

QUESTIONNAIRE REPORT SHEET



iPQ2AWS CCP-CLF.V1.1




AWS CERTIFIED CLOUD PRACTITIONER

CLF-C01 PRACTICE TESTS

iPQ.No	Correct Answer	Points	Explanation
2AWSCCP-CLF01	3	15.4	Answer = 3 Both A and B Explanation : AWS Auto Scaling: AWS Auto Scaling monitors your applications and adjusts the capacity automatically to maintain consistent, predictable performance at the lowest possible cost. Using AWS Auto Scaling, it is simple to scale an application across multiple resources and services in minutes. AWS offers auto-scaling at no additional cost. You only pay for the AWS resources required to run your applications, as well as the Amazon CloudWatch monitoring fees. AWS Identity and Access Management (AWS IAM) enables the secure management of access to AWS services and resources. Using IAM, you can create and manage AWS users and groups and grant or deny their access to AWS resources using permissions. IAM is a feature of your AWS account that is provided at no extra cost.

iPQ.No	Correct Answer	Points	Explanation
2AWSCCP-CLF02	3	15.4	Answer = 3 Both A and B Explanation : AWS Direct Connect is a cloud service that facilitates the establishment of a dedicated network connection between your premises and AWS. Using AWS Direct Connect, you can establish private connectivity between AWS and your datacenter, office, or colocation environment, which in many instances can reduce network costs, increase bandwidth throughput, and provide a more consistent network experience than Internet-based connections. Using AWS Direct Connect:  The diagram illustrates the five steps for setting up AWS Direct Connect: 1. Decide on an AWS Direct Connect location and connection size. 2. Create your connection request(s) on the AWS management console. 3. Download Letter of Authorization (LoA) from the AWS Management console. 4. Provide LoA to an APN Partner and ask them to establish the connection on your behalf. 5. Configure virtual interfaces to establish network connectivity. https://aws.amazon.com/directconnect/ AWS Virtual Private Network (VPN): AWS Virtual Private Network (VPN) solutions create secure connections between local networks, remote offices, client devices, and the AWS global network. AWS VPN consists of two services: AWS Site-to-Site VPN and AWS Client VPN. How AWS Client VPN Operates:  The diagram illustrates the six steps for setting up AWS Client VPN: 1. Configure the Client VPN endpoint. 2. Associate subnet(s) to the Client VPN endpoint. 3. Enable end user access to resources by adding authorization rules. 4. Enable connectivity to other networks by adding rules and routes. 5. Download and install the free VPN client on users' devices. 6. Users enter their corporate login credentials to connect to your network over AWS Client VPN. via https://aws.amazon.com/vpn/


iPQ.No	Correct Answer	Points	Explanation
2AWSCCP-CLF03	1	15.4	Answer = 1 Managing guest operating system patches on Amazon Elastic Compute Cloud (Amazon EC2) Explanation : AWS and the customer share responsibility for security and compliance. AWS operates, manages, and controls the components, from the host operating system and virtualization layer to the physical security of the service's operating facilities, in this shared model. In accordance with the AWS shared responsibility model, the customer is responsible for security "in" the cloud. Customers who deploy an Amazon EC2 instance are responsible for managing the guest operating system (including updates and security patches), any application software or utilities installed by the customer on the instances, and the configuration of the AWS-provided firewall (referred to as a security group) on each instance. Exam Alert: Please review the shared responsibility model in depth, as there will be multiple questions on it on the exam.  The diagram illustrates the AWS Shared Responsibility Model. It is divided into two main sections: CUSTOMER (Responsibility for security 'in' the cloud) and AWS (Responsibility for security 'of' the cloud). The CUSTOMER section includes: CUSTOMER DATA, PLATFORM, APPLICATIONS, IDENTITY & ACCESS MANAGEMENT, OPERATING SYSTEM, NETWORK & FIREWALL CONFIGURATION, CLIENT-SIDE DATA ENCRYPTION & DATA INTEGRITY AUTHENTICATION, SERVER-SIDE ENCRYPTION (FILE SYSTEM AND/OR DATA), and NETWORKING TRAFFIC PROTECTION (ENCRYPTION, INTEGRITY, IDENTITY). The AWS section includes: SOFTWARE (COMPUTE, STORAGE, DATABASE, NETWORKING), and HARDWARE/AWS GLOBAL INFRASTRUCTURE (REGIONS, AVAILABILITY ZONES, EDGE LOCATIONS). via - https://aws.amazon.com/compliance/shared-responsibility-model/

iPQ.No	Correct Answer	Points	Explanation
2AWSCCP-CLF04	3	15.4	Answer = 3 Both A and B Explanation : Amazon CloudWatch is a monitoring and observability service designed for DevOps engineers, developers, and site reliability engineers (SREs). CloudWatch provides data and actionable insights for monitoring applications, responding to system-wide performance changes, optimizing resource utilization, and obtaining a consolidated view of operational health. AWS Lambda: AWS Lambda enables the execution of code without the need to provision or manage servers. You only pay for the amount of computing time you use. The maximum execution time of the lambda is 15 minutes, so it can be used to execute this log backup process. To develop a solution for the given use case, you can create a CloudWatch Events rule with a cron expression that triggers on a schedule. The lambda can then be set as the target for this rule.


iPQ.No	Correct Answer	Points	Explanation
2AWSCCP-CLF05	2	15.4	Answer = 2 AWS Local Zones Explanation : AWS Local Zones enable the use of select AWS services, such as compute and storage services, closer to more end-users, thereby providing them with extremely low-latency access to locally-running applications. AWS Local Zones are also connected to the parent region via Amazon's redundant and extremely high-bandwidth private network, providing applications running in AWS Local Zones with quick, secure, and seamless access to the remaining AWS services. For low-latency requirements, you should use AWS Local Zones to deploy workloads closer to your end users. AWS Local Zones have their own connection to the internet and support AWS Direct Connect, allowing resources created in a Local Zone to communicate with local end-users with minimal latency. AWS Local Zones provide access to a variety of AWS services, including Amazon Elastic Compute Cloud (EC2), Amazon Virtual Private Cloud (VPC), Amazon Elastic Block Store (EBS), Amazon FSx, Amazon Elastic Load Balancing, Amazon EMR, Amazon ElastiCache, and Amazon Relational Database Service (RDS). You can also use Amazon EC2 Auto Scaling, Amazon EKS clusters, Amazon ECS clusters, Amazon EC2 Systems Manager, Amazon CloudWatch, AWS CloudTrail, and AWS CloudFormation to orchestrate or work with local services. AWS Local Zones also provide a high-bandwidth, secure connection to the AWS Region, enabling you to connect seamlessly to the complete set of services in the AWS Region using the same APIs and tools.


iPQ.No	Correct Answer	Points	Explanation
2AWSCCP-CLF06	3	15.4	Answer = 3 Both A and B Explanation : AWS Lambda is a compute service that enables code execution without server provisioning or management. AWS Lambda executes your code only when necessary and automatically scales from a few daily requests to thousands per second. When your Lambda function is not executing, you incur no costs. AWS Lambda enables the execution of code for virtually any type of application or backend service with no administration required. AWS Lambda executes your code on a high-availability compute infrastructure and manages the compute resources, including server and operating system maintenance, capacity provisioning and automatic scaling, code monitoring and logging. How AWS Lambda Operates:  aws.amazon.com/lambda

iPQ.No	Correct Answer	Points	Explanation																																								
2AWS CCP-CLF07	1	15.4	Answer = 1 AWS Enterprise Support Explanation : AWS provides its customers with four distinct support plans: AWS Developer Support, AWS Business Support, AWS Enterprise On-Ramp Support, and AWS Enterprise Support. Every AWS customer receives a basic support plan. AWS Enterprise Support offers customers a concierge-like service with a primary focus on assisting customers in achieving their goals and achieving success in the cloud. You receive 24x7 technical support from high-quality engineers, tools and technology to automatically manage the health of your environment, consultative architectural guidance delivered in the context of your applications and use-cases, and a designated Technical Account Manager (TAM) to coordinate access to proactive and preventative programs and AWS subject matter experts. Exam Alert: Please review the differences between the AWS Developer Support, AWS Business Support, AWS Enterprise On-Ramp Support, and AWS Enterprise Support plans, as you can anticipate at least two questions on the exam: <table><thead><tr><th></th><th>Developer</th><th>Business</th><th>Enterprise On-Ramp</th><th>Enterprise</th></tr></thead><tbody><tr><td></td><td>Recommended if you are experimenting or testing in AWS.</td><td>Minimum recommended tier if you have production workloads in AWS</td><td>Recommended if you have production and/or business critical workloads in AWS.</td><td>Recommended if you have business and/or mission critical workloads in AWS.</td></tr><tr><td>AWS Trusted Advisor Best Practice Checks</td><td>Service Quota and basic Security checks</td><td>Full set of checks</td><td>Full set of checks</td><td>Full set of checks</td></tr><tr><td>AWS Trusted Advisor Priority</td><td></td><td></td><td></td><td>Prioritized recommendations curated by your AWS account team</td></tr><tr><td>Enhanced Technical Support</td><td>Business hours** web access to Cloud Support Associates Unlimited cases with 1 primary contact Prioritized responses on AWS re:Post</td><td>24/7 phone, web, and chat access to Cloud Support Engineers Unlimited cases and unlimited contacts (IAM supported) Prioritized responses on AWS re:Post Access to AWS Support App in Slack</td><td>24/7 phone, web, and chat access to Cloud Support Engineers Unlimited cases and unlimited contacts (IAM supported) Prioritized responses on AWS re:Post Access to AWS Support App in Slack</td><td>24/7 phone, web, and chat access to Cloud Support Engineers Unlimited cases and unlimited contacts (IAM supported) Prioritized responses on AWS re:Post Access to AWS Support App in Slack</td></tr><tr><td>Case Severity / Response Times*</td><td>General guidance: < 24 hours** System impaired: < 12 hours**</td><td>General guidance: < 24 hours System impaired: < 12 hours Production system impaired: < 4 hours Production system down: < 1 hour</td><td>General guidance: < 24 hours System impaired: < 12 hours Production system impaired: < 4 hours Production system down: < 1 hour Business-critical system down: < 30 minutes</td><td>General guidance: < 24 hours System impaired: < 12 hours Production system impaired: < 4 hours Production system down: < 1 hour Business/Mission-critical system down: < 15 minutes</td></tr><tr><td>Architectural Guidance</td><td>General</td><td>Contextual to your use-cases</td><td>Consultative review and guidance based on your applications (one-per-year)</td><td>Consultative review and guidance based on your applications</td></tr><tr><td>Programmatic Case Management</td><td></td><td>AWS Support API</td><td>AWS Support API</td><td>AWS Support API</td></tr></tbody></table>		Developer	Business	Enterprise On-Ramp	Enterprise		Recommended if you are experimenting or testing in AWS.	Minimum recommended tier if you have production workloads in AWS	Recommended if you have production and/or business critical workloads in AWS.	Recommended if you have business and/or mission critical workloads in AWS.	AWS Trusted Advisor Best Practice Checks	Service Quota and basic Security checks	Full set of checks	Full set of checks	Full set of checks	AWS Trusted Advisor Priority				Prioritized recommendations curated by your AWS account team	Enhanced Technical Support	Business hours** web access to Cloud Support Associates Unlimited cases with 1 primary contact Prioritized responses on AWS re:Post	24/7 phone, web, and chat access to Cloud Support Engineers Unlimited cases and unlimited contacts (IAM supported) Prioritized responses on AWS re:Post Access to AWS Support App in Slack	24/7 phone, web, and chat access to Cloud Support Engineers Unlimited cases and unlimited contacts (IAM supported) Prioritized responses on AWS re:Post Access to AWS Support App in Slack	24/7 phone, web, and chat access to Cloud Support Engineers Unlimited cases and unlimited contacts (IAM supported) Prioritized responses on AWS re:Post Access to AWS Support App in Slack	Case Severity / Response Times*	General guidance: < 24 hours** System impaired: < 12 hours**	General guidance: < 24 hours System impaired: < 12 hours Production system impaired: < 4 hours Production system down: < 1 hour	General guidance: < 24 hours System impaired: < 12 hours Production system impaired: < 4 hours Production system down: < 1 hour Business-critical system down: < 30 minutes	General guidance: < 24 hours System impaired: < 12 hours Production system impaired: < 4 hours Production system down: < 1 hour Business/Mission-critical system down: < 15 minutes	Architectural Guidance	General	Contextual to your use-cases	Consultative review and guidance based on your applications (one-per-year)	Consultative review and guidance based on your applications	Programmatic Case Management		AWS Support API	AWS Support API	AWS Support API
	Developer	Business	Enterprise On-Ramp	Enterprise																																							
	Recommended if you are experimenting or testing in AWS.	Minimum recommended tier if you have production workloads in AWS	Recommended if you have production and/or business critical workloads in AWS.	Recommended if you have business and/or mission critical workloads in AWS.																																							
AWS Trusted Advisor Best Practice Checks	Service Quota and basic Security checks	Full set of checks	Full set of checks	Full set of checks																																							
AWS Trusted Advisor Priority				Prioritized recommendations curated by your AWS account team																																							
Enhanced Technical Support	Business hours** web access to Cloud Support Associates Unlimited cases with 1 primary contact Prioritized responses on AWS re:Post	24/7 phone, web, and chat access to Cloud Support Engineers Unlimited cases and unlimited contacts (IAM supported) Prioritized responses on AWS re:Post Access to AWS Support App in Slack	24/7 phone, web, and chat access to Cloud Support Engineers Unlimited cases and unlimited contacts (IAM supported) Prioritized responses on AWS re:Post Access to AWS Support App in Slack	24/7 phone, web, and chat access to Cloud Support Engineers Unlimited cases and unlimited contacts (IAM supported) Prioritized responses on AWS re:Post Access to AWS Support App in Slack																																							
Case Severity / Response Times*	General guidance: < 24 hours** System impaired: < 12 hours**	General guidance: < 24 hours System impaired: < 12 hours Production system impaired: < 4 hours Production system down: < 1 hour	General guidance: < 24 hours System impaired: < 12 hours Production system impaired: < 4 hours Production system down: < 1 hour Business-critical system down: < 30 minutes	General guidance: < 24 hours System impaired: < 12 hours Production system impaired: < 4 hours Production system down: < 1 hour Business/Mission-critical system down: < 15 minutes																																							
Architectural Guidance	General	Contextual to your use-cases	Consultative review and guidance based on your applications (one-per-year)	Consultative review and guidance based on your applications																																							
Programmatic Case Management		AWS Support API	AWS Support API	AWS Support API																																							

iPQ.No	Correct Answer	Points	Explanation
2AWSCCP-CLF08	2	15.4	Answer = 2 Amazon MQ Explanation : Amazon MQ is a managed message broker service for Apache ActiveMQ and RabbitMQ that simplifies the configuration and management of message brokers on AWS. Amazon MQ reduces your operational obligations by handling the provisioning, configuration, and maintenance of message brokers. Amazon MQ's compatibility with your existing applications via industry-standard APIs and protocols allows you to migrate to AWS without having to rewrite any code. Amazon MQ supports industry-standard APIs and protocols, allowing you to switch from any standards-based message broker to Amazon MQ without rewriting your applications' messaging code. AWS recommends considering Amazon SQS and Amazon SNS when developing cloud-based applications from scratch. How Amazon MQ works:  via https://aws.amazon.com/amazon-mq/
2AWSCCP-CLF09	1	15.4	Answer = 1 High Availability Explanation : A system that is available at a given time is capable of delivering the intended functionality. Highly available systems are those that can tolerate a certain amount of degradation while remaining accessible. You can run application instances in a multi-AZ deployment on AWS Cloud to achieve high availability.

iPQ.No	Correct Answer	Points	Explanation
2AWSCCP-CLF10	3	15.4	Answer = 3 Routing based on timing Explanation : Amazon Route 53 is a highly available and scalable DNS web service in the cloud. It is intended to provide developers and businesses with a highly dependable and cost-effective method for directing end users to Internet applications by translating domain names such as <code>www.example.com</code> into the numeric IP addresses such as <code>192.0.2.1</code> that computers use to communicate with one another. If your application is hosted in multiple AWS regions, you can improve performance for your users by serving their requests from the AWS region with the lowest latency using a latency-based routing policy. To utilize latency-based routing, you must generate latency records for your resources across multiple AWS regions. When Amazon Route 53 receives a DNS query for your domain or subdomain (<code>example.com</code> or <code>acme.com</code>), it determines which AWS regions you've created latency records for, then selects a latency record for the region with the lowest latency. Amazon Route 53 responds with the selected record's value, such as a web server's IP address. Overview of Amazon Route 53 Routing Policy: Choosing a routing policy PDF Kindle RSS When you create a record, you choose a routing policy, which determines how Amazon Route 53 responds to queries: • Simple routing policy – Use for a single resource that performs a given function for your domain, for example, a web server that serves content for the <code>example.com</code> website. • Failover routing policy – Use when you want to configure active-passive failover. • Geolocation routing policy – Use when you want to route traffic based on the location of your users. • Geoproximity routing policy – Use when you want to route traffic based on the location of your resources and, optionally, shift traffic from resources in one location to resources in another. • Latency routing policy – Use when you have resources in multiple AWS Regions and you want to route traffic to the region that provides the best latency. • Multivalue answer routing policy – Use when you want Route 53 to respond to DNS queries with up to eight healthy records selected at random. • Weighted routing policy – Use to route traffic to multiple resources in proportions that you specify.


iPQ.No	Correct Answer	Points	Explanation
2AWS CCP-CLF11	2	15.4	Answer = 2 Amazon Elastic Container Service (Amazon ECS): Launch type for Fargate containers Explanation : AWS Fargate is a serverless container compute engine. Amazon Elastic Container Service (Amazon ECS) and Amazon Elastic Kubernetes Service (Amazon EKS) are both supported. AWS Fargate makes it simple to concentrate on developing applications. AWS Fargate eliminates the need to provision and manage servers, enables you to specify and pay for resources per application, and enhances security by isolating applications by design. AWS Fargate allots the appropriate amount of compute resources, eliminating the need to select instances and scale cluster capacity. There is no overprovisioning or payment for additional servers, as you only pay for the resources required to run your containers. AWS Fargate executes each task or pod in its own kernel, giving each task or pod its own isolated computing environment. This allows your application to be designed with workload isolation and enhanced security. Using AWS Fargate:  https://aws.amazon.com/fargate/

iPQ.No	Correct Answer	Points	Explanation
2AWSCCP-CLF12	1	15.4	Answer = 1 Amazon Web Services Web Application Firewall (AWS WAF) Explanation : AWS Web Application Firewall (AWS WAF) is a web application firewall that helps protect web applications from attacks by enabling the configuration of rules that permit, block, or monitor (count) web requests based on conditions that you define. IP addresses, HTTP headers, HTTP bodies, URI strings, SQL injection, and cross-site scripting are examples of these conditions. The IP address-based match rule can be used to block specific geographies. The IP address to country lookup database's precision varies by region. AWS reports that based on recent tests, the IP address-to-country mapping accuracy is 99.8% overall. How the AWS Web Application Firewall (AWS WAF) Operates:  https://aws.amazon.com/waf/


iPQ.No	Correct Answer	Points	Explanation
2AWSCCP-CLF13	4	15.4	Answer = 4 AWS auto-scaling Explanation : AWS Auto Scaling monitors applications and adjusts capacity automatically to maintain consistent, predictable performance at the lowest cost possible. Using AWS Auto Scaling, it is simple to scale an application across multiple resources and services in minutes. The service offers a simple and powerful user interface that enables the creation of scaling plans for Amazon EC2 instances and Spot Fleets, Amazon ECS tasks, Amazon DynamoDB tables and indexes, and Amazon Aurora Replicas. AWS Auto Scaling simplifies scaling by providing recommendations to optimize performance, costs, or a balance of the two. How AWS Auto Scaling Operates:  aws.amazon.com/autoscaling

iPQ.No	Correct Answer	Points	Explanation
2AWSCCP-CLF14	3	15.4	Answer = 3 Both A and B Explanation : The majority of AWS's services are region-specific. However, by definition, few services must have a global reach due to the nature of the service they provide. Some of the global services include AWS Identity and Access Management (AWS IAM), Amazon CloudFront, Amazon Route 53, and AWS Web Application Firewall (AWS WAF). AWS Identity and Access Management (AWS IAM) enables the secure administration of access to AWS services and resources. Using AWS Identity and Access Management (AWS IAM), you can create and manage IAM users and IAM user groups, as well as grant or deny their access to AWS resources, using permissions. Amazon CloudFront is a developer-friendly content delivery network (CDN) service that securely delivers data, videos, applications, and APIs to customers worldwide with low latency and high transfer speeds.

iPQ.No	Correct Answer	Points	Explanation
2AWS CCP-CLF15	4	15.4	Answer = 4 Inspector for Amazon.com. Explanation : Amazon Inspector is an automated security assessment service that enhances the security and compliance of AWS-deployed applications by evaluating them automatically for exposure, vulnerabilities, and deviations from best practices. Amazon Inspector Overview: IDENTIFY APPLICATION SECURITY ISSUES Amazon Inspector helps you to identify security vulnerabilities as well as deviations from security best practices in applications, both before they are deployed, and while they are running in a production environment. This helps improve the overall security posture of your applications deployed on AWS. INTEGRATE SECURITY INTO DEVOPS Amazon Inspector is an API-driven service that analyzes network configurations in your AWS account and uses an optional agent for visibility into your Amazon EC2 instances. This makes it easy for you to build Inspector assessments right into your existing DevOps process, decentralizing and automating vulnerability assessments, and empowering your development and operations teams to make security assessments an integral part of the deployment process. INCREASE DEVELOPMENT AGILITY Amazon Inspector helps you reduce the risk of introducing security issues during development and deployment by automating the security assessment of your applications and proactively identifying vulnerabilities. This allows you to develop and iterate on new applications quickly and assess compliance with best practices and policies. LEVERAGE AWS SECURITY EXPERTISE The AWS security organization is continuously assessing the AWS environment and updating a knowledge base of security best practices and rules. Amazon Inspector makes this expertise available to you in the form of a service that simplifies the process of establishing and enforcing best practices within your AWS environment. STREAMLINE SECURITY COMPLIANCE Amazon Inspector gives security teams and auditors visibility into the security testing that is being performed during development of applications on AWS. This streamlines the process of validating and demonstrating that security and compliance standards and best practices are being followed throughout the development process. ENFORCE SECURITY STANDARDS Amazon Inspector allows you to define standards and best practices for your applications and validate adherence to these standards. This simplifies enforcement of your organization's security standards and best practices, and helps to proactively manage security issues before they impact your production application. http://aws.amazon.com/inspector/
2AWS CCP-CLF16	1	15.4	Answer = 1 Hybrid deployment Explanation : The most common method of hybrid deployment is between the cloud and existing on-premises infrastructure to extend an organization's infrastructure into the cloud while connecting cloud resources to internal systems. Cloud Computing Deployment Models: Cloud Computing Deployment Models  Cloud A cloud-based application is fully deployed in the cloud and all parts of the application run in the cloud. Applications in the cloud have either been created in the cloud or have been migrated from an existing infrastructure to take advantage of the benefits of cloud computing. Cloud-based applications can be built on low-level infrastructure pieces or can use higher level services that provide abstraction from the management, architecting, and scaling requirements of core infrastructure.  Hybrid A hybrid deployment is a way to connect infrastructure and applications between cloud-based resources and existing resources that are not located in the cloud. The most common method of hybrid deployment is between the cloud and existing on-premises infrastructure to extend, and grow, an organization's infrastructure into the cloud while connecting cloud resources to internal system. For more information on how AWS can help you with your hybrid deployment, please visit our hybrid page.  On-premises Deploying resources on-premises, using virtualization and resource management tools, is sometimes called "private cloud". On-premises deployment does not provide many of the benefits of cloud computing but is sometimes sought for its ability to provide dedicated resources. In most cases this deployment model is the same as legacy IT infrastructure while using application management and virtualization technologies to try and increase resource utilization. via - https://aws.amazon.com/types-of-cloud-computing/

iPQ.No	Correct Answer	Points	Explanation
2AWSCCP-CLF17	2	15.4	Answer = 2 Edge Location Management Explanation : Security and Compliance is a shared responsibility between AWS and the customer. The AWS Shared Responsibility Model can help relieve the customer's operational burden as AWS operates, manages, and controls the components from the host operating system and virtualization layer down to the physical security of the facilities in which the service operates. AWS is responsible for security "of" the cloud. This covers their global infrastructure elements including Regions, Availability Zones (AZ), and Edge Locations.
2AWSCCP-CLF18	3	15.4	Answer = 3 Both A and B Explanation : Amazon CloudWatch is a monitoring and observability service designed for DevOps engineers, developers, site reliability engineers (SREs), and IT managers. It provides data and actionable insights to monitor applications, respond to system-wide performance changes, optimize resource utilization, and get a unified view of operational health. Amazon Simple Notification Service (SNS) is a highly available, durable, secure, fully managed pub/sub messaging service that allows you to decouple microservices, distributed systems, and serverless applications. Using SNS:  https://aws.amazon.com/sns/

iPQ.No	Correct Answer	Points	Explanation
2AWSCCP-CLF19	1	15.4	Answer = 1 Compute, Store, and Transfer Outbound Data Explanation : AWS has three primary cost drivers: compute, storage, and outbound data transfer. Inbound data transfer and data transfer between AWS services in the same region are typically free. AWS Cloud Pricing Fundamentals: Understand the fundamentals of pricing There are three fundamental drivers of cost with AWS: compute, storage, and outbound data transfer. These characteristics vary somewhat, depending on the AWS product and pricing model you choose. In most cases, there is no charge for inbound data transfer or for data transfer between other AWS services within the same region. There are some exceptions, so be sure to verify data transfer rates before beginning. Outbound data transfer is aggregated across services and then charged at the outbound data transfer rate. This charge appears on the monthly statement as AWS Data Transfer Out. The more data you transfer, the less you pay per GB. For compute resources, you pay hourly from the time you launch a resource until the time you terminate it, unless you have made a reservation for which the cost is agreed upon beforehand. For data storage and transfer, you typically pay per GB. https://d0.awsstatic.com/whitepapers/aws_pricing_overview.pdf

iPQ.No	Correct Answer	Points	Explanation																																																								
2AWS CCP-CLF20	4	15.4	Answer = 4 AWS Budgets Explanation : AWS Budgets enables you to set custom budgets and receive alerts when your costs or usage exceed (or are forecast to exceed) your budgeted amount. You can also use AWS Budgets to set reservation utilization or coverage targets and receive alerts when your utilization drops below the threshold you define. AWS Budgets Overview: AWS Budgets Features Create and manage budgets Set custom cost and usage budgets to more easily manage your AWS spend. Monitor your budget status from the AWS Budgets dashboard or AWS Budgets reports. Refine your budget using filters Track your cost or usage across multiple dimensions by adding filters related to Service, Linked Account(s), Region, Tag, and more. Add notifications to your budget Set up to five alert thresholds for each budget. Each alert can notify up to ten email recipients as well as publish updates to a Slack channel, Amazon Chime room, or Amazon SNS topic of your choice. Getting Started AWS Budgets Dashboard The AWS Budgets Dashboard is your hub for creating, tracking, and inspecting your budgets. From the AWS Budgets Dashboard, you can create, edit, and manage your budgets, as well as view the status of each of your budgets. You can also view additional details about your budgets, such as a high-level variance analysis and a budget criteria summary. Budgets can be created at the monthly, quarterly, or yearly level, and you can customize the start and end dates. You can further refine your budget to track costs associated with multiple dimensions, such as AWS service, linked account, tag, and others. Budget alerts can be sent via email and/or Amazon Simple Notification Service (SNS) topic.  <table><thead><tr><th>Budget name</th><th>Budget type</th><th>Current</th><th>Budgeted</th><th>Forecasted</th><th>Current vs. budgeted</th><th>Forecasted vs. budgeted</th></tr></thead><tbody><tr><td>Project Name Cost Budget</td><td>Cost</td><td>\$45.85</td><td>\$45.00</td><td>\$45.33</td><td>8.17% over</td><td>10.17% over</td></tr><tr><td>Project ID Region Budget</td><td>Cost</td><td>\$60.21</td><td>\$100.00</td><td>\$100.00</td><td>60.21% over</td><td>100.00% over</td></tr><tr><td>Usage Monthly Cost Budget</td><td>Cost</td><td>\$141.56</td><td>\$175.00</td><td>\$167.08</td><td>80.89% over</td><td>95.82% over</td></tr><tr><td>Usage Q3 Cost Budget</td><td>Cost</td><td>\$100.00</td><td>\$100.00</td><td>\$100.00</td><td>0% over</td><td>0.00% over</td></tr><tr><td>400 Usage Budget</td><td>Usage</td><td>4,000</td><td>4,000</td><td>4,000.00</td><td>0% over</td><td>0.00% over</td></tr><tr><td>Monthly Usage Budget</td><td>Usage</td><td>3,000.00</td><td>4,000</td><td>3,000.00</td><td>25.00% over</td><td>75.00% over</td></tr><tr><td>Quarterly Budget</td><td>Cost</td><td>\$100.00</td><td>\$100.00</td><td>\$100.00</td><td>0.00% over</td><td>0.00% over</td></tr></tbody></table>	Budget name	Budget type	Current	Budgeted	Forecasted	Current vs. budgeted	Forecasted vs. budgeted	Project Name Cost Budget	Cost	\$45.85	\$45.00	\$45.33	8.17% over	10.17% over	Project ID Region Budget	Cost	\$60.21	\$100.00	\$100.00	60.21% over	100.00% over	Usage Monthly Cost Budget	Cost	\$141.56	\$175.00	\$167.08	80.89% over	95.82% over	Usage Q3 Cost Budget	Cost	\$100.00	\$100.00	\$100.00	0% over	0.00% over	400 Usage Budget	Usage	4,000	4,000	4,000.00	0% over	0.00% over	Monthly Usage Budget	Usage	3,000.00	4,000	3,000.00	25.00% over	75.00% over	Quarterly Budget	Cost	\$100.00	\$100.00	\$100.00	0.00% over	0.00% over
Budget name	Budget type	Current	Budgeted	Forecasted	Current vs. budgeted	Forecasted vs. budgeted																																																					
Project Name Cost Budget	Cost	\$45.85	\$45.00	\$45.33	8.17% over	10.17% over																																																					
Project ID Region Budget	Cost	\$60.21	\$100.00	\$100.00	60.21% over	100.00% over																																																					
Usage Monthly Cost Budget	Cost	\$141.56	\$175.00	\$167.08	80.89% over	95.82% over																																																					
Usage Q3 Cost Budget	Cost	\$100.00	\$100.00	\$100.00	0% over	0.00% over																																																					
400 Usage Budget	Usage	4,000	4,000	4,000.00	0% over	0.00% over																																																					
Monthly Usage Budget	Usage	3,000.00	4,000	3,000.00	25.00% over	75.00% over																																																					
Quarterly Budget	Cost	\$100.00	\$100.00	\$100.00	0.00% over	0.00% over																																																					

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


Exam Alert:

It is important to distinguish between CloudWatch Billing and AWS Budgets:

CloudWatch Billing Alarms: Notifies you when actual costs exceed a specified threshold.

AWS Budgets: Sends an alert when the actual cost exceeds the budgeted amount, or when the budgeted amount is exceeded by the cost forecast.

iPQ.No	Correct Answer	Points	Explanation
2AWSCCP-CLF21	3	15.4	Answer = 3 Both A and B Explanation : Using machine learning to analyze historical utilization metrics, AWS Compute Optimizer identifies the optimal AWS resource configurations, such as Amazon EC2 instance types, Amazon EBS volume configurations, and AWS Lambda function memory sizes. AWS Compute Optimizer provides recommendations for particular EC2 instance types, EC2 Auto Scaling groups, Amazon EBS volumes, and AWS Lambda functions. AWS Compute Optimizer computes a performance risk score for each resource dimension of the recommended instance, such as CPU, memory, EBS throughput, EBS IOPS, disk throughput, disk throughput, network throughput, and network packets per second (PPS). AWS Compute Optimizer provides EC2 instance type and size recommendations for EC2 Auto Scaling groups with a fixed group size, meaning desired, minimum, and maximum are all set to the same value and there is no associated scaling policy. AWS Compute Optimizer provides IOPS and throughput recommendations for general purpose (SSD) (gp3) volumes, as well as IOPS recommendations for provisioned IOPS (io1 and io2) volumes. AWS Compute Optimizer facilitates the optimization of two classes of Lambda functions. The first category contains lambda functions with potentially excessive memory sizes. The second category includes Lambda functions that are computationally intensive and may benefit from additional CPU power.


iPQ.No	Correct Answer	Points	Explanation
2AWSCCP-CLF22	4	15.4	Answer = 4 Database security Explanation : AWS and the customer share responsibility for security and compliance. AWS operates, manages, and controls the components, from the host operating system and virtualization layer to the physical security of the service's operating facilities, in this shared model. Under the AWS Shared Responsibility Model, customers are responsible for managing their data, which includes encrypting it. Overview of the AWS Shared Responsibility Model, found at  The diagram illustrates the AWS Shared Responsibility Model. On the left, two large chevron-shaped boxes represent the primary responsibilities: a blue box for the CUSTOMER (Responsibility for security 'IN' the cloud) and an orange box for AWS (Responsibility for security 'OF' the cloud). To the right, a stack of layers represents the cloud architecture. The top four layers are blue and fall under the Customer's responsibility: CUSTOMER DATA, PLATFORM, APPLICATIONS, IDENTITY & ACCESS MANAGEMENT, OPERATING SYSTEM, NETWORK & FIREWALL CONFIGURATION, and a split layer for CLIENT-SIDE DATA ENCRYPTION & DATA INTEGRITY AUTHENTICATION, SERVER-SIDE ENCRYPTION (FILE SYSTEM AND/OR DATA), and NETWORKING TRAFFIC PROTECTION (ENCRYPTION, INTEGRITY, IDENTITY). The next two layers are orange and fall under AWS's responsibility: SOFTWARE (which includes COMPUTE, STORAGE, DATABASE, and NETWORKING) and HARDWARE/AWS GLOBAL INFRASTRUCTURE (which includes REGIONS, AVAILABILITY ZONES, and EDGE LOCATIONS). https://aws.amazon.com/compliance/shared-responsibility-model/.

iPQ.No	Correct Answer	Points	Explanation
2AWSCCP-CLF23	3	15.4	Answer = 3 Both A and B Explanation : Foundations are part of the AWS Well-Architected Framework's reliability pillar. Prior to designing a system, AWS recommends establishing foundational requirements that influence its dependability. Amazon VPC, AWS Trusted Advisor, and AWS Service Quotas (previously known as AWS Service Limits) are the services that comprise foundations. AWS Trusted Advisor is an online tool that provides real-time guidance to help you provision your resources in accordance with AWS best practices for cost optimization, security, fault tolerance, service limits, and performance enhancement. Whether establishing new workflows, developing applications, or as part of ongoing enhancements, Trusted Advisor's recommendations help maintain the optimal provisioning of your solutions. Service Quotas allows you to view and manage your AWS service quotas from a centralized location. In AWS, quotas, also known as limits, represent the maximum values for the resources, actions, and items in your account. Each AWS service defines its quotas and determines their default values.


iPQ.No	Correct Answer	Points	Explanation																																																																																								
2AWS CCP-CLF24	3	15.4	Answer = 3 To store the thumbnails, use Amazon S3 One Zone-Infrequent Access (S3 One Zone-IA) Explanation : Amazon S3 One Zone-Infrequent Access (S3 One Zone-IA) is designed to store data that is accessed infrequently but requires rapid access when required. Amazon S3 One Zone-Infrequent Access (S3 One Zone-IA) stores data in a single availability zone (AZ) and costs 20% less than Amazon S3 Standard-Infrequent Access (S3 Standard-IA). Other S3 storage classes store data in a minimum of three availability zones (AZs). Amazon S3 One Zone-Infrequent Access (S3 One Zone-IA) provides the same high durability, high throughput, and low latency as S3 Standard, in addition to a low per-GB storage cost and per-GB retrieval fee. Although Amazon S3 One Zone-Infrequent Access (S3 One Zone-IA) provides less availability than S3 Standard, this is not a concern for the specified use case because thumbnails can be easily regenerated. Amazon S3 One Zone-Infrequent Access (S3 One Zone-IA) is the best option for this use case because thumbnails are rarely accessed but must be accessed quickly when needed. Exam Alert: Please review the following comparison of S3 storage classes, as you can anticipate a few questions on this aspect of S3: <table><caption>Performance across the S3 Storage Classes</caption><thead><tr><th></th><th>S3 Standard</th><th>S3 Intelligent-Tiering*</th><th>S3 Standard-IA</th><th>S3 One Zone-IA†</th><th>S3 Glacier Instant Retrieval</th><th>S3 Glacier Flexible Retrieval</th><th>S3 Glacier Deep Archive</th></tr></thead><tbody><tr><td>Designed for durability</td><td>99.999999999% (11 9's)</td><td>99.999999999% (11 9's)</td><td>99.999999999% (11 9's)</td><td>99.999999999% (11 9's)</td><td>99.999999999% (11 9's)</td><td>99.999999999% (11 9's)</td><td>99.999999999% (11 9's)</td></tr><tr><td>Designed for availability</td><td>99.99%</td><td>99.9%</td><td>99.9%</td><td>99.5%</td><td>99.9%</td><td>99.99%</td><td>99.99%</td></tr><tr><td>Availability SLA</td><td>99.9%</td><td>99%</td><td>99%</td><td>99%</td><td>99%</td><td>99.1%</td><td>99.9%</td></tr><tr><td>Availability Zones</td><td>≥3</td><td>≥3</td><td>≥3</td><td>1</td><td>≥3</td><td>≥3</td><td>≥3</td></tr><tr><td>Minimum capacity charge per object</td><td>N/A</td><td>N/A</td><td>128 KB</td><td>128 KB</td><td>128 KB</td><td>40 KB</td><td>40 KB</td></tr><tr><td>Minimum storage duration charge</td><td>N/A</td><td>N/A</td><td>30 days</td><td>30 days</td><td>90 days</td><td>90 days</td><td>180 days</td></tr><tr><td>Retrieval charge</td><td>N/A</td><td>N/A</td><td>per GB retrieved</td><td>per GB retrieved</td><td>per GB retrieved</td><td>per GB retrieved</td><td>per GB retrieved</td></tr><tr><td>First byte latency</td><td>milliseconds</td><td>milliseconds</td><td>milliseconds</td><td>milliseconds</td><td>milliseconds</td><td>minutes or hours</td><td>hours</td></tr><tr><td>Storage type</td><td>Object</td><td>Object</td><td>Object</td><td>Object</td><td>Object</td><td>Object</td><td>Object</td></tr><tr><td>Lifecycle transitions</td><td>Yes</td><td>Yes</td><td>Yes</td><td>Yes</td><td>Yes</td><td>Yes</td><td>Yes</td></tr></tbody></table>		S3 Standard	S3 Intelligent-Tiering*	S3 Standard-IA	S3 One Zone-IA†	S3 Glacier Instant Retrieval	S3 Glacier Flexible Retrieval	S3 Glacier Deep Archive	Designed for durability	99.999999999% (11 9's)	99.999999999% (11 9's)	99.999999999% (11 9's)	99.999999999% (11 9's)	99.999999999% (11 9's)	99.999999999% (11 9's)	99.999999999% (11 9's)	Designed for availability	99.99%	99.9%	99.9%	99.5%	99.9%	99.99%	99.99%	Availability SLA	99.9%	99%	99%	99%	99%	99.1%	99.9%	Availability Zones	≥3	≥3	≥3	1	≥3	≥3	≥3	Minimum capacity charge per object	N/A	N/A	128 KB	128 KB	128 KB	40 KB	40 KB	Minimum storage duration charge	N/A	N/A	30 days	30 days	90 days	90 days	180 days	Retrieval charge	N/A	N/A	per GB retrieved	per GB retrieved	per GB retrieved	per GB retrieved	per GB retrieved	First byte latency	milliseconds	milliseconds	milliseconds	milliseconds	milliseconds	minutes or hours	hours	Storage type	Object	Object	Object	Object	Object	Object	Object	Lifecycle transitions	Yes	Yes	Yes	Yes	Yes	Yes	Yes
	S3 Standard	S3 Intelligent-Tiering*	S3 Standard-IA	S3 One Zone-IA†	S3 Glacier Instant Retrieval	S3 Glacier Flexible Retrieval	S3 Glacier Deep Archive																																																																																				
Designed for durability	99.999999999% (11 9's)	99.999999999% (11 9's)	99.999999999% (11 9's)	99.999999999% (11 9's)	99.999999999% (11 9's)	99.999999999% (11 9's)	99.999999999% (11 9's)																																																																																				
Designed for availability	99.99%	99.9%	99.9%	99.5%	99.9%	99.99%	99.99%																																																																																				
Availability SLA	99.9%	99%	99%	99%	99%	99.1%	99.9%																																																																																				
Availability Zones	≥3	≥3	≥3	1	≥3	≥3	≥3																																																																																				
Minimum capacity charge per object	N/A	N/A	128 KB	128 KB	128 KB	40 KB	40 KB																																																																																				
Minimum storage duration charge	N/A	N/A	30 days	30 days	90 days	90 days	180 days																																																																																				
Retrieval charge	N/A	N/A	per GB retrieved	per GB retrieved	per GB retrieved	per GB retrieved	per GB retrieved																																																																																				
First byte latency	milliseconds	milliseconds	milliseconds	milliseconds	milliseconds	minutes or hours	hours																																																																																				
Storage type	Object	Object	Object	Object	Object	Object	Object																																																																																				
Lifecycle transitions	Yes	Yes	Yes	Yes	Yes	Yes	Yes																																																																																				

iPQ.No	Correct Answer	Points	Explanation
2AWSCCP-CLF25	3	15.4	Answer = 3 Amazon EC2 (Amazon Elastic Compute Cloud) Explanation : Amazon Elastic Compute Cloud (Amazon EC2) is a web service that offers secure, scalable cloud computing capacity with per-second billing. It is the simplest method for provisioning servers in the AWS cloud and gaining access to the underlying OS. Amazon EC2 reduces the time required to obtain and boot new server instances to minutes, allowing you to quickly scale capacity both up and down in response to fluctuating computing demands. Overview of Amazon Elastic Compute Cloud (Amazon EC2), accessible at Amazon EC2 presents a true virtual computing environment, allowing you to use web service interfaces to launch instances with a variety of operating systems, load them with your custom application environment, manage your network's access permissions, and run your image using as many or few systems as you desire. To use Amazon EC2, you simply: • Select a pre-configured, templated Amazon Machine Image (AMI) to get up and running immediately. Or create an AMI containing your applications, libraries, data, and associated configuration settings. • Configure security and network access on your Amazon EC2 instance. • Choose which instance type(s) you want, then start, terminate, and monitor as many instances of your AMI as needed, using the web service APIs or the variety of management tools provided. • Determine whether you want to run in multiple locations, utilize static IP endpoints, or attach persistent block storage to your instances. • Pay only for the resources that you actually consume, like instance-hours or data transfer. https://aws.amazon.com/ec2/.
2AWSCCP-CLF26	1	15.4	Answer = 1 Amazon Web Services Health Dashboard: Service Health Explanation : AWS Health Dashboard: Service Health is the central location for monitoring the availability and performance of AWS services. You can view the overall status of AWS services and sign in to view communications specific to your AWS account or organization. Current status information can be obtained by visiting https://health.aws.amazon.com/health/status.

iPQ.No	Correct Answer	Points	Explanation
2AWS CCP-CLF27	1	15.4	Answer = 1 Network Access Control List (Network ACL) acts as a firewall at the subnet level, whereas Security Group acts as a firewall at the instance level Explanation : A security group controls incoming and outgoing traffic for your instance by acting as a virtual firewall. When launching an instance in a VPC, you can assign the instance up to five security groups. Not the subnet level, but the instance level is where security groups operate. A network access control list (network ACL) is an optional security layer for your VPC that functions as a subnet-level firewall for controlling traffic into and out of one or more subnets. Overview of Security Groups: Security group basics The following are the basic characteristics of security groups for your VPC: - There are quotas on the number of security groups that you can create per VPC, the number of rules that you can add to each security group, and the number of security groups that you can associate with a network interface. For more information, see Amazon VPC quotas. - You can specify allow rules, but not deny rules. - You can specify separate rules for inbound and outbound traffic. - When you create a security group, it has no inbound rules. Therefore, no inbound traffic originating from another host to your instance is allowed until you add inbound rules to the security group. - By default, a security group includes an outbound rule that allows all outbound traffic. You can remove the rule and add outbound rules that allow specific outbound traffic only. If your security group has no outbound rules, no outbound traffic originating from your instance is allowed. - Security groups are stateful — if you send a request from your instance, the response traffic for that request is allowed to flow in regardless of inbound security group rules. Responses to allowed inbound traffic are allowed to flow out, regardless of outbound rules. Note Some types of traffic are tracked differently from other types. For more information, see Connection tracking in the Amazon EC2 User Guide for Linux Instances. - Instances associated with a security group can't talk to each other unless you add rules allowing the traffic (exception: the default security group has these rules by default). - Security groups are associated with network interfaces. After you launch an instance, you can change the security groups that are associated with the instance, which changes the security groups associated with the primary network interface (eth0). You can also specify or change the security groups associated with any other network interface. By default, when you create a network interface, it's associated with the default security group for the VPC, unless you specify a different security group. For more information about network interfaces, see Elastic network interfaces. - When you create a security group, you must provide it with a name and a description. The following rules apply: - Names and descriptions can be up to 255 characters in length. - Names and descriptions are limited to the following characters: a-z, A-Z, 0-9, spaces, and . _ - / ! # , @ [] + = & ; { } \$ % ' . - A security group name cannot start with sg- as these indicate a default security group. - A security group name must be unique within the VPC. - A security group can only be used in the VPC that you specify when you create the security group. https://docs.aws.com/vpc/latest/userguide/VPC_SecurityGroups.html Overview of Network Access Control List: Network ACL basics The following are the basic things that you need to know about network ACLs: - Your VPC automatically comes with a modifiable default network ACL. By default, it allows all inbound and outbound IPv4 traffic and, if applicable, IPv6 traffic. - You can create a custom network ACL and associate it with a subnet. By default, each custom network ACL denies all inbound and outbound traffic until you add rules.


iPQ.No	Correct Answer	Points	Explanation
2AWSCCP-CLF28	4	15.4	Answer = 4 Management of capacity and performance Explanation : The AWS Cloud Adoption Framework (AWS CAF) leverages AWS experience and best practices to assist you in digitally transforming and accelerating your business outcomes through innovative AWS usage. Utilize the AWS CAF to identify and prioritize transformation opportunities, evaluate and enhance your cloud readiness, and evolve your transformation roadmap iteratively. AWS CAF categorizes its capabilities as follows: business, people, governance, platform, security, and operations. In your cloud transformation journey, each perspective consists of a collection of capabilities that functionally related stakeholders own or manage. A perspective on operations ensures that your cloud services are delivered at a level that meets your organization's requirements. The AWS Cloud Adoption Framework (AWS CAF) includes performance and capacity management from the perspective of operations. Foundational Capabilities:  The diagram illustrates the Foundational Capabilities of the AWS Cloud Adoption Framework (AWS CAF), organized into six pillars: Business (Green): Strategy Management, Portfolio Management, Innovation Management, Product Management, Strategic Partnership, Data Monetization, Business Insights, Data Science. People (Orange): Culture Evolution, Transformational Leadership, Cloud Fluency, Workforce Transformation, Change Acceleration, Organization Design, Organizational Alignment. Governance (Blue): Program & Project Management, Benefits Management, Risk Management, Cloud Financial Management, Application Portfolio Mgmt, Data Governance, Data Curation. Platform (Yellow): Platform Architecture, Data Architecture, Platform Engineering, Data Engineering, Provisioning and Orchestration, Modern App. Development, CI/CD. Security (Grey): Security Governance, Security Assurance, Identity & Access Management, Threat Detection, Vulnerability Management, Infrastructure Protection, Data Protection, Application Security, Incident Response. Operations (Dark Blue): Observability, Event Mgmt (AIOps), Incident and Problem Mgmt, Change & Release Management, Performance and Capacity Mgmt, Configuration Management, Patch Management, Availability and Continuity Mgmt, Application Management. http://docs.aws.amazon.com/whitepapers/latest/overview-aws-cloud-adoption-framework/foundational-

iPQ.No	Correct Answer	Points	Explanation
2AWS CCP-CLF29	2	15.4	Answer = 2 Demand-Based Instance Explanation : On-Demand Instance: An on-demand instance is one that can be utilized on demand. You decide when to launch, stop, hibernate, start, restart, or terminate it. There is no requirement for a long-term contract when purchasing on-demand instances. There is no upfront cost, and you only pay for the seconds your on-demand instances are active. The cost per second to run an instance on demand is fixed. Instances created on demand cannot be interrupted. Overview of EC2 Pricing Options: On-Demand With On-Demand instances, you pay for compute capacity by the hour or the second depending on which instances you run. No longer-term commitments or upfront payments are needed. You can increase or decrease your compute capacity depending on the demands of your application and only pay the specified per hourly rates for the instance you use. On-Demand instances are recommended for: - Users that prefer the low cost and flexibility of Amazon EC2 without any up-front payment or long-term commitment - Applications with short-term, spiky, or unpredictable workloads that cannot be interrupted - Applications being developed or tested on Amazon EC2 for the first time See On-Demand pricing » Spot instances Amazon EC2 Spot instances allow you to request spare Amazon EC2 computing capacity for up to 90% off the On-Demand price. Learn More. Spot instances are recommended for: - Applications that have flexible start and end times - Applications that are only feasible at very low compute prices - Users with urgent computing needs for large amounts of additional capacity See Spot pricing » Savings Plans Savings Plans are a flexible pricing model that offer low prices on EC2 and Fargate usage, in exchange for a commitment to a consistent amount of usage (measured in \$/hour) for a 1 or 3 year term. Dedicated Hosts A Dedicated Host is a physical EC2 server dedicated for your use. Dedicated Hosts can help you reduce costs by allowing you to use your existing server-bound software licenses, including Windows Server, SQL Server, and SUSE Linux Enterprise Server (subject to your license terms), and can also help you meet compliance requirements. Learn more. - Can be purchased On-Demand (hourly). - Can be purchased as a Reservation for up to 70% off the On-Demand price. See Dedicated pricing » Reserved Instances Reserved Instances provide you with a significant discount (up to 75%) compared to On-Demand instance pricing. In addition, when Reserved Instances are assigned to a specific Availability Zone, they provide a capacity reservation, giving you additional confidence in your ability to launch instances when you need them. For applications that have steady state or predictable usage, Reserved Instances can provide significant savings compared to using On-Demand instances. See How to Purchase Reserved Instances for more information. Reserved Instances are recommended for: - Applications with steady state usage - Applications that may require reserved capacity - Customers that can commit to using EC2 over a 1 or 3 year term to reduce their total computing costs https://aws.amazon.com/ec2/pricing/


iPQ.No	Correct Answer	Points	Explanation
2AWSCCP-CLF30	3	15.4	Answer = 3 Amazon Web Services Web Application Firewall (AWS WAF) Explanation : AWS Web Application Firewall (AWS WAF) is a web application firewall that helps protect your web applications or APIs from common web exploits that can compromise availability, security, or consume excessive resources. AWS Web Application Firewall (AWS WAF) gives you control over the traffic that reaches your applications by allowing you to create security rules that block common attack patterns like SQL injection and cross-site scripting. You can also mitigate Web layer DDoS attacks with rate-based rules. How the AWS Web Application Firewall (AWS WAF) Works:  via - https://aws.amazon.com/waf/ An SQL injection attack exploits any of the known SQL flaws that allow the SQL server to execute malicious code. For instance, if a SQL server is susceptible to an injection attack, an attacker may be able to go to a website's search box and enter code that forces the site's SQL server to dump all of its stored usernames and passwords. Similar to a SQL injection attack, a cross-site scripting attack involves injecting malicious code into a website; however, the website itself is not under attack. Instead, the malicious code injected by the attacker runs only in the user's browser when they visit the compromised website, and it targets the visitor directly, not the website.

iPQ.No	Correct Answer	Points	Explanation
2AWS CCP-CLF31	3	15.4	Answer = 3 Both A and B Explanation : The AWS Well-Architected Framework provides guidance for constructing infrastructure for cloud-based applications that is secure, high-performing, resilient, and efficient. Based on six pillars — operational excellence, security, reliability, performance efficiency, cost optimization, and sustainability — the framework offers customers and partners a consistent method for evaluating scalable architectures and implementing designs.
2AWS CCP-CLF32	3	15.4	Answer = 3 Both A and B Explanation : The AWS Service Catalog enables organizations to create and manage catalogs of AWS-approved IT services. These IT services may encompass everything from virtual machine images, servers, software, and databases to multi-tier application architectures. AWS Partner Network (APN): Organizations can use the AWS Partner Network (APN) to identify the appropriate AWS services for developing solutions on AWS Cloud. The AWS Partner Network (APN) is a global partner program for technology and consulting firms that use Amazon Web Services to develop customer solutions and services.

iPQ.No	Correct Answer	Points	Explanation
2AWSCCP-CLF33	3	15.4	Answer = 3 Both A and B Explanation : The email address and password used to register for AWS services are the credentials for the AWS root user account. Therefore, the root user account has complete access to all AWS resources associated with that account. Access restrictions to the root user account are not possible. Multi-factor authentication (MFA) should be enabled on the root user account as a best practice. The password for the root user account can be changed after account creation. Create individual user accounts using AWS IAM for each employee performing administrative tasks, and assign administrative permissions as needed. AWS Root User Account Best Practices for Security: The AWS Account Root User PDF Kindle RSS When you first create an Amazon Web Services (AWS) account, you begin with a single sign-in identity that has complete access to all AWS services and resources in the account. This identity is called the AWS account <i>root user</i> and is accessed by signing in with the email address and password that you used to create the account. ⚠ Important We strongly recommend that you do not use the root user for your everyday tasks, even the administrative ones. Instead, adhere to the best practice of using the root user only to create your first IAM user. Then securely lock away the root user credentials and use them to perform only a few account and service management tasks. To view the tasks that require you to sign in as the root user, see AWS Tasks That Require Root User. For a tutorial on how to set up an administrator for daily use, see Creating Your First IAM Admin User and Group. You can create, rotate, disable, or delete access keys (access key IDs and secret access keys) for your AWS account root user. You can also change your root user password. Anyone who has root user credentials for your AWS account has unrestricted access to all the resources in your account, including billing information. When you create access keys, you create the access key ID and secret access key as a set. During access key creation, AWS gives you one opportunity to view and download the secret access key part of the access key. If you don't download it or if you lose it, you can delete the access key and then create a new one. You can create root user access keys with the IAM console, AWS CLI, or AWS API. A newly created access key has the status of <i>active</i>, which means that you can use the access key for CLI and API calls. You are limited to two access keys for each IAM user, which is useful when you want to rotate the access keys. You can also assign up to two access keys to the root user. When you disable an access key, you can't use it for API calls, and inactive keys do count toward your limit. You can create or delete an access key any time. However, when you delete an access key, it's gone forever and can't be retrieved. You can change the email address and password on the Security Credentials page. You can also choose Forgot password? on the AWS sign-in page to reset your password. via: https://docs.aws.amazon.com/IAM/latest/UserGuide/id_root-user.html Lock Away Your AWS Account Root User Access Keys You use an access key (an access key ID and secret access key) to make programmatic requests to AWS. However, do not use your AWS account root user access key. The access key for your AWS account root user gives full access to all your resources for all AWS services, including your billing information. You cannot reduce the permissions associated with your AWS account root user access key.


iPQ.No	Correct Answer	Points	Explanation
2AWSACP-CLF34	1	15.4	Answer = 1 Amazon RDS Multi-AZ improves the availability of databases Explanation : Amazon RDS multi-AZ deployments provide improved availability and durability for Amazon Relational Database Service (Amazon RDS) instances, making them an obvious choice for production database workloads. Amazon RDS automatically creates a primary DB instance and synchronously replicates the data to a standby instance in a different Availability Zone (AZ) when you provision an Amazon RDS Multi-AZ Instance with one standby. Amazon RDS performs an automatic failover to the standby in the event of an infrastructure failure, allowing you to resume database operations as soon as the failover is complete. How Amazon RDS Multi-AZ Works  via - https://aws.amazon.com/rds/features/multi-az/ Exam Alert: Please compare Multi-AZ, Multi-Region, and Read Replica deployments for RDS at Read replicas, Multi-AZ deployments, and multi-region deployments Amazon RDS read replicas complement Multi-AZ deployments. While both features maintain a second copy of your data, there are differences between the two: Multi-AZ deployments Multi-Region deployments Read replicas

iPQ.No	Correct Answer	Points	Explanation
2AWSCCP-CLF35	3	15.4	Answer = 3 IAM is AWS Identity and Access Management Explanation : AWS Identity and Access Management (IAM) permits secure management of access to AWS services and resources. Using IAM, you can create and manage AWS users and groups and grant or deny their access to AWS resources using permissions. IAM enables security best practices by permitting the assignment of unique security credentials to users and groups in order to specify which AWS service APIs and resources they can access. These features make IAM a crucial service for the overall security of your AWS account's resources. IAM is secure by default; users do not have access to AWS resources unless explicit permissions are granted.

IPQ.No	Correct Answer	Points	Explanation
2AWSCCP-CLF36	3	15.4	Answer = 3 Trusted Advisor on AWS Explanation : AWS Trusted Advisor is an online tool that provides real-time guidance to help you provision your resources in accordance with AWS best practices for cost optimization, security, fault tolerance, service limits, and performance enhancement. Whether establishing new workflows, developing applications, or as part of ongoing enhancements, Trusted Advisor's recommendations help maintain the optimal provisioning of your solutions. All AWS clients have access to the seven core Trusted Advisor checks, which enhance the security and performance of the AWS environment. How Trusted Advisor Works:  via - https://aws.amazon.com/premiumsupport/technology/trusted-advisor/ AWS Trusted Advisor Recommendations: Like your customized cloud expert, AWS Trusted Advisor analyzes your AWS environment and provides best practice recommendations in five categories: Cost Optimization  0  9  0  \$7,516.85 Potential monthly savings Performance  3  7  0  Security  2  4  11  Fault Tolerance  0  15  5  Service Limits  37  0  1  Core Checks & Recommendations All AWS customers get access to the seven core Trusted Advisor checks to help increase the security and performance of the AWS environment. Checks include: Security - S3 Bucket Permissions - Security Groups - Specific Ports Unrestricted - IAM Use - MFA on Root Account - EBS Public Snapshots - RDS Public Snapshots Service Limits via - https://aws.amazon.com/premiumsupport/technology/trusted-advisor/


iPQ.No	Correct Answer	Points	Explanation
2AWSCCP-CLF37	1	15.4	Answer = 1 Volume Gateway, Tape Gateway, and File Gateway Explanation : AWS Storage Gateway is a hybrid cloud storage service that links your on-premises environments to the AWS Cloud. Storage Gateway is used by customers to simplify storage management and reduce costs for key hybrid cloud storage use cases. These include the migration of tape backups to the cloud, the reduction of on-premises storage with cloud-backed file shares, the provision of low-latency access to data in AWS for on-premises applications, and a variety of migration, archiving, processing, and disaster recovery use cases. The AWS Storage Gateway service offers three distinct gateway types: tape gateway, file gateway, and volume gateway. These gateways connect on-premises applications to cloud storage and cache data locally for low-latency access. Overview of Gateway Storage Types: Gateway Types File Gateway The File Gateway presents a file interface that enables you to store files as objects in Amazon S3 using the industry-standard NFS and SMB file protocols, and access those files via NFS and SMB from your datacenter or Amazon EC2, or access those files as objects with the S3 API. POSIX-style metadata, including ownership, permissions, and timestamps are durably stored in Amazon S3 in the user-metadata of the object associated with the file. Once objects are transferred to S3, they can be managed as native S3 objects, and bucket policies such as versioning, lifecycle management, and cross-region replication and apply directly to objects stored in your bucket. Customers use the File Gateway to store file data into S3 for use by object-based workloads including data analytics or machine learning, as a cost-effective storage target for backups, and as a repository or tier in the cloud for application file storage. Learn More » Tape Gateway The Tape Gateway presents itself to your existing backup application as an industry-standard iSCSI-based virtual tape library (VTL), consisting of a virtual media changer and virtual tape drives. You can continue to use your existing backup applications and workflows while writing to a nearly limitless collection of virtual tapes. Each virtual tape is stored in Amazon S3. When you no longer require immediate or frequent access to data contained on a virtual tape, you can have your backup application move it from the Storage Gateway Virtual Tape Library into an archive tier that sits on top of Amazon S3 Glacier cloud storage, further reducing storage costs. Storage Gateway is currently compatible with most leading backup applications. The Tape Gateway's VTL interface eliminates large upfront tape automation capital expenses, multi-year maintenance contract commitments, and ongoing media costs. You pay only for the capacity you use and scale as your needs grow. The need to transport storage media to offsite facilities and handle tape media manually goes away, and your archives benefit from the design and durability of the AWS Cloud platform. Learn More » Volume Gateway The Volume Gateway presents your applications block storage volumes using the iSCSI protocol. Data written to these volumes can be asynchronously backed up as point-in-time snapshots of your volumes, and stored in the cloud as Amazon EBS snapshots. You can back up your on-premises Volume Gateway volumes using the service's native snapshot scheduler or the AWS Backup service. In both cases, volume backups are stored as Amazon EBS snapshots in AWS. These snapshots are incremental backups that capture only changed blocks. All snapshot storage is also compressed to minimize your storage charges. When connecting to the Volume Gateway with the iSCSI block interface, you can run the gateway in two modes: cached and stored. In cached mode, you store your primary data in Amazon S3 and retain your frequently accessed data locally in cache. With this mode, you can achieve substantial cost savings on primary storage, minimizing the need to scale your storage on-premises, while retaining low-latency access to your frequently accessed data.


iPQ.No	Correct Answer	Points	Explanation
2AWS CCP-CLF38	2	15.4	Answer = 2 The AWS Market Explanation : AWS Marketplace is a digital catalog containing thousands of software listings from independent software vendors, making it simple to discover, test, purchase, and deploy software that runs on AWS. AWS Marketplace contains thousands of software listings in popular categories such as security, networking, storage, machine learning, Internet of Things, business intelligence, databases, and DevOps. You may utilize AWS Marketplace as either a buyer (subscriber) or a seller (provider), or both. Anyone with an AWS account can use AWS Marketplace as a buyer or seller.
2AWS CCP-CLF39	1	15.4	Answer = 1 You can deploy AWS Shield automatically when a DDoS attack is detected Explanation : Auto-scaling on AWS is useful during DDoS attacks because it can rapidly scale out resources. However, it cannot deploy the AWS Shield service automatically to its group of resources.


iPQ.No	Correct Answer	Points	Explanation
2AWSCCP-CLF40	3	15.4	Answer = 3 Amazon Elastic Container Registry (Amazon ECR) is a registry for container images Explanation : The Amazon Elastic Container Registry (Amazon ECR) can be utilized to store, manage, and deploy Docker container images. Amazon Elastic Container Registry (Amazon ECR) eliminates the requirement to operate container repositories. Then, you can retrieve Docker images from Amazon Elastic Container Registry (Amazon ECR) and run them with Amazon Elastic Container Service (Amazon ECS). To understand how Amazon Elastic Container Registry (Amazon ECR) operates, please refer to this diagram:  https://aws.amazon.com/ecr/
2AWSCCP-CLF41	3	15.4	Answer = 3 Only outgoing data transfers incur fees Explanation : One of the primary advantages of cloud services is the ability to optimize costs to meet changing business requirements. AWS services have no complex dependencies or licensing requirements, allowing you to build innovative, cost-effective solutions with the most recent technology. AWS has three primary cost drivers: computation, storage, and outbound data transfer. Depending on the AWS product and pricing model that you select, these characteristics may vary. Outbound data from all AWS regions to the Internet is billed at region-specific, tiered data transfer rates. Transferring data from the internet into all AWS regions is free.

iPQ.No	Correct Answer	Points	Explanation
2AWS CCP-CLF42	4	15.4	Answer = 4 AWS Elastic Beanstalk Explanation : AWS Elastic Beanstalk is an easy-to-use service for deploying and scaling web applications and services built with Java,.NET, PHP, Node.js, Python, Ruby, Go, and Docker on well-known servers including Apache, Nginx, Passenger, and IIS. AWS Elastic Beanstalk handles the deployment, including capacity provisioning, load balancing, auto-scaling, and application health monitoring, automatically. In addition, you retain complete control over the AWS resources that power your application and can always access the underlying resources. AWS Elastic Beanstalk is free; you only pay for the AWS resources required to store and execute your applications. Advantages of AWS Elastic Beanstalk, accessed from: Fast and simple to begin Elastic Beanstalk is the fastest and simplest way to deploy your application on AWS. You simply use the AWS Management Console, a Git repository, or an integrated development environment (IDE) such as Eclipse or Visual Studio to upload your application, and Elastic Beanstalk automatically handles the deployment details of capacity provisioning, load balancing, auto-scaling, and application health monitoring. Within minutes, your application will be ready to use without any infrastructure or resource configuration work on your part. Developer productivity Elastic Beanstalk provisions and operates the infrastructure and manages the application stack (platform) for you, so you don't have to spend the time or develop the expertise. It will also keep the underlying platform running your application up-to-date with the latest patches and updates. Instead, you can focus on writing code rather than spending time managing and configuring servers, databases, load balancers, firewalls, and networks. Impossible to outgrow Elastic Beanstalk automatically scales your application up and down based on your application's specific need using easily adjustable Auto Scaling settings. For example, you can use CPU utilization metrics to trigger Auto Scaling actions. With Elastic Beanstalk, your application can handle peaks in workload or traffic while minimizing your costs. Complete resource control You have the freedom to select the AWS resources, such as Amazon EC2 instance type, that are optimal for your application. Additionally, Elastic Beanstalk lets you "open the hood" and retain full control over the AWS resources powering your application. If you decide you want to take over some (or all) of the elements of your infrastructure, you can do so seamlessly by using Elastic Beanstalk's management capabilities. https://aws.amazon.com/elasticbeanstalk/

iPQ.No	Correct Answer	Points	Explanation
2AWS CCP-CLF43	4	15.4	Answer = 4 Google Athena Explanation : Amazon Athena is an interactive query service that facilitates standard SQL analysis of Amazon S3 data. Amazon Athena is serverless, meaning there is no infrastructure to manage, and you only pay for the queries you execute. Amazon Athena's key features can be found at: Start querying instantly Serverless, no ETL Athena is serverless. You can quickly query your data without having to setup and manage any servers or data warehouses. Just point to your data in Amazon S3, define the schema, and start querying using the built-in query editor. Amazon Athena allows you to tap into all your data in S3 without the need to set up complex processes to extract, transform, and load the data (ETL). Pay per query Only pay for data scanned With Amazon Athena, you pay only for the queries that you run. You are charged \$5 per terabyte scanned by your queries. You can save from 30% to 90% on your per-query costs and get better performance by compressing, partitioning, and converting your data into columnar formats. Athena queries data directly in Amazon S3. There are no additional storage charges beyond S3. Open, powerful, standard Built on Presto, runs standard SQL Amazon Athena uses Presto with ANSI SQL support and works with a variety of standard data formats, including CSV, JSON, ORC, Avro, and Parquet. Athena is ideal for quick, ad-hoc querying but it can also handle complex analysis, including large joins, window functions, and arrays. Amazon Athena is highly available; and executes queries using compute resources across multiple facilities and multiple devices in each facility. Amazon Athena uses Amazon S3 as its underlying data store, making your data highly available and durable. Fast, really fast Interactive performance even for large datasets With Amazon Athena, you don't have to worry about having enough compute resources to get fast, interactive query performance. Amazon Athena automatically executes queries in parallel, so most results come back within seconds. https://aws.amazon.com/athena/ Simply point Amazon Athena to your Amazon Simple Storage Service (Amazon S3) data, define the schema, and begin standard SQL queries to use Amazon Athena. The majority of results are delivered in seconds. There is no need for complex ETL jobs to prepare your data for analysis with Amazon Athena. This makes it simple for anyone with SQL knowledge to analyze large datasets efficiently.



iPQ.No	Correct Answer	Points	Explanation
2AWSCCP-CLF44	4	15.4	Answer = 4 Amazon Elastic File System (Amazon EFS) Explanation : Amazon Elastic File System (Amazon EFS) is an elastic NFS file system that is simple, scalable, and fully managed. It is designed to scale on-demand to petabytes without disrupting applications, expanding and contracting automatically as files are added or removed, removing the need to provision and manage capacity to accommodate growth. Amazon EFS is designed to provide massively parallel shared access to thousands of Amazon EC2 instances, allowing your applications to achieve high levels of aggregate throughput and IOPS with low latencies. Amazon Elastic File System (Amazon EFS):  via - https://aws.amazon.com/efs/

iPQ.No	Correct Answer	Points	Explanation
2AWSCCP-CLF45	3	15.4	Answer = 3 AWS Global Accelerator Explanation : AWS Global Accelerator is a service that enhances the performance and availability of your applications for local and global users. It provides static IP addresses that function as a fixed entry point to your application endpoints in a single or multiple AWS regions, including your application load balancers, network load balancers, and Amazon EC2 instances. AWS Global Accelerator uses the AWS global network to optimize the path from your users to your applications, thereby enhancing the performance of your traffic by up to 60%. AWS Global Accelerator enhances the performance of a variety of TCP or UDP applications by proxying packets at the edge to applications running in one or more AWS regions. AWS Global Accelerator is suitable for non-HTTP use cases, such as gaming (UDP), IoT (MQTT), and voice over IP, as well as HTTP use cases that require static IP addresses or fast, deterministic regional failover. How AWS Global Accelerator Works: How it works  Directly access web applications Without AWS Global Accelerator  It can take many networks to reach the application. Paths to and from the application may differ. Each hop impacts performance and can introduce risks. With AWS Global Accelerator  Adding AWS Global Accelerator removes these inefficiencies. It leverages the Global AWS Network, resulting in improved performance. aws.amazon.com/global-accelerator Exam Alert:

iPQ.No	Correct Answer	Points	Explanation
2AWS CCP-CLF46	3	15.4	Answer = 3 Amazon CloudHSM Explanation : AWS CloudHSM is a cloud-based Hardware Security Module (HSM) that allows you to generate and use encryption keys in the AWS Cloud with ease. With CloudHSM, you can manage your encryption keys using FIPS 140-2 Level 3 HSMs that have been validated. It is a fully managed service that automates tedious administrative tasks like hardware provisioning, software patching, high availability, and backups. Please see this detailed CloudHSM description at  The diagram illustrates the AWS CloudHSM architecture. It shows an AWS Cloud icon at the top left. Below it, a box labeled 'VPC' contains an 'AWS CLOUDHSM' icon (a blue padlock with a key) and an 'Application HSM Client' icon (a blue box with a key). The 'Application HSM Client' is connected to a 'VPC INSTANCE' icon (a blue box with a key) via a dashed line labeled 'SSL'. The 'VPC INSTANCE' is also connected to the 'AWS CLOUDHSM' icon via a dashed line labeled 'SSL'. The entire VPC box is connected to the AWS Cloud icon via a dashed line labeled 'VPC'. AWS CloudHSM runs in your own Amazon Virtual Private Cloud (VPC), enabling you to easily use your HSMs with applications running on your Amazon EC2 instances. With CloudHSM, you can use standard VPC security controls to manage access to your HSMs. Your applications connect to your HSMs using mutually authenticated SSL channels established by your HSM client software. Since your HSMs are located in Amazon datacenters near your EC2 instances, you can reduce the network latency between your applications and HSMs versus an on-premises HSM. A: AWS manages the hardware security module (HSM) appliance, but does not have access to your keys B: You control and manage your own keys C: Application performance improves (due to close proximity with AWS workloads) D: Secure key storage in tamper-resistant hardware available in multiple Availability Zones (AZs) E: Your HSMs are in your Virtual Private Cloud (VPC) and isolated from other AWS networks. Separation of duties and role-based access control is inherent in the design of the AWS CloudHSM. AWS monitors the health and network availability of your HSMs but is not involved in the creation and management of the key material stored within your HSMs. You control the HSMs and the generation and use of your encryption keys. https://aws.amazon.com/cloudhsm/.

iPQ.No	Correct Answer	Points	Explanation
2AWSCCP-CLF47	4	15.4	Answer = 4 AWS Corporations Explanation : AWS Organizations enables the centralized management of billing, access, compliance, and security, as well as the sharing of resources across multiple AWS accounts. Using AWS Organizations, you can automate account creation, create groups of accounts that reflect your business requirements, and apply policies for governance to these groups. Additionally, you can streamline billing by configuring a single payment method for all of your AWS accounts. AWS Organizations is available at no additional cost to all AWS customers. Key AWS Organizations Features: CENTRALLY MANAGE POLICIES ACROSS MULTIPLE AWS ACCOUNTS To improve control over your AWS environment, you can use AWS Organizations to create groups of accounts, and then attach policies to a group to ensure the correct policies are applied across the accounts without requiring custom scripts and manual processes. GOVERN ACCESS TO AWS SERVICES, RESOURCES, AND REGIONS AWS Organizations allows you to restrict what services and actions are allowed in your accounts. You can use Service Control Policies (SCPs) to apply permission guardrails on AWS Identity and Access Management (IAM) users and roles. For example, you can apply an SCP that restricts users in accounts in your organization from launching any resources in regions that you do not explicitly allow. AUTOMATE AWS ACCOUNT CREATION AND MANAGEMENT AWS Organizations helps you simplify IT operations by automating AWS account creation and management. The Organizations APIs enable you to create new accounts programmatically, and to add the new accounts to a group. The policies attached to the group are automatically applied to the new account. For example, you can automate the creation of new accounts for workload or application isolation and grant entities in those accounts access only to the necessary AWS services. CONFIGURE AWS SERVICES ACROSS MULTIPLE ACCOUNTS AWS Organizations helps you configure AWS services and share resources across accounts in your organization. For example, Organizations integrates with AWS Single Sign-on to enable you to easily provision access for all of your developers to accounts in your organization from a single place. You can make central changes to access permissions and have them automatically updated on accounts in your organization. CONSOLIDATE BILLING ACROSS MULTIPLE AWS ACCOUNTS You can use AWS Organizations to set up a single payment method for all the AWS accounts in your organization through consolidated billing. With consolidated billing, you can see a combined view of charges incurred by all your accounts, as well as take advantage of pricing benefits from aggregated usage, such as volume discounts for Amazon EC2 and Amazon S3.

iPQ.No	Correct Answer	Points	Explanation
2AWS CCP-CLF48	2	15.4	Answer = 2 Amazon Simple Storage Service (Amazon S3) is object-based storage, while Amazon Elastic Block Store (Amazon EBS) and Amazon Elastic File System (Amazon EFS) are block-based and file-based storage, respectively Explanation : Amazon Elastic File System (Amazon EFS) is an elastic NFS file system that is simple, scalable, and fully managed. Amazon Elastic Block Store (Amazon EBS) is a high-performance, easy-to-use block storage service designed for use with Amazon Elastic Compute Cloud (Amazon EC2) for throughput- and transaction-intensive workloads at any scale. Amazon S3 is an object storage service with industry-leading scalability, data availability, security, and performance.

iPQ.No	Correct Answer	Points	Explanation
2AWS CCP-CLF49	1	15.4	Answer = 1 Resize images to create thumbnails quickly Explanation : Amazon Rekognition cannot be used to resize photos to create thumbnails. You can identify objects, people, text, scenes, and activities in images and videos, as well as detect any inappropriate content, using Amazon Rekognition. Amazon Rekognition also offers highly accurate facial analysis and facial search capabilities that can be used for a variety of user verification, people counting, and public safety applications. Amazon Recognition Use-Cases: Key features  Labels With Amazon Rekognition, you can identify thousands of objects (such as bike, telephone, building), and scenes (such as parking lot, beach, city). When analyzing video, you can also identify specific activities such as "delivering a package" or "playing soccer". Learn more  Custom labels With Amazon Rekognition Custom Labels, you can extend the detection capabilities of Amazon Rekognition to extract information from images that is uniquely helpful to your business. For example, you can find your corporate logo in social media, identify your products on store shelves, classify your machine parts in an assembly line, or detect your animated characters in videos. Learn more  Content moderation Amazon Rekognition helps you identify potentially unsafe or inappropriate content across both image and video assets and provides you with detailed labels that allow you to accurately control what you want to allow based on your needs. Use Amazon A2I to enhance the accuracy of Amazon Rekognition image moderation predictions using human review. Learn more  Text detection In photos and videos, text appears very differently than neat words on a printed page. Amazon Rekognition can read skewed and distorted text to capture information like store names, forced narratives overlaid on media, street signs, and text on product packaging. Learn more via: https://aws.amazon.com/rekognition/  

iPQ.No	Correct Answer	Points	Explanation
2AWSCCP-CLF50	4	15.4	Answer = 4 Amazon Rekognition Explanation : Three types of cloud computing exist: infrastructure as a service (IaaS), platform as a service (PaaS), and software as a service (SaaS). IaaS comprises the fundamental components of cloud IT. Typically, it offers access to networking capabilities, computers (virtual or on dedicated hardware), and data storage space. IaaS offers the greatest degree of IT resource flexibility and management control. Amazon EC2 (on AWS), GCP, Azure, Rackspace, Digital Ocean, and Linode are some examples. PaaS eliminates the need to manage underlying infrastructure (typically hardware and operating systems) and enables you to focus on application deployment and management. You do not have to be concerned with resource acquisition, capacity planning, software maintenance, patching, or any of the other undifferentiated heavy lifting required to run your application. Elastic Beanstalk (on AWS), Heroku, Google App Engine (GCP), and Windows Azure (Microsoft) are examples of such platforms. SaaS offers a comprehensive product that is operated and managed by the service provider. With a SaaS offering, you are not responsible for maintaining the service or managing the underlying infrastructure. You only need to consider how you will utilize the software in question. Amazon Rekognition, Google Apps (Gmail), Dropbox, and Zoom are a few examples. Types of Cloud Computing: Cloud Computing Models There are three main models for cloud computing. Each model represents a different part of the cloud computing stack.

iPQ.No	Correct Answer	Points	Explanation
2AWSCCP-CLF51	4	15.4	Answer = 4 Amazon FSx File Server for Windows Explanation : Amazon FSx for Windows File Server offers fully managed, highly reliable, and scalable file storage that is accessible via the industry-standard Service Message Block (SMB) protocol. It is built on Windows Server and offers a wide range of administrative features including user quotas, end-user file restore, and Microsoft Active Directory (AD) integration. Amazon FSx provides high levels of throughput, IOPS, and consistent sub-millisecond latencies to support a wide range of workloads. Amazon FSx is accessible from Windows, Linux, and macOS compute instances and devices. For Windows-based applications, Amazon FSx offers fully managed Windows file servers with features and performance optimized for "lift-and-shift" business-critical application workloads, such as home directories (user shares), media workflows, and ERP applications.
2AWSCCP-CLF52	1	15.4	Answer = 1 IAM Client Explanation : Access keys are long-term credentials for an IAM user or the AWS account root user. Access keys consist of two parts: an access key ID and a secret access key.
2AWSCCP-CLF53	3	15.4	Answer = 3 AWS Systems Manager Session Manager Explanation : It helps provide secure and auditable instance management without the need to open inbound ports, maintain bastion hosts, or manage SSH keys. AWS Systems Manager Session Manager helps to enable compliance with corporate policies that require controlled access to instances, increase security and auditability of access to the instance, and improve security and auditability of access to the instance.

iPQ.No	Correct Answer	Points	Explanation
2AWSCCP-CLF54	1	15.4	Answer = 1 Amazon EMR Explanation : Amazon EMR is the industry-leading cloud big data platform for processing vast amounts of data using open source tools such as Hadoop, Apache Spark, Apache Hive, Apache HBase, Apache Flink, Apache Hudi, and Presto.

iPQ.No	Correct Answer	Points	Explanation
2AWSCCP-CLF55	4	15.4	Answer = 4 All of the above Explanation : AWS Shield: AWS Shield is a managed distributed denial of service (DDoS) protection service that protects applications running on AWS. AWS Shield offers always-on detection and automatic inline mitigations that minimize application downtime and latency. AWS Web Application Firewall (AWS WAF): Using AWS Web Application Firewall (AWS WAF), you can configure web access control lists (Web ACLs) on your Amazon CloudFront distributions or application load balancers to filter and block requests based on request signatures. In addition, by using the AWS Web Application Firewall's (AWS WAF) rate-based rules, you can automatically block the IP addresses of bad actors when requests matching a rule exceed a threshold that you Amazon CloudFront with Amazon Route 53: AWS hosts Amazon CloudFront and Amazon Route 53 services on a distributed network of proxy servers in data centers called edge locations. Using the global Amazon network of edge locations for application delivery and DNS service is crucial for building a comprehensive defense against DDoS attacks against your dynamic web applications. AWS Shield, AWS Web Application Firewall (AWS WAF), and Amazon CloudFront with Amazon Route 53 mitigate DDoS attacks. How AWS Shield, CloudFront, and Route 53 work to help protect against DDoS attacks To help keep your dynamic web applications available when they are under DDoS attack, the steps in this post enable AWS Shield Standard by configuring your applications behind CloudFront and Route 53. AWS Shield Standard protects your resources from common, frequently occurring network and transport layer DDoS attacks. Attack traffic can be geographically isolated and absorbed using the capacity in edge locations close to the source. Additionally, you can configure geographical restrictions to help block attacks originating from specific countries. The request-routing technology in CloudFront connects each client to the nearest edge location, as determined by continuously updated latency measurements. HTTP and HTTPS requests sent to CloudFront can be monitored, and access to your application resources can be controlled at edge locations using AWS WAF. Based on conditions that you specify in AWS WAF, such as the IP addresses that requests originate from or the values of query strings, traffic can be allowed, blocked, or allowed and counted for further investigation or remediation. The following diagram shows how static and dynamic web application content can originate from endpoint resources within AWS or your corporate data center. For more details, see How CloudFront Delivers Content and How CloudFront Works with Regional Edge Caches. 

iPQ.No	Correct Answer	Points	Explanation
2AWSCCP-CLF56	4	15.4	Answer = 4 All of the above Explanation : AWS services can be accessed through three distinct methods: This is a straightforward web-based interface for accessing AWS services. AWS Command Line Interface (AWS CLI): You can automate service management with scripts and access AWS services from the command line. You can also access the AWS Software Development Kit (SDK), which provides language-specific abstracted APIs for AWS services.


iPQ.No	Correct Answer	Points	Explanation																																								
2AWS CCP-CLF57	2	15.4	Answer = 2 AWS Business Support Explanation : AWS provides its customers with four distinct support plans: AWS Developer Support, AWS Business Support, AWS Enterprise On-Ramp Support, and AWS Enterprise Support. Every AWS customer receives a basic support plan. AWS Business Support: Use the AWS Business Support plan if you have production workloads on AWS and require technical support and architectural guidance in the context of your specific use cases 24x7 via phone, email, and chat. For the given use case, the AWS Business Support plan is the most cost-effective option. Exam Alert: Please review the differences between the AWS Developer Support, AWS Business Support, AWS Enterprise On-Ramp Support, and AWS Enterprise Support plans, as you can anticipate at least two questions on the exam: <table><thead><tr><th></th><th>Developer</th><th>Business</th><th>Enterprise On-Ramp</th><th>Enterprise</th></tr><tr><td></td><td>Recommended if you are experimenting or testing in AWS.</td><td>Minimum recommended tier if you have production workloads in AWS</td><td>Recommended if you have production and/or business critical workloads in AWS.</td><td>Recommended if you have business and/or mission critical workloads in AWS.</td></tr></thead><tbody><tr><td>AWS Trusted Advisor Best Practice Checks</td><td>Service Quota and basic Security checks</td><td>Full set of checks</td><td>Full set of checks</td><td>Full set of checks</td></tr><tr><td>AWS Trusted Advisor Priority</td><td></td><td></td><td></td><td>Prioritized recommendations curated by your AWS account team</td></tr><tr><td>Enhanced Technical Support</td><td>Business hours** web access to Cloud Support Associates Unlimited cases with 1 primary contact Prioritized responses on AWS re:Post</td><td>24/7 phone, web, and chat access to Cloud Support Engineers Unlimited cases and unlimited contacts (IAM supported) Prioritized responses on AWS re:Post Access to AWS Support App in Slack</td><td>24/7 phone, web, and chat access to Cloud Support Engineers Unlimited cases and unlimited contacts (IAM supported) Prioritized responses on AWS re:Post Access to AWS Support App in Slack</td><td>24/7 phone, web, and chat access to Cloud Support Engineers Unlimited cases and unlimited contacts (IAM supported) Prioritized responses on AWS re:Post Access to AWS Support App in Slack</td></tr><tr><td>Case Severity / Response Times*</td><td>General guidance: < 24 hours** System impaired: < 12 hours**</td><td>General guidance: < 24 hours System impaired: < 12 hours Production system impaired: < 4 hours Production system down: < 1 hour</td><td>General guidance: < 24 hours System impaired: < 12 hours Production system impaired: < 4 hours Production system down: < 1 hour Business-critical system down: < 30 minutes</td><td>General guidance: < 24 hours System impaired: < 12 hours Production system impaired: < 4 hours Production system down: < 1 hour Business/Mission-critical system down: < 15 minutes</td></tr><tr><td>Architectural Guidance</td><td>General</td><td>Contextual to your use-cases</td><td>Consultative review and guidance based on your applications (one-per-year)</td><td>Consultative review and guidance based on your applications</td></tr><tr><td>Programmatic Case Management</td><td></td><td>AWS Support API</td><td>AWS Support API</td><td>AWS Support API</td></tr></tbody></table>		Developer	Business	Enterprise On-Ramp	Enterprise		Recommended if you are experimenting or testing in AWS.	Minimum recommended tier if you have production workloads in AWS	Recommended if you have production and/or business critical workloads in AWS.	Recommended if you have business and/or mission critical workloads in AWS.	AWS Trusted Advisor Best Practice Checks	Service Quota and basic Security checks	Full set of checks	Full set of checks	Full set of checks	AWS Trusted Advisor Priority				Prioritized recommendations curated by your AWS account team	Enhanced Technical Support	Business hours** web access to Cloud Support Associates Unlimited cases with 1 primary contact Prioritized responses on AWS re:Post	24/7 phone, web, and chat access to Cloud Support Engineers Unlimited cases and unlimited contacts (IAM supported) Prioritized responses on AWS re:Post Access to AWS Support App in Slack	24/7 phone, web, and chat access to Cloud Support Engineers Unlimited cases and unlimited contacts (IAM supported) Prioritized responses on AWS re:Post Access to AWS Support App in Slack	24/7 phone, web, and chat access to Cloud Support Engineers Unlimited cases and unlimited contacts (IAM supported) Prioritized responses on AWS re:Post Access to AWS Support App in Slack	Case Severity / Response Times*	General guidance: < 24 hours** System impaired: < 12 hours**	General guidance: < 24 hours System impaired: < 12 hours Production system impaired: < 4 hours Production system down: < 1 hour	General guidance: < 24 hours System impaired: < 12 hours Production system impaired: < 4 hours Production system down: < 1 hour Business-critical system down: < 30 minutes	General guidance: < 24 hours System impaired: < 12 hours Production system impaired: < 4 hours Production system down: < 1 hour Business/Mission-critical system down: < 15 minutes	Architectural Guidance	General	Contextual to your use-cases	Consultative review and guidance based on your applications (one-per-year)	Consultative review and guidance based on your applications	Programmatic Case Management		AWS Support API	AWS Support API	AWS Support API
	Developer	Business	Enterprise On-Ramp	Enterprise																																							
	Recommended if you are experimenting or testing in AWS.	Minimum recommended tier if you have production workloads in AWS	Recommended if you have production and/or business critical workloads in AWS.	Recommended if you have business and/or mission critical workloads in AWS.																																							
AWS Trusted Advisor Best Practice Checks	Service Quota and basic Security checks	Full set of checks	Full set of checks	Full set of checks																																							
AWS Trusted Advisor Priority				Prioritized recommendations curated by your AWS account team																																							
Enhanced Technical Support	Business hours** web access to Cloud Support Associates Unlimited cases with 1 primary contact Prioritized responses on AWS re:Post	24/7 phone, web, and chat access to Cloud Support Engineers Unlimited cases and unlimited contacts (IAM supported) Prioritized responses on AWS re:Post Access to AWS Support App in Slack	24/7 phone, web, and chat access to Cloud Support Engineers Unlimited cases and unlimited contacts (IAM supported) Prioritized responses on AWS re:Post Access to AWS Support App in Slack	24/7 phone, web, and chat access to Cloud Support Engineers Unlimited cases and unlimited contacts (IAM supported) Prioritized responses on AWS re:Post Access to AWS Support App in Slack																																							
Case Severity / Response Times*	General guidance: < 24 hours** System impaired: < 12 hours**	General guidance: < 24 hours System impaired: < 12 hours Production system impaired: < 4 hours Production system down: < 1 hour	General guidance: < 24 hours System impaired: < 12 hours Production system impaired: < 4 hours Production system down: < 1 hour Business-critical system down: < 30 minutes	General guidance: < 24 hours System impaired: < 12 hours Production system impaired: < 4 hours Production system down: < 1 hour Business/Mission-critical system down: < 15 minutes																																							
Architectural Guidance	General	Contextual to your use-cases	Consultative review and guidance based on your applications (one-per-year)	Consultative review and guidance based on your applications																																							
Programmatic Case Management		AWS Support API	AWS Support API	AWS Support API																																							

iPQ.No	Correct Answer	Points	Explanation
2AWS CCP-CLF58	3	15.4	Answer = 3 Both A and B Explanation : A "horizontally scalable" system is one whose capacity can be increased by adding more computers. This is in contrast to a "vertically scalable" system, which is limited to running its processes on a single computer; in such systems, the only way to improve performance is to add faster (or additional) CPUs, memory, or storage to a single computer. Horizontally scalable systems frequently outperform vertically scalable systems due to their ability to execute workloads in parallel and distribute them across multiple computers. ELB stands for Elastic Load Balancing. ELB distributes incoming application traffic across multiple targets, including Amazon EC2 instances, containers, IP addresses, and Lambda functions. It can manage the fluctuating application traffic load within a single availability zone (AZ) or across multiple AZs. This pertains to horizontal scaling. Amazon Relational Database Service (Amazon RDS) read replicas. Amazon Relational Database Service (Amazon RDS) simplifies the deployment, management, and scalability of cloud-based relational databases. Read replicas allow you to create synchronized read-only copies of your master database. Additionally, you can position your read replica in a different AWS region closer to your users to improve performance. Read replicas are an illustration of horizontal resource scaling.

iPQ.No	Correct Answer	Points	Explanation
2AWSCCP-CLF59	3	15.4	Answer = 3 Acceptable Use Policy for AWS Explanation : The Acceptable Use Policy outlines prohibited uses of the web services provided by Amazon Web Services, Inc. and its affiliates (the "Services") and the website accessible at http://aws.amazon.com (the "AWS Site"). This policy is available at https://aws.amazon.com/aup/ and is updated as necessary by AWS.
2AWSCCP-CLF60	4	15.4	Answer = 4 AWS Config Explanation : AWS Config is a service that enables you to assess, audit, and evaluate your AWS resource configurations. Config continuously monitors and records the configurations of your AWS resources and automates the comparison of recorded configurations to desired configurations. Consider resource-specific history, auditing, and compliance, and consider configuration. You can do the following using AWS Config: 1. Examine the configurations of your AWS resources for the desired settings. 2. Obtain a snapshot of the current configurations of supported AWS resources associated with your account. 3. Retrieve the configurations of one or more account-managed resources. 4. Get the configuration history of one or more resources. 5. Receive an alert when a resource is created, updated, or deleted. 6. View the connections between resources. For instance, you may wish to locate all the resources that utilize a specific security group.

iPQ.No	Correct Answer	Points	Explanation
2AWSCCP-CLF61	2	15.2	Answer = 2 Your Account Health Dashboard for AWS Explanation : AWS Health: Your Account Health Dashboard provides notifications and remediation guidance when AWS is experiencing events that may have an effect on you. With AWS Health: Your Account Health Dashboard, alerts are triggered by changes in the health of your AWS resources, providing visibility into events and guidance for diagnosing and resolving problems quickly. You can view the current status on this page at https://phd.aws.amazon.com/phd/home.

IPQ.No	Correct Answer	Points	Explanation				
2AWS CCP-CLF62	1	15.2	Answer = 1 A subnet spans only one availability zone (AZ) in the region, whereas an Amazon Virtual Private Cloud (Amazon VPC) spans all availability zones (AZ) in the region Explanation : Amazon Virtual Private Cloud (Amazon VPC) is a logically isolated section of the AWS Cloud in which you can launch AWS resources in a custom-defined virtual network. You have full control over your virtual networking environment, which includes the selection of your IP address range, the creation of subnets, and the configuration of route tables and network gateways. An Amazon Virtual Private Cloud (Amazon VPC) spans across all of the region's availability zones (AZ). Within your Amazon Virtual Private Cloud (Amazon VPC), a subnet is a range of IP addresses. A subnet only spans one availability zone (AZ) within a region. Overview of Amazon Virtual Private Cloud (Amazon VPC) and Subnets The following diagram shows a new VPC with an IPv4 CIDR block.  The main route table has the following routes. <table><thead><tr><th>Destination</th><th>Target</th></tr></thead><tbody><tr><td>10.0.0.0/16</td><td>local</td></tr></tbody></table> A VPC spans all of the Availability Zones in the Region. After creating a VPC, you can add one or more subnets in each Availability Zone. You can optionally add subnets in a Local Zone, which is an AWS infrastructure deployment that places compute, storage, database, and other select services closer to your end users. A Local Zone enables your end users to run applications that require single-digit millisecond latencies. For information about the Regions that support Local Zones, see Available Regions in the <i>Amazon EC2 User Guide for Linux Instances</i>. When you create a subnet, you specify the CIDR block for the subnet, which is a subset of the VPC CIDR block. Each subnet must reside entirely within one Availability Zone and cannot span zones. Availability Zones are distinct locations that are engineered to be isolated from failures in other Availability Zones. By launching instances in separate Availability Zones, you can protect your applications from the failure of a single location. We assign a unique ID to each subnet. via - https://docs.aws.amazon.com/vpc/latest/userguide/VPC_Subnets.html	Destination	Target	10.0.0.0/16	local
Destination	Target						
10.0.0.0/16	local						

iPQ.No	Correct Answer	Points	Explanation
2AWS CCP-CLF63	2	15.2	Answer = 2 Elasticity Explanation : The ability to acquire resources as needed and release them when they are no longer required is referred to as the "elasticity of the cloud." With cloud computing, you do not need to overprovision resources in advance to accommodate future peaks in business activity. Instead, you provide the required quantity of resources. You can instantly scale these resources up or down to increase or reduce capacity as your business's needs change. What exactly is Elasticity?  via Amazon: https://aws.amazon.com/what-is-cloud-computing/
2AWS CCP-CLF64	3	15.2	Answer = 3 Amazon DynamoDB Explanation : Amazon DynamoDB is a NoSQL database that supports key-value and document data models and enables developers to build serverless, modern applications that can start small and scale to support petabytes of data and tens of millions of read and write requests per second. Amazon DynamoDB supports both document and key-value data models. This allows Amazon DynamoDB to have a flexible schema, allowing each row to contain any number of columns at any given time. This allows you to easily adapt the table schema to changing business requirements, unlike relational databases.

iPQ.No	Correct Answer	Points	Explanation
2AWSCCP-CLF65	2	15.2	Answer = 2 Allows a single contact to open an unlimited number of cases Explanation : The AWS Developer Support plan permits an unlimited number of cases to be opened by a single primary contact.