

## CLF-C02 Dumps

### AWS Certified Cloud Practitioner

<https://www.certleader.com/CLF-C02-dumps.html>



**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)

Which abilities are benefits of the AWS Cloud? (Select TWO.)

- A. Trade variable expenses for capital expenses.
- B. Deploy globally in minutes.
- C. Plan capacity in advance of deployments.
- D. Take advantage of economies of scale.
- E. Reduce dependencies on network connectivity.

**Answer:** AB

**Explanation:**

The AWS Cloud offers many benefits, such as:

? Trade variable expenses for capital expenses: You can pay only for the resources you use, instead of investing in fixed costs upfront. This reduces the risk and complexity of planning and managing your IT infrastructure<sup>4</sup>

? Deploy globally in minutes: You can leverage the global infrastructure of AWS to deploy your applications and data in multiple regions and availability zones. This enables you to reach your customers faster, improve performance, and increase reliability<sup>5</sup>

**NEW QUESTION 3**

- (Topic 3)

Which AWS service provides this functionality?

- A. AWS IAM Identity Center (AWS Single Sign-On)
- B. AWS Systems Manager
- C. AWS Config
- D. AWS Control Tower

**Answer:** D

**Explanation:**

AWS Control Tower is a service that provides an easy way to set up and govern a secure, multi-account AWS environment. It automates the creation of accounts, organizational units, policies, and best practices based on the AWS Well-Architected Framework. AWS IAM Identity Center (AWS Single Sign-On) is a service that enables users to centrally manage access to multiple AWS accounts and business applications using a single sign-on experience. AWS Systems Manager is a service that provides operational management for AWS resources and applications. AWS Config is a service that enables users to assess, audit, and evaluate the configurations of AWS resources.

**NEW QUESTION 4**

- (Topic 3)

A company wants to make an upfront commitment for continued use of its production Amazon EC2 instances in exchange for a reduced overall cost.

Which pricing options meet these requirements with the LOWEST cost? (Select TWO.)

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

**Answer:** CD

**Explanation:**

Reserved Instances (RIs) are a pricing model that allows you to reserve EC2 instances for a specified period of time (one or three years) and receive a significant discount compared to On-Demand pricing. RIs are suitable for workloads that have predictable usage patterns and require a long-term commitment. You can choose between three payment options: All Upfront, Partial Upfront, or No Upfront. The more you pay upfront, the greater the discount<sup>1</sup>.

Savings Plans are a flexible pricing model that can help you reduce your EC2 costs by up to 72% compared to On-Demand pricing, in exchange for a commitment to a consistent amount of usage (measured in \$/hour) for a one or three year term. Savings Plans apply to usage across EC2, AWS Lambda, and AWS Fargate. You can choose between two types of Savings Plans: Compute Savings Plans and EC2 Instance Savings Plans. Compute Savings Plans offer the most flexibility and apply to any instance family, size, OS, tenancy, or region. EC2 Instance Savings Plans offer the highest discount and apply to a specific instance family within a region<sup>2</sup>.

Spot Instances are a pricing model that allows you to bid for unused EC2 capacity in the AWS cloud and are available at a discount of up to 90% compared to On-Demand pricing. Spot Instances are suitable for fault-tolerant or stateless workloads that can run on heterogeneous hardware and have flexible start and end times. However, Spot Instances are not guaranteed and can be interrupted by AWS at any time if the demand for capacity increases or your bid price is lower than the current Spot price<sup>3</sup>.

On-Demand Instances are a pricing model that allows you to pay for compute capacity by the hour or second with no long-term commitments. On-Demand Instances are suitable for short-term, spiky, or unpredictable workloads that cannot be interrupted, or for applications that are being developed or tested on EC2 for the first time. However, On-Demand Instances are the most expensive option among the four pricing models<sup>4</sup>.

Dedicated Hosts are physical EC2 servers fully dedicated for your use. Dedicated Hosts can help you reduce costs by allowing you to use your existing server-bound software licenses, such as Windows Server, SQL Server, and SUSE Linux Enterprise Server. Dedicated Hosts can be purchased On-Demand or as part of Savings Plans. Dedicated Hosts are suitable for workloads that need to run on dedicated physical servers or have strict licensing requirements. However, Dedicated Hosts are not the lowest cost option among the four pricing models.

**NEW QUESTION 5**

- (Topic 3)

Which option is a customer responsibility under the AWS shared responsibility model?

- A. Maintenance of underlying hardware of Amazon EC2 instances
- B. Application data security
- C. Physical security of data centers
- D. Maintenance of VPC components

**Answer:** B

**Explanation:**

The option that is a customer responsibility under the AWS shared responsibility model is B. Application data security.

According to the AWS shared responsibility model, AWS is responsible for the security of the cloud, while the customer is responsible for the security in the cloud. This means that AWS manages the security of the underlying infrastructure, such as the hardware, software, networking, and facilities that run the AWS services, while the customer manages the security of their applications, data, and resources that they use on top of AWS<sup>12</sup>. Application data security is one of the customer responsibilities under the AWS shared responsibility model. This means that the customer is responsible for protecting their application data from unauthorized access, modification, deletion, or leakage. The customer can use various AWS services and features to help with application data security, such as encryption, key management, access control, logging, and auditing<sup>12</sup>. Maintenance of underlying hardware of Amazon EC2 instances is not a customer responsibility under the AWS shared responsibility model. This is part of the AWS responsibility to secure the cloud. AWS manages the physical servers that host the Amazon EC2 instances and ensures that they are updated, patched, and replaced as needed<sup>13</sup>.

Physical security of data centers is not a customer responsibility under the AWS shared responsibility model. This is also part of the AWS responsibility to secure the cloud. AWS operates and controls the facilities where the AWS services are hosted and ensures that they are protected from unauthorized access, environmental hazards, fire, and theft<sup>14</sup>. Maintenance of VPC components is not a customer responsibility under the AWS shared responsibility model. This is a shared responsibility between AWS and the customer. AWS provides the VPC service and ensures that it is secure and reliable, while the customer configures and manages their own VPCs and related components, such as subnets, route tables, security groups, network ACLs, gateways, and endpoints<sup>15</sup>.

References:

1: Shared Responsibility Model - Amazon Web Services (AWS) 2: AWS Cloud Computing - W3Schools 3: [Amazon EC2 FAQs - Amazon Web Services] 4: [AWS Security - Amazon Web Services] 5: [Amazon Virtual Private Cloud (VPC) - Amazon Web Services]

**NEW QUESTION 6**

- (Topic 3)

A company is migrating its applications from on-premises to the AWS Cloud. The company wants to ensure that the applications are assigned only the minimum permissions that are needed to perform all operations.

Which AWS service will meet these requirements'?

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

**Answer:** A

**Explanation:**

AWS Identity and Access Management (IAM) is a service that helps you securely control access to AWS resources for your users. You use IAM to control who can use your AWS resources (authentication) and what resources they can use and in what ways (authorization). IAM also enables you to follow the principle of least privilege, which means granting only the permissions that are necessary to perform a task<sup>1</sup>. References: AWS Identity and Access Management (IAM) - AWS Documentation

**NEW QUESTION 7**

- (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 resources<sup>56</sup>. 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 8**

- (Topic 3)

A company uses AWS Organizations. The company wants to apply security best practices from the AWS Well-Architected Framework to all of its AWS accounts. Which AWS service will meet these requirements?

- A. Amazon Macie
- B. Amazon Detective
- C. AWS Control Tower
- D. AWS Secrets Manager

**Answer:** C

**Explanation:**

AWS Control Tower is the easiest way to set up and govern a secure, multi-account AWS environment based on best practices established through AWS's experience working with thousands of enterprises as they move to the cloud. With AWS Control Tower, builders can provision new AWS accounts in a few clicks, while you have peace of mind knowing your accounts conform to your organization's policies. AWS Control Tower automates the setup of a baseline environment, or landing zone, that is a secure, well-architected multi-account AWS environment<sup>1</sup>. AWS Control Tower helps you apply security best practices from the AWS Well-Architected Framework to all of your AWS accounts<sup>2</sup>.

**NEW QUESTION 9**

- (Topic 3)

A company that has multiple business units wants to centrally manage and govern its AWS Cloud environments. The company wants to automate the creation of AWS accounts, apply service control policies (SCPs), and simplify billing processes. Which AWS service or tool should the company use to meet these requirements?

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

**Answer:** A

**Explanation:**

AWS Organizations is an AWS service that enables you to centrally manage and govern your AWS Cloud environments across multiple business units. AWS Organizations allows you to create an organization that consists of AWS accounts that you create or invite to join. You can group your accounts into organizational units (OUs) and apply service control policies (SCPs) to them. SCPs are a type of policy that specify the maximum permissions for the accounts in your organization, and can help you enforce compliance and security requirements. AWS Organizations also simplifies billing processes by enabling you to consolidate and pay for all member accounts with a single payment method. You can also use AWS Organizations to automate the creation of AWS accounts by using APIs or AWS CloudFormation templates. References: What is AWS Organizations?, Policy-Based Management - AWS Organizations

**NEW QUESTION 10**

- (Topic 3)

A company needs to engage third-party consultants to help maintain and support its AWS environment and the company's business needs. Which AWS service or resource will meet these requirements?

- A. AWS Support
- B. AWS Organizations
- C. AWS Service Catalog
- D. AWS Partner Network (APN)

**Answer:** D

**Explanation:**

The AWS service or resource that will meet these requirements is D. AWS Partner Network (APN).

AWS Partner Network (APN) is a global community of consulting and technology partners that offer a wide range of services and solutions for AWS customers. APN partners can help customers design, architect, build, migrate, and manage their workloads and applications on AWS. APN partners have access to various resources, training, tools, and support to enhance their AWS expertise and deliver value to customers<sup>12</sup>.

AWS Support is a service that provides technical assistance and guidance for AWS customers. AWS Support offers different plans with varying levels of response time, access channels, and features. AWS Support does not directly engage third-party consultants, but rather connects customers with AWS experts and resources<sup>3</sup>.

AWS Organizations is a service that allows customers to manage multiple AWS accounts within a single organization. AWS Organizations enables customers to create groups of accounts, apply policies, automate account creation, and consolidate billing. AWS Organizations does not directly engage third-party consultants, but rather helps customers simplify and optimize their AWS account management<sup>4</sup>.

AWS Service Catalog is a service that allows customers to create and manage catalogs of IT services that are approved for use on AWS. AWS Service Catalog enables customers to control the configuration, deployment, and governance of their IT services. AWS Service Catalog does not directly engage third-party consultants, but rather helps customers standardize and streamline their IT service delivery.

References:

1: AWS Partner Network (APN) - Amazon Web Services (AWS) 2: Find an APN Partner - Amazon Web Services (AWS) 3: AWS Support – Amazon Web Services 4: AWS Organizations – Amazon Web Services 5: AWS Service Catalog – Amazon Web Services

**NEW QUESTION 10**

- (Topic 3)

A company wants to receive a notification when a specific AWS cost threshold is reached.

Which AWS services or tools can the company use to meet this requirement? (Select TWO.)

- A. Amazon Simple Queue Service (Amazon SQS)
- B. AWS Budgets
- C. Cost Explorer
- D. Amazon CloudWatch
- E. AWS Cost and Usage Report

**Answer:** BD

**Explanation:**

AWS Budgets and Amazon CloudWatch are two AWS services or tools that the company can use to receive a notification when a specific AWS cost threshold is reached. AWS Budgets allows users to set custom budgets to track their costs and usage, and respond quickly to alerts received from email or Amazon Simple Notification Service (Amazon SNS) notifications if they exceed their threshold. Users can create cost budgets with fixed or variable target amounts, and configure their notifications for actual or forecasted spend. Users can also set up custom actions to run automatically or through an approval process when a budget target is exceeded. For example, users could automatically apply a custom IAM policy that denies them the ability to provision additional resources within an account. Amazon CloudWatch is a service that monitors applications, responds to performance changes, optimizes resource use, and provides insights into operational health. Users can use CloudWatch to collect and track metrics, which are variables they can measure for their resources and applications. Users can create alarms that watch metrics and send notifications or automatically make changes to the resources they are monitoring when a threshold is breached. Users can use CloudWatch to monitor their AWS costs and usage by creating billing alarms that send notifications when their estimated charges exceed a specified threshold amount. Users can also use CloudWatch to monitor their Reserved Instance (RI) or Savings Plans utilization and coverage, and receive notifications when they fall below a certain level.

References: Cloud Cost And Usage Budgets - AWS Budgets, What is Amazon CloudWatch?, Creating a billing alarm - Amazon CloudWatch

**NEW QUESTION 15**

- (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 18**

- (Topic 3)

Which benefits can customers gain by using AWS Marketplace? (Select TWO.)

- A. Speed of business
- B. Fewer legal objections
- C. Ability to pay with credit cards
- D. No requirement for product licenses for any products
- E. Free use of all services for the first hour

**Answer:** AB

**Explanation:**

AWS Marketplace is a digital catalog that offers thousands of software products and solutions from independent software vendors (ISVs) and AWS partners. Customers can use AWS Marketplace to find, buy, and deploy software on AWS. Some of the benefits of using AWS Marketplace are:

? Speed of business: You can quickly and easily discover and deploy software that meets your business needs, without having to go through lengthy procurement processes. You can also use AWS Marketplace to test and compare different solutions before making a purchase decision.

? Fewer legal objections: You can benefit from standardized contract terms and conditions that are pre-negotiated between AWS and the ISVs. This reduces the time and effort required to review and approve legal agreements.

**NEW QUESTION 21**

- (Topic 3)

In the AWS shared responsibility model, which tasks are the responsibility of AWS? (Select TWO.)

- A. Patch an Amazon EC2 instance operating system.
- B. Configure a security group.
- C. Monitor the health of an Availability Zone.
- D. Protect the infrastructure that runs Amazon EC2 instances.

E. Manage access to the data in an Amazon S3 bucket

**Answer:** CD

**Explanation:**

According to the AWS shared responsibility model, AWS is responsible for the security of the cloud, which includes the tasks of monitoring the health of an Availability Zone and protecting the infrastructure that runs Amazon EC2 instances. An Availability Zone is a physically isolated location within an AWS Region that has its own power, cooling, and network connectivity. AWS monitors the health and performance of each Availability Zone and notifies customers of any issues or disruptions. AWS also protects the infrastructure that runs AWS services, such as Amazon EC2, by implementing physical, environmental, and operational security measures. AWS is not responsible for patching an Amazon EC2 instance operating system, configuring a security group, or managing access to the data in an Amazon S3 bucket. These are the customer's responsibilities for security in the cloud. The customer must ensure that the operating system and applications on their EC2 instances are up to date and secure. The customer must also configure the security group rules that control the inbound and outbound traffic for their EC2 instances. The customer must also manage the access permissions and encryption settings for their S3 buckets and objects<sup>2</sup>

**NEW QUESTION 25**

- (Topic 3)

A company has deployed an application in the AWS Cloud. The company wants to ensure that the application is highly resilient. Which component of AWS infrastructure can the company use to meet this requirement?

- A. Content delivery network (CDN)
- B. Edge locations
- C. Wavelength Zones
- D. Availability Zones

**Answer:** D

**Explanation:**

Availability Zones are components of AWS infrastructure that can help the company ensure that the application is highly resilient. Availability Zones are multiple, isolated locations within each AWS Region. Each Availability Zone has independent power, cooling, and physical security, and is connected to the other Availability Zones in the same Region via low-latency, high-throughput, and highly redundant networking. Availability Zones allow you to operate production applications and databases that are more highly available, fault tolerant, and scalable than would be possible from a single data center.

**NEW QUESTION 29**

- (Topic 3)

A company needs to deploy applications in the AWS Cloud as quickly as possible. The company also needs to minimize the complexity that is related to the management of AWS resources.

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

- A. AWS config
- B. AWS Elastic Beanstalk
- C. Amazon EC2
- D. Amazon Personalize

**Answer:** B

**Explanation:**

AWS Elastic Beanstalk is the AWS service that allows customers to deploy applications in the AWS Cloud as quickly as possible. AWS Elastic Beanstalk automatically handles the deployment, from capacity provisioning, load balancing, and auto-scaling to application health monitoring. Customers can upload their code and Elastic Beanstalk will take care of the rest<sup>1</sup>. AWS Elastic Beanstalk also minimizes the complexity that is related to the management of AWS resources. Customers can retain full control of the underlying AWS resources powering their applications and adjust the settings to suit their needs<sup>1</sup>. Customers can also use the AWS Management Console, the AWS Command Line Interface (AWS CLI), or APIs to manage their applications<sup>1</sup>.

AWS Config is the AWS service that enables customers to assess, audit, and evaluate the configurations of their AWS resources. AWS Config continuously monitors and records the configuration changes of the resources and evaluates them against desired configurations or best practices<sup>2</sup>. AWS Config does not help customers deploy applications in the AWS Cloud as quickly as possible or minimize the complexity that is related to the management of AWS resources.

Amazon EC2 is the AWS service that provides secure, resizable compute capacity in the cloud. Customers can launch virtual servers called instances and choose from various configurations of CPU, memory, storage, and networking resources<sup>3</sup>. Amazon EC2 does not automatically handle the deployment or management of AWS resources for customers. Customers have to manually provision, configure, monitor, and scale their instances and other related resources.

Amazon Personalize is the AWS service that enables customers to create personalized recommendations for their users based on their behavior and preferences. Amazon Personalize uses machine learning to analyze data and deliver real-time recommendations<sup>4</sup>. Amazon Personalize does not help customers deploy applications in the AWS Cloud as quickly as possible or minimize the complexity that is related to the management of AWS resources.

**NEW QUESTION 33**

- (Topic 3)

A software engineer wants to launch a virtual machine (VM) and MySQL database on AWS.

Which AWS service will meet these requirements with the LEAST operational effort?

- A. Amazon Elastic Container Service (Amazon ECS)
- B. AWS Elastic Beanstalk
- C. Amazon Lightsail
- D. Amazon EC2

**Answer:** B

**Explanation:**

AWS Elastic Beanstalk is a service that enables you to quickly deploy and manage applications in the AWS Cloud without worrying about the infrastructure that runs those applications. You simply upload your application, and Elastic Beanstalk automatically handles the details of capacity provisioning, load balancing, scaling, and application health monitoring. Elastic Beanstalk supports several platform configurations for Java, .NET, PHP, Node.js, Python, Ruby, Go, and Docker web applications that can run on familiar servers such as Apache, Nginx, Passenger, and IIS. You can also use Elastic Beanstalk to launch a virtual machine (VM) and MySQL database on AWS with the least operational effort. Amazon Elastic Container Service (Amazon ECS) is a fully managed container orchestration service that enables you to easily run, scale, and secure Docker containerized applications on AWS. However, it requires more operational effort than Elastic

Beanstalk, as you need to define your application architecture and the specifications of the containers that run it. Amazon Lightsail is an easy-to-use cloud platform that offers everything you need to build an application or website, plus a cost-effective, monthly plan. It is designed for developers who have little or no prior cloud experience and want to launch and manage applications on AWS with minimal complexity. However, it does not support MySQL databases, and it requires more operational effort than Elastic Beanstalk, as you need to configure your VM and database settings. Amazon EC2 is a web service that provides secure, resizable compute capacity in the cloud. It allows you to launch a virtual machine (VM) and MySQL database on AWS, but it requires the most operational effort, as you need to provision, monitor, and manage your EC2 instances and database.

**NEW QUESTION 37**

- (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 39**

- (Topic 3)

A team of researchers is going to collect data at remote locations around the world. Many locations do not have internet connectivity. The team needs to capture the data in the field, and transfer it to the AWS Cloud later. Which AWS service will support these requirements?

- A. AWS Outposts
- B. AWS Transfer Family
- C. AWS Snow Family
- D. AWS Migration Hub

**Answer: C**

**Explanation:**

AWS Snow Family is a group of devices that transport data in and out of AWS. AWS Snow Family devices are physical devices that can transfer up to exabytes of data. One exabyte is 1 000 000 000 000 megabytes. AWS Snow Family devices are designed for use in remote locations where internet connectivity is limited or unavailable. You can use these devices to collect and process data at the edge, and then ship them back to AWS for data upload. AWS Snow Family consists of three types of devices: AWS Snowcone, AWS Snowball, and AWS Snowmobile. References: 1: Edge Computing Devices, Secure Data Transfer - AWS Snow Family - AWS, 2: AWS Snow Family Documentation, 3: AWS Snow Family - W3Schools, 4: AWS Snow Family: Data Storage, Migration, and Computation

**NEW QUESTION 42**

- (Topic 3)

A company wants to integrate natural language processing (NLP) into business intelligence (BI) dashboards. The company wants to ask questions and receive answers with relevant visualizations. Which AWS service or tool will meet these requirements?

- A. Amazon Macie
- B. Amazon Rekognition
- C. Amazon QuickSight Q
- D. Amazon Lex

**Answer: C**

**Explanation:**

Amazon QuickSight Q is a natural language query feature that allows users to ask questions about their data and receive answers in the form of relevant visualizations. Amazon Macie is a data security and data privacy service that uses machine learning and pattern matching to discover and protect sensitive data in AWS. Amazon Rekognition is a computer vision service that can analyze images and videos for faces, objects, scenes, text, and more. Amazon Lex is a service for building conversational interfaces using voice and text.

**NEW QUESTION 45**

- (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 files<sup>1</sup>. 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 48**

- (Topic 3)

Which AWS service can identify when an Amazon EC2 instance was terminated?

- A. AWS Identity and Access Management (IAM)
- B. AWS CloudTrail
- C. AWS Compute Optimizer
- D. Amazon EventBridge

**Answer:** B

**Explanation:**

AWS CloudTrail is the AWS service that can identify when an Amazon EC2 instance was terminated. AWS CloudTrail is a service that records API calls and events for AWS accounts and resources. AWS CloudTrail can capture the `TerminateInstances` event, which is triggered when an EC2 instance is terminated by a user or an AWS service. The event contains information such as the instance ID, the user identity, the source IP address, the time, and the reason for the termination<sup>2</sup>. Customers can use the CloudTrail console, the AWS CLI, or the AWS SDKs to view and search for the `TerminateInstances` events in their event history or in their S3 buckets where they store their CloudTrail logs<sup>3</sup>.

**NEW QUESTION 51**

- (Topic 3)

A company wants to migrate its high-performance computing (HPC) application to Amazon EC2 instances. The application has multiple components. The application must have fault tolerance and must have the ability to fail over automatically.

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

- A. Multiple AWS Regions
- B. Multiple edge locations
- C. Multiple Availability Zones
- D. Regional edge caches

**Answer:** C

**Explanation:**

Using EC2 instances in multiple Availability Zones is an AWS infrastructure solution that meets the requirements of migrating a high performance computing (HPC) application to AWS with fault tolerance and failover capabilities, and with the least latency between components. An Availability Zone is a physically isolated location within an AWS Region that has its own power, cooling, and network connectivity. EC2 instances within the same Region can communicate with each other using low-latency private IP addresses. By using EC2 instances in multiple Availability Zones, the company can achieve fault tolerance and failover for their HPC application, because they can distribute the workload and data across different locations that are independent of each other. If one Availability Zone becomes unavailable or impaired, the company can redirect the traffic and data to another Availability Zone without affecting the performance and availability of the application<sup>5</sup>.

**NEW QUESTION 56**

- (Topic 3)

Which AWS service or feature enables users to encrypt data at rest in Amazon S3?

- A. IAM policies
- B. Server-side encryption
- C. Amazon GuardDuty
- D. Client-side encryption

**Answer:** B

**Explanation:**

Server-side encryption is an encryption option that Amazon S3 provides to encrypt data at rest in Amazon S3. With server-side encryption, Amazon S3 encrypts an object before saving it to disk in its data centers and decrypts it when you download the objects. You have three server-side encryption options to choose from: SSE-S3, SSE-C, and SSE-KMS. SSE-S3 uses keys that are managed by Amazon S3. SSE-C allows you to manage your own encryption keys. SSE-KMS uses keys that are managed by AWS Key Management Service (AWS KMS)<sup>5</sup>.

**NEW QUESTION 58**

- (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 62**

- (Topic 3)

A development team wants to deploy multiple test environments for an application in a fast repeatable manner. Which AWS service should the team use?

- A. Amazon EC2
- B. AWS CloudFormation
- C. Amazon QuickSight
- D. Amazon Elastic Container Service (Amazon ECS)

**Answer:** B

**Explanation:**

AWS CloudFormation is a service that allows you to model and provision your AWS resources using templates. You can define your infrastructure as code and automate the creation and update of your resources. AWS CloudFormation also supports nested stacks, change sets, and rollback features to help you manage complex and dynamic environments<sup>34</sup>. References:

? AWS CloudFormation

? AWS Certified Cloud Practitioner Exam Guide

**NEW QUESTION 66**

- (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<sup>1</sup>. 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<sup>23</sup>. References: <sup>2</sup>: AWS Lambda - Amazon Web Services (AWS), <sup>3</sup>: AWS Lambda Documentation, <sup>1</sup>: Amazon CLF-C02: What is customer responsibility under AWS ... - PUPUWEB

**NEW QUESTION 71**

- (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<sup>4567</sup>. References: <sup>4</sup>: Compute Savings Plans – Amazon Web Services, <sup>5</sup>: What are Savings Plans? - Savings Plans, <sup>6</sup>: How To Cut Your AWS Bill With Savings Plans (and avoid some common ...), <sup>7</sup>: AWS Savings Plans vs Reserved Instances

- GorillaStack

**NEW QUESTION 72**

- (Topic 3)

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

- A. Configure the AWS provided security group firewall.
- B. Classify company assets in the AWS Cloud.
- C. Determine which Availability Zones to use for Amazon S3 buckets.

- D. Patch or upgrade Amazon DynamoDB.
- E. Select Amazon EC2 instances to run AWS Lambda on.
- F. AWS Config

**Answer:** AB

**Explanation:**

According to the AWS shared responsibility model, the customer is responsible for security in the cloud, which includes the tasks of configuring the AWS provided security group firewall and classifying company assets in the AWS Cloud. A security group is a virtual firewall that controls the inbound and outbound traffic for one or more EC2 instances. The customer must configure the security group rules to allow or deny traffic based on protocol, port, or source and destination IP address<sup>2</sup> Classifying company assets in the AWS Cloud means identifying the types, categories, and sensitivity levels of the data and resources that the customer stores and processes on AWS. The customer must also determine the applicable compliance requirements and regulations that apply to their assets, and implement the appropriate security controls and measures to protect them

**NEW QUESTION 76**

- (Topic 3)

A company needs an automated vulnerability management service that continually scans AWS workloads for software vulnerabilities. Which AWS service will meet these requirements?

- A. Amazon GuardDuty
- B. Amazon Inspector
- C. AWS Security Hub
- D. AWS Shield

**Answer:** B

**Explanation:**

The correct answer is B. Amazon Inspector.

Amazon Inspector is an automated vulnerability management service that continually scans AWS workloads for software vulnerabilities and unintended network exposure. Amazon Inspector automatically discovers workloads, such as Amazon EC2 instances, containers, and Lambda functions, and scans them for software vulnerabilities and unintended network exposure<sup>12</sup>.

Amazon GuardDuty is a threat detection service that monitors your AWS accounts and workloads for malicious or unauthorized activity. Amazon GuardDuty does not scan for software vulnerabilities, but rather analyzes AWS CloudTrail, Amazon VPC Flow Logs, and DNS logs to detect threats such as compromised credentials, backdoors, or crypto mining<sup>3</sup>.

AWS Security Hub is a security and compliance service that aggregates and prioritizes security findings from multiple AWS services and partner solutions. AWS Security Hub does not scan for software vulnerabilities, but rather provides a comprehensive view of your security posture across your AWS accounts<sup>4</sup>.

AWS Shield is a managed service that protects your web applications and network resources from distributed denial-of-service (DDoS) attacks. AWS Shield does not scan for software vulnerabilities, but rather provides detection and mitigation of DDoS attacks at the network and application layers<sup>5</sup>.

References:

1: Automated Software Vulnerability Management - Amazon Inspector - AWS 3: [Amazon GuardDuty – Intelligent Threat Detection Made Easy] 2: AWS Re-Launches Amazon Inspector with New Architecture and Features - InfoQ 4: [AWS Security Hub – Unified Security and Compliance Center] 5: [AWS Shield – Managed DDoS Protection]

**NEW QUESTION 78**

- (Topic 3)

A company is moving to the AWS Cloud to reduce operational overhead for its application infrastructure. Which IT operation will the company still be responsible for after the migration to AWS?

- A. Security patching of AWS Elastic Beanstalk
- B. Backups of data that is stored in Amazon Aurora
- C. Termination of Amazon EC2 instances that are managed by AWS Auto Scaling
- D. Configuration of IAM access controls

**Answer:** D

**Explanation:**

AWS Elastic Beanstalk, Amazon Aurora, and AWS Auto Scaling are managed services that reduce the operational overhead for the customers. AWS is responsible for security patching, backups, and termination of these services. However, the customers are still responsible for configuring IAM access controls to manage the permissions and policies for their AWS resources. This is part of the AWS shared responsibility model, which defines the security and compliance responsibilities of AWS and the customers. You can learn more about the AWS shared responsibility model from this whitepaper or this digital course.

**NEW QUESTION 80**

- (Topic 3)

A company hosts a large amount of data in AWS. The company wants to identify if any of the data should be considered sensitive. Which AWS service will meet the requirement?

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

**Answer:** B

**Explanation:**

Amazon Macie is a fully managed service that uses machine learning and pattern matching to help you detect, classify, and better protect your sensitive data stored in the AWS Cloud<sup>1</sup>. Macie can automatically discover and scan your Amazon S3 buckets for sensitive data such as personally identifiable information (PII), financial information, healthcare information, intellectual property, and credentials<sup>1</sup>. Macie also provides you with a dashboard that shows the type, location, and volume of sensitive data in your AWS environment, as well as alerts and findings on potential security issues<sup>1</sup>.

The other options are not suitable for identifying sensitive data in AWS. Amazon Inspector is a service that helps you find security vulnerabilities and deviations from best practices in your Amazon EC2 instances<sup>2</sup>. AWS Identity and Access Management (IAM) is a service that helps you manage access to your AWS

resources by creating users, groups, roles, and policies<sup>3</sup>. Amazon CloudWatch is a service that helps you monitor and troubleshoot your AWS resources and applications by collecting metrics, logs, events, and alarms<sup>4</sup>. References:

? 1: What Is Amazon Macie? - Amazon Macie

? 2: What Is Amazon Inspector? - Amazon Inspector

? 3: What Is IAM? - AWS Identity and Access Management

? 4: What Is Amazon CloudWatch? - Amazon CloudWatch

**NEW QUESTION 84**

- (Topic 3)

Which AWS Cloud Adoption Framework (AWS CAF) capability belongs to the people perspective?

- A. Data architecture
- B. Event management
- C. Cloud fluency
- D. Strategic partnership

**Answer: C**

**Explanation:**

Cloud fluency is a capability that belongs to the people perspective of the AWS Cloud Adoption Framework (AWS CAF). Cloud fluency is the ability of the workforce to understand the benefits, challenges, and best practices of cloud computing, and to apply them to their roles and responsibilities. Cloud fluency helps the organization to adopt a cloud mindset, culture, and skills, and to leverage the full potential of the cloud. Cloud fluency can be achieved through various methods, such as training, certification, mentoring, coaching, and hands-on experience. Cloud fluency is one of the four capabilities of the people perspective, along with culture, organizational structure, and leadership. The other three capabilities belong to different perspectives of the AWS CAF. Data architecture is a capability of the platform perspective, which helps you design and implement data solutions that meet your business and technical requirements. Event management is a capability of the operations perspective, which helps you monitor and respond to events that affect the availability, performance, and security of your cloud resources. Strategic partnership is a capability of the business perspective, which helps you establish and maintain relationships with external stakeholders, such as customers, partners, suppliers, and regulators, to create value and achieve your business goals. References: AWS Cloud Adoption Framework: People Perspective, AWS CAF - Cloud Adoption Framework - W3Schools

**NEW QUESTION 89**

- (Topic 3)

A company is running an Amazon EC2 instance in a VPC.

An ecommerce company is using Amazon EC2 Auto Scaling groups to manage a fleet of web servers running on Amazon EC2.

This architecture follows which AWS Well-Architected Framework best practice?

- A. Secure the workload
- B. Decouple infrastructure components
- C. Design for failure
- D. Think parallel

**Answer: C**

**Explanation:**

Design for failure is one of the best practices of the AWS Well-Architected Framework. It means that the architecture should be resilient and fault-tolerant, and able to handle failures without impacting the availability and performance of the applications. By using Amazon EC2 Auto Scaling groups, the ecommerce company can design for failure by automatically scaling the number of EC2 instances up or down based on demand or health status. Amazon EC2 Auto Scaling groups can also distribute the EC2 instances across multiple Availability Zones, which are isolated locations within an AWS Region that have independent power, cooling, and network connectivity. This way, the company can ensure that their web servers can handle traffic spikes, recover from failures, and provide a consistent user experience

**NEW QUESTION 91**

- (Topic 3)

A company is hosting an application in the AWS Cloud. The company wants to verify that underlying AWS services and general AWS infrastructure are operating normally.

Which combination of AWS services can the company use to gather the required information? (Select TWO.)

- A. AWS Personal Health Dashboard
- B. AWS Systems Manager
- C. AWS Trusted Advisor
- D. AWS Service Health Dashboard
- E. AWS Service Catalog

**Answer: AD**

**Explanation:**

AWS Personal Health Dashboard and AWS Service Health Dashboard are two AWS services that can help the company to verify that underlying AWS services and general AWS infrastructure are operating normally. AWS Personal Health Dashboard provides a personalized view into the performance and availability of the AWS services you are using, as well as alerts that are automatically triggered by changes in the health of those services. In addition to event-based alerts, Personal Health Dashboard provides proactive notifications of scheduled activities, such as any changes to the infrastructure powering your resources, enabling you to better plan for events that may affect you. These notifications can be delivered to you via email or mobile for quick visibility, and can always be viewed from within the AWS Management Console. When you get an alert, it includes detailed information and guidance, enabling you to take immediate action to address AWS events impacting your resources<sup>3</sup>. AWS Service Health Dashboard provides a general status of AWS services, and the Service health view displays the current and historical status of all AWS services. This page shows reported service events for services across AWS Regions. You don't need to sign in or have an AWS account to access the AWS Service Health Dashboard – Service health page. You can also subscribe to RSS feeds for specific services or regions to receive notifications about service events<sup>4</sup>. References: Getting started with your AWS Health Dashboard – Your account health, Introducing AWS Personal Health Dashboard

**NEW QUESTION 92**

- (Topic 3)

A company wants to allow users to authenticate and authorize multiple AWS accounts by using a single set of credentials. Which AWS service or resource will meet this requirement?

- A. AWS Organizations
- B. IAM user
- C. AWS IAM Identity Center (AWS Single Sign-On)
- D. AWS Control Tower

**Answer: C**

**Explanation:**

AWS IAM Identity Center (AWS Single Sign-On) is a cloud-based service that makes it easy to centrally manage single sign-on (SSO) access to multiple AWS accounts and business applications. You can use AWS SSO to enable your users to sign in to the AWS Management Console or the AWS Command Line Interface (AWS CLI) with their existing corporate credentials<sup>2</sup>. You can also manage SSO access and user permissions across all your AWS accounts in AWS Organizations<sup>3</sup>. References: AWS Single Sign-On - AWS Documentation, AWS Organizations - AWS Documentation

**NEW QUESTION 97**

- (Topic 3)

Which pillar of the AWS Well-Architected Framework includes the AWS shared responsibility model?

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

**Answer: D**

**Explanation:**

The AWS Well-Architected Framework is a set of best practices and guidelines for designing and operating reliable, secure, efficient, and cost-effective systems in the cloud. The framework consists of five pillars: operational excellence, performance efficiency, reliability, security, and cost optimization. The security pillar covers the AWS shared responsibility model, which defines the security and compliance responsibilities of AWS and the customers. You can learn more about the AWS Well-Architected Framework from [this whitepaper] or [this digital course].

**NEW QUESTION 101**

- (Topic 3)

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

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

**Answer: A**

**Explanation:**

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

**NEW QUESTION 106**

- (Topic 3)

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

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

**Answer: A**

**Explanation:**

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

A multi-region deployment model is a cloud computing architecture that distributes an application and its data across multiple geographic regions. A multi-region deployment model enables a company to achieve global reach, high availability, disaster recovery, and performance optimization. By deploying an application in multiple regions, a company can serve customers from the nearest region, reduce latency, increase redundancy, and comply with data sovereignty regulations<sup>12</sup>. A single-region deployment model is a cloud computing architecture that runs an application and its data within a single geographic region. A single-region deployment model is simpler and cheaper than a multi-region deployment model, but it has limited scalability, availability, and performance. A single-region deployment model may not be suitable for a company that wants to deploy an application globally, as it may face challenges such as network latency, regional outages, or regulatory compliance<sup>12</sup>.

A multi-AZ (Availability Zone) deployment model is a cloud computing architecture that distributes an application and its data across multiple isolated locations within a single region. An Availability Zone is a physically separate location within an AWS Region that has independent power, cooling, and networking. A multi-AZ deployment model enhances the availability and durability of an application by providing redundancy and fault tolerance within a region<sup>34</sup>.

A single-AZ deployment model is a cloud computing architecture that runs an application and its data within a single Availability Zone. A single-AZ deployment

model is the simplest and most cost-effective option, but it has no redundancy or fault tolerance. A single-AZ deployment model may not be suitable for a company that wants to deploy an application globally, as it may face challenges such as network latency, regional outages, or regulatory compliance<sup>34</sup>.

References:

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

**NEW QUESTION 107**

- (Topic 3)

A company wants to migrate to AWS and use the same security software it uses on premises. The security software vendor offers its security software as a service on AWS.

Where can the company purchase the security solution?

- A. AWS Partner Solutions Finder
- B. AWS Support Center
- C. AWS Management Console
- D. AWS Marketplace

**Answer:** D

**Explanation:**

AWS Marketplace is an online store that helps customers find, buy, and immediately start using the software and services that run on AWS. Customers can choose from a wide range of software products in popular categories such as security, networking, storage, machine learning, business intelligence, database, and DevOps. Customers can also use AWS Marketplace to purchase software as a service (SaaS) solutions that are integrated with AWS. Customers can benefit from simplified procurement, billing, and deployment processes, as well as flexible pricing options and free trials. Customers can also leverage AWS Marketplace to discover and subscribe to solutions offered by AWS Partners, such as the security software vendor mentioned in the question. References: AWS Marketplace, [AWS Marketplace: Software as a Service (SaaS)], [AWS Cloud Practitioner Essentials: Module 6 - AWS Pricing, Billing, and Support]

**NEW QUESTION 108**

- (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 model<sup>2</sup>. 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 111**

- (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 MySQL<sup>4567</sup>. 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 112**

- (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 outcomes<sup>1</sup>.

? 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 activities<sup>1</sup>.

? 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 production<sup>1</sup>.

? 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 sustained<sup>1</sup>.

**NEW QUESTION 113**

- (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 behalf<sup>3</sup>

? 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 AWS<sup>4</sup>

**NEW QUESTION 115**

- (Topic 3)

A company wants to migrate its on-premises relational databases to the AWS Cloud. The company wants to use infrastructure as close to its current geographical location as possible.

Which AWS service or resource should the company use to select its Amazon RDS deployment area?

- A. Amazon Connect
- B. AWS Wavelength
- C. AWS Regions
- D. AWS Direct Connect

**Answer:** C

**Explanation:**

AWS Regions are the AWS service or resource that the company should use to select its Amazon RDS deployment area. AWS Regions are separate geographic areas where AWS clusters its data centers. Each AWS Region consists of multiple, isolated, and physically separate Availability Zones within a geographic area. Each AWS Region is designed to be isolated from the other AWS Regions to achieve the highest possible fault tolerance and stability. AWS provides a more extensive global footprint than any other cloud provider, and to support its global footprint and ensure customers are served across the world, AWS opens new Regions rapidly. AWS maintains multiple geographic Regions, including Regions in North America, South America, Europe, China, Asia Pacific, South Africa, and the Middle East. Amazon RDS is available in several AWS Regions worldwide. To create or work with an Amazon RDS DB instance in a specific AWS Region, you must use the corresponding regional service endpoint. You can choose the AWS Region that meets your latency or legal requirements. You can also use multiple AWS Regions to design a disaster recovery solution or to distribute your read workload. References: Global Infrastructure Regions & AZs - [aws.amazon.com](https://aws.amazon.com), Regions, Availability Zones, and Local Zones - Amazon Relational Database Service

**NEW QUESTION 119**

- (Topic 3)

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

Which solution achieves this goal?

- A. Use EC2 instances in a single Availability Zone.
- B. Use EC2 instances in multiple AWS Regions.
- C. Use EC2 instances in multiple Availability Zones in the same AWS Region.
- D. Use EC2 instances in the same edge location and the same AWS Region.

**Answer:** C

**Explanation:**

The solution that achieves the goal of having Amazon EC2 instances share the same geographic area but use multiple independent underlying power sources is to use EC2 instances in multiple Availability Zones in the same AWS Region. An Availability Zone is a physically isolated location within an AWS Region that has its own power, cooling, and network connectivity. An AWS Region is a geographical area that consists of two or more Availability Zones. By using multiple Availability Zones, users can increase the fault tolerance and resilience of their applications, as well as reduce latency for end users<sup>3</sup>.

Using EC2 instances in a single Availability Zone, multiple AWS Regions, or the same edge location and the same AWS Region would not meet the requirement of having multiple independent power sources.

**NEW QUESTION 123**

- (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 protocols<sup>1</sup>. AWS Security Hub is a service that provides a comprehensive view of the security posture of AWS accounts and resources<sup>2</sup>. Security groups are features that act as firewalls for controlling traffic at the instance level<sup>3</sup>. AWS WAF is a web application firewall that helps protect web applications from common web exploits<sup>4</sup>.

**NEW QUESTION 128**

- (Topic 3)

A company wants to use the latest technologies and wants to minimize its capital investment. Instead of upgrading on-premises infrastructure, the company wants to move to the AWS Cloud.

Which AWS Cloud benefit does this scenario describe?

- A. Increased speed to market
- B. The trade of infrastructure expenses for operating expenses
- C. Massive economies of scale
- D. The ability to go global in minutes

**Answer: B**

**Explanation:**

The trade of infrastructure expenses for operating expenses is one of the benefits of the AWS Cloud. By moving to the AWS Cloud, the company can avoid the upfront costs of purchasing and maintaining on-premises infrastructure, such as servers, storage, network, and software. Instead, the company can pay only for the AWS resources and services that they use, as they use them. This reduces the risk and complexity of planning and managing IT infrastructure, and allows the company to focus on innovation and growth. Increased speed to market, massive economies of scale, and the ability to go global in minutes are also benefits of the AWS Cloud, but they are not the best ones to describe this scenario. Increased speed to market means that the company can launch new products and services faster by using AWS services and tools. Massive economies of scale means that the company can benefit from the lower costs and higher performance that AWS achieves by operating at a large scale. The ability to go global in minutes means that the company can deploy their applications and data in multiple regions and availability zones around the world to reach their customers faster and improve performance and reliability<sup>5</sup>.

**NEW QUESTION 132**

- (Topic 3)

A company wants to define a central data protection policy that works across AWS services for compute, storage, and database resources.

Which AWS service will meet this requirement?

- A. AWS Batch
- B. AWS Elastic Disaster Recovery
- C. AWS Backup
- D. Amazon FSx

**Answer: C**

**Explanation:**

The AWS service that will meet this requirement is C. AWS Backup.

AWS Backup is a service that allows you to define a central data protection policy that works across AWS services for compute, storage, and database resources. You can use AWS Backup to create backup plans that specify the frequency, retention, and lifecycle of your backups, and apply them to your AWS resources using tags or resource IDs. AWS Backup supports various AWS services, such as Amazon EC2, Amazon EBS, Amazon RDS, Amazon DynamoDB, Amazon EFS, Amazon FSx, and AWS Storage Gateway<sup>12</sup>. AWS Batch is a service that allows you to run batch computing workloads on AWS. AWS Batch does not provide a central data protection policy, but rather enables you to optimize the allocation and utilization of your compute resources<sup>3</sup>.

AWS Elastic Disaster Recovery is a service that allows you to prepare for and recover from disasters using AWS. AWS Elastic Disaster Recovery does not provide a central data protection policy, but rather helps you minimize downtime and data loss by replicating your applications and data to AWS<sup>4</sup>.

Amazon FSx is a service that provides fully managed file storage for Windows and Linux applications. Amazon FSx does not provide a central data protection policy, but rather offers features such as encryption, snapshots, backups, and replication to protect your file systems<sup>5</sup>.

References:

1: AWS Backup – Centralized backup across AWS services 3: AWS Batch – Run Batch Computing Jobs on AWS 2: Data Protection Reference Architectures with AWS Backup 4: AWS Elastic Disaster Recovery – Prepare for and recover from disasters using AWS 5: Amazon FSx – Fully managed file storage for Windows and Linux applications

**NEW QUESTION 134**

- (Topic 3)

A company wants to migrate its PostgreSQL database to AWS. The company does not use the database frequently.

Which AWS service or resource will meet these requirements with the LEAST management overhead?

- A. PostgreSQL on Amazon EC2
- B. Amazon RDS for PostgreSQL
- C. Amazon Aurora PostgreSQL-Compatible Edition
- D. Amazon Aurora Serverless

**Answer: D**

**Explanation:**

Amazon Aurora Serverless is an on-demand, auto-scaling configuration for Amazon Aurora PostgreSQL-Compatible Edition. It is a fully managed service that automatically scales up and down based on the application's actual needs. Amazon Aurora Serverless is suitable for applications that have infrequent, intermittent, or unpredictable database workloads, and that do not require the full power and range of options provided by provisioned Aurora clusters. Amazon Aurora Serverless eliminates the need to provision and manage database instances, and reduces the management overhead associated with database administration tasks such as scaling, patching, backup, and recovery. References: Amazon Aurora Serverless, Choosing between Aurora Serverless and provisioned Aurora DB clusters, [AWS Cloud Practitioner Essentials: Module 4 - Databases in the Cloud]

**NEW QUESTION 138**

- (Topic 3)

A company wants an automated process to continuously scan its Amazon EC2 instances for software vulnerabilities. Which AWS service will meet these requirements?

- A. Amazon GuardDuty
- B. Amazon Inspector
- C. Amazon Detective
- D. Amazon Cognito

**Answer:** B

**Explanation:**

Amazon Inspector is the AWS service that can be used to perform vulnerability scans on AWS EC2 instances for software vulnerabilities automatically in a periodic fashion. Amazon Inspector automatically discovers EC2 instances and scans them for software vulnerabilities and unintended network exposure. Amazon Inspector uses AWS Systems Manager (SSM) and the SSM Agent to collect information about the software application inventory of the EC2 instances. This data is then scanned by Amazon Inspector for software vulnerabilities<sup>12</sup>. Amazon Inspector also integrates with other AWS services, such as Amazon EventBridge and AWS Security Hub, to automate discovery, expedite vulnerability routing, and shorten mean time to remediate (MTTR) vulnerabilities<sup>2</sup>.

**NEW QUESTION 142**

- (Topic 3)

Which capabilities are in the platform perspective of the AWS Cloud Adoption Framework (AWS CAF)? (Select TWO.)

- A. Performance and capacity management
- B. Data engineering
- C. Continuous integration and continuous delivery (CI/CD)
- D. Infrastructure protection
- E. Change and release management

**Answer:** BC

**Explanation:**

The platform perspective of the AWS Cloud Adoption Framework (AWS CAF) helps you build an enterprise-grade, scalable, hybrid cloud platform, modernize existing workloads, and implement new cloud-native solutions<sup>1</sup>. It comprises seven capabilities, two of which are data engineering and CI/CD<sup>1</sup>.

? Data engineering: This capability helps you design and evolve a fit-for-purpose data and analytics architecture that can reduce complexity, cost, and technical debt while enabling you to gain actionable insights from exponentially growing data volumes<sup>1</sup>. It involves selecting key technologies for each of your architectural layers, such as ingestion, storage, catalog, processing, and consumption. It also involves supporting real-time data processing and adopting a Lake House architecture to facilitate data movements between data lakes and purpose-built data stores<sup>1</sup>.

? CI/CD: This capability helps you automate the delivery of your cloud solutions using a set of practices and tools that enable faster and more reliable deployments<sup>1</sup>. It involves establishing a pipeline that can build, test, and deploy your code across multiple environments. It also involves adopting a DevOps culture that fosters collaboration, feedback, and continuous improvement among your development and operations teams<sup>1</sup>.

References:

? 1: Platform perspective: infrastructure and applications - An Overview of the AWS Cloud Adoption Framework

**NEW QUESTION 147**

- (Topic 3)

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

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

**Answer:** A

**Explanation:**

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

**NEW QUESTION 148**

- (Topic 3)

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

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

**Answer:** C

**Explanation:**

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

**NEW QUESTION 153**

- (Topic 3)

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

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

**Answer:** C

**Explanation:**

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

**NEW QUESTION 154**

- (Topic 3)

Which perspective in the AWS Cloud Adoption Framework (AWS CAF) includes a capability for well-designed data and analytics architecture?

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

**Answer:** D

**Explanation:**

The correct answer is D. Platform.

The Platform perspective in the AWS Cloud Adoption Framework (AWS CAF) includes a capability for well-designed data and analytics architecture. This capability helps you design, implement, and optimize your data and analytics solutions on AWS, using services such as Amazon S3, Amazon Redshift, Amazon EMR, Amazon Kinesis, Amazon Athena, and Amazon QuickSight. A well-designed data and analytics architecture enables you to collect, store, process, analyze, and visualize data from various sources, and derive insights that can drive your business decisions<sup>12</sup>.

The Security perspective does not include a capability for data and analytics architecture, but it does include a capability for data protection, which helps you secure your data at rest and in transit using encryption, key management, access control, and auditing<sup>13</sup>.

The Governance perspective does not include a capability for data and analytics architecture, but it does include a capability for data governance, which helps you manage the quality, availability, usability, integrity, and security of your data assets<sup>14</sup>.

The Operations perspective does not include a capability for data and analytics architecture, but it does include a capability for data operations, which helps you monitor, troubleshoot, and optimize the performance and availability of your data pipelines and workloads<sup>1</sup>.

References:

1: Foundational capabilities - An Overview of the AWS Cloud Adoption Framework 2: [AWS Cloud Adoption Framework: Platform Perspective] 3: [AWS Cloud Adoption Framework: Security Perspective] 4: [AWS Cloud Adoption Framework: Governance Perspective] : [AWS Cloud Adoption Framework: Operations Perspective]

**NEW QUESTION 156**

- (Topic 3)

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

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

**Answer:** AB

**Explanation:**

These are two scenarios that represent the concept of elasticity on AWS. Elasticity means the ability to adjust the resources and capacity of the system in response to changes in demand or environment. Scaling the number of Amazon EC2 instances based on traffic means using services such as AWS Auto Scaling or Elastic Load Balancing to add or remove instances as the traffic increases or decreases. Resizing Amazon RDS instances as business needs change means using the Amazon RDS console or API to modify the instance type, storage type, or storage size of the database as the workload grows or shrinks. You can learn more about the concept of elasticity on AWS from [this webpage] or [this digital course].

**NEW QUESTION 161**

- (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](#)<sup>1</sup>

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 162**

- (Topic 3)

A company is using Amazon DynamoDB.

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

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

**Answer: C**

**Explanation:**

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

**NEW QUESTION 167**

- (Topic 3)

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

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

**Answer: C**

**Explanation:**

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

**NEW QUESTION 169**

- (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 173**

- (Topic 3)

A developer who has no AWS Cloud experience wants to use AWS technology to build a web application. Which AWS service should the developer use to start building the application?

- A. Amazon SageMaker
- B. AWS Lambda
- C. Amazon Lightsail
- D. Amazon Elastic Container Service (Amazon ECS)

**Answer: C**

**Explanation:**

Amazon Lightsail is an easy-to-use cloud platform that offers everything you need to build an application or website, plus a cost-effective, monthly plan<sup>1</sup>. It is designed for developers who have little or no prior cloud experience and want to launch and manage applications on AWS with minimal complexity<sup>2</sup>. Amazon SageMaker is a service for building, training, and deploying machine learning models<sup>3</sup>. AWS Lambda is a service that lets you run code without provisioning or managing servers<sup>4</sup>. Amazon Elastic Container Service (Amazon ECS) is a fully managed container orchestration service.

**NEW QUESTION 177**

- (Topic 3)

A company is planning to host its workloads on AWS.

Which AWS service requires the company to update and patch the guest operating system?

- A. Amazon DynamoDB
- B. Amazon S3
- C. Amazon EC2
- D. Amazon Aurora

**Answer: C**

**Explanation:**

Amazon EC2 is an AWS service that provides scalable, secure, and resizable compute capacity in the cloud. Amazon EC2 allows customers to launch and manage virtual servers, called instances, that run a variety of operating systems and applications. Customers have full control over the configuration and management of their instances, including the guest operating system. Therefore, customers are responsible for updating and patching the guest operating system on their EC2 instances, as well as any other software or utilities installed on the instances. AWS provides tools and services, such as AWS Systems Manager and AWS OpsWorks, to help customers automate and simplify the patching process. References: Shared Responsibility Model, Shared responsibility model, [Amazon EC2]

**NEW QUESTION 180**

- (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 181**

- (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 183**

- (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 185**

- (Topic 3)

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

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

**Answer: BD**

**Explanation:**

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

**NEW QUESTION 190**

- (Topic 3)

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

Which AWS service will meet this requirement?

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

**Answer: B**

**Explanation:**

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

**NEW QUESTION 194**

- (Topic 3)

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

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

**Answer: A**

**Explanation:**

VPC Flow Logs is a feature that enables you to capture information about the IP traffic going to and from network interfaces in your VPC. Flow log data can be published to Amazon CloudWatch Logs, Amazon S3, or Amazon Kinesis Data Firehose. You can use VPC Flow Logs to diagnose network issues, monitor traffic patterns, detect security anomalies, and comply with auditing requirements<sup>34</sup>. References: Logging IP traffic using VPC Flow Logs - Amazon Virtual Private Cloud, New – VPC Traffic Mirroring – Capture & Inspect Network Traffic | AWS News Blog

**NEW QUESTION 199**

- (Topic 3)

Which database engines does Amazon Aurora support? (Select TWO.)

- A. Oracle
- B. Microsoft SQL Server
- C. MySQL
- D. PostgreSQL
- E. MongoDB

**Answer:** CD

**Explanation:**

Amazon Aurora is a relational database service that is compatible with MySQL and PostgreSQL engines. It delivers up to five times the performance of MySQL and up to three times the performance of PostgreSQL. It also provides high availability, scalability, security, and durability<sup>1</sup>

**NEW QUESTION 203**

- (Topic 3)

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

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

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

**Answer:** D

**Explanation:**

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

**NEW QUESTION 207**

- (Topic 3)

A customer runs an On-Demand Amazon Linux EC2 instance for 3 hours, 5 minutes, and 6 seconds.

For how much time will the customer be billed?

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

**Answer:** C

**Explanation:**

Amazon EC2 usage is calculated by either the hour or the second based on the size of the instance, operating system, and the AWS Region where the instances are launched. Pricing is per instance-hour consumed for each instance, from the time an instance is launched until it's terminated or stopped. Each partial instance-hour consumed is billed per-second for Linux instances and as a full hour for all other instance types<sup>1</sup>. Therefore, the customer will be billed for 3 hours and 6 minutes for running an On-Demand Amazon Linux EC2 instance for 3 hours, 5 minutes, and 6 seconds. References: Understand Amazon EC2 instance-hours billing

**NEW QUESTION 208**

- (Topic 3)

Which AWS Cloud deployment model uses AWS Outposts as part of the application deployment infrastructure?

- A. On-premises
- B. Serverless
- C. Cloud-native
- D. Hybrid

**Answer:** D

**Explanation:**

AWS Outposts is a fully managed service that extends AWS infrastructure, services, APIs, and tools to customer premises. By providing local access to AWS managed infrastructure, AWS Outposts enables customers to build and run applications on premises using the same programming interfaces as in AWS Regions, while using local compute and storage resources for lower latency and local data processing needs. An Outpost is a pool of AWS compute and storage capacity deployed at a customer site. AWS operates, monitors, and manages this capacity as part of an AWS Region. You can create subnets on your Outpost and specify them when you create AWS resources such as EC2 instances, EBS volumes, ECS clusters, and RDS instances. Instances in Outpost subnets communicate with other instances in the AWS Region using private IP addresses, all within the same VPC. Outposts solutions allow you to extend and run native AWS services on premises, and is available in a variety of form factors, from 1U and 2U Outposts servers to 42U Outposts racks, and multiple rack deployments. With AWS Outposts, you can run some AWS services locally and connect to a broad range of services available in the local AWS Region<sup>2</sup>. AWS Outposts is a hybrid cloud

deployment model that uses AWS Outposts as part of the application deployment infrastructure. Hybrid cloud is a cloud computing environment that uses a mix of on-premises, private cloud, and public cloud services with orchestration between the platforms. Hybrid cloud provides businesses with greater flexibility, more deployment options, and optimized costs. By using AWS Outposts, customers can benefit from the fully managed infrastructure, services, APIs, and tools of AWS on premises, while still having access to the full range of AWS services available in the Region for a truly consistent hybrid experience<sup>3</sup>. References: On-Premises Private Cloud - AWS Outposts Family - AWS, What is AWS Outposts? - AWS Outposts

**NEW QUESTION 211**

- (Topic 3)

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

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

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

**Answer: C**

**Explanation:**

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

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

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

Reference:

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

**NEW QUESTION 213**

- (Topic 3)

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

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

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

**Answer: B**

**Explanation:**

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

**NEW QUESTION 217**

- (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 221**

- (Topic 2)

A company wants to securely store Amazon RDS database credentials and automatically rotate user passwords periodically.

Which AWS service or capability will meet these requirements?

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

**Answer: C**

**Explanation:**

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

**NEW QUESTION 225**

- (Topic 2)

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

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

**Answer:** B

**Explanation:**

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

**NEW QUESTION 227**

- (Topic 2)

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

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

**Answer:** AE

**Explanation:**

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

**NEW QUESTION 228**

- (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 233**

- (Topic 2)

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

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

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

**Answer: B**

**Explanation:**

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

**NEW QUESTION 237**

- (Topic 2)

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

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

**Answer: B**

**Explanation:**

The AWS service that is used to temporarily provide federated security credentials to a user is AWS Security Token Service (AWS STS). AWS STS is a service that enables customers to request temporary, limited-privilege credentials for AWS Identity and Access Management (IAM) users or for users that they authenticate (federated users). The company can use AWS STS to grant federated users access to AWS resources without creating permanent IAM users or sharing long-term credentials. AWS STS helps customers manage and secure access to their AWS resources for federated users. Amazon GuardDuty, AWS Secrets Manager, and AWS Certificate Manager are not the best services to use for this purpose. Amazon GuardDuty is a threat detection service that monitors for malicious activity and unauthorized behavior across the AWS accounts and resources. AWS Secrets Manager is a service that helps customers manage and rotate secrets, such as database credentials, API keys, and passwords. AWS Certificate Manager is a service that helps customers provision, manage, and deploy public and private Secure Sockets Layer/Transport Layer Security (SSL/TLS) certificates for use with AWS services and internal connected resources. These services are more useful for different types of security and compliance tasks, rather than providing temporary federated security credentials to a user.

**NEW QUESTION 242**

- (Topic 2)

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

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

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

**Answer: C**

**Explanation:**

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

**NEW QUESTION 247**

- (Topic 2)

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

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

**Answer: B**

**Explanation:**

The AWS service or tool that provides recommendations to help users get rightsized Amazon EC2 instances based on historical workload usage data is AWS Compute Optimizer. AWS Compute Optimizer is a service that analyzes the configuration and performance of the AWS resources, such as Amazon EC2 instances, and provides recommendations for optimal resource types and sizes based on the workload patterns and metrics. AWS Compute Optimizer helps users improve the performance, availability, and cost efficiency of their AWS resources. AWS Pricing Calculator, AWS App Runner, and AWS Systems Manager are not the best services or tools to use for this purpose. AWS Pricing Calculator is a tool that helps users estimate the cost of using AWS services based on their requirements and preferences. AWS App Runner is a service that helps users easily and quickly deploy web applications and APIs without managing any infrastructure. AWS Systems Manager is a service that helps users automate and manage the configuration and operation of their AWS resources and applications<sup>34</sup>

**NEW QUESTION 250**

- (Topic 2)

Which AWS service provides the SIMPLEST way for the company to establish a website on AWS?

- A. Amazon Elastic File System (Amazon EFS)
- B. AWS Elastic Beanstalk
- C. AWS Lambda
- D. Amazon Lightsail

**Answer:** D

**Explanation:**

Amazon Lightsail is an easy-to-use cloud platform that offers you everything needed to build an application or website, plus a cost-effective, monthly plan. Whether you're new to the cloud or looking to get on the cloud quickly with AWS infrastructure you trust, we've got you covered. Lightsail provides the simplest way for the company to establish a website on AWS.

**NEW QUESTION 253**

- (Topic 2)

A user is moving a workload from a local data center to an architecture that is distributed between the local data center and the AWS Cloud. Which type of migration is this?

- A. On-premises to cloud native
- B. Hybrid to cloud native
- C. On-premises to hybrid
- D. Cloud native to hybrid

**Answer:** C

**Explanation:**

C is correct because moving a workload from a local data center to an architecture that is distributed between the local data center and the AWS Cloud is an example of an on-premises to hybrid migration. A hybrid cloud is a cloud computing environment that uses a mix of on-premises, private cloud, and public cloud services with orchestration between the platforms. A is incorrect because on-premises to cloud native migration is the process of moving a workload from a local data center to an architecture that is fully hosted and managed on the AWS Cloud. B is incorrect because hybrid to cloud native migration is the process of moving a workload from an architecture that is distributed between the local data center and the AWS Cloud to an architecture that is fully hosted and managed on the AWS Cloud. D is incorrect because cloud native to hybrid migration is the process of moving a workload from an architecture that is fully hosted and managed on the AWS Cloud to an architecture that is distributed between the local data center and the AWS Cloud.

**NEW QUESTION 254**

- (Topic 2)

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

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

**Answer:** B

**Explanation:**

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

**NEW QUESTION 256**

.....

## Thank You for Trying Our Product

\* 100% Pass or Money Back

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

\* One year free update

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

\* Trusted by Millions

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

\* Shop Securely

All transactions are protected by VeriSign!

**100% Pass Your CLF-C02 Exam with Our Prep Materials Via below:**

<https://www.certleader.com/CLF-C02-dumps.html>