

# Amazon-Web-Services

## Exam Questions CLF-C02

AWS Certified Cloud Practitioner



### NEW QUESTION 1

- (Topic 3)

A company wants to monitor for misconfigured security groups that are allowing unrestricted access to specific ports. Which AWS service will meet this requirement?

- A. AWS Trusted Advisor
- B. Amazon CloudWatch
- C. Amazon GuardDuty
- D. AWS Health Dashboard

**Answer:** A

#### Explanation:

AWS Trusted Advisor is an online tool that provides you real time guidance to help you provision your resources following AWS best practices, including security and performance. It can help you monitor for misconfigured security groups that are allowing unrestricted access to specific ports. Amazon CloudWatch is a service that monitors your AWS resources and the applications you run on AWS. Amazon GuardDuty is a threat detection service that continuously monitors for malicious activity and unauthorized behavior. AWS Health Dashboard provides relevant and timely information to help you manage events in progress, and provides proactive notification to help you plan for scheduled activities.

### 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 is migrating to the AWS Cloud to meet storage needs. The company wants to optimize costs based on the amount of storage that the company uses. Which AWS offering or benefit will meet these requirements MOST cost-effectively?

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

**Answer:** D

#### Explanation:

Volume-based discounts are an AWS offering or benefit that can help the company optimize costs based on the amount of storage that the company uses.

Volume-based discounts are discounts that AWS provides for some storage services, such as Amazon S3 and Amazon EBS, when the company stores a large amount of data. The more data the company stores, the lower the price per GB. For example, Amazon S3 offers six storage classes, each with a different price per GB. The price per GB decreases as the amount of data stored in each storage class increases

### NEW QUESTION 5

- (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 6**

- (Topic 3)

Which AWS service is a cloud security posture management (CSPM) service that aggregates alerts from various AWS services and partner products in a standardized format?

- A. AWS Security Hub
- B. AWS Trusted Advisor
- C. Amazon EventBridge
- D. Amazon GuardDuty

**Answer:** A

**Explanation:**

AWS Security Hub is a cloud security posture management (CSPM) service that performs security best practice checks, aggregates alerts, and enables automated remediation. Security Hub collects findings from the security services enabled across your AWS accounts, such as intrusion detection findings from Amazon GuardDuty, vulnerability scans from Amazon Inspector, and sensitive data identification findings from Amazon Macie. Security Hub also collects findings from partner security products using a standardized AWS Security Finding Format, eliminating the need for time-consuming data parsing and normalization efforts. Customers can designate an administrator account that can access all findings across their accounts. References: AWS Security Hub Overview, AWS Security Hub FAQs

**NEW QUESTION 7**

- (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 8

- (Topic 3)

What is the LEAST expensive AWS Support plan that provides the full set of AWS Trusted Advisor best practice checks for cost optimization?

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

**Answer: B**

#### Explanation:

AWS Business Support is the least expensive AWS Support plan that provides the full set of AWS Trusted Advisor best practice checks for cost optimization. AWS Trusted Advisor is a service that provides best practices and recommendations for cost optimization, performance, security, and fault tolerance. AWS Business Support also provides other benefits, such as 24/7 technical support, unlimited cases, and faster response times. AWS Enterprise Support is the most expensive AWS Support plan that provides the same benefits as AWS Business Support, plus additional benefits, such as a technical account manager and enterprise concierge support. AWS Developer Support and AWS Basic Support are cheaper AWS Support plans that provide only a limited set of AWS Trusted Advisor best practice checks for cost optimization .

### NEW QUESTION 9

- (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 10

- (Topic 3)

Which of the following services can be used to block network traffic to an instance? (Select TWO.)

- A. Security groups
- B. Amazon Virtual Private Cloud (Amazon VPC) flow logs
- C. Network ACLs
- D. Amazon CloudWatch
- E. AWS CloudTrail

**Answer: AC**

#### Explanation:

Security groups and network ACLs are two AWS services that can be used to block network traffic to an instance. Security groups are virtual firewalls that control the inbound and outbound traffic for your instances at the instance level. You can specify which protocols, ports, and source or destination IP addresses are allowed or denied for each instance. Security groups are stateful, which means that they automatically allow return traffic for any allowed inbound or outbound traffic<sup>123</sup>. Network ACLs are virtual firewalls that control the inbound and outbound traffic for your subnets at the subnet level. You can create rules to allow or deny traffic based on protocols, ports, and source or destination IP addresses. Network ACLs are stateless, which means that you have to explicitly allow return traffic for any allowed inbound or outbound traffic<sup>456</sup>. References: 1: Security groups for your VPC - Amazon Virtual Private Cloud, 2: Security Groups for Your VPC - Amazon Elastic Compute Cloud, 3: AWS Security Groups: Everything You Need to Know, 4: Network ACLs - Amazon Virtual Private Cloud, 5: Control traffic to subnets using network ACLs - Amazon Virtual Private Cloud, 6: AWS Network ACLs: Everything You Need to Know

### NEW QUESTION 10

- (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 12

- (Topic 3)

A company wants to ensure that all of its Amazon EC2 instances have compliant operating system patches. Which AWS service will meet these requirements?

- A. AWS Compute Optimizer
- B. AWS Elastic Beanstalk
- C. AWS AppSync
- D. AWS Systems Manager

**Answer: D**

#### Explanation:

AWS Systems Manager gives you visibility and control of your infrastructure on AWS. Systems Manager provides a unified user interface so you can view operational data from multiple AWS services and allows you to automate operational tasks across your AWS resources. You can use Systems Manager to apply OS patches, create system images, configure Windows and Linux operating systems, and execute PowerShell commands<sup>5</sup>. Systems Manager can help you ensure that all of your Amazon EC2 instances have compliant operating system patches by using the Patch Manager feature.

#### NEW QUESTION 13

- (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<sup>5</sup>.

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 17

- (Topic 3)

What is a benefit of using AWS serverless computing?

- A. Application deployment and management are not required
- B. Application security will be fully managed by AWS
- C. Monitoring and logging are not needed
- D. Management of infrastructure is offloaded to AWS

**Answer: D**

#### Explanation:

AWS serverless computing is a way of building and running applications without thinking about servers. AWS manages the infrastructure for you, so you don't have to provision, scale, patch, or monitor servers. You only pay for the compute time you consume, and you can focus on your application logic instead of managing servers<sup>12</sup>. References: Serverless Computing – Amazon Web Services, AWS Serverless Computing, Benefits, Architecture and Use-cases - XenonStack

#### NEW QUESTION 22

- (Topic 3)

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

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

**Answer: B**

**Explanation:**

AWS Artifact is a service that provides on-demand access to security and compliance reports from AWS and Independent Software Vendors (ISVs) who sell their products on AWS Marketplace. You can use AWS Artifact to download auditor-issued reports, certifications, accreditations, and other third-party attestations of AWS compliance with various standards and regulations, such as PCI-DSS, HIPAA, FedRAMP, GDPR, and more<sup>1234</sup>. You can also use AWS Artifact to review, accept, and manage your agreements with AWS and apply them to current and future accounts within your organization<sup>2</sup>. References: 1: Cloud Compliance - Amazon Web Services (AWS), 2: Security Compliance Management - AWS Artifact - AWS, 3: AWS Compliance Contact Us - Amazon Web Services, 4: AWS SECURITY AND COMPLIANCE QUICK REFERENCE GUIDE

**NEW QUESTION 23**

- (Topic 3)

A company needs to migrate a PostgreSQL database from on-premises to Amazon RDS. Which AWS service or tool should the company use to meet this requirement?

- A. Cloud Adoption Readiness Tool
- B. AWS Migration Hub
- C. AWS Database Migration Service (AWS DMS)
- D. AWS Application Migration Service

**Answer: C**

**Explanation:**

AWS Database Migration Service (AWS DMS) is a managed and automated 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 DMS supports migration between 20-plus database and analytics engines, such as PostgreSQL, Oracle, MySQL, SQL Server, MongoDB, Amazon Aurora, Amazon RDS, Amazon Redshift, and Amazon S3. AWS DMS also provides schema conversion and validation tools, as well as monitoring and security features. AWS DMS is a cost-effective and reliable solution for database migration, as you only pay for the compute resources and additional log storage used during the migration process, and you can minimize the downtime and data loss with

Multi-AZ and ongoing replication<sup>12</sup>

To migrate a PostgreSQL database from on-premises to Amazon RDS using AWS DMS, you need to perform the following steps:

? Create an AWS DMS replication instance in the same AWS Region as your target Amazon RDS PostgreSQL DB instance. The replication instance is a server that runs the AWS DMS replication software and connects to your source and target endpoints. You can choose the instance type, storage, and network settings based on your migration requirements<sup>3</sup>

? Create a source endpoint that points to your on-premises PostgreSQL database.

You need to provide the connection details, such as the server name, port, database name, user name, and password. You also need to specify the engine name as postgres and the SSL mode as required<sup>4</sup>

? Create a target endpoint that points to your Amazon RDS PostgreSQL DB instance. You need to provide the connection details, such as the server name, port, database name, user name, and password. You also need to specify the engine name as postgres and the SSL mode as verify-full.

? Create a migration task that defines the migration settings and options, such as the replication instance, the source and target endpoints, the migration type (full load, full load and change data capture, or change data capture only), the table mappings, the task settings, and the task monitoring role. You can also use the AWS Schema Conversion Tool (AWS SCT) to convert your source schema to the target schema and apply it to the target endpoint before or after creating the migration task.

? Start the migration task and monitor its progress and status using the AWS DMS console, the AWS CLI, or the AWS DMS API. You can also use AWS CloudFormation to automate the creation and execution of the migration task.

The other options are not suitable for migrating a PostgreSQL database from on-premises to Amazon RDS. Cloud Adoption Readiness Tool is a tool that helps you assess your readiness for cloud adoption based on six dimensions: business, people, process, platform, operations, and security. It does not perform any database migration tasks. AWS Migration Hub is a service that helps you track and manage the progress of your application migrations across multiple AWS and partner services, such as AWS DMS, AWS Application Migration Service, AWS Server Migration Service, and CloudEndure Migration. It does not perform any database migration tasks itself, but rather integrates with other migration services. 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. It does not support database migration, but rather replicates your servers as Amazon Machine Images (AMIs) and launches them as EC2 instances on AWS. References: AWS Database Migration Service, What is AWS Database Migration Service?, Working with an AWS DMS replication instance, Creating source and target endpoints for PostgreSQL, [Creating a target endpoint for Amazon RDS for PostgreSQL], [Creating a migration task for AWS DMS], [AWS Schema Conversion Tool], [Starting a migration task for AWS DMS], [AWS CloudFormation], [Cloud Adoption Readiness Tool], [AWS Migration Hub], [AWS Application Migration Service]

**NEW QUESTION 28**

- (Topic 3)

A company processes personally identifiable information (PII) and must keep data in the country where it was generated. The company wants to use Amazon EC2 instances for these workloads.

Which AWS service will meet these requirements?

- A. AWS Outposts
- B. AWS Storage Gateway
- C. AWS DataSync
- D. AWS OpsWorks

**Answer: A**

**Explanation:**

AWS Outposts is an AWS service that extends AWS infrastructure, services, APIs, and tools to virtually any datacenter, co-location space, or on-premises facility. AWS Outposts enables you to run Amazon EC2 instances and other AWS services locally, while maintaining a consistent and seamless connection to the AWS Cloud. AWS Outposts is ideal for workloads that require low latency, local data processing, or data residency. By using AWS Outposts, the company can process personally identifiable information (PII) and keep data in the country where it was generated, while leveraging the benefits of AWS

**NEW QUESTION 30**

- (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 32**

- (Topic 3)

A company needs to run a workload for several batch image rendering applications. It is acceptable for the workload to experience downtime. Which Amazon EC2 pricing model would be MOST cost-effective in this situation?

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

**Answer:** D

**Explanation:**

Amazon EC2 Spot Instances are instances that use spare EC2 capacity that is available at up to a 90% discount compared to On-Demand prices. You can use Spot Instances for various stateless, fault-tolerant, or flexible applications such as big data, containerized workloads, high-performance computing (HPC), and test & development workloads. Spot Instances are ideal for workloads that can be interrupted, such as batch image rendering applications<sup>1</sup>. On-Demand Instances are instances that let you pay for compute capacity by the hour or second (minimum of 60 seconds) with no long-term commitments. This frees you from the costs and complexities of planning, purchasing, and maintaining hardware and transforms what are commonly large fixed costs into much smaller variable costs<sup>2</sup>. Reserved Instances are instances that provide you with a significant discount (up to 75%) compared to On-Demand Instance pricing. In exchange, you select a term and make an upfront payment to reserve a certain amount of compute capacity for that term<sup>3</sup>. Dedicated Instances are instances that run in a VPC on hardware that's dedicated to a single customer. Your Dedicated Instances are physically isolated at the host hardware level from instances that belong to other AWS accounts<sup>4</sup>.

**NEW QUESTION 37**

- (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 40**

- (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 45

- (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 47

- (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 49

- (Topic 3)

A company wants to integrate its online shopping website with social media login credentials. Which AWS service can the company use to make this integration?

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

**Answer: C**

#### Explanation:

Amazon Cognito is a service that enables you to add user sign-up and sign-in features to your web and mobile applications. Amazon Cognito also supports social and enterprise identity federation, which means you can allow your users to sign in with their existing credentials from identity providers such as Google, Facebook, Apple, and Amazon. Amazon Cognito integrates with OpenID Connect (OIDC) and Security Assertion Markup Language (SAML) 2.0 protocols to facilitate the authentication and authorization process. Amazon Cognito also provides advanced security features, such as adaptive authentication, user verification, and multi-factor authentication (MFA). References: Amazon Cognito, What is Amazon Cognito?

#### NEW QUESTION 53

- (Topic 3)

Which AWS service or feature is associated with a subnet in a VPC and is used to control inbound and outbound traffic?

- A. Amazon Inspector
- B. Network ACLs
- C. AWS Shield
- D. VPC Flow Logs

**Answer:** B

**Explanation:**

Network ACLs (network access control lists) are an optional layer of security for your VPC that act as a firewall for controlling traffic in and out of one or more subnets. You can use network ACLs to allow or deny traffic based on protocol, port, or source and destination IP address. Network ACLs are stateless, meaning that they do not track the traffic that flows through them. Therefore, you must create rules for both inbound and outbound traffic.

**NEW QUESTION 56**

- (Topic 3)

A company needs to store infrequently used data for data archives and long-term backups.

A company needs a history report about how its Amazon EC2 instances were modified last month.

Which AWS service can be used to meet this requirement?

- A. AWS Service Catalog
- B. AWS Config
- C. Amazon CloudWatch
- D. AWS Artifact

**Answer:** B

**Explanation:**

AWS Config is a service that enables you to assess, audit, and evaluate the configurations of your AWS resources. AWS Config continuously monitors and records

your AWS resource configurations and allows you to automate the evaluation of recorded configurations against desired configurations. AWS Config can also track changes to your EC2 instances over time and provide a history report of the modifications. AWS Service Catalog, Amazon CloudWatch, and AWS Artifact are not the best services to meet this requirement. AWS Service Catalog is a service that allows you to create and manage catalogs of IT services that are approved for use on AWS. Amazon CloudWatch is a service that monitors your AWS resources and applications and provides metrics, alarms, dashboards, and logs. AWS Artifact is a service that provides on-demand access to AWS security and compliance reports and online agreements

**NEW QUESTION 58**

- (Topic 3)

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

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

**Answer:** D

**Explanation:**

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

**NEW QUESTION 60**

- (Topic 3)

Which company needs to apply security rules to a subnet for Amazon EC2 instances. Which AWS service or feature provides this functionality?

- A. Network ACLs
- B. Security groups
- C. AWS Certificate Manager (ACM)
- D. AWS Config

**Answer:** A

**Explanation:**

Network ACLs (network access control lists) are an AWS service or feature that provides the functionality of applying security rules to a subnet for EC2 instances. A subnet is a logical partition of an IP network within a VPC (virtual private cloud). A VPC is a logically isolated section of the AWS Cloud where the company can launch AWS resources in a virtual network that they define. A network ACL is a virtual firewall that controls the inbound and outbound traffic for one or more subnets. The company can use network ACLs to allow or deny traffic based on protocol, port, or source and destination IP address. Network ACLs are stateless, meaning that they do not track the traffic that flows through them. Therefore, the company must create rules for both inbound and outbound traffic

**NEW QUESTION 61**

- (Topic 3)

A company needs a fully managed file server that natively supports Microsoft workloads and file systems The file server must also support the SMB protocol.

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

- A. Amazon Elastic File System (Amazon EFS)
- B. Amazon FSx for Lustre
- C. Amazon FSx for Windows File Server
- D. Amazon Elastic Block Store (Amazon EBS)

**Answer:** C

**Explanation:**

Amazon FSx for Windows File Server is a fully managed file server that supports Microsoft workloads and file systems, including the SMB protocol. It provides features such as user quotas, end-user file restore, and Microsoft Active Directory integration. Amazon EFS is a fully managed file system that supports the NFS protocol, not SMB. Amazon FSx for Lustre is a fully managed file system that supports high-performance computing workloads, not Microsoft workloads. Amazon EBS is a block storage service that does not provide a file system or SMB support. References: Amazon FSx for Windows File Server, Amazon FSx for Lustre, Amazon EFS, Amazon EBS

**NEW QUESTION 63**

- (Topic 3)

A company wants to set AWS spending targets and track costs against those targets. Which AWS tool or feature should the company use to meet these requirements?

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

**Answer: B**

**Explanation:**

AWS Budgets is a tool that allows users to set AWS spending targets and track costs against those targets. Users can create budgets for various dimensions, such as service, linked account, tag, and more. Users can also receive alerts when the actual or forecasted costs exceed or are projected to exceed the budgeted amount. AWS Cost Explorer, AWS Cost and Usage Report, and Savings Plans are other AWS tools or features that can help users manage and optimize their AWS costs, but they do not enable users to set and track spending targets.

**NEW QUESTION 67**

- (Topic 3)

A company wants to launch its web application in a second AWS Region. The company needs to determine which services must be regionally configured for this launch.

Which AWS services can be configured at the Region level? (Select TWO.)

- A. Amazon EC2
- B. Amazon Route 53
- C. Amazon CloudFront
- D. AWS WAF
- E. Amazon DynamoDB

**Answer: BD**

**Explanation:**

Amazon Route 53 and AWS WAF are AWS services that can be configured at the Region level. Amazon Route 53 is a highly available and scalable cloud Domain Name System (DNS) web service that lets you register domain names, route traffic to resources, and check the health of your resources. 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. Amazon EC2, Amazon CloudFront, and Amazon DynamoDB are AWS services that can be configured at the global level or the Availability Zone level.

**NEW QUESTION 71**

- (Topic 3)

Which AWS service provides the ability to manage infrastructure as code?

- A. AWS CodePipeline
- B. AWS CodeDeploy
- C. AWS Direct Connect
- D. AWS CloudFormation

**Answer: D**

**Explanation:**

The AWS service that provides the ability to manage infrastructure as code is AWS CloudFormation. 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 you 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 you to monitor and control the changes to your infrastructure.

**NEW QUESTION 73**

- (Topic 3)

Which AWS services can be used to store files? (Select TWO.)

- A. Amazon S3
- B. AWS Lambda
- C. Amazon Elastic Block Store (Amazon EBS)
- D. Amazon SageMaker
- E. AWS Storage Gateway

**Answer: AC**

**Explanation:**

Amazon S3 and Amazon EBS are two AWS services that can be used to store files. Amazon S3 is an object storage service that offers high scalability, durability,

availability, and performance. Amazon EBS is a block storage service that provides persistent and low-latency storage volumes for Amazon EC2 instances. AWS Lambda, Amazon SageMaker, and AWS Storage Gateway are other AWS services that have different purposes, such as serverless computing, machine learning, and hybrid cloud storage .

#### NEW QUESTION 75

- (Topic 3)

Which AWS service provides storage that can be mounted across multiple Amazon EC2 instances?

- A. Amazon Workspaces
- B. Amazon Elastic File System (Amazon EFS)
- C. AWS Database Migration Service (AWS DMS)
- D. AWS Snowball Edge

**Answer: B**

#### Explanation:

Amazon EFS is a fully managed service that provides scalable and elastic file storage for multiple Amazon EC2 instances. Amazon EFS supports the Network File System (NFS) protocol, which allows multiple EC2 instances to access the same file system concurrently. You can learn more about Amazon EFS from this webpage or this digital course.

#### NEW QUESTION 76

- (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<sup>1</sup>. 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<sup>2</sup>. Amazon Rekognition is a computer vision service that can analyze images and videos for faces, objects, scenes, text, and more<sup>3</sup>. Amazon Lex is a service for building conversational interfaces using voice and text<sup>4</sup>.

#### NEW QUESTION 81

- (Topic 3)

A company has created an AWS Cost and Usage Report and wants to visualize the report. Which AWS service should the company use to ingest and display this information?

- A. Amazon QuickSight
- B. Amazon Pinpoint
- C. Amazon Neptune
- D. Amazon Kinesis

**Answer: A**

#### Explanation:

Amazon QuickSight is an AWS service that provides business intelligence and data visualization capabilities. Amazon QuickSight enables you to ingest, analyze, and display data from various sources, such as AWS Cost and Usage Reports, Amazon S3, Amazon Athena, Amazon Redshift, and Amazon RDS. You can use Amazon QuickSight to create interactive dashboards and charts that show insights and trends from your data. You can also share your dashboards and charts with other users or embed them into your applications.

#### NEW QUESTION 86

- (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 88

- (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>12</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>13</sup>.

**NEW QUESTION 92**

- (Topic 3)

A company is running its application in the AWS Cloud. The company wants to periodically review its AWS account for cost optimization opportunities. Which AWS service or tool can the company use to meet these requirements?

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

**Answer: A**

**Explanation:**

AWS Cost Explorer is an AWS service or tool that the company can use to periodically review its AWS account for cost optimization opportunities. AWS Cost Explorer is a tool that enables the company to visualize, understand, and manage their AWS costs and usage over time. The company can use AWS Cost Explorer to access interactive graphs and tables that show the breakdown of their costs and usage by service, region, account, tag, and more. The company can also use AWS Cost Explorer to forecast their future costs, identify trends and anomalies, and discover potential savings by using Reserved Instances or Savings Plans.

**NEW QUESTION 95**

- (Topic 3)

Which AWS services or features give users the ability to create a network connection between two VPCs? (Select TWO.)

- A. VPC endpoints
- B. Amazon Route 53
- C. VPC peering
- D. AWS Direct Connect
- E. AWS Transit Gateway

**Answer: CE**

**Explanation:**

VPC peering and AWS Transit Gateway are two AWS services or features that give users the ability to create a network connection between two VPCs. VPC peering is a networking connection between two VPCs that enables you to route traffic between them privately. You can create a VPC peering connection between your own VPCs, with a VPC in another AWS account, or with a VPC in a different AWS Region. Traffic between peered VPCs never traverses the public internet. VPC peering does not support transitive peering relationships, which means that if VPC A is peered with VPC B, and VPC B is peered with VPC C, then VPC A and VPC C are not automatically peered<sup>789</sup>. AWS Transit Gateway is a networking service that acts as a regional router for your VPCs and on-premises networks. You can attach up to 5,000 VPCs and VPN connections to a single transit gateway and route traffic between them. AWS Transit Gateway simplifies the management and scalability of your network architecture, as you only need to create and manage a single connection from the central transit gateway to each connected network. AWS Transit Gateway supports transitive routing, which means that any network that is attached to the transit gateway can communicate with any other network that is attached to the same transit gateway . References: 7: VPC peering - Amazon Virtual Private Cloud, 8: Connect VPCs using VPC peering - Amazon Virtual Private Cloud, 9: Amazon VPC-to-Amazon VPC connectivity options - Amazon Virtual Private Cloud, : [AWS Transit Gateway - Amazon Web Services], : [Connect VPCs using AWS Transit Gateway - Amazon Virtual Private Cloud], : [AWS Transit Gateway: Simplify Your Network Architecture]

**NEW QUESTION 100**

- (Topic 3)

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

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

**Answer: BE**

**Explanation:**

Elasticity in the AWS Cloud refers to the ability to acquire resources as you need them and release resources when you no longer need them. In the cloud, you want to do this automatically<sup>1</sup>. This means that you can rightsized resources as demand shifts, and you can easily procure resources when they are needed. Elasticity is not related to how quickly an Amazon EC2 instance can be restarted, the maximum amount of RAM an Amazon EC2 instance can use, or the pay-as-you-go billing model. These are aspects of scalability, performance, and cost, respectively<sup>2</sup>.

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

? Elasticity - AWS Well-Architected Framework

- ? Elastic - Reactive Systems on AWS
- ? What is the difference between scalability and elasticity?

#### NEW QUESTION 105

- (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.

#### NEW QUESTION 110

- (Topic 3)

An IT engineer needs to access AWS services from an on-premises application. Which credentials or keys does the application need for authentication?

- A. AWS account user name and password
- B. IAM access key and secret
- C. Amazon EC2 key pairs
- D. AWS Key Management Service (AWS KMS) keys

**Answer: B**

#### Explanation:

IAM access keys are long-term credentials that consist of an access key ID and a secret access key. You use access keys to sign programmatic requests that you make to AWS. If you need to access AWS services from an on-premises application, you can use IAM access keys to authenticate your requests. AWS account user name and password are used to sign in to the AWS Management Console. Amazon EC2 key pairs are used to connect to your EC2 instances using SSH. AWS Key Management Service (AWS KMS) keys are used to encrypt and decrypt your data using the AWS Encryption SDK or the AWS CLI.

#### NEW QUESTION 114

- (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).

#### NEW QUESTION 119

- (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 121

- (Topic 3)

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

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

**Answer:** BE

#### Explanation:

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

#### NEW QUESTION 126

- (Topic 3)

A company has 5 TB of data stored in Amazon S3. The company plans to occasionally run queries on the data for analysis. Which AWS service should the company use to run these queries in the MOST cost-effective manner?

- A. Amazon Redshift
- B. Amazon Athena
- C. Amazon Kinesis
- D. Amazon RDS

**Answer:** B

#### Explanation:

Amazon Athena is a serverless, interactive analytics service that allows users to run SQL queries on data stored in Amazon S3. It is ideal for occasional queries on large datasets, as it does not require any server provisioning, configuration, or management. Users only pay for the queries they run, based on the amount of data scanned. Amazon Athena supports various data formats, such as CSV, JSON, Parquet, ORC, and Avro, and integrates with AWS Glue Data Catalog to create and manage schemas. Amazon Athena also supports querying data from other sources, such as on-premises or other cloud systems, using data connectors<sup>1</sup>.

Amazon Redshift is a fully managed data warehouse service that allows users to run complex analytical queries on petabyte-scale data. However, it requires users to provision and maintain clusters of nodes, and pay for the storage and compute capacity they use. Amazon Redshift is more suitable for frequent and consistent queries on structured or semi-structured data<sup>2</sup>.

Amazon Kinesis is a platform for streaming data on AWS, enabling users to collect, process, and analyze real-time data. It is not designed for querying data stored in Amazon S3. Amazon Kinesis consists of four services: Kinesis Data Streams, Kinesis Data Firehose, Kinesis Data Analytics, and Kinesis Video Streams<sup>3</sup>.

Amazon RDS is a relational database service that provides six database engines: Amazon Aurora, PostgreSQL, MySQL, MariaDB, Oracle Database, and SQL Server. It simplifies database administration tasks such as backup, patching, scaling, and replication. However, it is not optimized for querying data stored in Amazon S3. Amazon RDS is more suitable for transactional workloads that require high performance and availability<sup>4</sup>.

References:

? Interactive SQL - Serverless Query Service - Amazon Athena - AWS

? [Amazon Redshift – Data Warehouse Solution - AWS]

? [Amazon Kinesis - Streaming Data Platform - AWS]

? [Amazon Relational Database Service (RDS) – AWS]

#### NEW QUESTION 131

- (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 135

- (Topic 3)

A developer wants to deploy an application quickly on AWS without manually creating the required resources. Which AWS service will meet these requirements?

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

**Answer: B**

**Explanation:**

AWS Elastic Beanstalk is a service that allows you to deploy and manage applications on AWS without manually creating and configuring the required resources, such as EC2 instances, load balancers, security groups, databases, and more. AWS Elastic Beanstalk automatically handles the provisioning, scaling, load balancing, health monitoring, and updating of your application, while giving you full control over the underlying AWS resources if needed. AWS Elastic Beanstalk supports a variety of platforms and languages, such as Java, .NET, PHP, Node.js, Python, Ruby, Go, and Docker. You can use the AWS Management Console, the AWS CLI, the AWS SDKs, or the AWS Elastic Beanstalk API to create and manage your applications. You can also use AWS CodeStar, AWS CodeCommit, AWS CodeBuild, AWS CodeDeploy, and AWS CodePipeline to integrate AWS Elastic Beanstalk with your development and deployment workflows<sup>12</sup>

**NEW QUESTION 137**

- (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: 2: AWS Lambda - Amazon Web Services (AWS), 3: AWS Lambda Documentation, 1: Amazon CLF-C02: What is customer responsibility under AWS ... - PUPUWEB

**NEW QUESTION 142**

- (Topic 3)

A company has teams that have different job roles and responsibilities. The company's employees often change teams. The company needs to manage permissions for the employees so that the permissions are appropriate for the job responsibilities.

Which IAM resource should the company use to meet this requirement with the LEAST operational overhead?

- A. IAM user groups
- B. IAM roles
- C. IAM instance profiles
- D. IAM policies for individual users

**Answer: B**

**Explanation:**

IAM roles are a way of granting temporary permissions to entities that need to access AWS resources, such as users, applications, or services. IAM roles allow customers to assign permissions to entities without having to create or manage IAM users or credentials for them. IAM roles can be assumed by different entities depending on the trust policy attached to the role. For example, IAM roles can be assumed by IAM users in the same or different AWS accounts, AWS services such as EC2 or Lambda, or external identities such as federated users or web identities. IAM roles can also be switched by IAM users to temporarily change their permissions. IAM roles are recommended for managing permissions for employees who often change teams, because they allow customers to define permissions based on job roles and responsibilities, and easily assign or revoke them as needed. IAM roles also reduce the operational overhead of creating, updating, or deleting IAM users or credentials for each employee or team change.

**NEW QUESTION 145**

- (Topic 3)

A company is building a mobile app to provide shopping recommendations to its customers. The company wants to use a graph database as part of the shopping recommendation engine.

Which AWS database service should the company choose?

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

**Answer: C**

**Explanation:**

Amazon Neptune is a service that provides a fully managed graph database that supports property graphs and RDF graphs. It can be used to build applications that work with highly connected datasets, such as shopping recommendations, social networks, fraud detection, and knowledge graphs<sup>2</sup>. Amazon DynamoDB is a service that provides a fully managed NoSQL database that delivers fast and consistent performance at any scale. Amazon Aurora is a service that provides a fully managed relational database that is compatible with MySQL and PostgreSQL. Amazon DocumentDB (with MongoDB compatibility) is a service that provides a fully managed document database that is compatible with MongoDB.

#### NEW QUESTION 148

- (Topic 3)

A company wants an AWS service to provide product recommendations based on its customer data. Which AWS service will meet this requirement?

- A. Amazon Polly
- B. Amazon Personalize
- C. Amazon Comprehend
- D. Amazon Rekognition

**Answer: B**

#### Explanation:

Amazon Personalize is an AWS service that helps developers quickly build and deploy a custom recommendation engine with real-time personalization and user segmentation<sup>1</sup>. It uses machine learning (ML) to analyze customer data and provide relevant recommendations based on their preferences, behavior, and context. Amazon Personalize can be used for various use cases such as optimizing recommendations, targeting customers more accurately, maximizing the value of unstructured text, and promoting items using business rules<sup>1</sup>.

The other options are not suitable for providing product recommendations based on customer data. Amazon Polly is a service that converts text into lifelike speech. Amazon Comprehend is a service that uses natural language processing (NLP) to extract insights from text and documents. Amazon Rekognition is a service that uses computer vision (CV) to analyze images and videos for faces, objects, scenes, and activities.

References:

- ? 1: Cloud Products - Amazon Web Services (AWS)
- ? 2: Recommender System – Amazon Personalize – Amazon Web Services
- ? 3: Top 25 AWS Services List 2023 - GeeksforGeeks
- ? 4: AWS to Azure services comparison - Azure Architecture Center
- ? 5: The 25+ Best AWS Cost Optimization Tools (Updated 2023) - CloudZero
- ? 6: Amazon Polly – Text-to-Speech Service - AWS
- ? 7: Natural Language Processing - Amazon Comprehend - AWS
- ? 8: Image and Video Analysis - Amazon Rekognition - AWS

#### NEW QUESTION 150

- (Topic 3)

What is the purpose of having an internet gateway within a VPC?

- A. To create a VPN connection to the VPC
- B. To allow communication between the VPC and the internet
- C. To impose bandwidth constraints on internet traffic
- D. To load balance traffic from the internet across Amazon EC2 instances

**Answer: B**

#### Explanation:

An internet gateway is a service that allows for internet traffic to enter into a VPC. Otherwise, a VPC is completely segmented off and then the only way to get to it is potentially through a VPN connection rather than through internet connection. An internet gateway is a logical connection between an AWS VPC and the internet. It supports IPv4 and IPv6 traffic. It does not cause availability risks or bandwidth constraints on your network traffic<sup>1</sup>. An internet gateway enables resources in your public subnets (such as EC2 instances) to connect to the internet if the resource has a public IPv4 address or an IPv6 address. Similarly, resources on the internet can initiate a connection to resources in your subnet using the public IPv4 address or IPv6 address<sup>2</sup>. An internet gateway also provides a target in your VPC route tables for internet-routable traffic. For communication using IPv4, the internet gateway also performs network address translation (NAT). For communication using IPv6, NAT is not needed because IPv6 addresses are public<sup>2</sup>. To enable access to or from the internet for instances in a subnet in a VPC using an internet gateway, you must create an internet gateway and attach it to your VPC, add a route to your subnet's route table that directs internet-bound traffic to the internet gateway, ensure that instances in your subnet have a public IPv4 address or an IPv6 address, and ensure that your network access control lists and security group rules allow the desired internet traffic to flow to and from your instance<sup>2</sup>. References: Connect to the internet using an internet gateway, AWS Internet Gateway and VPC Routing

#### NEW QUESTION 155

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

#### NEW QUESTION 158

- (Topic 3)

A company wants to monitor its workload performance. The company wants to ensure that the cloud services are delivered at a level that meets its business

needs.

Which AWS Cloud Adoption Framework (AWS CAF) perspective will meet these requirements?

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

**Answer: D**

**Explanation:**

The Operations perspective helps you monitor and manage your cloud workloads to ensure that they are delivered at a level that meets your business needs. Common stakeholders include chief operations officer (COO), cloud director, cloud operations manager, and cloud operations engineers<sup>1</sup>. The Operations perspective covers capabilities such as workload health monitoring, incident management, change management, release management, configuration management, and disaster recovery<sup>2</sup>. The Business perspective helps ensure that your cloud investments accelerate your digital transformation ambitions and business outcomes. Common stakeholders include chief executive officer (CEO), chief financial officer (CFO), chief information officer (CIO), and chief technology officer (CTO). The Business perspective covers capabilities such as business case development, value realization, portfolio management, and stakeholder management<sup>3</sup>.

The Governance perspective helps you orchestrate your cloud initiatives while maximizing organizational benefits and minimizing transformation-related risks. Common stakeholders include chief transformation officer, CIO, CTO, CFO, chief data officer (CDO), and chief risk officer (CRO). The Governance perspective covers capabilities such as governance framework, budget and cost management, compliance management, and data governance<sup>4</sup>.

The Platform perspective helps you build an enterprise-grade, scalable, hybrid cloud platform, modernize existing workloads, and implement new cloud-native solutions. Common stakeholders include CTO, technology leaders, architects, and engineers. The Platform perspective covers capabilities such as platform design and implementation, workload migration and modernization, cloud-native development, and DevOps<sup>5</sup>. References:

- ? AWS Cloud Adoption Framework: Operations Perspective
- ? AWS Cloud Adoption Framework - Operations Perspective
- ? AWS Cloud Adoption Framework: Business Perspective
- ? AWS Cloud Adoption Framework: Governance Perspective
- ? AWS Cloud Adoption Framework: Platform Perspective

**NEW QUESTION 163**

- (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 168**

- (Topic 3)

A company needs to implement identity management for a fleet of mobile apps that are running in the AWS Cloud.

Which AWS service will meet this requirement?

- A. Amazon Cognito
- B. AWS Security Hub
- C. AWS Shield
- D. AWS WAF

**Answer: A**

**Explanation:**

Amazon Cognito is a service that provides identity management for mobile and web applications, allowing users to sign up, sign in, and access AWS resources with different identity providers. AWS Security Hub is a service that provides a comprehensive view of the security posture of AWS accounts and resources. AWS Shield is a service that provides protection against distributed denial of service (DDoS) attacks. AWS WAF is a web application firewall that helps protect web applications from common web exploits.

**NEW QUESTION 169**

- (Topic 3)

A company runs a MySQL database in its on-premises data center. The company wants to run a copy of this database in the AWS Cloud.

Which AWS service would support this workload?

- A. Amazon RDS
- B. Amazon Neptune
- C. Amazon ElastiCache for Redis
- D. Amazon Quantum Ledger Database (Amazon QLDB)

**Answer: A**

**Explanation:**

Amazon Relational Database Service (Amazon RDS) is a web service that makes it easier to set up, operate, and scale a relational database in the cloud. It provides cost-efficient and resizable capacity, while automating time-consuming administration tasks such as hardware provisioning, database setup, patching, and backups. Amazon RDS supports six popular database engines: Amazon Aurora, PostgreSQL, MySQL, MariaDB, Oracle Database, and SQL Server. Amazon RDS can support running a copy of a MySQL database in the AWS Cloud, as it offers compatibility, scalability, and availability features.

**NEW QUESTION 172**

- (Topic 3)

Which AWS service provides a single location to track the progress of application migrations?

- A. AWS Application Discovery Service
- B. AWS Application Migration Service
- C. AWS Service Catalog
- D. AWS Migration Hub

**Answer: D**

**Explanation:**

AWS Migration Hub is a service that provides a single location to track the progress of application migrations across multiple AWS and partner solutions. It allows you to choose the AWS and partner migration tools that best fit your needs, while providing visibility into the status of migrations across your portfolio of applications<sup>1</sup>. AWS Migration Hub supports migration status updates from the following tools: AWS Application Migration Service, AWS Database Migration Service, CloudEndure Migration, Server Migration Service, and Migrate for Compute Engine<sup>1</sup>.

The other options are not correct for the following reasons:

? AWS Application Discovery Service is a service that helps you plan your migration projects by automatically identifying servers, applications, and dependencies in your on-premises data centers<sup>2</sup>. It does not track the progress of application migrations, but rather provides information to help you plan and scope your migrations.

? AWS Application Migration Service is a service that helps you migrate and modernize applications from any source infrastructure to AWS with minimal downtime and disruption<sup>3</sup>. It is one of the migration tools that can send status updates to AWS Migration Hub, but it is not the service that provides a single location to track the progress of application migrations.

? AWS Service Catalog is a service that allows you to create and manage catalogs of IT services that are approved for use on AWS<sup>4</sup>. It does not track the progress of application migrations, but rather helps you manage the provisioning and governance of your IT services.

References:

? 1: What Is AWS Migration Hub? - AWS Migration Hub

? 2: What Is AWS Application Discovery Service? - AWS Application Discovery Service

? 3: App Migration Tool - AWS Application Migration Service - AWS

? 4: What Is AWS Service Catalog? - AWS Service Catalog

**NEW QUESTION 174**

- (Topic 3)

A company wants to generate a list of IAM users. The company also wants to view the status of various credentials that are associated with the users, such as password, access keys, and multi-factor authentication (MFA) devices

Which AWS service or feature will meet these requirements?

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

**Answer: A**

**Explanation:**

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

**NEW QUESTION 179**

- (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:**

These are two of the seven capabilities that are in the platform perspective of the AWS Cloud Adoption Framework (AWS CAF). The platform perspective helps you build an enterprise-grade, scalable, hybrid cloud platform, modernize existing workloads, and implement new cloud-native solutions<sup>1</sup>. The other five capabilities are:

? Platform architecture – Establish and maintain guidelines, principles, patterns, and guardrails for your cloud environment.

? Platform engineering – Build a compliant multi-account cloud environment with enhanced security features, and packaged, reusable cloud products.

? Platform operations – Manage and optimize your cloud environment with automation, monitoring, and incident response.

? Application development – Develop and deploy cloud-native applications using modern architectures and best practices.

? Application migration – Migrate your existing applications to the cloud using proven methodologies and tools.

Performance and capacity management, infrastructure protection, and change and release management are not capabilities of the platform perspective. They are part of the operations perspective, which helps you achieve operational excellence in the cloud<sup>2</sup>. The operations perspective comprises six capabilities:

? Performance and capacity management – Monitor and optimize the performance and capacity of your cloud workloads.

- ? Infrastructure protection – Protect your cloud infrastructure from unauthorized access, malicious attacks, and data breaches.
- ? Change and release management – Manage changes and releases to your cloud workloads using automation and governance.
- ? Configuration management – Manage the configuration of your cloud resources and applications using automation and version control.
- ? Incident management – Respond to incidents affecting your cloud workloads using best practices and tools.
- ? Service continuity management – Ensure the availability and resilience of your cloud workloads using backup, recovery, and disaster recovery strategies.

#### NEW QUESTION 182

- (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 lets you ask questions about your data using everyday language and get answers in seconds. You can type questions such as “What are the total sales by region?” or “How did marketing campaign A perform?” and get answers in the form of relevant visualizations, such as charts or tables. You can also use Q to drill down into details, filter data, or perform calculations. Q uses machine learning to understand your data and your intent, and provides suggestions and feedback to help you refine your questions.

#### NEW QUESTION 186

- (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 191

- (Topic 3)

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

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

**Answer: A**

#### Explanation:

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

#### NEW QUESTION 195

- (Topic 3)

A company deployed an application on an Amazon EC2 instance. The application ran as expected for 6 months. In the past week, users have reported latency issues. A system administrator found that the CPU utilization was at 100% during business hours. The company wants a scalable solution to meet demand.

Which AWS service or feature should the company use to handle the load for its application during periods of high demand?

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

**Answer: A**

#### Explanation:

Auto Scaling groups are a feature that allows users to automatically scale the number of Amazon EC2 instances up or down based on demand or a predefined schedule. Auto Scaling groups can help improve the performance and availability of applications by adjusting the capacity in response to traffic fluctuations<sup>1</sup>. AWS Global Accelerator is a service that improves the availability and performance of applications by routing traffic through AWS edge locations<sup>2</sup>. Amazon Route 53 is a service that provides scalable and reliable domain name system (DNS) service<sup>3</sup>. An Elastic IP address is a static IPv4 address that can be associated with an Amazon EC2 instance<sup>4</sup>.

### NEW QUESTION 197

- (Topic 3)

A company wants a list of all users in its AWS account, the status of all of the users' access keys, and if multi-factor authentication (MFA) has been configured. Which AWS service or feature will meet these requirements?

- A. AWS Key Management Service (AWS KMS)
- B. IAM Access Analyzer
- C. IAM credential report
- D. Amazon CloudWatch

**Answer: C**

#### Explanation:

IAM credential report is a feature that allows you to generate and download a report that lists all IAM users in your AWS account and the status of their various credentials, including access keys and MFA devices. You can use this report to audit the security status of your IAM users and ensure that they follow the best practices for using AWS<sup>1</sup>.

AWS Key Management Service (AWS KMS) is a service that allows you to create and manage encryption keys to protect your data. It does not provide information about IAM users or their credentials<sup>2</sup>.

IAM Access Analyzer is a feature that helps you identify the resources in your AWS account, such as S3 buckets or IAM roles, that are shared with an external entity. It does not provide information about IAM users or their credentials<sup>3</sup>.

Amazon CloudWatch is a service that monitors and collects metrics, logs, and events from your AWS resources and applications. It does not provide information about IAM users or their credentials<sup>4</sup>.

References:

- ? Getting credential reports for your AWS account - AWS Identity and Access Management
- ? AWS Key Management Service - Amazon Web Services
- ? IAM Access Analyzer - AWS Identity and Access Management
- ? Amazon CloudWatch - Amazon Web Services

### NEW QUESTION 200

- (Topic 3)

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

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

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

**Answer: B**

#### Explanation:

AWS Business Support is the most cost-effective AWS Support plan that provides chat access to a cloud support engineer 24/7. AWS Business Support also offers phone and email support, as well as a response time of less than one hour for urgent issues. AWS Business Support does not include access to infrastructure event management, which is a feature of AWS Enterprise Support. AWS Enterprise Support is more expensive and provides additional benefits, such as a technical account manager, a support concierge, and a response time of less than 15 minutes for critical issues. AWS Developer Support and AWS Basic Support do not provide chat access to a cloud support engineer. AWS Developer Support provides email support and a response time of less than 12 hours for general guidance issues. AWS Basic Support provides customer service and account support, as well as access to forums and documentation<sup>1</sup>

### NEW QUESTION 203

- (Topic 3)

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

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

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

**Answer: D**

#### Explanation:

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

### NEW QUESTION 207

- (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 210**

- (Topic 3)

A company wants to use guidelines from the AWS Well-Architected Framework to limit human error and facilitate consistent responses to events. Which of the following is a Well-Architected design principle that will meet these requirements?

- A. Use AWS CodeDeploy.
- B. Perform operations as code.
- C. Migrate workloads to a Dedicated Host.
- D. Use AWS Compute Optimizer.

**Answer:** B

**Explanation:**

This is a design principle of the operational excellence pillar of the AWS Well-Architected Framework. Performing operations as code means using scripts, templates, or automation tools to perform routine tasks, such as provisioning, configuration, deployment, and monitoring. This reduces human error, increases consistency, and enables faster recovery from failures. You can learn more about the operational excellence pillar from this whitepaper or this digital course.

**NEW QUESTION 212**

- (Topic 3)

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

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

**Answer:** D

**Explanation:**

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

**NEW QUESTION 216**

- (Topic 3)

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

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

**Answer:** A

**Explanation:**

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

A multi-region deployment model is a cloud computing architecture that distributes an application and its data across multiple geographic regions. A multi-region deployment model enables a company to achieve global reach, high availability, disaster recovery, and performance optimization. By deploying an application in multiple regions, a company can serve customers from the nearest region, reduce latency, increase redundancy, and comply with data sovereignty regulations12.

A single-region deployment model is a cloud computing architecture that runs an application and its data within a single geographic region. A single-region deployment model is simpler and cheaper than a multi-region deployment model, but it has limited scalability, availability, and performance. A single-region deployment model may not be suitable for a company that wants to deploy an application globally, as it may face challenges such as network latency, regional outages, or regulatory compliance12.

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

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

References:

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

**NEW QUESTION 221**

- (Topic 3)

At what support level do users receive access to a support concierge?

- A. Basic Support
- B. Developer Support
- C. Business Support

D. Enterprise Support

**Answer:** D

**Explanation:**

Users receive access to a support concierge at the Enterprise Support level. A support concierge is a team of AWS billing and account experts that specialize in working with enterprise accounts. They can help users with billing and account inquiries, cost optimization, FinOps support, cost analysis, and prioritized answers to billing questions. The support concierge is included as part of the Enterprise Support plan, which also provides access to a Technical Account Manager (TAM), Infrastructure Event Management, AWS Trusted Advisor, and 24/7 technical support. References: AWS Support Plan Comparison, AWS Enterprise Support Plan, AWS Support Concierge

**NEW QUESTION 224**

- (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 229**

- (Topic 3)

A company wants a time-series database service that makes it easier to store and analyze trillions of events each day. Which AWS service will meet this requirement?

- A. Amazon Neptune
- B. Amazon Timestream
- C. Amazon Forecast
- D. Amazon DocumentDB (with MongoDB compatibility)

**Answer:** B

**Explanation:**

Amazon Timestream is a fast, scalable, and serverless time-series database service for IoT and other operational applications that makes it easy to store and analyze trillions of events per day up to 1,000 times faster and at as little as 1/10th the cost of relational databases<sup>1</sup>. Amazon Timestream saves you time and cost in managing the lifecycle of time series data, and its purpose-built query engine lets you access and analyze recent and historical data together with a single query<sup>1</sup>. Amazon Timestream has built-in time series analytics functions, helping you identify trends and patterns in near real time<sup>1</sup>. The other options are not suitable for storing and analyzing trillions of events per day. Amazon Neptune is a graph database service that supports highly connected data sets. Amazon Forecast is a machine learning service that generates accurate forecasts based on historical data. Amazon DocumentDB (with MongoDB compatibility) is a document database service that supports MongoDB workloads.

References:

? 1: Time Series Database – Amazon Timestream – Amazon Web Services

**NEW QUESTION 230**

- (Topic 3)

An ecommerce company wants to use Amazon EC2 Auto Scaling to add and remove EC2 instances based on CPU utilization. Which AWS service or feature can initiate an Amazon EC2 Auto Scaling action to achieve this goal?

- A. Amazon Simple Queue Service (Amazon SQS)
- B. Amazon Simple Notification Service (Amazon SNS)
- C. AWS Systems Manager
- D. Amazon CloudWatch alarm

**Answer:** D

**Explanation:**

Amazon CloudWatch alarm is an AWS service or feature that can initiate an Amazon EC2 Auto Scaling action based on CPU utilization. Amazon CloudWatch is a monitoring and observability service that collects and tracks metrics, logs, events, and alarms for your AWS resources and applications. Amazon CloudWatch alarms are actions that you can configure to send notifications or automatically make changes to the resources you are monitoring based on rules that you define<sup>67</sup>.

Amazon EC2 Auto Scaling is a service that helps you maintain application availability and allows you to automatically add or remove EC2 instances according to definable conditions. You can create dynamic scaling policies that track a specific CloudWatch metric, such as CPU utilization, and define what action to take when the associated CloudWatch alarm is in ALARM. When the policy is in effect, Amazon EC2 Auto Scaling adjusts the group's desired capacity up or down when the threshold of an alarm is

breached<sup>89</sup>. References: 6: Cloud Monitoring - Amazon CloudWatch - AWS, 7: Amazon

CloudWatch Documentation, 8: Dynamic scaling for Amazon EC2 Auto Scaling, 9: Amazon EC2 Auto Scaling Documentation

**NEW QUESTION 235**

- (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:**

AWS Enterprise Support provides customers with a designated technical account manager (TAM) who is a single point of contact for all technical and operational issues. The TAM provides consultative architectural and operational guidance delivered in the context of the customer's applications and use-cases to help them achieve the greatest value from AWS. The TAM also helps customers with proactive services, such as strategic business reviews, security improvement programs, guided Well-Architected reviews, cost optimization workshops, and more<sup>1</sup>.

A full set of AWS Trusted Advisor checks is not an additional benefit of AWS Enterprise Support, as it is also included in the AWS Business Support plan<sup>2</sup>. AWS Trusted Advisor is a tool that provides best practice recommendations for cost optimization, performance, security, fault tolerance, and service limits.

Phone, email, and chat access to cloud support engineers 24 hours a day, 7 days a week is not an additional benefit of AWS Enterprise Support, as it is also included in the AWS Business Support plan<sup>2</sup>. Cloud support engineers can help customers with technical issues, such as troubleshooting, configuration, usage, and service features.

A consultative review and architecture guidance for the company's applications is not an additional benefit of AWS Enterprise Support, as it is also included in the AWS Business Support plan<sup>2</sup>. Customers can request a consultative review from a solutions architect who will provide best practices and recommendations based on the customer's use-cases and goals.

**NEW QUESTION 237**

- (Topic 3)

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

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

**Answer: C**

**Explanation:**

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

AWS, AWS Cloud Development Kit Documentation, AWS Cloud Development Kit - Wikipedia, AWS CDK Intro Workshop | AWS CDK Workshop

**NEW QUESTION 239**

- (Topic 3)

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

- A. Redundancy
- B. Operational excellence
- C. Availability
- D. Multi-Region

**Answer: B**

**Explanation:**

The AWS Well-Architected Framework helps cloud architects build secure, high-performing, resilient, and efficient infrastructure for their applications and workloads. Based on five pillars — operational excellence, security, reliability, performance efficiency, and cost optimization — the Framework provides a consistent approach for customers and partners to evaluate architectures, and implement designs that can scale over time. Operational excellence is one of the pillars of the Framework, and it focuses on running and monitoring systems to deliver business value, and continually improving processes and procedures.

**NEW QUESTION 244**

- (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 249

- (Topic 3)

A company is building an application on AWS. The application needs to comply with credit card regulatory requirements. The company needs proof that the AWS services and deployment are in compliance.

Which actions should the company take to meet these requirements? (Select TWO.)

- A. Use Amazon Inspector to submit the application for certification.
- B. Ensure that the application's underlying hardware components comply with requirements.
- C. Use AWS Artifact to access AWS documents about the compliance of the services.
- D. Get the compliance of the application certified by a company assessor.
- E. Use AWS Security Hub to certify the compliance of the application.

**Answer:** CD

#### Explanation:

Using AWS Artifact to access AWS documents about the compliance of the services, and getting the compliance of the application certified by a company assessor are actions that the company should take to meet the requirements of complying with credit card regulatory requirements. AWS Artifact is a service that provides on-demand access to AWS security and compliance reports and select online agreements. Reports available in AWS Artifact include our Service Organization Control (SOC) reports, Payment Card Industry (PCI) reports, and certifications from accreditation bodies across geographies and compliance verticals that validate the implementation and operating effectiveness of AWS security controls. AWS Artifact can help you demonstrate compliance with credit card regulatory requirements by providing you with proof that the AWS services and deployment are in compliance. Getting the compliance of the application certified by a company assessor is an action that the company should take to ensure that the application meets the specific requirements of the credit card industry. A company assessor is an independent third-party entity that is qualified to assess the compliance of the application with the relevant standards and regulations. Using Amazon Inspector to submit the application for certification is not an action that the company should take, because 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, but it does not provide certification for the applications. Ensuring that the application's underlying hardware components comply with requirements is not an action that the company should take, because the application is deployed on AWS, and AWS is responsible for the security and compliance of the underlying hardware components. This is part of the shared responsibility model, where AWS is responsible for security of the cloud, and customers are responsible for security in the cloud. Using AWS Security Hub to certify the compliance of the application is not an action that the company should take, because AWS Security Hub is a service that gives you a comprehensive view of your security posture across your AWS accounts and helps you check your environment against security industry standards and best practices, but it does not provide certification for the applications.

#### NEW QUESTION 253

- (Topic 3)

Which AWS services are connectivity services for a VPC? (Select TWO.)

- A. AWS Site-to-Site VPN
- B. AWS Direct Connect
- C. Amazon Connect
- D. AWS Key Management Service (AWS KMS)
- E. AWS Identity and Access Management (IAM)

**Answer:** A

#### Explanation:

AWS Site-to-Site VPN and AWS Direct Connect are AWS services that are connectivity services for a VPC. AWS Site-to-Site VPN is a service that enables you to securely connect your on-premises network or branch office site to your Amazon Virtual Private Cloud (Amazon VPC). You can establish VPN connections over the internet or over AWS Direct Connect<sup>1</sup>. AWS Direct Connect is a service that lets you establish a dedicated network connection between your network and one of the AWS Direct Connect locations. Using AWS Direct Connect, you can create a private connection between AWS and your datacenter, office, or colocation environment, which can reduce your network costs, increase bandwidth throughput, and provide a more consistent network experience than internet-based connections<sup>2</sup>. Amazon Connect is a service that lets you set up and manage a contact center in the cloud, but it does not provide network connectivity between the VPC and your on-premises network. AWS Key Management Service (AWS KMS) is a service that makes it easy for you to create and manage cryptographic keys and control their use across a wide range of AWS services and in your applications, but it does not provide network connectivity between the VPC and your on-premises network. AWS Identity and Access Management (IAM) is a service that enables you to manage access to AWS services and resources securely, but it does not provide network connectivity between the VPC and your on-premises network.

#### NEW QUESTION 257

- (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 261

- (Topic 3)

A company is building an application in the AWS Cloud. The company wants to use temporary credentials for the application to access other AWS resources. Which AWS service will meet these requirements?

- A. AWS Key Management Service (Aws KMS)
- B. AWS CloudHSM
- C. Amazon Cognito
- D. AWS Security Token Service (Aws STS)

**Answer:** D

**Explanation:**

AWS Security Token Service (AWS STS) is a service that provides temporary security credentials to users or applications that need to access AWS resources. The temporary credentials have a limited lifetime and can be configured to last from a few minutes to several hours. The credentials are not stored with the user or application, but are generated dynamically and provided on request. The credentials work almost identically to long-term access key credentials, but have the advantage of not requiring distribution, rotation, or revocation<sup>1</sup>.

AWS Key Management Service (AWS KMS) is a service that provides encryption and decryption services for data and keys. It does not provide temporary security credentials<sup>2</sup>. AWS CloudHSM is a service that provides hardware security modules (HSMs) for cryptographic operations and key management. It does not provide temporary security credentials<sup>3</sup>.

Amazon Cognito is a service that provides user authentication and authorization for web and mobile applications. It can also provide temporary security credentials for authenticated users, but not for applications<sup>4</sup>.

**NEW QUESTION 266**

- (Topic 3)

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

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

**Answer:** A

**Explanation:**

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

**NEW QUESTION 268**

- (Topic 3)

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

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

**Answer:** B

**Explanation:**

Amazon EC2 instance store provides temporary block-level storage for your EC2 instance. This storage is located on disks that are physically attached to the host computer. Instance store is ideal for temporary storage of information that changes frequently, such as buffers, caches, scratch data, and other temporary content. It can also be used to store temporary data that you replicate across a fleet of instances, such as a load-balanced pool of web servers. An instance store consists of one or more instance store volumes exposed as block devices. The size of an instance store as well as the number of devices available varies by instance type and instance size. The virtual devices for instance store volumes are ephemeral<sup>[0-23]</sup>. Instance types that support one instance store volume have ephemeral<sup>0</sup>. Instance types that support two or more instance store volumes have ephemeral<sup>0</sup>, ephemeral<sup>1</sup>, and so on. Instance store pricing Instance store volumes are included as part of the instance's usage cost. The data on an instance store volume persists even if the instance is rebooted. However, the data does not persist if the instance is stopped, hibernated, or terminated. When the instance is stopped, hibernated, or terminated, every block of the instance store volume is cryptographically erased. Therefore, do not rely on instance store volumes for valuable, long-term data. If you need to retain the data stored on an instance store volume beyond the lifetime of the instance, you need to manually copy that data to more persistent storage, such as an Amazon EBS volume, an Amazon S3 bucket, or an Amazon EFS file system. There are some events that can result in your data not persisting throughout the lifetime of the instance. The following table indicates whether data on instance store volumes is persisted during specific events, for both virtualized and bare metal instances<sup>1</sup>. References: Amazon EC2 instance store - Amazon Elastic Compute Cloud

**NEW QUESTION 273**

- (Topic 3)

Which of the following is a benefit of operating in the AWS Cloud?

- A. The ability to migrate on-premises network devices to the AWS Cloud
- B. The ability to expand compute, storage, and memory when needed
- C. The ability to host custom hardware in the AWS Cloud
- D. The ability to customize the underlying hypervisor layer for Amazon EC2

**Answer:** B

**Explanation:**

One of the benefits of operating in the AWS Cloud is the ability to expand compute, storage, and memory when needed, which enables users to scale their

applications and resources up or down based on demand. This also helps users optimize their costs and performance. The ability to migrate on-premises network devices to the AWS Cloud, the ability to host custom hardware in the AWS Cloud, and the ability to customize the underlying hypervisor layer for Amazon EC2 are not benefits of operating in the AWS Cloud, as they are either not possible or not recommended by AWS .

#### NEW QUESTION 275

- (Topic 3)

A company wants to create a globally accessible ecommerce platform for its customers. The company wants to use a highly available and scalable DNS web service to connect users to the platform.

Which AWS service will meet these requirements?

- A. Amazon EC2
- B. Amazon VPC
- C. Amazon Route 53
- D. Amazon RDS

**Answer: C**

#### Explanation:

Amazon Route 53 is a highly available and scalable Domain Name System (DNS) web service that can route internet traffic to the company's ecommerce platform<sup>1</sup>. Route 53 can also register domain names, check the health of resources, and provide global DNS features<sup>2</sup>. Route 53 can connect users to the platform by translating human-readable names like `www.example.com` into the numeric IP addresses that computers use to communicate with each other<sup>2</sup>.

References: 1: Amazon Route 53 | DNS Service | AWS; 2: What is Amazon Route 53? - Amazon Route 53

#### NEW QUESTION 279

- (Topic 3)

A company has deployed an Amazon EC2 instance.

Which option is an AWS responsibility under the AWS shared responsibility model?

- A. Managing and encrypting application data
- B. Installing updates and security patches of guest operating system
- C. Configuration of infrastructure devices
- D. Configuration of security groups on each instance

**Answer: C**

#### Explanation:

According to the AWS shared responsibility model, AWS is responsible for protecting the infrastructure that runs all of the services offered in the AWS Cloud, such as data centers, hardware, software, networking, and facilities<sup>1</sup>. This includes the configuration of infrastructure devices, such as routers, switches, firewalls, and load balancers<sup>2</sup>. Customers are responsible for managing their data, applications, operating systems, security groups, and other aspects of their AWS environment<sup>1</sup>. Therefore, options A, B, and D are customer responsibilities, not AWS responsibilities. References: 1: AWS Well-Architected Framework - Elasticity; 2: Reactive Systems on AWS - Elastic

#### NEW QUESTION 284

.....

## **Thank You for Trying Our Product**

### **We offer two products:**

1st - We have Practice Tests Software with Actual Exam Questions

2nd - Questions and Answers in PDF Format

### **CLF-C02 Practice Exam Features:**

- \* CLF-C02 Questions and Answers Updated Frequently
- \* CLF-C02 Practice Questions Verified by Expert Senior Certified Staff
- \* CLF-C02 Most Realistic Questions that Guarantee you a Pass on Your First Try
- \* CLF-C02 Practice Test Questions in Multiple Choice Formats and Updates for 1 Year

**100% Actual & Verified — Instant Download, Please Click**  
**[Order The CLF-C02 Practice Test Here](#)**