

complete your programming course

about resources, doubts and more!

MYEXAM.FK

Amazon

(AWS Certified Cloud Practitioner CLF-C02)

AWS Certified Cloud Practitioner CLF-C02

Total: **718 Questions**
Link:

Question: 1

A company plans to use an Amazon Snowball Edge device to transfer files to the AWS Cloud. Which activities related to a Snowball Edge device are available to the company at no cost?

- A. Use of the Snowball Edge appliance for a 10-day period
- B. The transfer of data out of Amazon S3 and to the Snowball Edge appliance
- C. The transfer of data from the Snowball Edge appliance into Amazon S3
- D. Daily use of the Snowball Edge appliance after 10 days

Answer: C

Explanation:

The correct answer is **C: The transfer of data from the Snowball Edge appliance into Amazon S3**. This is because AWS charges for the Snowball Edge appliance usage and for data transferred out of AWS (egress). Transferring data into Amazon S3 from the Snowball Edge is generally free. Let's break down why the other options are incorrect and solidify why option C is the definitive answer.

Option A is incorrect because AWS charges for the use of the Snowball Edge appliance, irrespective of whether it is used for the standard 10-day period or an extended duration. The cost is typically based on a daily rate, and a free period is not usually offered for this service.

Option B is incorrect because transferring data out of Amazon S3 and to the Snowball Edge appliance incurs data transfer charges. AWS charges for data egress, which refers to data leaving the AWS network. Moving data from S3 to a Snowball Edge involves data egress.

Option D is incorrect because AWS continues to charge a daily rate for using the Snowball Edge appliance even after the initial 10-day period. There's no "free" daily usage after the initial period.

Therefore, only the transfer of data into Amazon S3 from the Snowball Edge is offered at no cost. This is a key benefit of using Snowball Edge for data migration – you are only billed for the appliance rental and, importantly, egress bandwidth when data is transferred out of AWS, not into it. AWS promotes data import to encourage customers to store data within the AWS ecosystem.

For further research, refer to the official AWS Snowball documentation and pricing pages:

AWS Snowball Documentation: <https://aws.amazon.com/snowball/>
AWS Snowball Pricing: <https://aws.amazon.com/snowball/pricing/>

These resources clearly outline the costs associated with using Snowball Edge, including data transfer and appliance usage fees, which support the explanation provided. They specifically highlight the "no cost" nature of importing data into S3 from the appliance.

Question: 2

A company has deployed applications on Amazon EC2 instances. The company needs to assess application vulnerabilities and must identify infrastructure deployments that do not meet best practices.

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

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

Answer: B

Explanation:

The correct answer is B, Amazon Inspector. Amazon Inspector is a vulnerability management service that automatically assesses applications for vulnerabilities and deviations from best practices. It inspects EC2 instances and container images for software vulnerabilities and unintended network accessibility, helping improve the security posture of deployed applications. It provides findings that highlight potential security issues and recommends remediation steps.

AWS Trusted Advisor (Option A) provides recommendations across cost optimization, performance, security, fault tolerance, and service limits, but it does not specifically focus on application vulnerability assessments.

AWS Config (Option C) enables you to assess, audit, and evaluate the configurations of your AWS resources. It tracks resource configurations over time, but it doesn't perform vulnerability scanning or application-level security assessments.

Amazon GuardDuty (Option D) is a threat detection service that continuously monitors for malicious activity and unauthorized behavior to protect your AWS accounts and workloads. It focuses on identifying threats based on log analysis and network activity but doesn't directly assess application vulnerabilities.

Therefore, Amazon Inspector is the most suitable service for the company's stated requirement of assessing application vulnerabilities and identifying infrastructure deployments that do not meet best practices from a security perspective.

References:

Amazon Inspector: <https://aws.amazon.com/inspector/>

Question: 3

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

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

A. Create an Amazon S3 bucket for each user. Mount each bucket by using an S3 file system mounting utility.

B. Configure and deploy an AWS Storage Gateway file gateway. Connect each user's workstation to the file gateway.

C. Move each user's working environment to Amazon WorkSpaces. Set up an Amazon WorkDocs account for each user.

D. Deploy an Amazon EC2 instance and attach an Amazon Elastic Block Store (Amazon EBS) Provisioned IOPS volume. Share the EBS volume directly with the users.

Answer: B

Explanation:

The correct answer is B: Configure and deploy an AWS Storage Gateway file gateway. Connect each user's workstation to the file gateway.

Here's why this is the most operationally efficient solution:

AWS Storage Gateway, specifically the File Gateway type, is designed to solve the problem of extending on-premises storage to the cloud while maintaining local access performance. A File Gateway caches frequently accessed data locally, minimizing latency for users. This addresses the requirement of retaining performance benefits while offloading less frequently used files to Amazon S3 in the cloud.

Option A, using S3 buckets per user and mounting them individually, is operationally complex. Managing numerous S3 buckets, access policies, and file system mounts introduces significant overhead. S3 file system mounting utilities can also introduce performance limitations compared to a dedicated gateway appliance.

Option C, migrating users to Amazon WorkSpaces and using Amazon WorkDocs, is a more drastic solution that likely involves significant cost and effort associated with migrating user environments. This approach addresses the storage problem but also requires a change in the way users interact with their files and potentially requires retraining. It is much more than a storage extension.

Option D, deploying an EC2 instance with an EBS volume and sharing it directly, introduces complexity in managing the EC2 instance, the EBS volume, and the file sharing mechanism. It requires manually configuring file sharing permissions and dealing with potential performance bottlenecks. The user must handle all data management tasks and will pay for unused data.

In contrast, the File Gateway solution offers a managed service that simplifies the connection between on-premises users and cloud storage, reducing the operational burden. The file gateway handles the caching and transfer of data to S3, allowing the company to scale its storage capacity without major changes to user workflows or IT infrastructure.

Relevant Links:

AWS Storage Gateway: <https://aws.amazon.com/storagegateway/>

File Gateway documentation:

<https://docs.aws.amazon.com/storagegateway/latest/userguide/WhatIsStorageGateway.html>

Question: 4

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

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

Answer: C

Explanation:

The correct answer is C: Have the EC2 instance assume a role to obtain the privileges to upload the file. Here's why:

IAM roles provide a secure and manageable way to grant permissions to AWS services and applications. When an EC2 instance is assigned an IAM role, it can temporarily assume the role's permissions without needing to store long-term credentials directly on the instance. This is achieved through temporary security credentials provided by the AWS Security Token Service (STS).

Option A is extremely insecure. Hardcoding credentials directly into applications is a major security risk. If the application or code repository is compromised, the credentials are exposed, granting unauthorized access to the S3 bucket.

Option B is marginally better than option A, but still introduces significant security vulnerabilities. Storing credentials in a text file on the EC2 instance makes them susceptible to unauthorized access if the instance is compromised. Anyone gaining access to the instance could potentially read the file and obtain the

credentials.

Option D violates the principle of least privilege and introduces a severe security risk. Modifying the S3 bucket policy to allow any service to upload at any time effectively removes all access control, potentially leading to data breaches and unauthorized modifications.

Using IAM roles aligns with security best practices by avoiding long-term credentials stored directly on the EC2 instance. The EC2 instance assumes a role only when it needs to access the S3 bucket, and the temporary credentials provided by STS expire after a set period, minimizing the risk of unauthorized access. Roles also centralize permission management through IAM, making it easier to audit and control access to AWS resources.

Further Reading:

IAM Roles for EC2:https://docs.aws.amazon.com/IAM/latest/UserGuide/id_roles_use_ec2.html

Security Best Practices in IAM:<https://docs.aws.amazon.com/IAM/latest/UserGuide/best-practices.html> **AWS STS (Security Token Service):**<https://docs.aws.amazon.com/STS/latest/APIReference/welcome.html>

Question: 5

Which option is a customer responsibility when using Amazon DynamoDB under the AWS Shared Responsibility Model?

- A. Physical security of DynamoDB
- B. Patching of DynamoDB
- C. Access to DynamoDB tables
- D. Encryption of data at rest in DynamoDB

Answer: C

Explanation:

The correct answer is C: Access to DynamoDB tables. Let's examine why, using the AWS Shared Responsibility Model as a guide.

The AWS Shared Responsibility Model dictates that AWS is responsible for the security of the cloud, while the customer is responsible for security in the cloud. AWS manages the underlying infrastructure that supports DynamoDB, including hardware, software, networking, and facilities. This encompasses the physical security of the data centers where DynamoDB runs (eliminating option A) and patching the DynamoDB service itself (eliminating option B). The data-at-rest encryption is also managed by Amazon, which handles the underlying hardware and software used to protect the data.

However, customers are responsible for managing access to their data. The customer defines who can access their DynamoDB tables, what actions they can perform (read, write, update, delete), and from where they can access the service. This is achieved by using AWS Identity and Access Management (IAM) to control who can authenticate and authorize with the DynamoDB resources. Customers must configure IAM policies to grant appropriate permissions to users, groups, and roles accessing their DynamoDB tables. This includes implementing principles of least privilege, regularly reviewing and updating access policies, and monitoring access logs to detect and respond to unauthorized activity. Therefore, determining who has access to the DynamoDB tables is the responsibility of the customer according to the AWS Shared Responsibility model.

For more information, refer to the AWS Shared Responsibility Model documentation:

<https://aws.amazon.com/compliance/shared-responsibility-model/> and the DynamoDB security best practices: <https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/security.html>.

Question: 6

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

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

Answer: C

Explanation:

The correct answer is C, Governance. The AWS Cloud Adoption Framework (AWS CAF) is designed to help organizations develop and execute efficient and effective plans for cloud adoption. It provides a structured approach to cloud adoption by grouping capabilities into perspectives. These perspectives help different stakeholders within an organization understand and manage the changes required to adopt the cloud successfully.

Governance is one of the six perspectives of the AWS CAF. This perspective focuses on skills and processes to manage and measure cloud investments, evaluating business risk, and complying with internal and external regulations. It addresses key organizational considerations such as compliance, risk management, data integrity, and security controls. It includes elements like policies, roles, responsibilities, and procedures related to cloud governance.

While options like Sustainability, Performance Efficiency, and Reliability are important aspects of cloud computing, they are not foundational perspectives within the AWS CAF itself. Performance Efficiency is a pillar of the AWS Well-Architected Framework, not the CAF. Reliability is an aspect considered across different perspectives but isn't a standalone perspective within the AWS CAF. Sustainability is gaining prominence but isn't formally integrated as a core pillar within the established CAF structure (though sustainable practices can certainly be incorporated into cloud strategies influenced by the CAF). Therefore, Governance is the most accurate answer as it represents a foundational capability and perspective that aligns with the core principles and framework of the AWS CAF.

For more information on the AWS Cloud Adoption Framework, refer to the following resources:

AWS CAF Whitepaper:<https://docs.aws.amazon.com/whitepapers/latest/aws-cloud-adoption-framework/aws-cloud-adoption-framework.pdf>

AWS Cloud Adoption Framework (AWS CAF) Overview:<https://aws.amazon.com/professional-services/CAF/>

Question: 7

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

Which AWS service meets these requirements?

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

Answer: C

Explanation:

The question asks for an AWS service that can manage cluster size, scheduling, and environment maintenance for a Docker environment currently running on EC2.

AWS Fargate is a serverless compute engine for containers that works with both Amazon ECS (Elastic Container Service) and Amazon EKS (Elastic Kubernetes Service). It allows you to run containers without managing servers or clusters. This directly addresses the company's need to avoid managing cluster size and environment maintenance, as Fargate automatically scales and manages the underlying infrastructure. Fargate also handles scheduling containers across its infrastructure.

AWS Lambda is a serverless compute service that runs code without provisioning or managing servers. While it's serverless, it's designed for event-driven functions, not long-running containerized applications. Therefore, it is not a suitable replacement for managing a Docker environment.

Amazon RDS (Relational Database Service) is a managed database service that simplifies setting up, operating, and scaling relational databases in the cloud. It does not manage containers or Docker environments.

Amazon Athena is an interactive query service that makes it easy to analyze data in Amazon S3 using standard SQL. It has no role in container management.

Therefore, AWS Fargate (Option C) is the correct answer because it directly addresses the requirements of managing cluster size, scheduling, and environment maintenance for containerized applications, removing the need to manage EC2 instances directly.

<https://aws.amazon.com/fargate/>

Question: 8

A company wants to run a NoSQL database on Amazon EC2 instances. Which task is the responsibility of AWS in this scenario?

- A. Update the guest operating system of the EC2 instances.
- B. Maintain high availability at the database layer.
- C. Patch the physical infrastructure that hosts the EC2 instances.
- D. Configure the security group firewall.

Answer: C

Explanation:

The correct answer is **C. Patch the physical infrastructure that hosts the EC2 instances.**

Here's why:

In the Shared Responsibility Model of AWS, responsibilities are divided between AWS and the customer. AWS is responsible for the security of the cloud, while the customer is responsible for security in the cloud. When a company runs a NoSQL database on EC2, it's utilizing Infrastructure as a Service (IaaS).

AWS manages the physical infrastructure that underlies the cloud services. This includes patching and maintaining the hardware, networking, and facilities that host EC2 instances. This responsibility ensures the core infrastructure remains secure and reliable.

Let's analyze the other options:

A. Update the guest operating system of the EC2 instances: Updating the guest OS is the customer's responsibility. The customer has full control over the OS running on their EC2 instances and must manage its patching and updates.

B. Maintain high availability at the database layer: While AWS provides services and features to achieve high availability, configuring and maintaining high availability for the database itself (e.g., database replication, failover mechanisms) falls under the customer's responsibility, especially when running the database on EC2. The customer is responsible for configuring their NoSQL database to be highly available.

D. Configure the security group firewall: Security groups are virtual firewalls associated with EC2 instances. Configuring these firewalls to control inbound and outbound traffic is the customer's responsibility.

Therefore, only patching the physical infrastructure, which includes the underlying hardware and facilities, is definitively and exclusively AWS's responsibility in this scenario.

Further Reading:

AWS Shared Responsibility Model: <https://aws.amazon.com/compliance/shared-responsibility-model/>
Amazon EC2: <https://aws.amazon.com/ec2/>

Question: 9

Which AWS services or tools can identify rightsizing opportunities for Amazon EC2 instances? (Choose two.)

- A. AWS Cost Explorer
- B. AWS Billing Conductor
- C. Amazon CodeGuru
- D. Amazon SageMaker
- E. AWS Compute Optimizer

Answer: AE

Explanation:

The correct answer is A. AWS Cost Explorer and E. AWS Compute Optimizer. Both services provide capabilities to analyze EC2 instance utilization and recommend right-sizing opportunities.

AWS Cost Explorer helps visualize, understand, and manage AWS costs and usage over time. It can identify underutilized EC2 instances by analyzing metrics like CPU utilization, network I/O, and disk I/O, allowing you to pinpoint instances that are larger than necessary for their workload. Cost Explorer's Savings Plans recommendations also factor in EC2 instance usage and can suggest switching to different instance types based on usage patterns. While Cost Explorer doesn't directly "right-size" instances, it provides the insights needed to make informed decisions about instance type adjustments. <https://aws.amazon.com/aws-cost-management/aws-cost-explorer/>

AWS Compute Optimizer is specifically designed to analyze the configuration and utilization metrics of AWS resources, including EC2 instances. It leverages machine learning to recommend optimal AWS resources for your workloads, reducing costs and improving performance. For EC2 instances, Compute Optimizer analyzes metrics such as CPU, memory, and network utilization to suggest instances that are a better fit for the observed workload. It can recommend down-sizing, up-sizing, or switching to different instance families based on the workload's needs. These recommendations are based on historical performance data and projections of future performance, providing valuable guidance for rightsizing. <https://aws.amazon.com/compute-optimizer/>

Why the other options are incorrect:

B. AWS Billing Conductor: AWS Billing Conductor is used to customize AWS billing data to reflect your internal business structure and cost allocation needs. It's about cost allocation and reporting, not rightsizing analysis.

C. Amazon CodeGuru: Amazon CodeGuru is a developer tool that provides intelligent recommendations for improving code quality and identifying an application's most expensive lines of code. It doesn't analyze EC2 instance utilization for rightsizing purposes.

D. Amazon SageMaker: Amazon SageMaker is a fully managed machine learning service. While it consumes EC2 resources, it does not directly provide rightsizing recommendations for other EC2 instances.

Question: 10

Which of the following are benefits of using AWS Trusted Advisor? (Choose two.)

- A. Providing high-performance container orchestration
- B. Creating and rotating encryption keys
- C. Detecting underutilized resources to save costs
- D. Improving security by proactively monitoring the AWS environment
- E. Implementing enforced tagging across AWS resources

Answer: CD

Explanation:

The correct answer is CD because AWS Trusted Advisor focuses on cost optimization and security improvements. Option C, detecting underutilized resources, aligns with Trusted Advisor's ability to identify services like EC2 instances or EBS volumes that are not being fully utilized, leading to potential cost savings by rightsizing or terminating them. This directly contributes to AWS's well-architected framework pillar of cost optimization. Option D, improving security by proactively monitoring the AWS environment, reflects Trusted Advisor's capability to check for security vulnerabilities such as open security groups, exposed access keys, or outdated software versions, which can then be remediated to enhance the security posture of the AWS environment. This aligns with the security pillar of the AWS well-architected framework.

Option A is incorrect because container orchestration is primarily handled by services like Amazon ECS, EKS, and Fargate, not Trusted Advisor. Option B is incorrect because key creation and rotation are functions of services such as AWS Key Management Service (KMS) or AWS CloudHSM, not Trusted Advisor. Option E is incorrect because implementing enforced tagging is usually managed via AWS Tag Policies through AWS Organizations or custom scripting, not directly by Trusted Advisor. Trusted Advisor simply checks resource tagging but doesn't enforce tag policies. Therefore, C and D are the most relevant benefits Trusted Advisor provides.

Further research:

AWS Trusted Advisor: <https://aws.amazon.com/premiumsupport/technology/trusted-advisor/> AWS Well-Architected Framework: <https://aws.amazon.com/well-architected/>

Question: 11

Which of the following is an advantage that users experience when they move on-premises workloads to the AWS Cloud?

- A. Elimination of expenses for running and maintaining data centers
- B. Price discounts that are identical to discounts from hardware providers

C.Distribution of all operational controls to AWS

D.Elimination of operational expenses

Answer: A

Explanation:

The correct answer is A: Elimination of expenses for running and maintaining data centers.

Moving on-premises workloads to AWS allows users to significantly reduce or completely eliminate the costs associated with managing their own physical data centers. These costs include expenses related to hardware procurement and maintenance, power consumption, cooling, physical security, and staffing for tasks like system administration and facility management. By leveraging AWS's infrastructure, users shift these responsibilities and associated expenses to AWS.

Option B is incorrect because price discounts from AWS are not necessarily identical to those offered by hardware providers. While AWS offers various pricing models like reserved instances and spot instances to reduce costs, these are specific to AWS services and not directly comparable to hardware vendor discounts.

Option C is incorrect because AWS does not take over all operational controls. While AWS manages the underlying infrastructure, users retain control over their applications, operating systems, and data residing on AWS. This concept is often referred to as the "shared responsibility model."

Option D, Elimination of operational expenses, is too broad. While many operational expenses are reduced or eliminated, some remain. For example, the user is still responsible for operating system patching, database administration, and application monitoring on virtual machines they create. Therefore, eliminating all operational expenses is an overstatement. The most significant cost advantage from the customer perspective is the removal of data center expenses.

For further research, refer to the AWS Shared Responsibility Model:

<https://aws.amazon.com/compliance/shared-responsibility-model/> and AWS Pricing:

<https://aws.amazon.com/pricing/>.

Question: 12

A company wants to manage deployed IT services and govern its infrastructure as code (IaC) templates. Which AWS service will meet this requirement?

A.AWS Resource Explorer

B.AWS Service Catalog

C.AWS Organizations

D.AWS Systems Manager

Answer: B

Explanation:

The correct answer is **B. AWS Service Catalog**. Here's why:

AWS Service Catalog is designed specifically for managing deployed IT services and governing infrastructure as code (IaC) templates within an organization. It allows organizations to create and manage catalogs of IT services that are approved for use. These services are defined as products, which can be infrastructure as code templates created using services like AWS CloudFormation.

Service Catalog provides a centralized repository for these approved products, ensuring consistency and

compliance across deployments. It enables administrators to control which services are available to end users, specify parameters and configurations, and track usage. Users can then easily self-provision approved services through a consistent and governed process, eliminating the need for manual configuration and reducing the risk of errors. This streamlines the provisioning process, making it faster and easier for users to access the IT services they need. It also provides centralized governance for the entire lifecycle of the service.

AWS Resource Explorer (A) helps discover and explore AWS resources across regions but doesn't directly govern IaC or manage service catalogs. AWS Organizations (C) manages AWS accounts centrally but does not handle the specific task of service catalog management or IaC governance. AWS Systems Manager (D) helps manage existing resources and automate operational tasks but it doesn't provide the service catalog functionality needed for organizing and governing the deployment of IT services.

Therefore, AWS Service Catalog directly addresses the requirements of managing deployed IT services and governing IaC templates, making it the most suitable choice.

Further reading:

AWS Service Catalog: <https://aws.amazon.com/servicecatalog/>

AWS Service Catalog Documentation: <https://docs.aws.amazon.com/servicecatalog/latest/adminguide/what-is.html>

Question: 13

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

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

Answer: C

Explanation:

The correct answer is C: AWS Cost Explorer. AWS Cost Explorer is specifically designed to help users visualize, understand, and manage their AWS spending and usage patterns over time. It provides interactive graphs and reports that enable users to analyze their cost and usage data at a granular level. Users can filter data by service, region, account, tag, and other dimensions to identify cost drivers and optimization opportunities.

AWS Cost Explorer offers features like cost forecasting, which predicts future spending based on historical trends. This allows users to proactively manage their budgets and avoid unexpected cost overruns. It also supports creating custom cost reports that can be tailored to specific needs, such as tracking the cost of a particular project or application. The service also allows users to set budgets and receive alerts when spending approaches or exceeds predefined thresholds.

The other options are incorrect. AWS Organizations is a service for managing multiple AWS accounts, not for cost visualization. AWS Pricing Calculator is a tool for estimating the cost of AWS services before deployment. AWS Service Catalog allows organizations to create and manage catalogs of IT services that are approved for use. These are not directly related to analyzing past spending or usage.

In summary, AWS Cost Explorer's features for visualizing, understanding, and managing spending and usage data directly address the question's requirements, making it the most appropriate answer. <https://aws.amazon.com/aws-cost-management/aws-cost-explorer/>

Question: 14

A company is using a central data platform to manage multiple types of data for its customers. The company wants to use AWS services to discover, transform, and visualize the data.

Which combination of AWS services should the company use to meet these requirements? (Choose two.)

- A. AWS Glue
- B. Amazon Elastic File System (Amazon EFS)
- C. Amazon Redshift
- D. Amazon QuickSight
- E. Amazon Quantum Ledger Database (Amazon QLDB)

Answer: AD

Explanation:

The requirement is to discover, transform, and visualize data. AWS Glue (A) is a serverless data integration service that makes it easy to discover, prepare, and combine data for analytics, machine learning, and application development. It provides capabilities to crawl data sources to infer schema, transform data using ETL (Extract, Transform, Load) jobs, and catalog the metadata. Thus, it directly addresses the data discovery and transformation requirements.

Amazon QuickSight (D) is a scalable, serverless, embeddable, machine learning-powered business intelligence (BI) service built for the cloud. It allows users to easily create and publish interactive BI dashboards, enabling visualization of the processed data from the data platform. QuickSight fulfills the visualization requirement.

Amazon EFS (B) is a fully managed, elastic NFS file system for use with AWS Cloud services and on-premises resources. While useful for storage, it doesn't directly contribute to the data discovery, transformation, or visualization needs. Amazon Redshift (C) is a data warehouse that can store large volumes of data and is excellent for data warehousing and complex querying but is not a primary tool for initial data discovery and transformation phases. It is more typically used after data has been transformed. Amazon QLDB (E) is a fully managed, serverless, transparent, and immutable ledger database. It's not suited for data discovery, transformation, or visualization; its primary purpose is maintaining an immutable transaction log.

Therefore, AWS Glue and Amazon QuickSight are the most suitable combination to address the company's needs.

Authoritative Links:

AWS Glue: <https://aws.amazon.com/glue/>

Amazon QuickSight: <https://aws.amazon.com/quicksight/>

Question: 15

A global company wants to migrate its third-party applications to the AWS Cloud. The company wants help from a global team of experts to complete the migration faster and more reliably in accordance with AWS internal best practices. Which AWS service or resource will meet these requirements?

- A. AWS Support
- B. AWS Professional Services
- C. AWS Launch Wizard

Answer: B

Explanation:

The correct answer is **B. AWS Professional Services**. Here's why:

AWS Professional Services is a global team of AWS experts that partners with customers to accelerate their cloud adoption journey. They offer services like migration support, application modernization, and operational excellence, all tailored to AWS best practices. Their expertise is crucial for companies looking to migrate complex third-party applications reliably and quickly. They can provide hands-on assistance, architectural guidance, and knowledge transfer to ensure a smooth transition.

Option A, AWS Support, provides technical assistance and troubleshooting for AWS services but doesn't typically offer proactive migration planning and execution.

Option C, AWS Launch Wizard, helps deploy specific workloads on AWS, but it's limited in scope and doesn't offer the comprehensive support needed for a full migration of multiple third-party applications.

Option D, AWS Managed Services (AMS), provides ongoing operational support for your AWS infrastructure, but it's more suitable for post-migration operations rather than the initial migration phase itself. While AMS can assist with post-migration management, Professional Services provides the experts and resources needed to ensure a successful and timely migration. Professional services helps to ensure the migration is performed by AWS experts.

Therefore, AWS Professional Services is the most suitable choice for a global company seeking expert assistance to migrate third-party applications to AWS quickly and reliably, adhering to AWS best practices.

Further research:

AWS Professional Services: <https://aws.amazon.com/professional-services/>

Question: 16

An e-learning platform needs to run an application for 2 months each year. The application will be deployed on Amazon EC2 instances. Any application downtime during those 2 months must be avoided.

Which EC2 purchasing option will meet these requirements MOST cost-effectively?

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

Answer: D

Explanation:

Here's a detailed justification for why On-Demand Instances are the most cost-effective option for the given scenario:

The key requirement is running an application for only two months a year with zero downtime. Let's analyze each option:

On-Demand Instances: These allow you to pay only for the compute time you use, by the hour or second. You start them when needed and stop them when done. For two months of use annually, this eliminates any cost

for the remaining ten months. This makes it immediately appealing when only requiring compute resources for short durations each year. <https://aws.amazon.com/ec2/pricing/on-demand/>

Reserved Instances: These offer significant cost savings (up to 75%) compared to On-Demand, but require a commitment of 1 or 3 years. Paying for a full year (or more) when only needing the instances for 2 months makes them a poor choice from a cost perspective. <https://aws.amazon.com/ec2/pricing/reserved-instances/>

Spot Instances: These offer the largest discounts, but come with the risk of interruption. AWS can terminate them with a two-minute warning if the spot price exceeds your bid. The requirement for no downtime eliminates Spot Instances as a viable option. <https://aws.amazon.com/ec2/spot/>

Dedicated Hosts: These are physical servers dedicated to your use. They are the most expensive option and are primarily used for compliance or licensing reasons where you need hardware isolation. They provide no cost benefit for this scenario and overcomplicate the deployment process.
<https://aws.amazon.com/ec2/dedicated-hosts/>

Therefore, On-Demand instances are the most cost-effective because you only pay for what you use during the specific 2-month period each year. It aligns with the short-term operational needs without long-term commitments or the risk of interruption.

Question: 17

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:

The correct answer is B, AWS Elastic Beanstalk. Here's why:

AWS Elastic Beanstalk is a Platform-as-a-Service (PaaS) offering that simplifies application deployment and management on AWS. It allows developers to upload their application code, and Elastic Beanstalk automatically handles the provisioning, deployment, load balancing, auto-scaling, and health monitoring of the application.

Option A, Amazon EC2, provides virtual servers in the cloud. While you can deploy an application on EC2, it requires manual configuration of the operating system, web server, application dependencies, and networking, which doesn't align with the requirement of quickly deploying an application without manual resource creation.

Option C, AWS CodeBuild, is a fully managed continuous integration service that compiles source code, runs tests, and produces software packages that are ready to deploy. It's part of a CI/CD pipeline but doesn't handle the deployment and infrastructure provisioning aspects required in the question.

Option D, Amazon Personalize, is a machine learning service that enables developers to create individualized recommendations for their customers. It is irrelevant to application deployment.

Elastic Beanstalk abstracts away the underlying infrastructure details, enabling developers to focus on writing code and deploying applications quickly. It supports various programming languages and platforms,

including Java, .NET, PHP, Node.js, Python, Ruby, Go, and Docker. Elastic Beanstalk handles tasks like creating EC2 instances, configuring load balancers, and setting up auto-scaling groups, reducing the operational burden on the developer. This capability addresses the requirement of deploying applications quickly without manual resource creation.

Further reading:

AWS Elastic Beanstalk: <https://aws.amazon.com/elasticbeanstalk/>

Question: 18

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 the most suitable feature to protect data in an S3 bucket from accidental deletion or overwriting. When enabled on an S3 bucket, versioning automatically keeps multiple versions of an object. If an object is accidentally deleted, the previous version can be easily restored. Similarly, if an object is overwritten, the original version is retained, preventing permanent data loss.

S3 Lifecycle rules automate object transitioning and deletion, but they are not designed for preventing accidental deletions; they are more for cost optimization based on data access patterns. S3 bucket policies control access to the bucket and its objects, defining who can perform what actions. While they are important for security, they don't prevent accidental actions by authorized users. S3 server-side encryption protects data at rest, preventing unauthorized access, but it doesn't address the problem of accidental deletion or overwriting.

Versioning is the only option that specifically provides a mechanism for recovery from unintended modifications or deletions, ensuring data durability and availability in the face of human error. It essentially creates a historical record of changes, allowing you to revert to any previous state. Each time an object is modified or deleted, a new version is created or a delete marker is added, respectively, without physically removing the older version.

Therefore, S3 Versioning directly addresses the requirement of protecting data from accidental deletion or overwriting by enabling the recovery of previous object versions.

Further Reading:

[Amazon S3 Versioning](#)
[Using versioning in S3 buckets](#)

Question: 19

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 correct answer is D. AWS CloudFormation enables Infrastructure as Code (IaC) on AWS. IaC is the practice of managing and provisioning infrastructure through code, rather than through manual processes. CloudFormation allows you to define your AWS resources in a template file (typically in YAML or JSON) and then automatically provision and configure those resources. This enables automation, version control, repeatability, and improved consistency in infrastructure deployments.

AWS CloudFormation uses these templates to create, update, and delete AWS resources in a safe and predictable manner. It provides a single source of truth for your infrastructure, making it easier to track changes and roll back deployments if necessary. Using IaC with CloudFormation reduces the risk of human error, speeds up deployment times, and helps enforce compliance.

The other options are incorrect because they do not directly provide IaC capabilities. AWS CodePipeline is a continuous integration and continuous delivery (CI/CD) service for automating your release pipelines. AWS CodeDeploy automates application deployments to various compute services. AWS Direct Connect establishes a dedicated network connection from your on-premises environment to AWS. While CodePipeline and CodeDeploy can be integrated with CloudFormation, they do not, by themselves, offer IaC functionality. CloudFormation is specifically designed to define and manage infrastructure through code.

Further resources:

AWS CloudFormation: <https://aws.amazon.com/cloudformation/>
Infrastructure as Code: <https://aws.amazon.com/devops/infrastructure-as-code/>

Question: 20

An online gaming company needs to choose a purchasing option to run its Amazon EC2 instances for 1 year. The web traffic is consistent, and any increases in traffic are predictable. The EC2 instances must be online and available without any disruption.

Which EC2 instance purchasing option will meet these requirements MOST cost-effectively?

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

Answer: B

Explanation:

The correct answer is B, Reserved Instances. Let's break down why.

The scenario describes consistent, predictable traffic and a need for continuous uptime. This immediately rules out Spot Instances (C and D), as they are subject to interruption if the spot price exceeds the bid price.

Spot Fleets also utilize Spot Instances and therefore inherit the same risk of interruption, making them unsuitable for applications requiring constant availability.

On-Demand Instances (A) offer flexibility and no commitment, but they are generally the most expensive option for long-term, predictable workloads.

Reserved Instances (B), on the other hand, offer a significant discount (up to 72% compared to On-Demand) in exchange for a 1-year or 3-year commitment. Since the workload is predictable and needs to be online for a year, the gaming company can leverage Reserved Instances to achieve substantial cost savings without compromising availability. They provide capacity reservation, which guarantees that the instances will be available when needed. Because the traffic is predictable, the company can accurately determine the required instance capacity and purchase the appropriate number of Reserved Instances. This provides the optimal balance of cost-effectiveness and reliability.

Therefore, Reserved Instances meet the requirements of consistent availability and cost optimization for a predictable, year-long workload.

Authoritative Links:

Amazon EC2 Instance Purchasing Options:<https://aws.amazon.com/ec2/pricing/>
Reserved Instances:<https://aws.amazon.com/ec2/reserved-instances/>

Question: 21

Which AWS service or feature allows a user to establish a dedicated network connection between a company's on-premises data center and the AWS Cloud?

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

Answer: A

Explanation:

The correct answer is A, AWS Direct Connect. Let's break down why:

AWS Direct Connect provides a dedicated network connection between your on-premises data center or office and AWS. This dedicated connection offers several advantages over traditional internet-based connections, primarily in terms of network performance, security, and cost. Direct Connect bypasses the public internet, resulting in more consistent and predictable network performance, lower latency, and enhanced security. This is crucial for applications requiring low latency or handling sensitive data.

VPC Peering (option B) enables you to connect two VPCs (Virtual Private Clouds) within AWS, allowing resources in those VPCs to communicate with each other as if they were part of the same network. However, VPC Peering doesn't bridge the gap between on-premises infrastructure and the AWS Cloud. It's strictly for intra-AWS connectivity.

AWS VPN (option C) also allows you to connect your on-premises network to AWS, but it uses the public internet as the underlying transport. While VPN provides a secure connection, it's subject to the inherent variability and potential security risks associated with internet traffic. Direct Connect, on the other hand, offers a private, dedicated circuit.

Amazon Route 53 (option D) is a highly available and scalable Domain Name System (DNS) web service. It's used to translate domain names into IP addresses, effectively directing users to your applications and services. While essential for managing your online presence, it doesn't establish a dedicated network connection to AWS.

In summary, AWS Direct Connect is designed specifically to create a private, dedicated network connection between your on-premises infrastructure and the AWS Cloud, addressing the requirements of the question directly. The other options focus on different aspects of networking and DNS management but don't offer the same dedicated connection capability.

For further research, refer to the official AWS documentation:

AWS Direct Connect:<https://aws.amazon.com/directconnect/>
VPC Peering:<https://docs.aws.amazon.com/vpc/latest/peering/what-is-vpc-peering.html> **AWS**
VPN:<https://aws.amazon.com/vpn/>
Amazon Route 53:<https://aws.amazon.com/route53/>

Question: 22

Which option is a physical location of the AWS global infrastructure?

- A. AWS DataSync
- B. AWS Region
- C. Amazon Connect
- D. AWS Organizations

Answer: B

Explanation:

The correct answer is AWS Region because it directly relates to the physical locations of AWS infrastructure.

AWS Regions are geographically distinct locations consisting of one or more Availability Zones. Each Availability Zone is a physically isolated and independent infrastructure housed within a distinct geographic location. These Availability Zones are designed to provide fault tolerance and high availability.

Option A, AWS DataSync, is a data transfer service that facilitates moving data between on-premises and AWS storage solutions, but it's not a physical location. Option C, Amazon Connect, is a cloud-based contact center service and not a physical location either. Option D, AWS Organizations, is an account management service that allows you to centrally manage multiple AWS accounts, but it doesn't represent a physical location. Therefore, only AWS Region accurately identifies a physical location within the AWS global infrastructure. The physical infrastructure including the data centers resides in the region.

Here are some authoritative links for further research:

AWS Global Infrastructure:<https://aws.amazon.com/about-aws/global-infrastructure/>
AWS Regions and Availability Zones:<https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/using-regions-availability-zones.html>

Question: 23

A company wants to protect its AWS Cloud information, systems, and assets while performing risk assessment and mitigation tasks.

Which pillar of the AWS Well-Architected Framework is supported by these goals?

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

Answer: B

Explanation:

The correct answer is **B. Security**. Here's a detailed justification:

The AWS Well-Architected Framework comprises five pillars: Operational Excellence, Security, Reliability, Performance Efficiency, and Cost Optimization. Each pillar focuses on a specific area of cloud architecture best practices.

The scenario emphasizes protecting information, systems, and assets within the AWS Cloud, along with performing risk assessment and mitigation. This aligns directly with the **Security pillar**. This pillar focuses on protecting information, systems, and assets to deliver business value through risk assessments and mitigation strategies. It emphasizes confidentiality, integrity, and availability of data while ensuring secure access controls and compliance. Key considerations include identity and access management, detective controls, infrastructure protection, data protection, and incident response. Addressing these considerations helps a company minimize risks and maintain a secure cloud environment.

Let's look at why the other options are less suitable:

A. Reliability: Reliability focuses on ensuring the system recovers from failures and meets demand through testing recovery scenarios and adapting your system as needed. While security contributes to overall reliability (e.g., preventing unauthorized changes that could cause failures), reliability itself does not primarily involve the protection of assets or risk assessment in the way the question describes.

C. Operational Excellence: Operational Excellence refers to the ability to run and monitor systems to deliver business value and to continually improve supporting processes. This pillar is about efficient operations, automation, and continuous improvement, not specifically security risk assessments and data protection.

D. Performance Efficiency: Performance Efficiency focuses on using computing resources efficiently to meet demands and maintaining that efficiency as demand changes and technologies evolve. It is about optimizing the use of resources for the best performance at the lowest cost, and doesn't specifically encompass security measures for protecting assets and mitigating risks.

Therefore, the scenario specifically highlights goals related to securing assets, assessing risks, and mitigating vulnerabilities, making the Security pillar the most relevant.

Authoritative Links for Further Research:

AWS Well-Architected Framework:<https://aws.amazon.com/well-architected/> **AWS Well-Architected Framework - Security Pillar:**
<https://docs.aws.amazon.com/wellarchitected/latest/security-pillar/security-pillar.html>

Question: 24

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:

The correct answer is B: To allow communication between the VPC and the internet.

An Internet Gateway (IGW) serves as a crucial component within a Virtual Private Cloud (VPC) by enabling instances within the VPC to connect to the internet, and conversely, allowing traffic from the internet to reach these instances. Without an IGW, a VPC is effectively isolated, meaning instances can communicate within the VPC but cannot access or be accessed by external networks.

The IGW performs two key functions: it provides a target in your VPC route tables for internet-routable traffic, and it performs network address translation (NAT) for instances that have been assigned public IPv4 addresses. When you create a route in your VPC's route table, directing traffic destined for the internet (0.0.0.0/0) to the IGW, instances in the associated subnets can access the internet, assuming they also have a public IP address or are behind a NAT gateway which itself routes through the IGW. The IGW is horizontally scaled, redundant, and highly available, ensuring it doesn't impose bandwidth constraints (contrary to option C) or require load balancing (as suggested by option D). It is not designed for creating VPN connections (option A); VPN connections require a Virtual Private Gateway or AWS VPN service.

Therefore, the fundamental purpose of the Internet Gateway is to provide a pathway for bi-directional communication between your VPC and the public internet. It is an essential element for hosting public-facing applications or enabling internet access for instances that require software updates, external services, or other internet-dependent functionalities.

[AWS Documentation - Internet Gateway](#)

Question: 25

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 correct answer is D, Implementing loosely coupled dependencies. The AWS Well-Architected Framework encourages building systems that are scalable, maintainable, and resilient. Decomposing a monolithic application into microservices directly aligns with the principle of loose coupling within the Operational Excellence, Reliability, and Performance Efficiency pillars.

Here's why:

Loose Coupling: Microservices architecture promotes independence between components. Each microservice operates autonomously and communicates with others through well-defined APIs. This reduces dependencies and allows individual services to be updated, scaled, or even replaced without impacting the entire application.

Scalability: Microservices enable independent scaling. If one part of the application experiences increased load, only the corresponding microservice needs to be scaled, rather than the entire monolith.

Maintainability: Smaller, independent codebases are easier to understand, test, and maintain. Developers can focus on specific microservices without having to navigate the complexity of a large monolithic code base.

Resilience: The failure of one microservice is less likely to bring down the entire application. Other microservices can continue to function, providing a more resilient system.

Options A, B, and C, while important best practices, don't directly address the problem of a monolithic application being difficult to scale and maintain. Functional testing (A) is a testing strategy, automation (B) improves deployment efficiency, and multi-location deployment (C) enhances availability but these are not the core solution to the monolith's inherent limitations. Loosely coupled microservices specifically addresses the problem presented in the scenario.

Reference:

AWS Well-Architected Framework: This document outlines the five pillars and best practices for designing and operating reliable, secure, efficient, and cost-effective systems in the cloud. Look under operational excellence, reliability, and performance efficiency pillars for specific benefits tied to loose coupling.

Question: 26

A company has an AWS account. The company wants to audit its password and access key rotation details for compliance purposes. Which AWS service or tool will meet this requirement?

- A. IAM Access Analyzer
- B. AWS Artifact
- C. IAM credential report
- D. AWS Audit Manager

Answer: C

Explanation:

The correct answer is C, IAM credential report. Here's a detailed justification:

IAM credential reports are a crucial feature within AWS Identity and Access Management (IAM) that provide a comprehensive, downloadable CSV file containing a snapshot of all users in an AWS account and their credential status. This report details information like password last used, password enabled status, access key age, access key status (active or inactive), and last key rotation date. This allows a company to audit password policies and access key rotation for compliance purposes directly within the AWS console, aligning with security best practices. https://docs.aws.amazon.com/IAM/latest/UserGuide/id_credentials_getting-report.html

IAM Access Analyzer (Option A) primarily identifies the resources in your organization and accounts, such as S3 buckets, IAM roles, or KMS keys, that are shared with an external entity. While security-focused, it doesn't provide the granular details of password and access key rotation.

<https://docs.aws.amazon.com/IAM/latest/UserGuide/what-is-access-analyzer.html>

AWS Artifact (Option B) is an on-demand resource for AWS compliance reports and agreements. It provides access to compliance reports like SOC reports and PCI DSS compliance documentation, and it allows you to accept agreements. It isn't directly used for auditing password and access key rotations.

<https://aws.amazon.com/artifact/>

AWS Audit Manager (Option D) helps you continuously audit your AWS usage to simplify how you assess risk and compliance with regulations and industry standards. While Audit Manager can assess IAM-related controls, the IAM credential report provides a direct, immediately available snapshot of user credential details for compliance. Audit Manager takes a broader, ongoing assessment approach.

<https://aws.amazon.com/audit-manager/>

Therefore, the IAM credential report is the most suitable and direct tool within AWS to meet the requirement of auditing password and access key rotation details for compliance.

Question: 27

A company wants to receive a notification when a specific AWS cost threshold is reached. Which AWS services or tools can the company use to meet this requirement? (Choose two.)

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

Answer: BD

Explanation:

The correct answer is B. AWS Budgets and D. Amazon CloudWatch.

AWS Budgets allows you to set custom cost and usage budgets and receive alerts when those budgets are exceeded. It directly addresses the requirement of getting notified when a specific cost threshold is reached.

You can configure various thresholds and notification methods, including email and integration with other AWS services. This is specifically designed for cost management and monitoring.

[<https://aws.amazon.com/aws-cost-management/aws-budgets/>]

Amazon CloudWatch enables monitoring of AWS resources and applications in real time. While it doesn't directly manage budgets, it can be used to monitor AWS billing metrics. By creating CloudWatch alarms based on these metrics, you can receive notifications when certain cost thresholds are reached. For example, you can set an alarm to trigger when the estimated charges for an AWS service exceed a predefined amount.

[<https://aws.amazon.com/cloudwatch/>]

Amazon SQS (A) is a message queuing service used for decoupling and scaling microservices, distributed systems, and serverless applications. It doesn't provide cost management or notification capabilities related to cost thresholds.

Cost Explorer (C) is a tool used for visualizing, understanding, and managing your AWS costs and usage over time. While it allows you to analyze cost trends, it doesn't provide proactive notifications when thresholds are exceeded. It's more of an analytical tool than an alerting mechanism.

AWS Cost and Usage Report (E) provides detailed data about your AWS costs and usage, which can be used for in-depth analysis. However, it doesn't offer built-in alerting capabilities when costs exceed specific thresholds. You'd need to build custom solutions using this data, whereas AWS Budgets and CloudWatch offer the functionality natively.

Question: 28

Which AWS service or resource provides answers to the most frequently asked security-related questions that AWS receives from its users?

- A. AWS Artifact
- B. Amazon Connect

- C.AWS Chatbot
- D.AWS Knowledge Center

Answer: D

Explanation:

The correct answer is D, AWS Knowledge Center. Here's a detailed justification:

The AWS Knowledge Center is the central repository for a vast collection of articles, FAQs, how-to guides, and troubleshooting tips curated by AWS experts. A core function of the Knowledge Center is to address common customer questions, particularly those related to security best practices, compliance, and incident response.

Options A, B, and C are not primarily designed to provide general security FAQs:

AWS Artifact focuses on providing on-demand access to AWS' compliance reports and security certifications. While related to security, it doesn't answer frequently asked questions. (<https://aws.amazon.com/artifact/>) **Amazon Connect** is a cloud-based contact center service. While security is a factor in its design and operation, it's not where you'd find answers to general security questions. (<https://aws.amazon.com/connect/>) **AWS Chatbot** enables interaction with AWS services using chat platforms. Though it can be used for some security-related tasks, it doesn't serve as a primary source of security FAQs. (<https://aws.amazon.com/chatbot/>)

Therefore, the AWS Knowledge Center is the most suitable resource to provide answers to frequently asked security-related questions as it is designed as a central location for AWS-related information.

Question: 29

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

Answer: AB

Explanation:

The AWS shared responsibility model outlines the security responsibilities between AWS and the customer. AWS is responsible for the security of the cloud, handling the physical infrastructure, hardware, networking, and the core services. The customer is responsible for security in the cloud, focusing on securing their data, applications, operating systems, and identities.

Option A, configuring security group firewalls, falls under customer responsibility because security groups are a customer-configurable resource used to control network traffic to and from AWS resources like EC2 instances. The customer decides which ports and protocols are open, and from which IP addresses or CIDR blocks. This directly impacts the security in the cloud, under the customer's control.

Option B, classifying company assets in the AWS Cloud, is also the customer's responsibility. Data classification involves categorizing data based on its sensitivity and importance. AWS provides tools like AWS Security Hub and AWS Identity and Access Management (IAM) that support data classification, but the

responsibility for identifying and classifying the data itself, as well as defining security policies based on that classification, rests with the customer. This is crucial for compliance and data protection.

Option C, determining which Availability Zones to use for Amazon S3 buckets, is partially a customer responsibility. While S3 itself is a regional service offering high availability, the customer decides where the data resides by selecting the AWS region. While S3 manages the distribution across AZs within that region, the initial regional placement is on the customer. However, between these two selected answers, A and B are more clearly defined customer responsibilities regarding security.

Option D, patching or upgrading Amazon DynamoDB, is AWS's responsibility. DynamoDB is a fully managed NoSQL database service. AWS handles all the underlying infrastructure, patching, upgrades, and maintenance of the database.

Option E, selecting Amazon EC2 instances to run AWS Lambda on, is incorrect because AWS Lambda is a serverless compute service. The customer does not manage or provision any EC2 instances to run Lambda functions. AWS handles the underlying infrastructure and resource allocation.

Therefore, configuring security groups and classifying company assets are both prime examples of tasks that remain the customer's responsibility under the AWS shared responsibility model.

Relevant Links:

AWS Shared Responsibility Model: <https://aws.amazon.com/compliance/shared-responsibility-model/>

Question: 30

Which of the following are pillars of the AWS Well-Architected Framework? (Choose two.)

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

Answer: BE

Explanation:

The correct answer is B. Reliability and E. Operational Excellence. These are indeed two of the five pillars of the AWS Well-Architected Framework.

The AWS Well-Architected Framework helps cloud architects build secure, high-performing, resilient, and efficient infrastructure for their applications. It's based on five pillars:

1. **Operational Excellence:** Focuses on running and monitoring systems to deliver business value and continually improving processes and procedures. This includes automating changes, responding to events, and defining standards to manage operations.
2. **Security:** Encompasses protecting information, systems, and assets while delivering business value through risk assessments and mitigation strategies. This pillar highlights the importance of identity and access management, detection controls, infrastructure protection, data protection, and incident response.
3. **Reliability:** Concerns the ability of a system to recover from failures and meet demand, while avoiding disruptions. Considerations involve fault tolerance, recovery procedures, and scalability of

the system.

4. **Performance Efficiency:** Focuses on using computing resources efficiently to meet requirements and maintaining that efficiency as demand changes and technologies evolve. This involves selecting the right resource types and sizes, monitoring performance, and adapting to changing requirements.
5. **Cost Optimization:** Involves running systems at the lowest price point without sacrificing other architectural pillars such as performance efficiency or reliability. Key aspects are understanding spending, analyzing usage, avoiding unused resources, and selecting the optimal resource type.

Availability (A) and Scalability (C) are related to the Reliability and Performance Efficiency pillars respectively, but are not pillars themselves. Responsive design (D) is a design principle related to web applications and user interfaces rather than a pillar of a cloud architecture framework.

Therefore, Reliability and Operational Excellence are the two correct options representing pillars in the AWS Well-Architected Framework.

Further reading:

AWS Well-Architected Framework: <https://wa.aws.amazon.com/>

Question: 31

Which AWS service or feature is used to send both text and email messages from distributed applications?

- A. Amazon Simple Notification Service (Amazon SNS)
- B. Amazon Simple Email Service (Amazon SES)
- C. Amazon CloudWatch alerts
- D. Amazon Simple Queue Service (Amazon SQS)

Answer: A

Explanation:

The correct answer is **A. Amazon Simple Notification Service (Amazon SNS)**.

Amazon SNS is a fully managed messaging service for both application-to-application (A2A) and application-to-person (A2P) communication. It facilitates sending messages (including text messages via SMS and email) from distributed applications to a large number of subscribers. SNS supports various delivery protocols, including HTTP, HTTPS, email, SMS, mobile push, and SQS.

While Amazon SES (Simple Email Service) excels at sending emails, it does not natively handle SMS messaging. CloudWatch alerts are primarily for monitoring and triggering actions based on performance metrics, not for general-purpose messaging. Amazon SQS (Simple Queue Service) is a message queuing service, allowing components of distributed applications to communicate asynchronously. SQS is excellent for decoupling services but doesn't directly send emails or SMS messages to end users like SNS does.

SNS topics act as communication channels. Applications can publish messages to a topic, and SNS distributes those messages to all subscribers configured for that topic. This publish/subscribe (pub/sub) model is ideal for notifying many recipients simultaneously. For example, you could use SNS to send an email to all customers after a successful purchase and simultaneously send an SMS message to administrators about a critical system alert. The broad range of delivery options makes SNS a perfect tool for applications requiring both email and SMS messaging functionalities.

For further research, consider exploring the following resources:

Question: 32

A user needs programmatic access to AWS resources through the AWS CLI or the AWS API. Which option will provide the user with the appropriate access?

- A. Amazon Inspector
- B. Access keys
- C. SSH public keys
- D. AWS Key Management Service (AWS KMS) keys

Answer: B

Explanation:

The correct answer is **B. Access keys**. Here's why:

When a user requires programmatic access to AWS services (e.g., using the AWS CLI, SDKs, or APIs), they need a way to authenticate their requests. Access keys (Access Key ID and Secret Access Key) are the primary mechanism provided by AWS for this purpose. These keys allow the user to securely identify themselves and their permissions to the AWS environment.

Option A, Amazon Inspector, is a vulnerability management service that automatically assesses the security posture of EC2 instances and container images. It doesn't provide programmatic access credentials.

Option C, SSH public keys, are used to securely connect to EC2 instances. While SSH provides a secure connection, it's not designed for programmatic access to AWS services through APIs or the CLI. SSH keys are used for interactive shell access to compute instances.

Option D, AWS Key Management Service (AWS KMS) keys, are used to encrypt data at rest and in transit. While KMS keys are crucial for security, they don't directly grant a user programmatic access to AWS resources. Instead, KMS keys are used by services that the user already has access to.

Access keys are explicitly designed to facilitate secure programmatic interactions with AWS. They are generated and managed through IAM (Identity and Access Management), allowing administrators to control which users or services have access to specific AWS resources and actions. The access key ID identifies the user, and the secret access key is like a password that proves the user's identity. Best practices dictate rotating access keys regularly to maintain security.

For further reading, consult the official AWS documentation:

Managing Access Keys for IAM Users:

https://docs.aws.amazon.com/IAM/latest/UserGuide/id_credentials_access-keys.html

Question: 33

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

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

- A. Reserved Instances
- B. Spot Instances

- C. On-Demand Instances
- D. Dedicated Instances

Answer: B

Explanation:

The optimal pricing model for the company's AWS Batch simulations is Spot Instances (B). Here's why:

Stateless and Fault-Tolerant Workloads: Spot Instances are ideal for workloads that are stateless and fault-tolerant. If a Spot Instance is interrupted, the simulation can be automatically restarted on another instance without data loss or significant impact, given the fault-tolerant nature of the simulations.

Cost Optimization: Spot Instances offer significantly reduced prices compared to On-Demand Instances, often up to 90% lower. This is because the company is bidding on unused EC2 capacity.

AWS Batch Integration: AWS Batch seamlessly integrates with Spot Instances. Batch can automatically manage the bidding and provisioning of Spot Instances, ensuring that the simulations run cost-effectively.

3-Hour Runtime: The 3-hour maximum runtime falls within the typical lifecycle of Spot Instances. While Spot Instances can be interrupted, the simulation has a defined and relatively short runtime, increasing the likelihood of successful completion.

Not Reserved Instances: Reserved Instances (A) are better suited for long-term, predictable workloads, not for fluctuating simulation requirements.

Not On-Demand Instances: On-Demand Instances (C) are more expensive than Spot Instances and are not cost-effective for a large number of simulations.

Not Dedicated Instances: Dedicated Instances (D) are the most expensive and are typically used for compliance or licensing requirements, not for cost optimization of compute-intensive simulations.

In summary, Spot Instances provide the best balance of cost savings and availability for the company's fault-tolerant, stateless, and relatively short-duration simulations managed by AWS Batch.

Authoritative Links:

AWS Spot Instances: <https://aws.amazon.com/ec2/spot/>

AWS Batch Pricing: <https://aws.amazon.com/batch/pricing/>

Question: 34

What does the concept of agility mean in AWS Cloud computing? (Choose two.)

- A. The speed at which AWS resources are implemented
- B. The speed at which AWS creates new AWS Regions
- C. The ability to experiment quickly
- D. The elimination of wasted capacity
- E. The low cost of entry into cloud computing

Answer: AC

Explanation:

The correct answers are A and C. Agility in AWS Cloud computing refers to the ability to rapidly develop, test, and launch new applications or functionalities. Option A, "The speed at which AWS resources are

implemented," is correct because AWS allows for the fast provisioning of infrastructure. Services like EC2, S3, and databases can be spun up in minutes, enabling organizations to quickly react to market demands and deploy updates more frequently than with traditional infrastructure. Option C, "The ability to experiment quickly," is also correct. AWS provides a plethora of services that encourage experimentation without large upfront investment. Concepts like Infrastructure as Code (IaC) and automated deployment pipelines facilitate continuous integration and continuous delivery (CI/CD), empowering developers to rapidly iterate and test new ideas in production-like environments. Options B, D, and E are incorrect as they relate more to AWS internal scaling, cost optimization, and accessibility, respectively. Useful resources:

AWS Cloud Benefits: <https://aws.amazon.com/what-is-aws/benefits/>

AWS Cloud Adoption Framework: <https://aws.amazon.com/cloud-adoption-framework/>

Question: 35

A company needs to block SQL injection attacks.
Which AWS service or feature can meet this requirement?

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

Answer: A

Explanation:

The correct answer is A: AWS WAF (Web Application Firewall).

AWS WAF is a web application firewall that helps protect your web applications from common web exploits and bots that may affect availability, compromise security, or consume excessive resources. One of the core functionalities of AWS WAF is to filter malicious traffic and protect against common web attacks, including SQL injection.

SQL injection is a code injection technique used to attack data-driven applications, in which malicious SQL statements are inserted into an entry field for execution. AWS WAF can inspect incoming HTTP requests and, based on rules that you define, block, allow, or count requests. These rules can include pre-configured rulesets specifically designed to prevent SQL injection attacks.

Option B, AWS Shield, provides protection against DDoS (Distributed Denial of Service) attacks, focusing on infrastructure and network layer protection rather than application-level vulnerabilities like SQL injection. While Shield helps keep services available, it doesn't directly address SQL injection.

Option C, Network ACLs (Network Access Control Lists), act as firewalls at the subnet level, controlling traffic entering and exiting a subnet. They operate at layers 3 and 4 of the OSI model (network and transport layers), primarily dealing with IP addresses, protocols, and ports. Network ACLs cannot inspect the content of HTTP requests to identify and block SQL injection attempts.

Option D, Security groups, act as virtual firewalls for your EC2 instances and other resources, controlling inbound and outbound traffic at the instance level. Similar to Network ACLs, they operate at the network and transport layers and don't inspect application-layer content to identify SQL injection attempts.

In summary, AWS WAF is designed for application-layer protection and provides specific rulesets to mitigate SQL injection attacks, making it the appropriate solution for this requirement. AWS Shield protects against DDoS attacks, while Network ACLs and security groups operate at lower network layers and do not analyze HTTP request content to prevent SQL injection.

Further research:

AWS WAF Documentation:<https://aws.amazon.com/waf/>

SQL Injection Prevention with AWS WAF:<https://aws.amazon.com/blogs/security/how-to-protect-from-sql-injection-by-using-aws-waf/>

Question: 36

Which AWS service or feature identifies whether an Amazon S3 bucket or an IAM role has been shared with an external entity?

- A. AWS Service Catalog
- B. AWS Systems Manager
- C. AWS IAM Access Analyzer
- D. AWS Organizations

Answer: C

Explanation:

The correct answer is C, AWS IAM Access Analyzer. Here's why:

AWS IAM Access Analyzer is designed to help you identify unintended resource access to your AWS resources, including S3 buckets and IAM roles. It analyzes resource policies to determine which resources are accessible to external entities (accounts outside your organization). It does this by using automated reasoning to mathematically verify resource policies against your defined trust boundaries.

Specifically, Access Analyzer identifies S3 buckets that have been made public or shared with other AWS accounts, pinpointing external access. It also analyzes IAM policies attached to roles to see if those roles grant access to other AWS accounts or principals outside of your organization. When Access Analyzer finds a policy that allows access from outside your trust zone, it generates a finding.

AWS Service Catalog (A) allows organizations to create and manage catalogs of IT services that are approved for use on AWS, which is unrelated to identifying external access. AWS Systems Manager (B) helps you manage your AWS and on-premises infrastructure, automating tasks such as patching and configuration management. While useful for overall management, it doesn't directly analyze resource policies for external access. AWS Organizations (D) is used for managing multiple AWS accounts, but it does not by itself identify whether a resource in an individual account is shared with an external entity. It's more about organizational structure and centralized billing/governance across accounts.

IAM Access Analyzer focuses specifically on the permissions and access granted to AWS resources. The other options are broader service offerings. Therefore, only Access Analyzer directly identifies whether an S3 bucket or an IAM role is shared with an external entity through its analysis of resource policies.

For further research, refer to these authoritative links:

AWS IAM Access Analyzer Documentation:<https://docs.aws.amazon.com/IAM/latest/UserGuide/access-analyzer.html>

AWS Security Blog Post on Access Analyzer:<https://aws.amazon.com/blogs/security/simplify-permissions-management-using-aws-iam-access-analyzer/>

Question: 37

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 Macie.

Answer: B

Explanation:

The correct answer is B, downloading the reports from AWS Artifact. AWS Artifact is a service that provides on-demand access to AWS compliance reports, such as SOC reports, PCI DSS compliance packages, and ISO certifications. It serves as a central repository for these resources, eliminating the need to contact AWS support or compliance teams directly for common compliance documentation.

Option A, contacting the AWS Compliance team, is not the most efficient method. While AWS Compliance teams can provide assistance, AWS Artifact is the designed self-service resource. Option C, opening a case with AWS Support, is also less efficient. Support is typically used for troubleshooting issues and not for routine access to readily available compliance reports. Option D, generating the reports with Amazon Macie, is incorrect. Amazon Macie is a security service that uses machine learning to discover and protect sensitive data; it does not generate compliance reports. Macie helps maintain data security posture, which is related to compliance but distinct from obtaining compliance documentation.

Therefore, AWS Artifact provides the quickest and easiest way for a Cloud Practitioner to get the compliance reports needed before migrating to AWS. It's designed to provide transparency and self-service access to compliance documentation.

Authoritative links:

AWS Artifact: <https://aws.amazon.com/artifact/>

Question: 38

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

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

Answer: A

Explanation:

The correct answer is A: Cost of application software licenses.

Here's why:

In the AWS Cloud, the shared responsibility model dictates cost allocation. AWS takes responsibility for the "infrastructure as a service" (IaaS) layer. This encompasses the underlying hardware, physical security, and operational costs associated with running the AWS data centers. The customer is responsible for everything above that, including the operating system, applications, data, and in this scenario, software licenses.

Specifically:

A. Cost of application software licenses: When an e-commerce company migrates its applications (e.g., database software, e-commerce platform software) to AWS, it retains the responsibility for procuring and managing the licenses for that software. This applies whether the licenses are perpetual or subscription-based.

B. Cost of the hardware infrastructure on AWS: AWS manages and covers the costs related to its hardware infrastructure. The company pays for the usage of that infrastructure, but AWS owns and maintains it.

C. Cost of power for the AWS servers: The electricity needed to power the AWS servers is an operational expense borne by AWS.

D. Cost of physical security for the AWS data center: Maintaining the physical security of data centers is a direct responsibility and cost for AWS.

The company's direct responsibility involves any applications they deploy, their associated licenses, and the management of their data. The Shared Responsibility Model clarifies who is responsible for different security and management aspects when using cloud services. In short, AWS takes care of the underlying infrastructure, while the customer is responsible for what they put "in" the cloud.

Further Research:

AWS Shared Responsibility Model: <https://aws.amazon.com/compliance/shared-responsibility-model/>

Question: 39

A company is setting up AWS Identity and Access Management (IAM) on an AWS account. Which recommendation complies with IAM security best practices?

- A. Use the account root user access keys for administrative tasks.
- B. Grant broad permissions so that all company employees can access the resources they need.
- C. Turn on multi-factor authentication (MFA) for added security during the login process.
- D. Avoid rotating credentials to prevent issues in production applications.

Answer: C

Explanation:

The correct answer is C: Turn on multi-factor authentication (MFA) for added security during the login process.

IAM security best practices emphasize minimizing risk and adhering to the principle of least privilege. Option A is incorrect because using the root user access keys for administrative tasks exposes the entire AWS account to significant risk. The root user has unrestricted access, and compromising these keys would grant an attacker complete control.

Option B violates the principle of least privilege. Granting broad permissions allows users to access resources they don't need, increasing the potential attack surface and the risk of accidental misconfiguration or malicious actions.

Option D is incorrect because rotating credentials regularly is crucial for mitigating the risk of compromised credentials. If a credential is stolen, its lifespan is limited if rotation is in place.

Option C, enabling MFA, adds an extra layer of security beyond just a username and password. Even if a password is compromised, an attacker would still need the MFA device to gain access. This significantly

reduces the likelihood of unauthorized access, aligning directly with IAM security best practices. MFA is highly recommended by AWS to improve account security.

Supporting links:

AWS IAM Best Practices: <https://docs.aws.amazon.com/IAM/latest/UserGuide/best-practices.html>
Securing your AWS account: <https://aws.amazon.com/security/security-best-practices/>
AWS Multi-Factor Authentication: <https://aws.amazon.com/iam/features/mfa/>

Question: 40

Elasticity in the AWS Cloud refers to which of the following? (Choose two.)

- A. How quickly an Amazon EC2 instance can be restarted
- B. The ability to rightsize 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 focuses on the ability to dynamically adapt resources to meet fluctuating demands. This essentially means that you can easily scale resources up or down based on your current needs, optimizing performance and cost. Option B, "The ability to rightsize resources as demand shifts," directly aligns with this concept. Elasticity allows you to automatically increase or decrease the computing capacity, storage, or other resources as your application's workload changes. When demand is high, you can scale up to ensure optimal performance; when demand decreases, you can scale down to minimize expenses.

Option E, "How easily resources can be procured when they are needed," is also a crucial aspect of elasticity. AWS provides easy access to a vast range of services and resources that can be provisioned quickly. This speed and ease of procurement enable you to respond swiftly to changing demands, which is vital for maintaining a seamless user experience.

Option A relates more to the recovery or maintenance of individual EC2 instances rather than the overall dynamic scaling concept. Option C describes the maximum RAM capacity of a specific instance, which is a fixed attribute, not a dynamic adjustment. Option D, the pay-as-you-go model, is related to cost optimization but is not directly defining elasticity. While pay-as-you-go enables cost savings achieved through elasticity, it is a separate benefit derived from cloud infrastructure. Elasticity is about scaling resources, while pay-as-you-go is about how you pay for them.

In summary, Elasticity on AWS is about dynamic resource allocation based on demand (B) and the ease of acquiring resources when required (E).

For further research, consider:

AWS Documentation on Elasticity: <https://docs.aws.amazon.com/whitepapers/latest/aws-overview/elastic-cloud-computing.html>

AWS Certified Cloud Practitioner Exam Guide: This document elaborates on the core concepts tested in the exam, including elasticity.

Question: 41

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

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

Answer: A**Explanation:**

The correct answer is A. AWS CloudTrail is a service that enables auditing and governance by recording API calls made on your AWS account. It delivers log files to an Amazon S3 bucket that you specify. These logs contain information such as the identity of the caller, the time of the API call, the source IP address of the caller, the request parameters, and the response elements returned by the AWS service. This makes it indispensable for security analysis, resource change tracking, compliance auditing, and troubleshooting.

AWS Trusted Advisor provides best practice recommendations across several categories, including cost optimization, security, fault tolerance, and performance, but it does not record API calls. Amazon Inspector is an automated security assessment service that helps improve the security and compliance of applications deployed on AWS. It identifies potential security vulnerabilities and deviations from security best practices. AWS X-Ray is a distributed tracing system that helps developers analyze and debug production, distributed applications, such as those built using a microservices architecture, but it is not primarily an auditing tool for API calls. CloudTrail's primary function is exactly to track API usage within the AWS environment, providing a trail of actions taken for accountability and compliance. Therefore, CloudTrail directly fulfills the requirement stated in the prompt.

<https://aws.amazon.com/cloudtrail/>

Question: 42

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:**

The correct answer is A: Managing the code within the Lambda function.

According to the AWS Shared Responsibility Model, AWS manages the security of the cloud, while the customer is responsible for security in the cloud. For AWS Lambda, this division of labor means AWS handles the underlying infrastructure, including the hardware, operating system, and runtime environment.

Customers, on the other hand, are responsible for the code they deploy within their Lambda functions.

Specifically, customers must ensure the code is secure, does not contain vulnerabilities, and adheres to best practices. They are responsible for managing any dependencies the code relies on and keeping those dependencies up to date. The security of the function itself, including its logic and data handling, falls under

the customer's responsibility.

Options B, C, and D are incorrect because they represent tasks handled by AWS. AWS maintains and secures the hardware and operating systems powering Lambda. The automatic scaling and pay-per-use nature of Lambda also eliminate the need for customers to manage the lifecycle (starting/stopping) of Lambda functions. These are managed by the AWS platform. In essence, customers focus solely on the function's code and configuration, while AWS handles everything else at the infrastructure layer. AWS documentation detailing the shared responsibility model and Lambda responsibilities is available here:

[AWS Shared Responsibility Model](#)
[Security Best Practices for AWS Lambda](#)

Question: 43

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:

The correct answer is Amazon Athena (B). Here's why:

Amazon Athena is a serverless query service that enables you to analyze data directly in Amazon S3 using standard SQL. Because it's serverless, you only pay for the queries you run. This makes it incredibly cost-effective for occasional data analysis.

Amazon Redshift (A) is a fully managed data warehouse service designed for complex analytical queries and large-scale data warehousing. While it offers powerful analytical capabilities, it requires provisioning and managing a cluster, which incurs costs even when you're not actively running queries. This makes it less suitable for occasional use cases.

Amazon Kinesis (C) is a platform for real-time data streaming and processing. It's designed for ingesting and processing high-volume, real-time data streams, not for running ad-hoc queries on data stored in S3.

Amazon RDS (D) is a relational database service that supports various database engines like MySQL, PostgreSQL, and SQL Server. It's designed for transactional workloads, not for analytical queries on data stored in S3.

Therefore, for a company with 5 TB of data in S3 that needs to occasionally run queries for analysis, Amazon Athena provides the most cost-effective solution because it avoids the costs associated with maintaining a dedicated data warehouse or database instance. The pay-per-query pricing model aligns perfectly with the infrequent usage pattern.

Further Research:

Amazon Athena: <https://aws.amazon.com/athena/>
Amazon Redshift: <https://aws.amazon.com/redshift/>
Amazon Kinesis: <https://aws.amazon.com/kinesis/>
Amazon RDS: <https://aws.amazon.com/rds/>

Question: 44

Which AWS service can be used at no additional cost?

- A. Amazon SageMaker
- B. AWS Config
- C. AWS Organizations
- D. Amazon CloudWatch

Answer: C

Explanation:

The correct answer is C, AWS Organizations. While many AWS services incur costs based on usage and resource consumption, AWS Organizations itself is offered at no additional charge. You only pay for the AWS resources that your accounts within the organization utilize. AWS Organizations provides tools for managing and governing multiple AWS accounts, allowing you to centrally manage billing, control access, comply with regulations, and share resources across your accounts. Key features like consolidated billing, which simplifies payment for multiple AWS accounts, and organizational units (OUs) for grouping accounts are available without direct cost. Service Control Policies (SCPs), a powerful tool to establish guardrails for IAM permissions across all accounts in an OU or an entire organization, also do not add to the overall AWS Organizations costs.

Amazon SageMaker (A) is a machine learning service that incurs costs based on usage of compute instances, data storage, and other resources consumed during model training and deployment. AWS Config (B) is a configuration management service that charges based on the number of configuration items recorded. Amazon CloudWatch (D) is a monitoring and observability service with charges associated with metrics stored, logs ingested, and alarms created. Therefore, AWS Organizations stands out as the service offered without incurring additional costs, aside from the resources used within the managed AWS accounts.

For detailed information, refer to the AWS Organizations documentation:

<https://aws.amazon.com/organizations/> and the AWS Pricing Overview: <https://aws.amazon.com/pricing/>

Question: 45

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

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

Answer: C

Explanation:

The correct answer is C, Cloud fluency. The AWS Cloud Adoption Framework (AWS CAF) helps organizations develop and execute efficient and effective cloud adoption strategies. It's organized around six perspectives: Business, People, Governance, Platform, Security, and Operations. The People perspective focuses on skills, knowledge, and organizational structures necessary to thrive in a cloud environment.

Cloud fluency, specifically, directly addresses the training, education, and overall skill development required for individuals and teams to effectively utilize AWS services. It emphasizes building a common understanding of cloud concepts and establishing a shared vocabulary. This includes fostering a culture of learning and experimentation related to cloud technologies.

Data architecture falls under the Platform perspective, which deals with the technical aspects of building and deploying cloud solutions. Event management is generally considered part of the Operations perspective, focused on monitoring, managing, and troubleshooting cloud-based workloads. Strategic partnership, while important, is more aligned with the Business perspective, which focuses on aligning cloud adoption with business outcomes and leveraging partnerships to achieve strategic goals. Therefore, options A, B, and D relate to elements outside of the People perspective and its goal of developing relevant cloud skills.

Further research:

AWS Cloud Adoption Framework:<https://aws.amazon.com/professional-services/CAF/> **AWS CAF Perspectives:**<https://docs.aws.amazon.com/whitepapers/latest/aws-cloud-adoption-framework/aws-caf-perspectives.html>

Question: 46

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? (Choose two.)

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

Answer: CD

Explanation:

The correct answer is **C. Reserved Instances** and **D. Savings Plans**.

Reserved Instances offer a significant discount (up to 75%) compared to On-Demand instances in exchange for a one- or three-year commitment. This upfront commitment lowers the cost if the company knows its long-term EC2 needs. There are three types of Reserved Instances: Standard, Convertible, and Scheduled.

Standard provides the most significant discounts but less flexibility, while Convertible allows changing instance attributes.

Savings Plans also offer lower pricing in exchange for a commitment to a consistent amount of usage, measured in dollars per hour, for one or three years. There are two types: Compute Savings Plans and EC2 Instance Savings Plans. Compute Savings Plans apply to EC2, AWS Lambda, and AWS Fargate usage, providing flexibility across different compute services. EC2 Instance Savings Plans are specific to EC2 instance families within a region and offer the lowest prices.

Why the other options are not the lowest cost:

A. Spot Instances: Spot Instances offer deeply discounted pricing but are not suitable for production workloads requiring consistent availability. They can be terminated with a two-minute warning if the Spot price exceeds the company's bid, disrupting production. They are best suited for fault-tolerant workloads.

B. On-Demand Instances: On-Demand Instances provide flexibility without upfront commitment but are the

most expensive pricing option. They are suitable for short-term, unpredictable workloads.

E. Dedicated Hosts: Dedicated Hosts offer physical servers dedicated to a single customer. They are the most expensive option, providing dedicated hardware and are primarily used for regulatory or licensing requirements. They do not provide the lowest cost.

Reserved Instances and Savings Plans are ideal when the company has consistent EC2 usage and wants to reduce costs by making an upfront commitment. Choosing between these two largely depends on the specific workload. Savings Plans are generally more flexible than Reserved Instances because they can apply to different instance types or even other compute services. But in general, both offer lower cost than other options for consistent workloads.

Authoritative Links:

[Amazon EC2 Pricing](#)

[AWS Savings Plans](#)

[Amazon EC2 Reserved Instances](#)

Question: 47

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

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

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

Answer: C

Explanation:

The correct answer is C: AWS Regions. Here's why:

AWS Regions are geographically isolated locations where AWS deploys and operates its data centers. Each Region contains multiple Availability Zones, which are distinct locations within a Region that are engineered to be isolated from failures. When deploying Amazon RDS databases, you need to select an AWS Region. This choice directly impacts proximity to your on-premises infrastructure. The closer the AWS Region is to your current geographical location, the lower the latency and faster the data transfer speeds for migration and ongoing operations. This will provide the closest infrastructure mirroring the existing setup.

Amazon Connect (A) is a cloud-based contact center service, irrelevant to database deployment location. AWS Wavelength (B) provides ultra-low latency infrastructure for 5G devices, but this is more about mobile applications and doesn't help select a general deployment area. AWS Direct Connect (D) establishes a dedicated network connection from on-premises to AWS, improving bandwidth and reducing latency after the Region has been selected. It doesn't assist in choosing the initial deployment area. Therefore, selecting an appropriate AWS Region is the first and most direct step for achieving the desired proximity. The official AWS documentation explicitly mentions selecting a Region for deploying resources. The documentation outlines how to choose the best region based on your business needs, including latency and location.

Further research:

AWS Regions and Availability Zones: <https://aws.amazon.com/about-aws/global-infrastructure/>

Amazon RDS: <https://aws.amazon.com/rds/>

Question: 48

A company is exploring the use of the AWS Cloud, and needs to create a cost estimate for a project before the infrastructure is provisioned.

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

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

Answer: B**Explanation:**

The correct answer is B, AWS Pricing Calculator. The AWS Pricing Calculator is specifically designed to estimate the cost of AWS services for your use cases before you provision any resources. It allows you to model your solutions, explore different service configurations, and then calculate the estimated monthly and annual cost.

AWS Free Tier (A) provides free usage of certain AWS services up to specified limits, but it's not for comprehensive cost estimation of a whole project. While it can reduce costs, it doesn't give the high-level forecasting capability the calculator offers.

AWS Billing and Cost Management (C) is a suite of tools to analyze and manage your AWS costs after you are already using AWS services. It includes cost allocation tags, budgets, and cost anomaly detection. It doesn't help with pre-deployment estimation.

AWS Cost and Usage Report (D) provides detailed information about your AWS costs and usage after they have been incurred. It's a powerful tool for cost analysis and optimization but not useful for initial cost estimation.

Therefore, the AWS Pricing Calculator best fits the requirement of estimating costs before infrastructure is provisioned. The other options are more related to cost tracking and management after resources are in use. The core purpose of the AWS Pricing Calculator is proactive cost planning.

For further research, refer to the AWS Pricing Calculator documentation:

<https://aws.amazon.com/pricing/calculator/>

Question: 49

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

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

Answer: A**Explanation:**

The correct answer is A: Deliver the content through Amazon CloudFront.

CloudFront is a content delivery network (CDN) service that caches content in edge locations around the world. This ensures that users receive content from the nearest edge location, reducing latency and improving performance for geographically distributed users. This aligns perfectly with the requirement of delivering images and videos globally with minimal latency. CDNs are designed specifically to handle high-traffic content delivery with optimized performance and reduced load on origin servers.

Option B, storing the content on Amazon S3 and enabling S3 cross-region replication, replicates data to multiple AWS Regions, improving data durability and availability. However, it doesn't directly address latency for global content delivery as users still have to access the data from specific S3 Regions, potentially far away from their location. Replication is more focused on data redundancy and disaster recovery, not necessarily fast content delivery.

Option C, implementing a VPN across multiple AWS Regions, is generally used for secure connections between networks. While it could indirectly help with access to resources in other Regions, it would add complexity and overhead, and it is not a cost-effective solution for content delivery compared to a CDN. VPNs are not optimized for content caching and global distribution.

Option D, delivering the content through AWS PrivateLink, provides private connectivity between VPCs and AWS services or supported AWS Marketplace partner services without exposing traffic to the public internet.

This is primarily useful for secure internal communication and accessing services privately within the AWS network. It's not designed for general public content delivery.

Therefore, using Amazon CloudFront provides the most cost-effective and performance-optimized solution for globally delivering images and videos with minimal latency. CloudFront automatically handles the caching and distribution, ensuring that users always receive the content from the closest available server.

For more information, see:

[Amazon CloudFront: What is CloudFront?](#)

[Amazon S3 Cross-Region Replication:](#)

[AWS PrivateLink:](#)

Question: 50

Which option is a benefit of the economies of scale based on the advantages of cloud computing?

- A. The ability to trade variable expense for fixed expense
- B. Increased speed and agility
- C. Lower variable costs over fixed costs
- D. Increased operational costs across data centers

Answer: C

Explanation:

The correct answer is C: Lower variable costs over fixed costs. Economies of scale in cloud computing refer to the cost advantages that arise from increased production. Cloud providers like AWS operate massive data centers, distributing the costs of infrastructure, power, and cooling across a huge customer base.

This large scale allows them to purchase resources in bulk, negotiate better rates with vendors, and implement efficiencies that smaller organizations cannot achieve on their own. Consequently, these cost savings are passed on to customers, resulting in lower per-unit costs for services.

Option A is incorrect because cloud computing allows trading fixed expense for variable expense. Instead of investing heavily in owning and maintaining physical infrastructure (fixed expense), users pay for resources as they are consumed (variable expense).

Option B, Increased speed and agility, is a benefit of cloud computing in general, but not directly related to economies of scale. While economies of scale can contribute to faster innovation and deployment, they are not the primary driver. Agility comes from the ease of provisioning and configuring resources.

Option D, Increased operational costs across data centers, is the opposite of what economies of scale provide. Economies of scale are about reducing operational costs.

Therefore, the most direct and accurate answer is C, because the scale of cloud providers reduces the per-unit cost of services, resulting in lower variable costs for users compared to the costs of establishing and operating their own infrastructure.

[AWS Economics of the Cloud](#)[Understanding Cloud Economics](#)

Question: 51

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:

The correct answer is C: AWS Cloud Development Kit (AWS CDK).

Here's a detailed justification:

The question asks about a software development framework for defining cloud resources as code and provisioning them via AWS CloudFormation. This describes the core functionality of Infrastructure as Code (IaC).

AWS CDK is specifically designed to address this. It allows developers to define cloud infrastructure using familiar programming languages like TypeScript, Python, Java, and .NET. The AWS CDK then synthesizes these definitions into AWS CloudFormation templates. This enables developers to leverage the power and reliability of CloudFormation while using higher-level programming constructs.

Here's why the other options are incorrect:

A. AWS CLI (Command Line Interface): The AWS CLI is a tool for interacting with AWS services through commands. While you can use it to interact with CloudFormation, it doesn't provide a high-level framework for defining resources as code in the same way as AWS CDK. You would be directly writing and managing CloudFormation templates.

B. AWS Developer Center: This is a portal providing resources, documentation, and tools for developing on AWS, but it's not a specific framework for defining infrastructure as code. It's a central location for developers, but it doesn't inherently provision resources.

D. AWS CodeStar: AWS CodeStar is a cloud-based development service for quickly developing, building, and deploying applications on AWS. While it integrates with other AWS services like CloudFormation, it's primarily focused on application development lifecycle, not directly on defining cloud resources as code. CodeStar can use CloudFormation, but is not an IaC framework itself.

In summary, AWS CDK provides the abstraction needed to define infrastructure as code using familiar programming languages, which is then translated into CloudFormation for provisioning and management. This aligns perfectly with the question's requirements.

Authoritative Links for Further Research:

AWS CDK Official Documentation:<https://aws.amazon.com/cdk/>
AWS CloudFormation:<https://aws.amazon.com/cloudformation/>

Question: 52

A company is developing an application that uses multiple AWS services. The application needs to use temporary, limited-privilege credentials for authentication with other AWS APIs. Which AWS service or feature should the company use to meet these authentication requirements?

- A. Amazon API Gateway
- B. IAM users
- C. AWS Security Token Service (AWS STS)
- D. IAM instance profiles

Answer: C

Explanation:

The correct answer is C: AWS Security Token Service (AWS STS).

AWS STS is specifically designed for issuing temporary security credentials. These temporary credentials provide limited privileges and have an expiration time, making them ideal for applications requiring short-term access to AWS resources without relying on long-term IAM user credentials.

Here's why the other options are incorrect:

A. Amazon API Gateway: API Gateway is primarily used for creating, publishing, maintaining, monitoring, and securing APIs. While it can integrate with IAM for authentication, it doesn't inherently provide temporary credentials for application-level authentication to other AWS services.

B. IAM users: IAM users represent long-term identities within an AWS account. While they can be used for authentication, using long-term credentials directly within an application poses security risks (credential exposure). It's best practice to avoid embedding IAM user credentials directly in applications.

D. IAM instance profiles: IAM instance profiles provide temporary credentials to EC2 instances so applications running on those instances can make AWS API requests. While they provide temporary credentials, they are specifically for EC2 instances, not for general application-level authentication across different AWS services.

AWS STS addresses the need for temporary, least-privilege access by issuing tokens with specific permissions and a limited lifespan. This significantly reduces the risk of compromised credentials and improves overall security posture. Applications can assume roles defined in IAM using AWS STS to obtain these temporary credentials, allowing them to interact with other AWS services securely. In summary, AWS STS is the best choice because it's specifically built to handle the creation and management of short-lived, limited-privilege credentials for application authentication, fulfilling the requirements for temporary access

and enhanced security.

Supporting documentation:

[AWS Security Token Service \(STS\)](#): Overview of the AWS STS service.

[IAM Roles](#): Detailed explanation of IAM roles and how they relate to temporary security credentials.

[Using temporary security credentials to access AWS](#)

Question: 53

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:

The correct answer is A, AWS Security Hub. AWS Security Hub is specifically designed to be a cloud security posture management (CSPM) service. Its primary function is to aggregate security alerts and findings from multiple AWS services like Amazon GuardDuty, Amazon Inspector, and AWS IAM Access Analyzer, as well as from integrated third-party partner products. By consolidating these alerts into a single pane of glass, Security Hub provides a unified view of your security state within the AWS environment. This standardization allows security teams to prioritize and respond to the most pressing security issues more efficiently. Security Hub also performs automated security checks based on industry best practices and compliance standards (like CIS benchmarks and PCI DSS) to proactively identify potential vulnerabilities and misconfigurations in your AWS infrastructure. This makes it a cornerstone for improving overall security posture.

AWS Trusted Advisor, on the other hand, focuses on optimization across cost, performance, security, fault tolerance, and service limits, offering recommendations rather than aggregating alerts. Amazon EventBridge is an event bus service facilitating event-driven architectures. It doesn't inherently provide security posture management. Amazon GuardDuty is a threat detection service that continuously monitors for malicious activity and unauthorized behavior, but it is a source of security findings for Security Hub, rather than a CSPM service itself. Security Hub's aggregation, standardization, and compliance checking capabilities distinctly qualify it as a CSPM solution.

[AWS Security Hub Documentation](#) [CSPM Definition](#)

Question: 54

Which AWS service is always provided at no charge?

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

Answer: B

Explanation:

AWS Identity and Access Management (IAM) is provided at no charge, while other AWS services typically incur costs based on usage. IAM allows you to manage access to AWS services and resources securely. It enables you to create and manage AWS users and groups, and use permissions to allow and deny their access to AWS resources. Creating an IAM user or using its core functionalities such as multi-factor authentication, identity federation, and defining IAM roles does not incur a fee.

Amazon S3 charges for storage, data transfer, and requests. Elastic Load Balancers charge based on the amount of data processed and the time the load balancer is running. AWS WAF charges per web ACL and per rule, in addition to charges per request. IAM's free offering is crucial for establishing a secure foundation for AWS deployments and implementing the principle of least privilege. While using other AWS services protected by IAM will cost you, the IAM service itself is designed to be a fundamental building block at no cost. This allows all AWS customers to easily set up permissions to use the resources that are billed.

For further research, you can refer to the official AWS documentation on IAM pricing: <https://aws.amazon.com/iam/pricing/> and the general AWS Pricing Overview page: <https://aws.amazon.com/pricing/>. These resources provide more details on AWS service pricing models.

Question: 55

To reduce costs, a company is planning to migrate a NoSQL database to AWS. Which AWS service is fully managed and can automatically scale throughput capacity to meet database workload demands?

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

Answer: C

Explanation:

The correct answer is C, Amazon DynamoDB. The question focuses on a fully managed NoSQL database service that can automatically scale throughput capacity.

Amazon DynamoDB is a fully managed NoSQL database service offered by AWS. This means AWS handles the operational aspects like hardware provisioning, patching, and backups, relieving the company of these burdens.

DynamoDB excels at automatic scaling of throughput capacity. It can automatically adjust read and write capacity based on the workload demands, ensuring optimal performance without manual intervention. This feature helps the company reduce costs by only paying for the resources actually consumed.

Amazon Redshift (A) is a fully managed data warehouse service, suitable for analytical workloads involving large datasets and complex queries. It is not designed as a general-purpose NoSQL database.

Amazon Aurora (B) is a fully managed relational database engine compatible with MySQL and PostgreSQL. It offers high performance and availability but isn't a NoSQL database.

Amazon RDS (D) (Relational Database Service) supports relational databases like MySQL, PostgreSQL, Oracle, SQL Server, and MariaDB. While RDS simplifies database administration, it doesn't offer a NoSQL solution with the same automatic scaling capabilities as DynamoDB for NoSQL workloads.

Therefore, DynamoDB's NoSQL nature, fully managed status, and automatic scaling capabilities make it the ideal choice for the described scenario, aligning with the company's goal of cost reduction.

For further research, you can refer to the official AWS documentation:

Amazon DynamoDB:<https://aws.amazon.com/dynamodb/>

Question: 56

A company is using Amazon DynamoDB.

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

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

Answer: C

Explanation:

The correct answer is C, managing database access permissions, because under the AWS shared responsibility model, AWS manages the security of the cloud, while the customer is responsible for security in the cloud.

Let's break down why the other options are incorrect in the context of DynamoDB:

A. Patch the operating system: DynamoDB is a fully managed NoSQL database service. AWS handles the underlying infrastructure, including the operating system patching, eliminating this operational burden from the customer.

B. Provision hosts: As a managed service, AWS automatically provisions and manages the hosts required to run DynamoDB. Customers do not have direct access to the underlying infrastructure.

D. Secure the operating system: Again, because DynamoDB is a managed service, AWS is responsible for the security of the operating system on which DynamoDB runs.

In contrast, the customer is responsible for configuring and managing access control to their DynamoDB tables and data. This includes:

Defining IAM (Identity and Access Management) roles and policies that grant specific permissions to users and applications.

Managing authentication and authorization for accessing the database. Implementing fine-grained access control based on attributes or data content.

Securing database access permissions is a critical aspect of data security, aligning directly with the customer's responsibility for security in the cloud. This ensures that only authorized users and applications can access sensitive data stored in DynamoDB.

Authoritative Links:

AWS Shared Responsibility Model:<https://aws.amazon.com/compliance/shared-responsibility-model/> **Security in DynamoDB:**<https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/security.html>

Question: 57

A company has a test AWS environment. A company is planning on testing an application within AWS. The application testing can be interrupted and does not need to run continuously.

Which Amazon EC2 purchasing option will meet these requirements MOST cost-effectively?

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

Answer: C

Explanation:

The correct answer is C, Spot Instances. Here's why:

Spot Instances offer significant cost savings (up to 90% compared to On-Demand) because you bid on unused EC2 capacity. This is ideal for fault-tolerant, flexible applications.

The question states the application "can be interrupted" and "does not need to run continuously." This perfectly aligns with the nature of Spot Instances, which can be terminated if your bid price is lower than the current Spot price.

On-Demand Instances (A) are suitable for short-term, irregular workloads but are more expensive than Spot Instances. They offer no cost benefit for interruptible workloads.

Dedicated Instances (B) are physically isolated hardware dedicated to a single customer, making them the most expensive option. They are used when you have regulatory compliance.

Reserved Instances (D) are for predictable, long-term workloads that are not suitable for interruptible tasks.

Spot Instances are best for applications where the processing is flexible, and the application can be stopped and restarted.

Therefore, Spot Instances are the most cost-effective option for testing an application that can be interrupted and doesn't need to run continuously.

Reference: <https://aws.amazon.com/ec2/spot/>

Question: 58

Which AWS service gives users the ability to discover and protect sensitive data that is stored in Amazon S3 buckets?

- A. Amazon Macie
- B. Amazon Detective
- C. Amazon GuardDuty
- D. AWS IAM Access Analyzer

Answer: A

Explanation:

The correct answer is A, Amazon Macie. Macie is a fully managed data security and data privacy service that uses machine learning and pattern matching to discover and protect sensitive data stored in Amazon S3. It provides insights into your S3 data security posture, enabling you to identify and classify sensitive data, such as personally identifiable information (PII) or financial data.

Macie automates the process of discovering, classifying, and protecting sensitive data at scale. It continuously monitors S3 buckets for access control vulnerabilities and provides alerts when it detects potential security risks, such as publicly accessible buckets containing sensitive information. It offers features like automated sensitive data discovery jobs, custom data identifiers, and integration with AWS Security Hub for centralized security management. Macie can automatically alert you when sensitive data is stored in an S3 bucket, allowing you to take corrective action quickly.

Amazon Detective (B) analyzes log data to investigate security findings and conduct root cause analysis, focusing on identifying security incidents and suspicious activities. Amazon GuardDuty (C) provides threat detection by monitoring your AWS accounts and workloads for malicious activity. AWS IAM Access Analyzer (D) identifies the resources in your organization and accounts that are shared with an external entity. While these are important security tools, they don't provide the specific functionality of discovering and protecting sensitive data within S3 buckets like Amazon Macie does. Therefore, Macie is the only service among the options designed specifically for the stated purpose.

Further research:

Amazon Macie:<https://aws.amazon.com/macie/>

Amazon Detective:<https://aws.amazon.com/detective/>

Amazon GuardDuty:<https://aws.amazon.com/guardduty/>

AWS IAM Access Analyzer:<https://aws.amazon.com/iam/features/access-analyzer/>

Question: 59

Which of the following services can be used to block network traffic to an instance? (Choose 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:

The correct answer is A and C: Security Groups and Network ACLs (NACLs).

Security Groups act as virtual firewalls at the instance level, controlling inbound and outbound traffic. You associate security groups with EC2 instances, and they filter traffic based on rules you define that specify allowed protocols, ports, and source/destination IP addresses. These rules are stateful, meaning if you allow inbound traffic from a particular source, the return traffic is automatically allowed, regardless of outbound rules.

Network ACLs, on the other hand, operate at the subnet level. They control traffic entering and exiting subnets within a VPC. NACLs are stateless, meaning that rules must be explicitly defined for both inbound and outbound traffic. If you allow inbound traffic from a particular source, you must also define a rule allowing the return traffic. NACLs also use rules to determine whether traffic is allowed or denied based on protocol, port, and IP address ranges.

Option B, Amazon VPC flow logs, captures information about the IP traffic going to, from, and within your VPC. This is used for auditing and monitoring purposes, not for blocking traffic.

[<https://docs.aws.amazon.com/vpc/latest/userguide/flow-logs.html>]

Option D, Amazon CloudWatch, is a monitoring and observability service. It collects and tracks metrics, collects and monitors log files, and sets alarms. It does not block network traffic directly.

[<https://aws.amazon.com/cloudwatch/>]

Option E, AWS CloudTrail, records AWS API calls for your account and delivers log files to an Amazon S3 bucket. This is for auditing and governance, not for blocking network traffic.

[<https://aws.amazon.com/cloudtrail/>]

Therefore, Security Groups and Network ACLs are the two services designed to control and block network traffic to instances at the instance and subnet levels, respectively. Security Groups are stateful and operate at the instance level, while NACLs are stateless and operate at the subnet level.

Question: 60

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:

The correct answer is **B. AWS CloudTrail**. Here's a detailed justification:

AWS CloudTrail is a service that enables governance, compliance, operational auditing, and risk auditing of your AWS account. It does this by logging API calls made within your AWS environment. Whenever an action is taken, such as terminating an EC2 instance, CloudTrail records this event, including who performed the action, when it occurred, and from where it was initiated. This creates an auditable history of all activities within your AWS infrastructure.

Specifically, when an EC2 instance is terminated, the API call that initiates this termination is captured by CloudTrail as an event. The CloudTrail log will contain information detailing the instance ID that was terminated, the user or role that initiated the termination, the timestamp of the termination, and the AWS region where the instance was located. Therefore, by examining CloudTrail logs, one can definitively identify when an EC2 instance was terminated and gather details about the termination event.

Let's examine why the other options are incorrect:

A. AWS Identity and Access Management (IAM): IAM manages access to AWS services and resources. While IAM policies control who can terminate an EC2 instance, IAM itself does not record when an instance is terminated. It only deals with permissions.

C. AWS Compute Optimizer: Compute Optimizer analyzes your EC2 instance usage and provides recommendations for optimizing instance types to reduce costs and improve performance. It doesn't track instance termination events.

D. Amazon EventBridge: EventBridge is a serverless event bus that allows you to build event-driven applications. While you could configure EventBridge to react to an EC2 termination event (triggered via CloudTrail data events), EventBridge itself doesn't record the termination event. CloudTrail is the source of the termination event.

In summary, CloudTrail is the primary service designed to log and audit API calls, including EC2 instance

termination events, making it the correct choice for identifying when an instance was terminated.

Authoritative Links:

AWS CloudTrail:<https://aws.amazon.com/cloudtrail/>

Logging Amazon EC2 API calls with AWS CloudTrail:

<https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/logging-using-cloudtrail.html>

Question: 61

Which of the following is a fully managed MySQL-compatible database?

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

Answer: D

Explanation:

The correct answer is Amazon Aurora (D) because it is a fully managed, MySQL-compatible relational database engine. Amazon Aurora offers the performance and availability of commercial-grade databases with the simplicity and cost-effectiveness of open-source databases. Its architecture is designed for the cloud and integrates tightly with other AWS services, providing scalability, security, and reliability. Being "fully managed" means that AWS handles tasks like database setup, patching, backups, and recovery, freeing users from operational overhead.

Amazon S3 (A) is an object storage service, not a relational database. It's used for storing files and other unstructured data. <https://aws.amazon.com/s3/>

Amazon DynamoDB (B) is a NoSQL database service. While fully managed and highly scalable, it's not compatible with MySQL. It uses a different data model optimized for key-value and document storage. <https://aws.amazon.com/dynamodb/>

Amazon Redshift (C) is a data warehouse service optimized for analytical workloads, not transactional workloads like MySQL. It's based on PostgreSQL but designed for handling large datasets for business intelligence and data analysis. <https://aws.amazon.com/redshift/>

Amazon Aurora comes in two flavors: MySQL-compatible and PostgreSQL-compatible. Because the question specifically asks for a MySQL-compatible database, Aurora is the only valid choice among the options. The key here is understanding the distinction between different types of databases and the services that offer them in a managed capacity on AWS. Aurora's MySQL compatibility makes it a suitable option for applications that already utilize MySQL. <https://aws.amazon.com/rds/aurora/>

Question: 62

Which AWS service supports a hybrid architecture that gives users the ability to extend AWS infrastructure, AWS services, APIs, and tools to data centers, co-location environments, or on-premises facilities?

- A. AWS Snowmobile
- B. AWS Local Zones
- C. AWS Outposts

Answer: C

Explanation:

The correct answer is C, AWS Outposts.

AWS Outposts is specifically designed to extend AWS infrastructure and services to on-premises environments. It allows you to run AWS services, infrastructure, and management tools on your own hardware in your data center or co-location facility. This enables a truly consistent hybrid cloud experience where you can leverage AWS services while maintaining data residency, low latency, and regulatory compliance requirements on-premises. AWS Outposts offers pre-configured racks of compute and storage that AWS installs and manages for you.

Let's examine why the other options are incorrect:

A. AWS Snowmobile: AWS Snowmobile is an exabyte-scale data transfer service used to move massive amounts of data to AWS. While it's part of the AWS ecosystem, it's a data transfer service, not a hybrid cloud solution that extends AWS infrastructure.

B. AWS Local Zones: AWS Local Zones are extensions of AWS Regions that allow you to run latency-sensitive applications closer to your end-users. While Local Zones bring AWS services closer to users, they are still fully managed by AWS within AWS infrastructure and do not extend AWS infrastructure to on-premises locations.

D. AWS Fargate: AWS Fargate is a serverless compute engine for containers that works with both Amazon Elastic Container Service (ECS) and Amazon Elastic Kubernetes Service (EKS). It removes the need to provision and manage servers. Fargate is a compute option within AWS Regions and does not extend AWS infrastructure on-premises.

In essence, AWS Outposts is the only service among the options that allows you to run AWS infrastructure, services, and APIs directly within your own on-premises environment, creating a hybrid cloud architecture.

Further Research:

AWS Outposts: <https://aws.amazon.com/outposts/>
AWS Snowmobile: <https://aws.amazon.com/snowmobile/>
AWS Local Zones: <https://aws.amazon.com/localzones/>
AWS Fargate: <https://aws.amazon.com/fargate/>

Question: 63

Which AWS service can run a managed PostgreSQL database that provides online transaction processing (OLTP)?

- A. Amazon DynamoDB
- B. Amazon Athena
- C. Amazon RDS
- D. Amazon EMR

Answer: C

Explanation:

The correct answer is **C. Amazon RDS (Relational Database Service)**.

Amazon RDS is a managed database service that supports several database engines, including PostgreSQL. It is specifically designed for relational databases and is well-suited for online transaction processing (OLTP) workloads. OLTP involves managing transaction-oriented applications, typically for data entry and retrieval transaction processing. RDS simplifies database administration tasks like patching, backups, and scaling, allowing users to focus on application development.

Amazon DynamoDB (A) is a NoSQL database service. While DynamoDB is extremely scalable and fast, it's optimized for key-value and document data models, making it less suitable for traditional relational OLTP workloads.

Amazon Athena (B) is a serverless query service that allows you to analyze data in Amazon S3 using standard SQL. It is primarily designed for analytics and reporting, not OLTP. Athena excels at large-scale data processing and query execution, making it an Online Analytical Processing (OLAP) service.

Amazon EMR (D) is a managed Hadoop framework. It's used for big data processing and analysis, often involving batch processing of large datasets. EMR is not designed for OLTP database workloads.

Therefore, Amazon RDS is the only service listed that is a managed relational database service with PostgreSQL support optimized for OLTP applications.

Further research:

Amazon RDS:<https://aws.amazon.com/rds/>

Amazon RDS for PostgreSQL:<https://aws.amazon.com/rds/postgresql/>

OLTP vs OLAP:<https://www.databricks.com/glossary/oltp-vs-olap>

Question: 64

A company wants to provide managed Windows virtual desktops and applications to its remote employees over secure network connections.

Which AWS services can the company use to meet these requirements? (Choose two.)

- A. Amazon Connect
- B. Amazon AppStream 2.0
- C. Amazon WorkSpaces
- D. AWS Site-to-Site VPN
- E. Amazon Elastic Container Service (Amazon ECS)

Answer: CD

Explanation:

The correct answer is **C. Amazon WorkSpaces** and **D. AWS Site-to-Site VPN**.

Amazon WorkSpaces: Amazon WorkSpaces is a fully managed, secure Desktop-as-a-Service (DaaS) solution. It allows companies to provision virtual, cloud-based Windows desktops for their employees. This directly addresses the requirement of providing managed Windows virtual desktops and applications.<https://aws.amazon.com/workspaces/>

AWS Site-to-Site VPN: AWS Site-to-Site VPN establishes a secure connection between an on-premises network and an AWS VPC (Virtual Private Cloud). In the context of remote employees, this allows them to securely connect to the company's network resources within AWS, including the Amazon WorkSpaces, over an encrypted connection. This meets the need for secure network connections.<https://aws.amazon.com/vpn/site-to-site-vpn/>

In this scenario, employees likely use the internet to access AWS resources. A VPN is the best choice as it creates a secure, encrypted connection.

Here's why the other options are less suitable:

A. Amazon Connect: Amazon Connect is a cloud-based contact center service. While valuable for customer service operations, it's not directly involved in providing managed Windows virtual desktops and applications to remote employees.<https://aws.amazon.com/connect/>

B. Amazon AppStream 2.0: Amazon AppStream 2.0 allows you to stream desktop applications to users without installing them locally. While it can provide access to applications, it is not a complete Desktop-as-a-Service (DaaS) like WorkSpaces. Workspaces provides a persistent desktop for users while Appstream is aimed at streaming specific apps. WorkSpaces is the more holistic solution for the stated requirements.<https://aws.amazon.com/appstream2/>

E. Amazon Elastic Container Service (Amazon ECS): Amazon ECS is a container orchestration service. While ECS could potentially host parts of the infrastructure needed, it doesn't directly provide managed Windows virtual desktops and applications in the way WorkSpaces does.<https://aws.amazon.com/ecs/>

Therefore, Amazon WorkSpaces and AWS Site-to-Site VPN together provide the comprehensive solution required: managed Windows virtual desktops delivered securely to remote employees.

Question: 65

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:

The correct answer is **A. AWS Trusted Advisor**. Here's why:

AWS Trusted Advisor is a service that acts as a cloud expert, providing recommendations to optimize your AWS infrastructure for security, cost optimization, performance, and fault tolerance. A key function of Trusted Advisor is to identify security vulnerabilities, including overly permissive security group rules. It specifically checks for security groups that have rules allowing unrestricted (0.0.0.0/0 or ::/0) access to ports like 22 (SSH), 3389 (RDP), and other commonly targeted ports. Trusted Advisor then generates alerts and recommendations, enabling the company to remediate these misconfigurations and improve their security posture. This aligns directly with the scenario's requirement of monitoring for misconfigured security groups allowing unrestricted access to specific ports.

While the other options offer valuable services, they don't directly address the stated need in the same way as Trusted Advisor. Amazon CloudWatch (B) is primarily a monitoring and observability service focused on metrics, logs, and events. While you could potentially create custom metrics and alarms to detect security group changes, it wouldn't provide the same built-in, automated security checks as Trusted Advisor. Amazon GuardDuty (C) is a threat detection service that analyzes CloudTrail logs, VPC Flow Logs, and DNS logs to identify malicious activity. Although it can detect compromised EC2 instances potentially resulting from a

misconfigured security group, it's a reactive threat detection service rather than a proactive configuration monitoring tool. Finally, AWS Health Dashboard (D) provides information about the general health of AWS services and resources, not specific misconfigurations within your account. Therefore, Trusted Advisor is the most appropriate service for the company's requirement.

For further research:

AWS Trusted Advisor:<https://aws.amazon.com/premiumsupport/technology/trusted-advisor/>

Question: 66

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

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

Answer: A

Explanation:

Amazon DynamoDB is the correct answer because it is a fully managed NoSQL database service renowned for its speed and scalability. DynamoDB is explicitly designed as a key-value and document database, making it ideally suited for applications requiring ultra-fast access to data. It offers single-digit millisecond latency at any scale, making it perfect for use cases like session management, gaming leaderboards, and shopping carts. Its ability to handle high-traffic volumes and large datasets without performance degradation is a key differentiator.

Amazon Aurora is a MySQL and PostgreSQL-compatible relational database. While it offers high performance for relational workloads, it is not inherently a key-value store and doesn't provide the same sub-millisecond latency for key-value operations as DynamoDB. Amazon DocumentDB (with MongoDB compatibility) is a document database service that provides MongoDB compatibility, enabling you to use existing MongoDB drivers and tools. While document databases are a form of NoSQL database, DocumentDB is not specifically optimized for sub-millisecond key-value access in the same way that DynamoDB is. Amazon Neptune is a graph database service that is optimized for complex relationships and graph traversals. It's not a key-value database, and its focus is on querying relationships between data points rather than simple key-value lookups.

Therefore, given the requirement for a key-value database with sub-millisecond latency at scale, Amazon DynamoDB stands out as the most appropriate and efficient solution. Its architecture and design choices are specifically geared towards this type of performance and workload.

Further reading:

[Amazon DynamoDB](#)
[NoSQL Databases on AWS](#)

Question: 67

A company is deploying a machine learning (ML) research project that will require a lot of compute power over several months. The ML processing jobs do not need to run at specific times. Which Amazon EC2 instance purchasing option will meet these requirements at the lowest cost?

- A. On-Demand Instances

- B.Spot Instances
- C.Reserved Instances
- D.Dedicated Instances

Answer: B

Explanation:

The correct answer is B, Spot Instances. Here's a detailed justification:

Spot Instances offer the lowest cost for Amazon EC2 compute capacity compared to On-Demand, Reserved, or Dedicated Instances. They achieve this by allowing users to bid on unused EC2 capacity. Because the company's ML processing jobs don't need to run at specific times, the potential for interruptions with Spot Instances is acceptable. If a Spot Instance is terminated due to a higher bid from another user, the processing can be resumed when the Spot Instance price falls below the user's bid again.

On-Demand Instances (Option A) provide compute capacity by the hour or second with no long-term commitments. While flexible, they are the most expensive option, unsuitable when cost optimization is the primary concern and jobs are flexible in execution.

Reserved Instances (Option C) offer a discount in exchange for a commitment to use EC2 instances for 1 or 3 years. Since the project lasts only several months and the company is researching, the long-term commitment may not be ideal and upfront costs could be prohibitive. Reserved Instances are best when you have predictable usage patterns.

Dedicated Instances (Option D) run in a VPC on hardware dedicated to a single customer. They are the most expensive option and typically used for compliance or regulatory reasons or specific licensing agreements. They do not offer cost optimization benefits. Since no specific compliance requirements were mentioned, this option would be least preferable for an ML project.

For ML research projects requiring high compute power without time constraints, Spot Instances are the most cost-effective, allowing the workload to be interrupted and resumed as Spot prices fluctuate, while On-Demand, Reserved, and Dedicated instances either cost more, require time commitments, or dedicated hardware respectively.

Authoritative Links:

Amazon EC2 Spot Instances: <https://aws.amazon.com/ec2/spot/>
Amazon EC2 Pricing: <https://aws.amazon.com/ec2/pricing/>

Question: 68

Which AWS services or features provide disaster recovery solutions for Amazon EC2 instances? (Choose two.)

- A.EC2 Reserved Instances
- B.EC2 Amazon Machine Images (AMIs)
- C.Amazon Elastic Block Store (Amazon EBS) snapshots
- D.AWS Shield
- E.Amazon GuardDuty

Answer: BC

Explanation:

The correct answer is B and C: EC2 Amazon Machine Images (AMIs) and Amazon Elastic Block Store (Amazon

EBS) snapshots are critical components in a disaster recovery strategy for Amazon EC2 instances.

EC2 AMIs: AMIs serve as templates containing the operating system, application server, and applications needed to launch an instance. In a disaster scenario, you can quickly launch new EC2 instances from a recent AMI in a different AWS Region, effectively restoring your application. By regularly backing up your AMIs and storing them in multiple Regions, you ensure that you can recover your EC2 instances should an AWS Region become unavailable. An AMI encapsulates the configuration of an instance, enabling its rapid recreation in a disaster recovery scenario.

Amazon EBS Snapshots: EBS volumes are block storage devices attached to EC2 instances. EBS snapshots are incremental backups of your EBS volumes stored in Amazon S3. They allow you to restore the data on those volumes to a point in time. In a disaster recovery scenario, you can use EBS snapshots to recreate EBS volumes in a different AWS Region and then attach them to newly launched EC2 instances, restoring your data. Consistent backups via EBS snapshots are essential for maintaining data integrity and availability.

Why other options are incorrect:

A. EC2 Reserved Instances: Reserved Instances are a billing discount; they don't provide disaster recovery capabilities directly. They pre-commit to usage to achieve cost savings but do not offer backups or failover solutions.

D. AWS Shield: AWS Shield provides protection against Distributed Denial of Service (DDoS) attacks. While important for security, it's not a disaster recovery solution for EC2 instances. It protects from malicious traffic, not service outages.

E. Amazon GuardDuty: Amazon GuardDuty is a threat detection service that monitors for malicious activity and unauthorized behavior. While important for security, it doesn't provide disaster recovery capabilities. It identifies threats but doesn't handle restoration of services.

Authoritative Links for Further Research:

Amazon EC2 AMIs: <https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/AMIs.html>

Amazon EBS Snapshots: <https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/EBSSnapshots.html>

Question: 69

Which AWS service provides command line access to AWS tools and resources directly from a web browser?

- A. AWS CloudHSM
- B. AWS CloudShell
- C. Amazon WorkSpaces
- D. AWS Cloud Map

Answer: B

Explanation:

The correct answer is B, AWS CloudShell. AWS CloudShell provides a browser-based, pre-authenticated shell accessible from the AWS Management Console. This allows users to manage and interact with AWS resources directly through a command-line interface, without needing to install or manage separate command-line tools on their local machines. CloudShell comes pre-configured with commonly used tools like the AWS CLI, Python, and Git, making it a convenient and ready-to-use environment for tasks such as deploying applications, managing infrastructure, and exploring AWS services. The integrated AWS CLI is pre-authenticated with the user's IAM credentials, eliminating the need for manual configuration of AWS credentials. AWS CloudHSM (A) is a cloud-based hardware security module that allows you to generate, store, and manage encryption keys securely. Amazon WorkSpaces (C) is a fully managed desktop

virtualization service. AWS Cloud Map (D) is a cloud resource discovery service. Therefore, only AWS CloudShell offers direct command-line access via a web browser.

<https://aws.amazon.com/cloudshell/>

Question: 70

A network engineer needs to build a hybrid cloud architecture connecting on-premises networks to the AWS Cloud using AWS Direct Connect. The company has a few VPCs in a single AWS Region and expects to increase the number of VPCs to hundreds over time.

Which AWS service or feature should the engineer use to simplify and scale this connectivity as the VPCs increase in number?

- A.VPC endpoints
- B.AWS Transit Gateway
- C.Amazon Route 53
- D.AWS Secrets Manager

Answer: B

Explanation:

The correct answer is **B. AWS Transit Gateway**. Here's a detailed justification:

AWS Transit Gateway is a network transit hub that simplifies connecting multiple VPCs and on-premises networks. As the company expands from a few VPCs to hundreds, a direct connection from each on-premises network to every VPC becomes increasingly complex and unmanageable. This creates a full-mesh network, requiring numerous individual connections and routing configurations.

AWS Transit Gateway simplifies this by acting as a central hub. Instead of creating point-to-point connections between each VPC and the on-premises network, you connect each VPC and the AWS Direct Connect connection to the Transit Gateway. This effectively centralizes routing and management, reducing the operational overhead.

Option A, VPC endpoints, are used to privately connect to AWS services without traversing the public internet, and are not designed for connecting multiple VPCs or on-premises networks. Option C, Amazon Route 53, is a DNS service and does not provide network connectivity between VPCs and on-premises. Option D, AWS Secrets Manager, is for securely storing and rotating secrets and is not relevant to network connectivity.

Therefore, AWS Transit Gateway is the most suitable solution for simplifying and scaling the hybrid cloud architecture by providing a single, centralized point for connecting multiple VPCs and on-premises networks via AWS Direct Connect. It simplifies routing, reduces administrative overhead, and scales effectively as the number of VPCs increases, making it the ideal choice for this scenario.

For more information, refer to the official AWS documentation:

AWS Transit Gateway: <https://aws.amazon.com/transit-gateway/>
AWS Direct Connect: <https://aws.amazon.com/directconnect/>

Question: 71

A company wants to assess its operational readiness. It also wants to identify and mitigate any operational risks ahead of a new product launch.

Which AWS Support plan offers guidance and support for this kind of event at no additional charge?

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

Answer: D

Explanation:

The correct answer is AWS Enterprise Support (D). Here's why:

Operational Readiness Reviews (ORR): Enterprise Support includes access to ORRs, which are specifically designed to help customers proactively assess their operational readiness for new product launches or significant changes. This aligns perfectly with the company's goal.

Proactive Guidance: The Enterprise Support plan provides proactive guidance from designated Technical Account Managers (TAMs). TAMs work closely with customers to understand their business needs and identify potential operational risks.

Risk Mitigation: Through ORRs and TAM guidance, Enterprise Support helps customers identify and mitigate potential risks before they impact production environments. This proactive approach is crucial for a successful product launch.

Other Support Plans:

Basic Support: This plan offers limited support and doesn't include ORRs or proactive guidance.

Developer Support: While Developer Support offers technical assistance, it lacks the strategic guidance and proactive assessments of Enterprise Support.

Business Support: Business Support offers broader coverage than Developer, but Enterprise Support stands out with features like ORR and designated TAMs.

Cost Considerations: While the question states "at no additional charge," this implicitly refers to the features included within the different support plans. Enterprise Support has a higher subscription cost overall, but the ORR benefit is already part of that subscription, not an additional charge.

Authoritative Links:

AWS Support Plans: <https://aws.amazon.com/premiumsupport/plans/>

AWS Enterprise Support: <https://aws.amazon.com/premiumsupport/enterprise/>

Question: 72

A company wants to establish a schedule for rotating database user credentials.

Which AWS service will support this requirement with the LEAST amount of operational overhead?

- A.AWS Systems Manager
- B.AWS Secrets Manager
- C.AWS License Manager
- D.AWS Managed Services

Answer: B

Explanation:

The correct answer is **B. AWS Secrets Manager**.

Here's a detailed justification:

AWS Secrets Manager is specifically designed for managing secrets, including database credentials, API keys, and other sensitive information. It allows you to securely store, rotate, and retrieve secrets throughout their lifecycle. A key feature of Secrets Manager is its ability to automate the rotation of credentials on a schedule you define. This addresses the company's requirement for rotating database user credentials.

By automating rotation, Secrets Manager significantly reduces the operational overhead compared to manually rotating credentials or using other services not specifically designed for secret management. It integrates directly with various AWS services and databases, streamlining the rotation process.

AWS Systems Manager (option A) provides operational insights and management capabilities, including patch management, configuration management, and automation. While it can be used to script a rotation process, it requires significantly more manual configuration and operational overhead than Secrets Manager, which offers built-in rotation functionality.

AWS License Manager (option C) is for managing software licenses, and AWS Managed Services (option D) provides ongoing management of your AWS infrastructure. Neither of these services are directly relevant or designed for rotating database credentials.

Therefore, AWS Secrets Manager is the most appropriate choice because it minimizes operational overhead by providing a dedicated service for secure storage, management, and automated rotation of database credentials, directly fulfilling the given requirement.

For further research, refer to the official AWS documentation:

AWS Secrets Manager:<https://aws.amazon.com/secrets-manager/>
Rotating AWS Secrets Manager secrets:
<https://docs.aws.amazon.com/secretsmanager/latest/userguide/rotating-secrets.html>

Question: 73

Which AWS service or feature can be used to create a private connection between an on-premises workload and an AWS Cloud workload?

- A. Amazon Route 53
- B. Amazon Macie
- C. AWS Direct Connect
- D. AWS PrivateLink

Answer: C

Explanation:

The correct answer is C: AWS Direct Connect. Here's why:

AWS Direct Connect establishes a dedicated network connection between your on-premises environment (like a data center or office) and AWS. This dedicated connection bypasses the public internet, offering more consistent network performance, lower latency, and increased security when compared to internet-based connections. This makes it ideal for scenarios where you need to reliably and securely transfer large amounts of data or run latency-sensitive applications between your on-premises environment and AWS.

Amazon Route 53 (Option A) is a scalable DNS (Domain Name System) web service. While it can be used to manage domain names and route traffic to AWS resources, it doesn't provide a dedicated, private connection between on-premises and AWS.

Amazon Macie (Option B) is a fully managed data security and data privacy service that uses machine learning and pattern matching to discover and protect sensitive data in AWS. It doesn't provide a network connection.

AWS PrivateLink (Option D) provides private connectivity between VPCs, AWS services, and your on-premises networks, without exposing your traffic to the public internet. While it achieves a similar goal of private connectivity, PrivateLink focuses primarily on connecting to AWS services within the AWS network, rather than directly connecting your on-premises network. Also, it often requires Direct Connect or VPN as the underlying connection from on-prem. Direct Connect is a more fundamental way to establish a private dedicated network connection.

Therefore, AWS Direct Connect is the most appropriate service for creating a private and dedicated connection between an on-premises workload and an AWS Cloud workload.

Further research:

AWS Direct Connect:<https://aws.amazon.com/directconnect/>

AWS PrivateLink:<https://aws.amazon.com/privatelink/>

Question: 74

Which AWS service is used to provide encryption for Amazon EBS?

- A. AWS Certificate Manager
- B. AWS Systems Manager
- C. AWS KMS
- D. AWS Config

Answer: C

Explanation:

The correct answer is AWS Key Management Service (KMS). AWS KMS is a managed service that makes it easy for you to create and control the cryptographic keys used to encrypt your data. Specifically for Amazon Elastic Block Storage (EBS) volumes, AWS KMS is the recommended and commonly used service for encryption.

AWS KMS allows you to generate, store, and manage encryption keys. When you enable encryption for an EBS volume, you can specify a KMS key to use. This key can be an AWS managed key, a customer managed key (CMK) stored in KMS, or a key from an external key management system integrated with KMS via a custom key store.

Using KMS for EBS encryption provides several benefits:

1. **Simplified Key Management:** KMS handles the complexities of key generation, rotation, and storage.
2. **Security:** Keys are protected by hardware security modules (HSMs) that are FIPS 140-2 Level 3 validated.
3. **Control:** You have granular control over who can access and use the keys.
4. **Compliance:** Helps you meet compliance requirements related to data encryption.
5. **Integration:** Seamlessly integrates with other AWS services like EBS.

The other options are not suitable for EBS encryption:

AWS Certificate Manager (ACM): Primarily used for managing SSL/TLS certificates for secure communication.

AWS Systems Manager: Primarily for automating operational tasks across your AWS resources, not

encryption.

AWS Config: Primarily for assessing, auditing, and evaluating the configurations of your AWS resources, not encryption.

Therefore, AWS KMS is the only service listed that provides encryption capabilities for Amazon EBS.

References:

[AWS KMS Documentation](#)
[Amazon EBS Encryption](#)

Question: 75

A company wants to manage its AWS Cloud resources through a web interface. Which AWS service will meet this requirement?

- A. AWS Management Console
- B. AWS CLI
- C. AWS SDK
- D. AWS Cloud9

Answer: A

Explanation:

The correct answer is A, AWS Management Console. The question focuses on a company's need for a web interface to manage AWS resources. The AWS Management Console is a web-based interface specifically designed for this purpose. It allows users to access and manage all AWS services through a graphical user interface. Users can create, configure, monitor, and manage their AWS resources like EC2 instances, S3 buckets, and databases via the console.

Option B, AWS CLI (Command Line Interface), is a command-line tool that requires users to interact with AWS services through terminal commands, which isn't a web interface. Option C, AWS SDK (Software Development Kit), provides libraries for developers to interact with AWS services programmatically using various programming languages, but it's not a direct interface for general management. Option D, AWS Cloud9, is a cloud-based integrated development environment (IDE), primarily for coding and debugging applications, not for general resource management through a web interface. Thus, only the AWS Management Console directly fulfills the requirement of managing AWS cloud resources via a web-based interface. The AWS Management Console offers a centralized, user-friendly platform for overseeing the entire AWS infrastructure.

Reference:

AWS Management Console: <https://aws.amazon.com/console/>

Question: 76

Which of the following are advantages of the AWS Cloud? (Choose two.)

- A. Trade variable expenses for capital expenses
- B. High economies of scale
- C. Launch globally in minutes
- D. Focus on managing hardware infrastructure
- E. Overprovision to ensure capacity

Answer: BC

Explanation:

The correct answer is B and C because they reflect key advantages offered by the AWS Cloud.

B. High economies of scale: AWS achieves economies of scale by aggregating the demand from hundreds of thousands of customers. This massive scale allows AWS to reduce costs significantly and pass the savings on to customers in the form of lower prices. AWS can optimize hardware utilization, negotiate better deals with vendors, and automate many operational tasks, leading to lower costs per unit of compute, storage, and other resources. This enables customers to benefit from a lower total cost of ownership (TCO) compared to managing their own infrastructure.

<https://aws.amazon.com/economics/>

C. Launch globally in minutes: AWS has a global infrastructure consisting of Regions and Availability Zones. This allows customers to deploy their applications and services in multiple geographic locations quickly and easily. The ability to launch globally in minutes enables businesses to reach a wider audience, improve performance by reducing latency for users around the world, and ensure business continuity by distributing workloads across multiple regions. This global reach and speed of deployment is a significant advantage over traditional on-premises infrastructure.

<https://aws.amazon.com/about-aws/global-infrastructure/>

Now, let's look at why the other options are incorrect:

A. Trade variable expenses for capital expenses: This statement is incorrect. The AWS Cloud allows you to trade capital expenses (CapEx) for variable expenses (OpEx). Instead of investing heavily in data centers and servers before knowing how you're going to use them, you only pay when you consume computing resources.

D. Focus on managing hardware infrastructure: This is the opposite of what AWS offers. AWS handles the hardware infrastructure management, allowing customers to focus on their core business and innovation. This is a major benefit of cloud computing.

E. Overprovision to ensure capacity: A key advantage of AWS is the ability to scale resources on demand. Instead of overprovisioning, you can dynamically adjust resources as needed, optimizing cost and performance.

Question: 77

Which AWS Cloud benefit is shown by an architecture's ability to withstand failures with minimal downtime?

- A. Agility
- B. Elasticity
- C. Scalability
- D. High availability

Answer: D

Explanation:

The correct answer is D, High Availability. High availability in cloud computing refers to a system's ability to remain operational and accessible, even when faced with failures. It is achieved through redundancy, fault tolerance, and automated recovery mechanisms. The architecture's ability to withstand failures and minimize downtime directly aligns with the core principle of high availability.

Agility (A) refers to the ability to rapidly innovate and deploy new services, which is not the primary concern when addressing system failures. Elasticity (B) refers to the system's ability to automatically scale resources up or down based on demand, which, while helpful for handling load spikes, doesn't directly address failure

recovery. Scalability (C) refers to the ability to handle increased workloads by adding resources, which is also related to capacity but doesn't inherently guarantee uninterrupted service during failures.

High availability directly addresses the scenario described in the question, whereas the other options are related but not the most pertinent. A highly available system is designed to minimize single points of failure and automatically reroute traffic or switch to redundant resources in the event of a failure. This ensures that users experience minimal disruption, upholding service continuity.

For further research, refer to the AWS Well-Architected Framework for reliability, which focuses on building systems that can withstand failures. You can find more information at:

<https://wa.aws.amazon.com/wellarchitected/2020-07-02T19-33-23/reliability/rel01-design-for-high-availability.en.html> and the AWS documentation on high availability: <https://aws.amazon.com/reliability/>.

Question: 78

A developer needs to maintain a development environment infrastructure and a production environment infrastructure in a repeatable fashion.

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

- A. AWS Ground Station
- B. AWS Shield
- C. AWS IoT Device Defender
- D. AWS CloudFormation

Answer: D

Explanation:

The correct answer is D, AWS CloudFormation. CloudFormation allows developers to define infrastructure as code using templates. These templates can be version-controlled and reused to consistently provision and manage infrastructure across multiple environments, such as development and production. This ensures a repeatable and predictable infrastructure setup. Options A, B, and C are incorrect because they do not address the requirement of repeatable infrastructure provisioning. AWS Ground Station is for controlling satellite communication, AWS Shield provides DDoS protection, and AWS IoT Device Defender manages IoT device security. CloudFormation's infrastructure-as-code approach enables automated deployment, configuration, and management of AWS resources, addressing the need for repeatability, versioning, and consistent environments for development and production, directly aligning with DevOps best practices.

Utilizing CloudFormation promotes infrastructure consistency, reduces manual configuration errors, and accelerates the deployment process.

Further Research:

[AWS CloudFormation Documentation Infrastructure as Code \(IaC\)](#)

Question: 79

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

- A. Maintain the security of the AWS Cloud.
- B. Configure firewalls and networks.
- C. Patch the operating system of Amazon RDS instances.

D.Implement physical and environmental controls.

Answer: B

Explanation:

The correct answer is B: Configure firewalls and networks. The AWS Shared Responsibility Model delineates the security and operational responsibilities between AWS and the customer. AWS is responsible for the security of the cloud, encompassing the physical infrastructure, global networking, hardware, and software that supports AWS services. This includes securing data centers, managing underlying infrastructure, and ensuring the availability of services.

The customer is responsible for the security in the cloud. This translates to managing the security of their data, applications, operating systems, network configurations (including firewalls), identity and access management (IAM), and client-side data encryption. In essence, the customer controls what they put into the cloud and how they configure it.

Option A is incorrect because maintaining the security of the AWS Cloud is solely AWS's responsibility. Option C is also incorrect because patching the operating system of Amazon RDS instances is AWS's responsibility, as RDS is a managed database service. AWS manages the underlying OS and database software, while the customer manages database-specific configurations and security settings within the database itself. Option D is incorrect because implementing physical and environmental controls falls under AWS's purview, concerning the physical security of their data centers.

Configuring firewalls and networks (Option B) directly pertains to the customer defining the security parameters and access rules for their resources deployed within the AWS environment. They control the inbound and outbound traffic and must ensure that only authorized access is granted. This is a crucial aspect of securing their cloud environment and is squarely within the customer's sphere of responsibility.

Further reading on the AWS Shared Responsibility Model can be found here:

[AWS Shared Responsibility Model](#)
[AWS Documentation on Security](#)

Question: 80

Which AWS service helps deliver highly available applications with fast failover for multi-Region and Multi-AZ architectures?

- A.AWS WAF
- B.AWS Global Accelerator
- C.AWS Shield
- D.AWS Direct Connect

Answer: B

Explanation:

The correct answer is B, AWS Global Accelerator. Here's why:

AWS Global Accelerator is specifically designed to improve the availability and performance of applications for global users. It does this by directing traffic to optimal endpoints based on user location, health check results, and configured traffic policies. It uses the AWS global network to route traffic intelligently to the closest healthy application endpoint, providing low latency and high availability.

Global Accelerator provides static IP addresses that act as a fixed entry point to your applications. This abstraction enables seamless failover to healthy endpoints in different AWS Regions or Availability Zones (AZs) without requiring users to update their DNS settings or client configurations. If one endpoint fails, Global Accelerator automatically redirects traffic to another healthy endpoint within seconds. This is crucial for multi-region and multi-AZ architectures aiming for fast failover and minimal disruption.

AWS WAF (A) is a web application firewall that helps protect your web applications from common web exploits and bots. While important for security, it doesn't inherently provide fast failover or route traffic across multiple regions for availability purposes.

AWS Shield (C) provides protection against Distributed Denial of Service (DDoS) attacks. It helps keep your applications available during such attacks but isn't focused on enabling multi-region failover in the same way as Global Accelerator.

AWS Direct Connect (D) establishes a dedicated network connection from your on-premises environment to AWS. While it can improve network performance, it doesn't inherently provide the intelligent traffic routing and fast failover capabilities of Global Accelerator. It's used for establishing a private connection, not necessarily high availability and failover between multiple AWS Regions or AZs.

Therefore, AWS Global Accelerator is the most suitable service for delivering highly available applications with fast failover in multi-Region and multi-AZ architectures.

Here are some authoritative links for further research:

AWS Global Accelerator:<https://aws.amazon.com/global-accelerator/>

AWS Global Accelerator Documentation:<https://docs.aws.amazon.com/global-accelerator/index.html>

MY EXAM.FK