTION 78

- (Topic 3)

Which AWS service is always provided at no charge?

- A. Amazon S3
- B. AWS Identity and Access Management (IAM)
- C. Elastic Load Balancers
- D. AWS WAF

Answer: B

Explanation:

AWS Identity and Access Management (IAM) is a web service that helps you securely control access to AWS resources. You can use IAM to create and manage AWS users and groups, and use permissions to allow and deny their access to AWS resources. IAM is always provided at no charge¹². References: 1: AWS Identity and Access Management (IAM) - Amazon Web Services (AWS), 2: Which aws service is always provided at no charge? - Brainly.in

NEW QUESTION 81

- (Topic 3)

Which options are AWS Cloud Adoption Framework (AWS CAF) cloud transformation journey recommendations? (Select TWO.)

- A. Envision phase
- B. Align phase
- C. Assess phase
- D. Mobilize phase
- E. Migrate and modernize phase

Answer: AB

Explanation:

The AWS Cloud Adoption Framework (AWS CAF) cloud transformation journey is a four-phase process that helps customers plan and execute their cloud migration and digital transformation. The four phases are:

? Envision phase: This phase focuses on demonstrating how cloud will help accelerate the business outcomes of the customer. It involves identifying and prioritizing transformation opportunities across four domains: business, people, governance, and platform. It also involves associating the transformation initiatives with key stakeholders and measurable business outcomes¹.

? Align phase: This phase focuses on identifying capability gaps across six perspectives: business, people, governance, platform, security, and operations. It also involves identifying cross-organizational dependencies and surfacing stakeholder concerns and challenges. The goal of this phase is to create strategies for improving the cloud readiness, ensure stakeholder alignment, and facilitate relevant organizational change management activities¹.

? Launch phase: This phase focuses on delivering pilot initiatives in production and demonstrating incremental business value. Pilots should be highly impactful and influence future direction. The customer should learn from the pilots and adjust their approach before scaling to full production¹.

? Scale phase: This phase focuses on expanding production pilots and business value to the desired scale and ensuring that the business benefits associated with the cloud investments are realized and sustained¹.

NEW QUESTION 85

- (Topic 3)

Which of the following actions are controlled with AWS Identity and Access Management (IAM)? (Select TWO.)

- A. Control access to AWS service APIs and to other specific resources.
- B. Provide intelligent threat detection and continuous monitoring.
- C. Protect the AWS environment using multi-factor authentication (MFA).
- D. Grant users access to AWS data centers.
- E. Provide firewall protection for applications from common web attacks.

Answer: AC

Explanation:

AWS Identity and Access Management (IAM) is a service that enables you to manage access to AWS services and resources securely. You can use IAM to perform the following actions:

? Control access to AWS service APIs and to other specific resources: You can create users, groups, roles, and policies that define who can access which AWS resources and how. You can also use IAM to grant temporary access to users or applications that need to perform certain tasks on your behalf³

? Protect the AWS environment using multi-factor authentication (MFA): You can enable MFA for your IAM users and root user to add an extra layer of security to your AWS account. MFA requires users to provide a unique authentication code from an approved device or SMS text message, in addition to their user name and password, when they sign in to AWS⁴

NEW QUESTION 90

- (Topic 3)

Which of the following is a software development framework that a company can use to define cloud resources as code and provision the resources through AWS CloudFormation?

- A. AWS CLI
- B. AWS Developer Center
- C. AWS Cloud Development Kit (AWS CDK)
- D. AWS CodeStar

Answer: C

Explanation:

AWS Cloud Development Kit (AWS CDK) is a software development framework that allows you to define cloud resources as code using familiar programming languages, such as TypeScript, Python, Java, .NET, and Go (in Developer Preview). You can use AWS CDK to model your application resources using high-level constructs that provide sensible defaults and best practices, or use low-level constructs that provide full access to the underlying AWS CloudFormation resources. AWS CDK synthesizes your code into AWS CloudFormation templates that you can deploy using the AWS CDK CLI or the AWS Management Console. AWS CDK also integrates with other AWS services, such as AWS CodeCommit, AWS CodeBuild, AWS CodePipeline, AWS Lambda, Amazon EC2, Amazon S3, and more, to help you automate your development and deployment processes. AWS CDK is an open-source framework that you can extend and contribute to. References: Cloud Development Framework - AWS Cloud Development Kit - AWS, AWS Cloud Development Kit Documentation, AWS Cloud Development Kit - Wikipedia, AWS CDK Intro Workshop | AWS CDK Workshop

NEW QUESTION 95

- (Topic 3)

Which AWS service or feature offers security for a VPC by acting as a firewall to control traffic in and out of subnets?

- A. AWS Security Hub

- B. Security groups
- C. Network ACL
- D. AWSWAF

Answer: C

Explanation:

A network access control list (network ACL) is a feature that acts as a firewall for controlling traffic in and out of one or more subnets in a virtual private cloud (VPC). Network ACLs can be configured with rules that allow or deny traffic based on the source and destination IP addresses, ports, and protocols¹. AWS Security Hub is a service that provides a comprehensive view of the security posture of AWS accounts and resources². Security groups are features that act as firewalls for controlling traffic at the instance level³. AWS WAF is a web application firewall that helps protect web applications from common web exploits⁴.

NEW QUESTION 96

- (Topic 3)

Which benefit does AWS offer exclusively to users who have an AWS Enterprise Support plan?

- A. Access to a technical project manager
 - B. Access to a technical account manager (TAM)
 - C. Access to a cloud support engineer
 - D. Access to a solutions architect
- A company wants to automatically set up and govern a multi-account AWS environment.

Answer: B

Explanation:

AWS Enterprise Support plan is the highest level of support that AWS offers to its customers. One of the exclusive benefits of this plan is the access to a technical account manager (TAM), who is a dedicated point of contact for guidance, advocacy, and support². A technical project manager, a cloud support engineer, and a solutions architect are not exclusive benefits of the AWS Enterprise Support plan, as they are also available to customers with lower-tier support plans or through other AWS services or programs^{3,4,5}.

NEW QUESTION 98

- (Topic 3)

A company needs to search for text in documents that are stored in Amazon S3. Which AWS service will meet these requirements?

- A. Amazon Kendra
- B. Amazon Rekognition
- C. Amazon Polly
- D. Amazon Lex

Answer: A

Explanation:

Amazon Kendra is a highly accurate and easy to use intelligent search service powered by machine learning. It enables users to easily find the content they are looking for, even when it is scattered across multiple locations and content repositories within their organization. Amazon Kendra supports natural language queries, and can search for text in documents stored in Amazon S3, as well as other sources such as SharePoint, OneDrive, Salesforce, ServiceNow, and more¹. Amazon Rekognition is a computer vision service that makes it easy to add image and video analysis to applications. It can detect objects, faces, text, scenes, activities, and emotions in images and videos. However, it is not designed for searching for text in documents stored in Amazon S3².

Amazon Polly is a text-to-speech service that turns text into lifelike speech. It can create audio versions of books, articles, podcasts, and more. However, it is not designed for searching for text in documents stored in Amazon S3³.

Amazon Lex is a service for building conversational interfaces using voice and text. It can create chatbots that can interact with users using natural language. However, it is not designed for searching for text in documents stored in Amazon S3⁴.

References:

- ? Amazon Kendra – Intelligent Search Service Powered by Machine Learning
- ? Amazon Rekognition – Video and Image - AWS
- ? Amazon Polly – Text-to-Speech Service - AWS
- ? Amazon Lex – Build Conversation Bots - AWS

NEW QUESTION 103

- (Topic 3)

Which option is the default pricing model for Amazon EC2 instances?

- A. On-Demand Instances
- B. Savings Plans
- C. Spot Instances
- D. Reserved Instances

Answer: A

Explanation:

On-Demand Instances are the default pricing model for Amazon EC2 instances. They allow users to pay for compute capacity by the second, with no long-term commitments or upfront payments. They are suitable for applications with short-term, irregular, or unpredictable workloads that cannot be interrupted³. Savings Plans are a pricing model that offer significant savings on Amazon EC2 and AWS Fargate usage, in exchange for a commitment to a consistent amount of usage (measured in \$/hour) for a 1- year or 3-year term. Spot Instances are a pricing model that offer spare Amazon EC2 compute capacity at up to 90% discount compared to On-Demand prices, but they can be interrupted by AWS with a two-minute notice when the demand exceeds the supply. Reserved Instances are a pricing model that offer up to 75% discount compared to On- Demand prices, in exchange for a commitment to use a specific instance type and size in a specific region for a 1-year or 3-year term.

NEW QUESTION 105

- (Topic 3)

Which type of AWS storage is ephemeral and is deleted when an Amazon EC2 instance is stopped or terminated?

- A. Amazon Elastic Block Store (Amazon EBS)
- B. Amazon EC2 instance store
- C. Amazon Elastic File System (Amazon EFS)
- D. Amazon S3

Answer: B

Explanation:

Amazon EC2 instance store provides temporary block-level storage for your EC2 instance. This storage is located on disks that are physically attached to the host computer. Instance store is ideal for temporary storage of information that changes frequently, such as buffers, caches, scratch data, and other temporary content. It can also be used to store temporary data that you replicate across a fleet of instances, such as a load-balanced pool of web servers. An instance store consists of one or more instance store volumes exposed as block devices. The size of an instance store as well as the number of devices available varies by instance type and instance size. The virtual devices for instance store volumes are ephemeral[0-23]. Instance types that support one instance store volume have ephemeral0. Instance types that support two or more instance store volumes have ephemeral0, ephemeral1, and so on. Instance store pricing Instance store volumes are included as part of the instance's usage cost. The data on an instance store volume persists even if the instance is rebooted. However, the data does not persist if the instance is stopped, hibernated, or terminated. When the instance is stopped, hibernated, or terminated, every block of the instance store volume is cryptographically erased. Therefore, do not rely on instance store volumes for valuable, long-term data. If you need to retain the data stored on an instance store volume beyond the lifetime of the instance, you need to manually copy that data to more persistent storage, such as an Amazon EBS volume, an Amazon S3 bucket, or an Amazon EFS file system. There are some events that can result in your data not persisting throughout the lifetime of the instance. The following table indicates whether data on instance store volumes is persisted during specific events, for both virtualized and bare metal instances¹. References: Amazon EC2 instance store - Amazon Elastic Compute Cloud

NEW QUESTION 109

- (Topic 3)

A company website is experiencing DDoS attacks.

Which AWS service can help protect the company website against these attacks?

- A. AWS Resource Access Manager
- B. AWS Amplify
- C. AWS Shield
- D. Amazon GuardDuty

Answer: C

Explanation:

AWS Shield is a managed DDoS protection service that safeguards applications running on AWS from distributed denial of service (DDoS) attacks. DDoS attacks are malicious attempts to disrupt the normal functioning of a website or application by overwhelming it with a large volume of traffic from multiple sources. AWS Shield provides two tiers of protection: Standard and Advanced. AWS Shield Standard is automatically enabled for all AWS customers at no additional cost. It protects your AWS resources, such as Amazon CloudFront, AWS Global Accelerator, and Amazon Route 53, from the most common and frequently occurring network and transport layer DDoS attacks. AWS Shield Advanced is an optional paid service that provides additional protection for your AWS resources and applications, such as Amazon Elastic Compute Cloud (Amazon EC2), Elastic Load Balancing (ELB), Amazon Simple Storage Service (Amazon S3), Amazon Relational Database Service (Amazon RDS), and AWS Elastic Beanstalk. AWS Shield Advanced offers enhanced detection and mitigation capabilities, 24/7 access to the AWS DDoS Response Team (DRT), real-time visibility and reporting, and cost protection against DDoS-related spikes in your AWS bill¹². References: AWS Shield, What is a DDOS Attack & How to Protect Your Site Against One

NEW QUESTION 110

- (Topic 3)

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

Which AWS service will meet these requirements?

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

Answer: A

Explanation:

Amazon CloudWatch is a service that provides monitoring and observability for AWS resources and applications. Users can create custom dashboards to collect and visualize metrics, logs, alarms, and events from different sources⁵. AWS X-Ray is a service that provides distributed tracing and analysis for applications. AWS Systems Manager is a service that provides operational management for AWS resources and applications. AWS CloudTrail is a service that provides governance, compliance, and auditing for AWS account activity.

NEW QUESTION 114

- (Topic 3)

A company wants to store data with high availability, encrypt the data at rest, and have direct access to the data over the internet.

Which AWS service will meet these requirements MOST cost-effectively?

- A. Amazon Elastic Block Store (AmazonEBS)
- B. Amazon S3
- C. Amazon Elastic File System (Amazon EFS)
- D. AWS Storage Gateway

Answer: C

Explanation:

Amazon Elastic File System (Amazon EFS) provides a simple, scalable, fully managed elastic NFS file system for use with AWS Cloud services and on-premises resources. It is built to scale on demand to petabytes without disrupting applications, growing and shrinking automatically as you add and remove files, eliminating the need to provision and manage capacity to accommodate growth. Amazon EFS offers two storage classes: the Standard storage class, and the Infrequent Access storage class (EFS IA).

EFS IA provides price/performance that is cost-optimized for files not accessed every day. Amazon EFS encrypts data at rest and in transit, and supports direct access over the internet⁴.

NEW QUESTION 117

- (Topic 3)

A company seeks cost savings in exchange for a commitment to use a specific amount of an AWS service or category of AWS services for 1 year or 3 years. Which AWS pricing model or offering will meet these requirements?

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

Answer: B

Explanation:

Savings Plans are an AWS pricing model or offering that can meet the requirements of seeking cost savings in exchange for a commitment to use a specific amount of an AWS service or category of AWS services for 1 year or 3 years. Savings Plans are flexible plans that offer significant discounts on AWS compute usage, such as EC2, Lambda, and Fargate. The company can choose from two types of Savings Plans: Compute Savings Plans and EC2 Instance Savings Plans. Compute Savings Plans provide the most flexibility and apply to any eligible compute usage, regardless of instance family, size, region, operating system, or tenancy. EC2 Instance Savings Plans provide more savings and apply to a specific instance family within a region. The company can select the amount of compute usage per hour (e.g., \$10/hour) that they want to commit to for the duration of the plan (1 year or 3 years). The company will pay the discounted Savings Plan rate for the amount of usage that matches their commitment, and the regular on-demand rate for any usage beyond that

NEW QUESTION 120

- (Topic 3)

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

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

Answer: BC

Explanation:

Amazon CloudFront and AWS Global Accelerator are two AWS services that make use of global edge locations. Edge locations are AWS sites that are deployed worldwide in major cities and places with a high population. Edge locations are used to cache data and reduce latency for end-user access¹.

Amazon CloudFront is a content delivery network (CDN) service that securely delivers data, videos, applications, and APIs to customers globally with low latency and high

transfer speeds. Amazon CloudFront uses a global network of over 200 edge locations and 13 regional edge caches to cache your content closer to your viewers, improving performance and reducing costs²³.

AWS Global Accelerator is a networking service that improves the availability and performance of your applications with local or global users. AWS Global Accelerator uses the AWS global network to route user traffic to the optimal endpoint based on health, performance, and policies. AWS Global Accelerator uses over 100 edge locations to bring your application endpoints closer to your users, reducing network hops and improving user experience⁴⁵. References: 1: AWS for the Edge - Amazon Web Services

(AWS), 2: Content Delivery Network (CDN) - Amazon CloudFront - AWS, 3: Amazon CloudFront Documentation, 4: AWS Global Accelerator - Amazon Web Services, 5: AWS Global Accelerator Documentation

NEW QUESTION 124

- (Topic 3)

A company is expecting a short-term spike in internet traffic for its application. During the traffic increase, the application cannot be interrupted. The company also needs to minimize cost and maximize flexibility.

A company needs to use a serverless interactive query service to analyze data in Amazon S3. The query service must support standard SQL.

Which AWS service will meet these requirements?

- A. Amazon Redshift
- B. AWS Glue
- C. Amazon Athena
- D. Amazon Kinesis Data Streams

Answer: C

Explanation:

Amazon Athena is a serverless interactive query service that makes it easy to analyze data in Amazon S3 using standard SQL. Athena is ideal for quick, ad-hoc querying but it can also handle complex analysis, including large joins, window functions, and arrays. Athena scales automatically—executing queries in parallel—so results are fast, even with large datasets and complex queries. Amazon Redshift is a fully managed, petabyte-scale data warehouse service that can run complex analytic queries against structured and semi-structured data using standard SQL. However, it is not a serverless service and requires provisioning and managing clusters of nodes. AWS Glue is a fully managed extract, transform, and load (ETL) service that makes it easy to prepare and load your data for analytics. However, it is not a query service and does not support standard SQL. Amazon Kinesis Data Streams is a service that enables you to build custom applications that process or analyze streaming data for specialized needs. However, it is not a query service and does not support standard SQL.

NEW QUESTION 125

- (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](#)¹

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 129

- (Topic 3)

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

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

Answer: A

Explanation:

Amazon S3 Glacier Deep Archive is the lowest-cost storage class in the cloud. It is designed for long-term data archiving that is rarely accessed. It offers a retrieval time of 12 hours and a durability of 99.999999999% (11 9's). It is ideal for data that must be retained for 7 years or longer to meet regulatory compliance requirements.

NEW QUESTION 134

- (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 workforce¹. References: [People Perspective - AWS Cloud Adoption Framework](#)

NEW QUESTION 135

- (Topic 3)

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

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

Answer: B

Explanation:

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

NEW QUESTION 139

- (Topic 3)

Which AWS services can limit manual errors by consistently provisioning AWS resources in multiple environments?

- A. AWS Config
- B. AWS CodeStar
- C. AWS CloudFormation

- D. AWS Cloud Development Kit (AWS CDK)
- E. AWS CodeBuild

Answer: CD

Explanation:

AWS CloudFormation and AWS Cloud Development Kit (AWS CDK) are AWS services that can limit manual errors by consistently provisioning AWS resources in multiple environments. AWS CloudFormation is a service that enables you to model and provision AWS resources using templates. You can use AWS CloudFormation to define the AWS resources and their dependencies that you need for your applications, and to automate the creation and update of those resources across multiple environments, such as development, testing, and production. AWS CloudFormation helps you ensure that your AWS resources are configured consistently and correctly, and that you can easily replicate or modify them as needed. AWS Cloud Development Kit (AWS CDK) is a service that enables you to use familiar programming languages, such as Python, TypeScript, Java, and C#, to define and provision AWS resources. You can use AWS CDK to write code that synthesizes into AWS CloudFormation templates, and to leverage the existing libraries and tools of your preferred language. AWS CDK helps you reduce the complexity and errors of writing and maintaining AWS CloudFormation templates, and to apply the best practices and standards of software development to your AWS infrastructure.

NEW QUESTION 140

- (Topic 3)

A company's headquarters is located on a different continent from where the majority of the company's customers live. The company wants an AWS Cloud environment setup that will provide the lowest latency to the customers.

A company wants to automate the creation of new AWS accounts and automatically prevent all users from creating Amazon EC2 instances.

Which AWS service provides this functionality?

- A. AWS Service Catalog
- B. AWS Organizations
- C. EC2 Image Builder
- D. AWS Systems Manager

Answer: B

Explanation:

AWS Organizations is a service that enables you to create and manage multiple AWS accounts centrally. You can use AWS Organizations to automate account creation, apply policies to control access and permissions, and consolidate billing across your accounts. You can also use AWS Organizations to prevent users from creating Amazon EC2 instances in certain regions or with certain configurations²

NEW QUESTION 144

- (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 practices⁴.

NEW QUESTION 145

- (Topic 3)

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

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

Answer: AD

Explanation:

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

NEW QUESTION 146

- (Topic 3)

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

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

Answer: B

Explanation:

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

NEW QUESTION 149

- (Topic 3)

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

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

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

Answer: B

Explanation:

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

NEW QUESTION 150

- (Topic 3)

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

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

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

Answer: B

Explanation:

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

NEW QUESTION 154

- (Topic 3)

A company uses AWS for its web application. The company wants to minimize latency and perform compute operations for the application as close to end users as possible.

Which AWS service or infrastructure component will provide this functionality?

- A. AWS Regions
- B. Availability Zones
- C. Edge locations
- D. AWS Direct Connect

Answer: C

Explanation:

Edge locations are sites that Amazon CloudFront uses to cache copies of your content for faster delivery to users at any location. You can use Amazon CloudFront to deliver your entire website, including dynamic, static, streaming, and interactive content using a global network of edge locations. Requests for your content are automatically routed to the nearest edge location, so content is delivered with the best possible performance. Edge locations can also host AWS Lambda functions to perform compute operations for your web application as close to end users as possible.

NEW QUESTION 155

- (Topic 3)

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

Which AWS service meets these requirements?

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

Answer: D

Explanation:

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

NEW QUESTION 157

- (Topic 3)

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

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

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

Answer: D

Explanation:

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

NEW QUESTION 159

- (Topic 3)

A company wants to set up a high-speed connection between its data center and its applications that run on AWS. The company must not transfer data over the internet.

Which action should the company take to meet these requirements?

- A. Transfer data to AWS by using AWS Snowball.
- B. Transfer data to AWS by using AWS Storage Gateway.
- C. Set up a VPN connection between the data center and an AWS Region.
- D. Set up an AWS Direct Connect connection between the company network and AWS.

Answer: D

Explanation:

AWS Direct Connect is a cloud service solution that makes it easy to establish a dedicated network connection from a customer's premises to AWS. AWS Direct Connect does not involve the public internet, and therefore can reduce network costs, increase bandwidth throughput, and provide a more consistent network experience than internet-based connections. AWS Snowball is a petabyte-scale data transport service that uses secure devices to transfer large amounts of data into and out of the AWS Cloud. AWS Storage Gateway is a hybrid cloud storage service that gives customers on-premises access to virtually unlimited cloud storage. A VPN connection enables customers to establish a secure and private connection between their network and AWS.

NEW QUESTION 161

- (Topic 3)

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

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

Answer: A

Explanation:

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

NEW QUESTION 166

- (Topic 3)

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

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

- A. Hardware lifecycle management
- B. Application optimization
- C. Server maintenance

D. Power, network, and cooling provisioning

Answer: B

Explanation:

Amazon RDS is a managed database service that handles most of the common database administration tasks, such as hardware provisioning, server maintenance, backup and recovery, patching, scaling, and replication. However, Amazon RDS does not optimize the application that interacts with the database. The company is still responsible for tuning the performance, security, and availability of the application according to its business requirements and best practices¹². References:

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

NEW QUESTION 170

- (Topic 3)

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

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

Answer: C

Explanation:

Amazon EC2 usage is calculated by either the hour or the second based on the size of the instance, operating system, and the AWS Region where the instances are launched. Pricing is per instance-hour consumed for each instance, from the time an instance is launched until it's terminated or stopped. Each partial instance-hour consumed is billed per-second for Linux instances and as a full hour for all other instance types¹. Therefore, the customer will be billed for 3 hours and 6 minutes for running an On-Demand Amazon Linux EC2 instance for 3 hours, 5 minutes, and 6 seconds. References: Understand Amazon EC2 instance-hours billing

NEW QUESTION 171

- (Topic 3)

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

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

Answer: B

Explanation:

AWS Budgets allows you to set custom budgets that alert you when your costs or usage exceed (or are forecasted to exceed) your budgeted amount. You can also use AWS Budgets to set reservation utilization or coverage targets and receive alerts when your utilization drops below the threshold you define. AWS Budgets provides you with a comprehensive view of your cost and usage, as well as your reservation utilization and coverage¹.

NEW QUESTION 175

- (Topic 2)

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

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

Answer: A

Explanation:

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

NEW QUESTION 179

- (Topic 2)

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

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

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

Answer: B

Explanation:

Amazon EC2 On-Demand Instances are instances that you pay for by the second, with no long-term commitments or upfront payments⁴. This option is suitable

for applications that have unpredictable or intermittent workloads, such as the one described in the question. Amazon EC2 Standard Reserved Instances are instances that you purchase for a one-year or three-year term, and pay a lower hourly rate compared to On-Demand Instances. This option is suitable for applications that have steady state or predictable usage. AWS Wavelength is a service that enables developers to build applications that deliver ultra-low latency to mobile devices and users by deploying AWS compute and storage at the edge of the 5G network. This option is not relevant for the application described in the question. Application Load Balancer is a type of load balancer that operates at the application layer and distributes traffic based on the content of the request. This option is not a service or feature to host the application, but rather to balance the traffic among multiple instances.

NEW QUESTION 183

- (Topic 2)

A company needs help managing multiple AWS linked accounts that are reported on a consolidated bill. Which AWS Support plan includes an AWS concierge whom the company can ask for assistance?

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

Answer: B

Explanation:

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

NEW QUESTION 188

- (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.

NEW QUESTION 193

- (Topic 2)

Which controls are the responsibility of both AWS and AWS customers, according to the AWS shared responsibility model? (Select TWO.)

- A. Physical and environmental controls
- B. Patch management
- C. Configuration management
- D. Account structures
- E. Choice of the AWS Region where data is stored

Answer: BC

Explanation:

Patch management and configuration management are controls that are the responsibility of both AWS and AWS customers, according to the AWS shared responsibility model. Patch management is the process of applying updates to software and applications to fix vulnerabilities, bugs, or performance issues. Configuration management is the process of defining and maintaining the settings and parameters of systems and applications to ensure their consistency and reliability. AWS is responsible for patching and configuring the software and services that it manages, such as the AWS global infrastructure, the hypervisor, and the AWS managed services. The customer is responsible for patching and configuring the software and services that they manage, such as the guest operating system, the applications, and the AWS customer-managed services. Physical and environmental controls are the responsibility of AWS, according to the AWS shared responsibility model. Physical and environmental controls are the measures that protect the physical security and availability of the AWS global infrastructure, such as power, cooling, fire suppression, and access control. AWS is responsible for maintaining these controls and ensuring the resilience and reliability of the AWS Cloud. Account structures are the responsibility of the customer, according to the AWS shared responsibility model. Account structures are the ways that customers organize and manage their AWS accounts and resources, such as using AWS Organizations, IAM users and roles, resource tagging, and billing preferences. The customer is responsible for creating and configuring these structures and ensuring the security and governance of their AWS environment. Choice of the AWS Region where data is stored is the responsibility of the customer, according to the AWS shared responsibility model. AWS Regions are geographic areas that consist of multiple isolated Availability Zones. Customers can choose which AWS Region to store their data and run their applications, depending on their latency, compliance, and cost requirements. The customer is responsible for selecting the appropriate AWS Region and ensuring the data sovereignty and regulatory compliance of their data.

NEW QUESTION 196

- (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 troubleshooting¹.

NEW QUESTION 201

- (Topic 2)

A company does not want to rely on elaborate forecasting to determine its usage of compute resources. Instead, the company wants to pay only for the resources that it uses. The company also needs the ability to increase or decrease its resource usage to meet business requirements. Which pillar of the AWS Well-Architected Framework aligns with these requirements?

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

Answer: D

Explanation:

Cost optimization is the pillar of the AWS Well-Architected Framework that aligns with the requirements of not relying on elaborate forecasting and paying only for the resources that are used. The cost optimization pillar focuses on the ability of a system to deliver business value at the lowest price point. Cost optimization involves using the right AWS services and resources for the workload, measuring and monitoring the cost and usage, and continuously improving the cost efficiency. Cost optimization also leverages the benefits of the AWS Cloud, such as pay-as-you-go pricing, elasticity, and scalability. For more information, see [Cost Optimization Pillar] and [Cost Optimization].

NEW QUESTION 205

- (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 time¹²

NEW QUESTION 207

- (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 210

- (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 215

- (Topic 2)

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

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

Answer: A

Explanation:

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

NEW QUESTION 219

- (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 customer.

NEW QUESTION 221

- (Topic 2)

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

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

Answer: D

Explanation:

AWS Outposts is a service that offers fully managed and configurable compute and storage racks built with AWS-designed hardware that allow you to run your workloads on premises and seamlessly connect to AWS services in the cloud. AWS Outposts is ideal for workloads that require low latency, local data processing, or local data storage. With AWS Outposts, you can use the same AWS APIs, tools, and infrastructure across on premises and the cloud to deliver a truly consistent hybrid experience. Availability Zones are isolated locations within each AWS Region that are engineered to be fault-tolerant and provide high availability. AWS Local Zones are extensions of AWS Regions that are placed closer to large population, industry, and IT centers where no AWS Region exists today. AWS Wavelength is a service that enables developers to build applications that deliver ultra-low latency to mobile devices and users by deploying AWS

compute and storage at the edge of the 5G network. None of these services or features can help you host a critical application with minimum latency at a remote site that has a slow internet connection.

NEW QUESTION 224

- (Topic 2)

A company moves a workload to AWS to run on Amazon EC2 instances. The company needs to run the workload in the most cost-effective way. What can the company do to meet this requirement?

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

Answer: D

Explanation:

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

NEW QUESTION 226

- (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 cloud³⁴

NEW QUESTION 229

- (Topic 2)

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

Which AWS service should the company use?

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

Answer: B

Explanation:

The company should use AWS WAF to safeguard the website from SQL injection or cross-site scripting. AWS WAF is a web application firewall that helps protect web applications from common web exploits that could affect availability, compromise security, or consume excessive resources. The company can use AWS WAF to create custom rules that block malicious requests that match certain patterns, such as SQL injection or cross-site scripting. AWS WAF can be applied to web applications that are behind an Application Load Balancer, Amazon CloudFront, or Amazon API Gateway. Amazon GuardDuty, AWS Trusted Advisor, and Amazon Inspector are not the best services to use for this purpose. Amazon GuardDuty is a threat detection service that monitors for malicious activity and unauthorized behavior across the AWS accounts and resources. AWS Trusted Advisor is a service that provides best practice recommendations for cost optimization, performance, security, and fault tolerance. Amazon Inspector is a service that assesses the security and compliance of applications running on Amazon EC2 instances¹²

NEW QUESTION 232

- (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 236

- (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. 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 238

- (Topic 2)

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

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

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

Answer: D

Explanation:

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

NEW QUESTION 241

- (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. 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 242

- (Topic 2)

A company is using Amazon RDS.

A company is launching a critical business application in an AWS Region. How can the company increase resilience for this application?

- A. Deploy a copy of the application in another AWS account.
- B. Deploy the application by using multiple VPCs.

- C. Deploy the application by using multiple subnets.
- D. Deploy the application by using multiple Availability Zones.

Answer: D

Explanation:

Deploying the application by using multiple Availability Zones is the best way to increase resilience for the application. According to the Amazon RDS User Guide, "Amazon RDS provides high availability and failover support for DB instances using Multi-AZ deployments. In a Multi-AZ deployment, Amazon RDS automatically provisions and maintains a synchronous standby replica in a different Availability Zone. The primary DB instance is synchronously replicated across Availability Zones to a standby replica to provide data redundancy, eliminate I/O freezes, and minimize latency spikes during system backups."⁴ Deploying a copy of the application in another AWS account, using multiple VPCs, or using multiple subnets do not provide the same level of resilience as using multiple Availability Zones.

NEW QUESTION 246

- (Topic 2)

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

Which AWS service will meet this requirement?

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

Answer: A

Explanation:

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

NEW QUESTION 249

- (Topic 2)

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

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

Answer: C

Explanation:

AWS is responsible for maintaining the physical and environmental controls of the AWS Cloud, such as power, cooling, fire suppression, and physical security¹. The customer is responsible for managing the IAM user permissions, creating security group rules for outbound access, applying Amazon EC2 operating system patches, and other aspects of security in the cloud¹.

NEW QUESTION 254

- (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 255

- (Topic 2)

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

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

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

Answer: AC

Explanation:

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

NEW QUESTION 259

- (Topic 2)

A company wants to develop a shopping application that records customer orders. The application needs to use an AWS managed database service to store data. Which AWS service should the company use to meet these requirements?

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

Answer: A

Explanation:

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

NEW QUESTION 261

- (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 265

- (Topic 2)

A company is preparing to launch a redesigned website on AWS. Users from around the world will download digital handbooks from the website. Which AWS solution should the company use to provide these static files securely?

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

Answer: B

Explanation:

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

NEW QUESTION 267

- (Topic 2)

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

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

Answer: A

Explanation:

AWS Artifact is the AWS service or tool that provides on-demand access to AWS security and compliance reports and AWS online agreements. AWS Trusted Advisor is a tool that provides real-time guidance to help users provision their resources following AWS best practices. Amazon Inspector is a service that helps users improve the security and compliance of their applications. AWS Billing console is a tool that helps users manage their AWS costs and usage. These concepts are explained in the AWS Cloud Practitioner Essentials course3.

NEW QUESTION 270

- (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 271

- (Topic 1)

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

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

Answer: A

Explanation:

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

NEW QUESTION 274

- (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 275

.....

Relate Links

100% Pass Your CLF-C02 Exam with ExamBible Prep Materials

<https://www.exambible.com/CLF-C02-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/>