



QUESTIONNAIRE REPORT SHEET

iPQ4AWS CCP.V1.1

iPQ.No	Correct Answer	Points	Explanation
4AWS CCP01	2	15.4	Answer = 2 Device Farm on AWS Explanation : With AWS Device Farm, you can test your web and mobile applications across a wide range of desktop browsers and actual mobile devices to improve their quality without having to set up and maintain any testing infrastructure. To speed up the execution of your test suite, the service lets you run your tests simultaneously on various desktop browsers or actual devices. It also provides videos and logs to help you rapidly spot problems with your app. Device Farm is made for programmers, QA groups, and customer service agents that create, test, and support mobile apps to improve the quality of their applications. Due to the variety of device models, different OS and firmware versions, carrier and manufacturer modifications, and reliance on other apps and remote services, application quality is becoming more and more crucial while also becoming more complex. Developers, QA, and support staff can perform both automated tests and manual tasks like reproducing customer issues, exploring new functionality, and carrying out manual test plans thanks to AWS Device Farm, which speeds up the development process by executing tests on multiple devices. By doing away with the requirement for internal device labs, lab managers, and automated infrastructure development, AWS Device Farm also provides significant savings. According to https://aws.amazon.com/device-farm/, how AWS Device Farm functions Testing on real mobile devices Automated Testing Test your app in parallel against a massive collection of physical devices in the AWS Cloud. Use one of our built-in frameworks, to test your applications without having to write or maintain test scripts, or use one of our supported automation testing frameworks.  <pre> graph LR A[Choose app to test Choose to test a native, hybrid, or web app] --> B[AWS Device Farm AWS Device Farm tests your app against your choice of real phones and tablets and allows developers to access their private fleet devices via their local host machines] B --> C[Get results Get results in minutes that pinpoint bugs & performance problems] </pre>


iPQ.No	Correct Answer	Points	Explanation
4AWS CCP02	3	15.4	Answer = 3 Both A and B Explanation : An online application called AWS Trusted Advisor offers real-time advice to help you supply your resources in accordance with AWS best practices. Regular recommendations from Trusted Advisor help maintain your solutions provisioned at their best, whether for creating new workflows, creating applications, or as part of continuing improvement. AWS Trusted Advisor examines your AWS infrastructure and offers best practice advice in the following five areas: cost reduction, performance improvement, security enhancement, fault tolerance, and service limit adjustments. The Trusted Advisor examines your AWS infrastructure and offers suggestions when there are chances to cut costs, enhance system performance, or plug security holes. It provides alerts on several of the most typical security configuration errors that can happen, such as failing to create IAM accounts for your internal users, granting public access to Amazon S3 buckets, failing to enable user activity logging (AWS CloudTrail), and failing to enable MFA on your root AWS account. Using https://aws.amazon.com/premiumsupport/technology/trusted-advisor/, find out how AWS Trusted Advisor functions. 

iPQ.No	Correct Answer	Points	Explanation
4AWS CCP03	3	15.4	Answer = 3 Both A and B Explanation : All requests can be blocked by the AWS Web Application Firewall (AWS WAF), with the exception of those you approve. All requests can be blocked by the AWS Web Application Firewall (AWS WAF), with the exception of those you approve. This is helpful if you need to deliver content to users of a restricted website who can easily be identified from properties in web requests, including the IP addresses they use to access the website. AWS Web Application Firewall (AWS WAF) may look for SQL code that is probably malicious (sometimes referred to as SQL injection). With conditions that you set, the AWS Web Application Firewall (AWS WAF) provides extra defense against web threats. You can use web request characteristics to define conditions, such as the IP addresses from which requests come, the presence of potentially malicious scripts (known as cross-site scripting), the presence of potentially malicious SQL code (known as SQL injection), and many others. The AWS Web Application Firewall (AWS WAF) is a web application firewall that aids in defending your web applications or APIs from widespread web exploits that could damage availability, compromise security, or use up excessive resources. The AWS Web Application Firewall (AWS WAF) enables you to establish security rules that block typical attack patterns like SQL injection and cross-site scripting, giving you control over how traffic enters your applications. Rate-based controls can also be used to lessen Web-layer DDoS attacks. Use this link to learn how AWS Web Application Firewall (AWS WAF) operates: aws.amazon.com/waf  The diagram illustrates the AWS WAF architecture and workflow. It shows a sequence of steps: 1. AWS Firewall Manager (Manage multiple AWS WAF deployments) connects to AWS WAF (Protect your web applications from common web exploits). 2. AWS WAF connects to Amazon CloudFront, Application Load Balancer, and Amazon API Gateway. 3. A 'Create a policy' step (Build your own rules using the visual rule builder, code in JSON, or deploy managed rules maintained by AWS and/or sellers from AWS Marketplace) leads to a 'Block & Filter' step (Protect against exploits and vulnerabilities such as SQL/XSS attacks; filter out unwanted traffic by defining specific patterns or by IP address). 4. The 'Block & Filter' step leads to a 'Monitor' step (Use Amazon CloudWatch for incoming traffic metrics & Amazon Kinesis Firehose for request details; then tune rules based on metrics & log data).

iPQ.No	Correct Answer	Points	Explanation
4AWS CCP04	1	15.4	Answer = 1 EC2 Auto Scaling on Amazon Explanation : You can make sure that you have the appropriate number of Amazon EC2 instances available to handle the load for your application by using Amazon EC2 Auto Scaling. You can define the minimum number of instances in each "auto-scaling group," and Amazon EC2 auto-scaling makes sure that your group never falls below this amount. through https://docs.aws.amazon.com/autoscaling/ec2/userguide/what-is-amazon-ec2-auto-scaling.html Amazon EC2 Auto Scaling Overview What Is Amazon EC2 Auto Scaling? PDF Kindle RSS Amazon EC2 Auto Scaling helps you ensure that you have the correct number of Amazon EC2 instances available to handle the load for your application. You create collections of EC2 instances, called <i>Auto Scaling groups</i>. You can specify the minimum number of instances in each Auto Scaling group, and Amazon EC2 Auto Scaling ensures that your group never goes below this size. You can specify the maximum number of instances in each Auto Scaling group, and Amazon EC2 Auto Scaling ensures that your group never goes above this size. If you specify the desired capacity, either when you create the group or at any time thereafter, Amazon EC2 Auto Scaling ensures that your group has this many instances. If you specify scaling policies, then Amazon EC2 Auto Scaling can launch or terminate instances as demand on your application increases or decreases. For example, the following Auto Scaling group has a minimum size of one instance, a desired capacity of two instances, and a maximum size of four instances. The scaling policies that you define adjust the number of instances, within your minimum and maximum number of instances, based on the criteria that you specify.  The diagram, titled "Auto Scaling group", shows a horizontal row of four orange squares representing EC2 instances. The first two squares are solid orange, and the last two are outlined in orange. Below the squares, three horizontal brackets indicate scaling limits: the first bracket under the first two squares is labeled "Minimum size"; the second bracket under the last two squares is labeled "Scale out as needed"; and the third, longest bracket under all four squares is labeled "Maximum size". A fourth bracket, labeled "Desired capacity", is positioned between the "Minimum size" and "Maximum size" brackets, spanning the width of the first two squares.


iPQ.No	Correct Answer	Points	Explanation																														
4AWS CCP05	3	15.4	Answer = 3 AWS Health Dashboard: The condition of your account Explanation : Only the Business and Enterprise help subscriptions include the entire range of AWS Trusted Advisor best practice inspections, expanded technical help with unlimited cases and contacts, and third-party software support. AWS Health Dashboard: All support plans provide access to your account's health. AWS Health Dashboard: Your Account Health alerts you to occurrences that AWS is experiencing and offers advice on how to address them. AWS Health Dashboard: Your Account Health provides event visibility and information to help you rapidly identify and address problems. Alerts are generated by changes in the health of your AWS resources. To find out the most recent status, visit https://health.aws.amazon.com/health/home. Exam Warning: As you might anticipate at least a few questions on the exam, please review the distinctions between the AWS Developer Support, AWS Business Support, AWS Enterprise On-Ramp Support, and AWS Enterprise Support plans: <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></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</td><td>General guidance: < 24 hours System impaired: < 12 hours Production system impaired: < 4 hours</td><td>General guidance: < 24 hours System impaired: < 12 hours Production system impaired: < 4 hours</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		General guidance: < 24 hours** System impaired: < 12 hours**	General guidance: < 24 hours System impaired: < 12 hours Production system impaired: < 4 hours	General guidance: < 24 hours System impaired: < 12 hours Production system impaired: < 4 hours	General guidance: < 24 hours System impaired: < 12 hours Production system impaired: < 4 hours
	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																													
	General guidance: < 24 hours** System impaired: < 12 hours**	General guidance: < 24 hours System impaired: < 12 hours Production system impaired: < 4 hours	General guidance: < 24 hours System impaired: < 12 hours Production system impaired: < 4 hours	General guidance: < 24 hours System impaired: < 12 hours Production system impaired: < 4 hours																													

iPQ.No	Correct Answer	Points	Explanation
4AWS CCP06	3	15.4	Answer = 3 Both A and B Explanation : All objects are stored in Amazon S3 buckets and can be organized with shared names called prefixes; you can also append up to 10 key-value pairs known as Amazon S3 object tags to each object, which can be created, updated, and deleted. Amazon Simple Storage Service (Amazon S3) is an object storage service that offers industry-leading scalability, data availability, security, and performance.
4AWS CCP07	1	15.4	Answer = 1 VPC: Virtual Private Cloud Explanation : You have complete control over your virtual networking environment, including the selection of your IP address range, the creation of subnets, the configuration of route tables and network gateways, and the launch of AWS resources in a virtual network that you define. Amazon Virtual Private Cloud (Amazon VPC) is a logically isolated section of the AWS Cloud where you can launch AWS resources in a virtual network that you define.

iPQ.No	Correct Answer	Points	Explanation
4AWS CCP08	2	15.4	Answer = 2 Warm-standby technique Explanation : The benefits of a lower RTO (recovery time objective) and RPO (recovery point objective) vs. the costs of implementing and operating a strategy must be weighed when choosing your disaster recovery (DR) strategy, and the pilot light and warm standby strategies both offer a good balance of benefits and costs. This strategy replicates data from the primary region to data resources in the recovery region that are ready to serve requests, such as Amazon Relational Database Service (Amazon RDS) DB instances or Amazon DynamoDB tables. In addition to replication, this strategy necessitates the creation of a continuous backup in the recovery region because when "human action"-type disasters occur, data can be deleted or corrupted, and replication will replicate t The DR endpoint can handle requests, but the warm standby strategy deploys a functional stack at reduced capacity; it may be more, but it is always less than the full production deployment for cost savings. If the passive stack is deployed to the recovery region at full capacity, however, then this strategy is known as "hot standby." Warm standby deploys a functional stack to the recovery region, making it simpler to test. Disaster recovery (DR) techniques are described in the following article:  Figure 1. DR strategies https://aws.amazon.com/blogs/architecture/disaster-recovery-dr-architecture-on-aws-part-iii-pilot-light-and-warm-standby

iPQ.No	Correct Answer	Points	Explanation
4AWS CCP09	2	15.4	Answer = 2 To operate any interruptible containerized task with Amazon Elastic Container Service (Amazon ECS) Explanation : You can utilize unutilized EC2 capacity in the AWS cloud by using Amazon EC2 Spot Instances, which are available for up to 90% less than on-demand costs. Spot instances can be used with Elastic Container Service (ECS) or Elastic Container Service for Kubernetes (EKS) to run any containerized workload, from distributed parallel test systems to applications that map millions of miles per day. Spot instances offer the flexibility of ad-hoc provisioning for multiple instance types in different availability zones, with the option to hibernate, stop, or restart instances as needed. accessed from: Benefits of Using Spot Instances to Run Containers  LOWER COSTS Significant price savings of up to 90% over On-Demand EC2 Instances. Simply mix Spot Instances with On-Demand and/or Reserved Instances or run on 100% Spot Instances to optimize cost and performance.  FASTER RESULTS Easily run multiple projects simultaneously and speed up job flows to generate business results faster and innovate faster without breaking the bank. Run and scale to large numbers of parallel tasks via Spot Fleet or Spot Instance pool.  RESOURCE FLEXIBILITY Flexibility of ad-hoc provisioning for multiple instance types in different Availability Zones, with an option to hibernate, stop or terminate instances when EC2 needs the capacity back and Spot Instances are reclaimed.  EASE OF USE Easily launch a Spot Instance via an API call, EC2 Fleet and the AWS Management Console. EC2 Spot is also integrated with other AWS services such as ECS, EKS, EC2 Auto Scaling groups and CloudFormation. https://aws.amazon.com/ec2/spot/containers-for-less/


iPQ.No	Correct Answer	Points	Explanation																																																																																								
4AWS CCP10	2	15.4	Answer = 2 Amazon S3 Basic Explanation : Low latency and high throughput performance are offered by Amazon S3 Standard, which is designed for the durability of 99.999999999% of objects across multiple availability zones (AZ). Amazon S3 Standard has no restrictions on the minimum storage duration for objects. Amazon S3 Standard also offers high durability, availability, and performance for object storage for frequently accessed data. You don't need to recall the exact figures; just keep in mind that Amazon S3 Standard and Amazon S3 Intelligent-Tiering do not incur any retrieval fees: Performance across the S3 Storage Classes <table><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>99.9%</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> via: https://aws.amazon.com/s3/storage-classes/		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%	99.9%	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%	99.9%	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																																																																																				
4AWS CCP11	3	15.4	Answer = 3 Pricing Calculator for AWS Explanation : You can investigate AWS services and generate a cost projection for your use cases on AWS using the AWS Pricing Calculator. You can investigate the price ranges and calculations used to arrive at your estimate, model your solutions before they are built, and discover the instance kinds and contract terms that best suit your requirements. This gives you the information you need to use AWS wisely. You have two options: price out setting up a fresh set of instances and services or plan your AWS prices and use. Visit https://calculator.aws/#/ to use the AWS Pricing Calculator. Additionally, AWS provides a free service called Migration Evaluator (formerly TSO Logic) for building data-driven business reasons for AWS cloud migration and planning.																																																																																								

iPQ.No	Correct Answer	Points	Explanation
4AWS CCP12	3	15.4	Answer = 3 CodeDeploy AWS Explanation : Application deployments to various compute services, such as Amazon EC2, AWS Fargate, AWS Lambda, and on-premises instances, can be automated with the help of AWS CodeDeploy. Your application deployments are totally automated using CodeDeploy, negating the need for manual processes. Through rolling updates and deployment health monitoring, CodeDeploy guards against application downtime during deployments.
4AWS CCP13	1	15.4	Answer = 1 The AWS Artifact Explanation : Your go-to, central resource for compliance-related information that is important to your organization is AWS Artifact. It offers on-demand access to certain online agreements as well as security and compliance information from AWS. The Service Organization Control (SOC) reports, Payment Card Industry (PCI) reports, and certifications from accrediting bodies across geographies and compliance verticals that vouch for the implementation and operational efficacy of AWS security controls are among the reports that are accessible in AWS Artifact. It is a free self-service portal that provides on-demand access to AWS compliance reports; it is not a service.
4AWS CCP14	2	15.4	Answer = 2 (Amazon EFS) Amazon Elastic File System Explanation : For usage with AWS cloud services and on-premises resources, Amazon Elastic File System (Amazon EFS) offers a straightforward, scalable, fully managed elastic NFS file system. You mount an EFS file system on your on-premises Linux server using the normal Linux mount command for mounting a file system. You must have an AWS Direct Connect or AWS VPN connection between your on-premises datacenter and your Amazon VPC in order to access EFS file systems from on-premises. through https://aws.amazon.com/efs/faq/ How Amazon EFS Functions 


iPQ.No	Correct Answer	Points	Explanation
4AWS CCP15	4	15.4	Answer = 4 Qualifications Report Explanation : A report on the status of all the users in your account's database's credentials, such as passwords, access keys, and MFA devices, can be generated and downloaded. The credentials report can be used to support your auditing and compliance initiatives. The report can be used to audit the results of requirements for credential lifecycle management, like password and access key rotation. An external auditor may be given the report, or you may enable him or her access so they can download it themselves.
4AWS CCP16	3	15.4	Answer = 3 (Amazon EC2) Amazon Elastic Compute Cloud Explanation : All AWS clients have AWS Shield Standard active by default. You can get AWS Shield Advanced if you want stronger defenses against attacks. You also have exclusive access to advanced, real-time analytics and statistics with Shield Advanced, giving you a wealth of information about assaults on your AWS resources. AWS Shield Advanced features sophisticated DDoS attack detection and mitigation for attacks on the network layer (layer 3), transport layer (layer 4), and application layer (layer 7) with the help of the DRT (DDoS response team). For web applications operating on the following resources: Amazon Elastic Compute Cloud, Elastic Load Balancing (ELB), Amazon CloudFront, Amazon Route 53, and AWS Global Accelerator, AWS Shield Advanced offers increased DDoS attack prevention.

iPQ.No	Correct Answer	Points	Explanation
4AWS CCP17	2	15.4	Answer = 2 One or more data centers situated nearby Explanation : One or more distinct data centers with redundant power, networking, and connectivity make up an availability zone (AZ) in an AWS region. A geographic area where AWS clusters data centers is referred to as an AWS region. Customers now have the option to run production databases and applications that are more fault-tolerant, highly available, and scalable than would be conceivable in a single data center thanks to AZs. High-bandwidth, low-latency networking over fully redundant, dedicated metro fiber connects every AZ in an AWS region, enabling high-throughput, low-latency networking between AZs. According to https://aws.amazon.com/about-aws/global-infrastructure/regions_az/, AWS Regions and Availability Zones (AZs) are described. Regions AWS has the concept of a Region, which is a physical location around the world where we cluster data centers. We call each group of logical data centers an Availability Zone. Each AWS Region consists of multiple, isolated, and physically separate AZ's within a geographic area. Unlike other cloud providers, who often define a region as a single data center, the multiple AZ design of every AWS Region offers advantages for customers. Each AZ has independent power, cooling, and physical security and is connected via redundant, ultra-low-latency networks. AWS customers focused on high availability can design their applications to run in multiple AZ's to achieve even greater fault-tolerance. AWS infrastructure Regions meet the highest levels of security, compliance, and data protection. AWS provides a more extensive global footprint than any other cloud provider, and to support its global footprint and ensure customers are served across the world, AWS opens new Regions rapidly. AWS maintains multiple geographic Regions, including Regions in North America, South America, Europe, China, Asia Pacific, South Africa, and the Middle East. Availability Zones An Availability Zone (AZ) is one or more discrete data centers with redundant power, networking, and connectivity in an AWS Region. AZ's give customers the ability to operate production applications and databases that are more highly available, fault tolerant, and scalable than would be possible from a single data center. All AZ's in an AWS Region are interconnected with high-bandwidth, low-latency networking, over fully redundant, dedicated metro fiber providing high-throughput, low-latency networking between AZ's. All traffic between AZ's is encrypted. The network performance is sufficient to accomplish synchronous replication between AZ's. AZ's make partitioning applications for high availability easy. If an application is partitioned across AZ's, companies are better isolated and protected from issues such as power outages, lightning strikes, tornadoes, earthquakes, and more. AZ's are physically separated by a meaningful distance, many kilometers, from any other AZ, although all are within 100 km (60 miles) of each other.


iPQ.No	Correct Answer	Points	Explanation
4AWS CCP18	3	15.4	Answer = 3 Demand-Side Instance Explanation : An instance that is available on demand is known as an on-demand instance. Its lifespan is completely under your control; you get to choose when to start, stop, hibernate, start, reboot, or terminate it. When you purchase on-demand instances, you are not obligated to make a long-term commitment. You simply pay for the time that your on-demand instances are running; there is no upfront fee. No long-term purchase commitment is required. Running an on-demand instance has a fixed cost per second. Instances that are on demand cannot be stopped. Therefore, for short-term, spikey, and crucial workloads, on-demand instances are the ideal option. Overview of Amazon 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, spikey, 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
4AWS CCP19	1	15.4	Answer = 1 Systems Manager for AWS Explanation : You have visibility and control over your AWS infrastructure thanks to AWS Systems Manager. You can automate operational tasks like compiling software inventories, running commands, managing patches, and configuring servers across the AWS Cloud and on-premises infrastructure using Systems Manager, which offers a unified user interface so you can view operational data from various AWS services. AWS Systems Manager provides tools for patch management, configuration compliance, and running commands. retrieved from: https://aws.amazon.com/systems-manager/faq  retrieved from: https://aws.amazon.com/systems-manager
4AWS CCP20	2	15.4	Answer = 2 AIM role Explanation : You can create an IAM identity with specified permissions in your account called an IAM role. In that it is an AWS identity with authorization policies that specify what the identity can and cannot do in AWS, an IAM role is comparable to an IAM user. You are given temporary security credentials for your session when you accept a role.

iPQ.No	Correct Answer	Points	Explanation
4AWS CCP21	4	15.4	Answer = 4 Cost Explorer for AWS Explanation : You can easily monitor, comprehend, and control your AWS expenses and use over time with the help of AWS Cost Explorer's user-friendly interface. Your top five cost-accumulating AWS services are broken down in a default report provided by AWS Cost Explorer, along with a thorough breakdown of all services in the table view. You can change the time range in the reports to view historical data going back up to twelve months in order to better understand your expense trends. In order to better prepare for the future and have a better picture of your expenditures and use, AWS Cost Explorer also provides forecasting. Details about AWS Cost Explorer may be found at AWS Cost Explorer Features Get started quickly A set of default reports are included to help you quickly gain insight into your cost drivers and usage trends. Forecast future costs and usage Use forecasting to get a better idea of your what your costs and usage may look like in the future, so that you can plan ahead. Set time interval and granularity Set a custom time period, and determine whether you would like to view your data at a monthly or daily level of granularity. Save your progress Once you arrive at a helpful view, save your progress as a new report that you can refer back to in the future. Filter/Group your data Dig deeper into your data by taking advantage of filtering and grouping functionality, using a variety of available dimensions. Build custom applications Directly access the interactive, ad-hoc analytics engine that powers AWS Cost Explorer. https://aws.amazon.com/aws-cost-management/aws-cost-explorer/.


iPQ.No	Correct Answer	Points	Explanation
4AWS CCP22	2	15.4	Answer = 2 Trusted Advisor for AWS Explanation : An online application called AWS Trusted Advisor offers real-time advice to help you supply your resources in accordance with AWS best practices. Regular recommendations from Trusted Advisor help maintain your solutions provisioned at their best, whether for creating new workflows, creating applications, or as part of continuing improvement. AWS Trusted Advisor examines your AWS infrastructure and offers best practice advice in the following five areas: cost reduction, performance improvement, security enhancement, fault tolerance, and service limit adjustments. When volumes seem to be underutilized, AWS Trusted Advisor can analyze Amazon Elastic Block Store (Amazon EBS) volume parameters and issue a warning. When a volume is produced, charges start. A volume is presumably not being utilized if it is detached or has very little write activity (except from boot volumes) for an extended length of time. Using https://aws.amazon.com/premiumsupport/technology/trusted-advisor/, find out how AWS Trusted Advisor functions.  <pre> graph LR A[As an AWS customer, you want the most value from your investment. Trusted Advisor can help.] --> B[AWS Trusted Advisor Trusted Advisor scans your AWS infrastructure, compares it to AWS best practices in five categories, and provides recommended actions] B --> C[Cost Optimization] B --> D[Performance] B --> E[Security] B --> F[Fault Tolerance] B --> G[Service Limits] C --> H[Investigation recommended] D --> I[no problem detected] E --> J[action recommended] F --> K[no problem detected] G --> L[no problem detected] H --> M[Recommendations & action links] I --> M J --> M K --> M L --> M </pre>

iPQ.No	Correct Answer	Points	Explanation
4AWS CCP23	3	15.4	Answer = 3 Both A and B Explanation : The four support plans that AWS offers—AWS Developer Support, AWS Business Support, AWS Enterprise On-Ramp Support, and AWS Enterprise Support—are designed to meet the needs of each of its customers. All AWS clients receive a free basic support package. Enterprise Support for AWS The AWS Enterprise Support plan should be used to offer clients concierge-style assistance with a primary goal of assisting the client in achieving their objectives and experiencing success in the cloud. With the AWS Enterprise Support plan, you get consultative architectural guidance delivered in the context of your applications and use-cases, 24x7 technical support from high-quality engineers, tools, and technology to automatically manage the health of your environment, and a designated Technical Account Manager (TAM) to coordinate access to proactive and preventative programs and AWS subject matter experts. The AWS Enterprise help plan offers technical help through phone, email, and chat around-the-clock, although it is more expensive than the AWS Business Support plan. You can initiate, manage, and close support cases, as well as operationally manage your Trusted Advisor check requests and status, using programmatic access (API access) to AWS Support Center services. The AWS Enterprise On-Ramp Support package also includes programmatic access (API access) to the AWS Support Center, it should be noted. Business Support by AWS If you use AWS for production workloads and need access to technical support and architectural advice on a 24/7 basis through phone, email, and chat for your particular use cases, you should use the AWS Business Support subscription. You can initiate, manage, and close support cases, as well as operationally manage your Trusted Advisor check requests and status, using programmatic access (API access) to AWS Support Center services. Exam Warning:

iPQ.No	Correct Answer	Points	Explanation
4AWS CCP24	1	15.4	Answer = 1 VPC peering relationship Explanation : A networking link called a VPC peering connection allows you to privately route traffic between two VPCs. Both VPCs allow communication between instances just like if they were on the same network. Between your VPCs, with a VPC in another AWS account, or with a VPC in a different AWS region, you can establish a VPC peering connection. An overview of VPC peering may be found at What is VPC peering? PDF Amazon Virtual Private Cloud (Amazon VPC) enables you to launch AWS resources into a virtual network that you've defined. A VPC peering connection is a networking connection between two VPCs that enables you to route traffic between them using private IPv4 addresses or IPv6 addresses. Instances in either VPC can communicate with each other as if they are within the same network. You can create a VPC peering connection between your own VPCs, or with a VPC in another AWS account. The VPCs can be in different regions (also known as an inter-region VPC peering connection).  AWS uses the existing infrastructure of a VPC to create a VPC peering connection; it is neither a gateway nor a VPN connection, and does not rely on a separate piece of physical hardware. There is no single point of failure for communication or a bandwidth bottleneck. A VPC peering connection helps you to facilitate the transfer of data. For example, if you have more than one AWS account, you can peer the VPCs across those accounts to create a file sharing network. You can also use a VPC peering connection to allow other VPCs to access resources you have in one of your VPCs. https://docs.aws.amazon.com/vpc/latest/peering/what-is-vpc-peering.html.


iPQ.No	Correct Answer	Points	Explanation
4AWS CCP25	2	15.4	Answer = 2 Read replicas can be used for disaster recovery as well as better read performance Explanation : Setting up, running, and scaling a relational database in the cloud is simple with Amazon Relational Database Service (Amazon RDS). Create read-only copies that are synchronized with your master database using read replicas. For better read performance, read replicas are used. For better speed, you may also move your read replica to a different AWS region that is closer to your users. In the event of a disaster, using a cross-region read replica can also assist ensure that you are able to quickly resume operations if you encounter a problem with regional availability. Read replicas are an illustration of horizontal resource scalability. through - https://aws.amazon.com/rds/features/multi-az/ Read the replica overview Amazon RDS Read Replicas provide enhanced performance and durability for RDS database (DB) instances. They make it easy to elastically scale out beyond the capacity constraints of a single DB instance for read-heavy database workloads. You can create one or more replicas of a given source DB instance and serve high-volume application read traffic from multiple copies of your data, thereby increasing aggregate read throughput. Read replicas can also be promoted when needed to become standalone DB instances. Read replicas are available in Amazon RDS for MySQL, MariaDB, PostgreSQL, Oracle, and SQL Server as well as Amazon Aurora. For the MySQL, MariaDB, PostgreSQL, Oracle, and SQL Server database engines, Amazon RDS creates a second DB instance using a snapshot of the source DB instance. It then uses the engines' native asynchronous replication to update the read replica whenever there is a change to the source DB instance. The read replica operates as a DB instance that allows only read-only connections; applications can connect to a read replica just as they would to any DB instance. Amazon RDS replicates all databases in the source DB instance. Amazon Aurora further extends the benefits of read replicas by employing an SSD-backed virtualized storage layer purpose-built for database workloads. Amazon Aurora replicas share the same underlying storage as the source instance, lowering costs and avoiding the need to copy data to the replica nodes. For more information about replication with Amazon Aurora, see the online documentation. A natural fit for production database workloads, Amazon RDS multi-AZ deployments offer improved availability and durability for RDS database (DB) instances. Amazon RDS automatically builds a primary DB instance and synchronously replicates the data to a backup instance in a separate availability zone (AZ) when you provision an Amazon RDS Multi-AZ DB Instance. As soon as the failover is over, you can resume database operations since Amazon RDS executes an immediate failover to the backup. After a failover, your application can continue database operation without the need for human administrative intervention because the endpoint for your DB instance stays the same. Consider multi-AZ deployment as improving your system's availability and dependability; however, multi-AZ deployment cannot be used by itself for disaster recovery. According to https://aws.amazon.com/rds/features/multi-az/, this is how multi-AZ deployment works.

iPQ.No	Correct Answer	Points	Explanation
4AWS CCP26	2	15.4	Answer = 2 (Amazon S3) Amazon Simple Storage Service Explanation : By customizing your objects' Amazon S3 lifecycles, you can manage them such that they are cost-effectively stored over their entire lifespan. A set of guidelines that specify the activities that Amazon S3 should do in relation to a collection of items is known as a lifecycle configuration. Two categories of acts exist: Actions for transitioning objects to another storage class are specified. For instance, you may decide to archive things to the Amazon S3 Glacier storage class a year after creating them or migrate them to the Amazon S3 Standard-IA storage class after 30 days. Actions for defining an object's expiration date. On your part, Amazon S3 removes expired objects.


iPQ.No	Correct Answer	Points	Explanation
4AWS CCP27	3	15.4	Answer = 3 Both A and B Explanation : AWS and the customer share responsibility for security and compliance. With AWS running the host operating system, the virtualization layer, and even the physical security of the premises where the service is housed, the operational burden on the user can be lessened thanks to this shared approach. Multi-factor authentication (MFA) should be enabled for all accounts. Examine AWS Identity and Access Management (AWS IAM) permissions and analyze user access patterns. Customers are in charge of implementing multi-factor authentication (MFA) across all accounts, examining access patterns, and checking permissions for AWS Identity and Access Management (AWS IAM) entities under the AWS Shared Responsibility Model. through https://aws.amazon.com/compliance/shared-responsibility-model/. AWS Shared Responsibility Model Overview  The diagram illustrates the AWS Shared Responsibility Model. On the left, two large arrows represent the shared responsibilities: a blue arrow pointing right labeled 'CUSTOMER' with 'RESPONSIBILITY FOR SECURITY 'IN' THE CLOUD' and an orange arrow pointing right labeled 'AWS' with 'RESPONSIBILITY FOR SECURITY 'OF' THE CLOUD'. To the right, a stack of boxes represents the layers of responsibility. The top three layers (Customer's responsibility) are: 'CUSTOMER DATA', 'PLATFORM, APPLICATIONS, IDENTITY & ACCESS MANAGEMENT', and 'OPERATING SYSTEM, NETWORK & FIREWALL CONFIGURATION'. The next layer (Shared responsibility) is divided into three boxes: 'CLIENT-SIDE DATA ENCRYPTION & DATA INTEGRITY AUTHENTICATION', 'SERVER-SIDE ENCRYPTION (FILE SYSTEM AND/OR DATA)', and 'NETWORKING TRAFFIC PROTECTION (ENCRYPTION, INTEGRITY, IDENTITY)'. The 'SOFTWARE' layer (AWS's responsibility) includes 'COMPUTE', 'STORAGE', 'DATABASE', and 'NETWORKING'. The bottom layer (AWS's responsibility) is 'HARDWARE/AWS GLOBAL INFRASTRUCTURE', which includes 'REGIONS', 'AVAILABILITY ZONES', and 'EDGE LOCATIONS'.

iPQ.No	Correct Answer	Points	Explanation
4AWS CCP28	3	15.4	Answer = 3 Both A and B Explanation : Direct Connect AWS A dedicated network connection from your premises to AWS may be easily established with the help of AWS Direct Connect, a cloud service solution. Your datacenter, office, or colocation environment can be privately connected to AWS via AWS Direct Connect, which in many circumstances can lower network costs, boost bandwidth throughput, and offer a more reliable network experience than Internet-based connections. Transit Gateway by AWS Through a single hub, AWS Transit Gateway links on-premises networks and virtual private clouds (VPC). Your network becomes simpler as a result, and complicated peer relationships are eliminated. Every new connection is only established once; it functions as a cloud router. Inter-region peering links AWS Transit Gateways using the AWS global network as you grow internationally. Your information is always automatically encrypted and never sent over a public network. See https://aws.amazon.com/transit-gateway/ for information on how AWS Transit Gateway can streamline your network. The diagram is divided into two parts: 'Without AWS Transit Gateway' and 'With AWS Transit Gateway'. Without AWS Transit Gateway: This part shows a complex network where a central 'Customer gateway' is connected to multiple 'VPN connections'. Each 'VPN connection' is linked to an 'Amazon VPC'. These VPCs are further interconnected through 'VPC Peering' relationships. A 'Direct Connect Gateway' is also shown connected to the VPCs. Below this diagram, text states: 'Complexity increases with scale. You must maintain routing tables within each VPC and connect to each onsite location using separate network gateways.' With AWS Transit Gateway: This part shows a simplified network where all 'Amazon VPCs' and 'VPN connections' are connected to a single central hub labeled 'AWS Transit Gateway'. The 'Customer gateway' and 'Direct Connect Gateway' are also connected to this central hub. Below this diagram, text states: 'Your network is streamlined and scalable. AWS Transit Gateway routes all traffic to and from each VPC or VPN, and you have one place to manage and monitor it all.'
4AWS CCP29	1	15.4	Answer = 1 Amazon Translate Explanation : With its neural machine translation technology, Amazon Translate provides quick, accurate, and cost-effective language translation. You can localize content, such as websites and applications, for users around the world using Amazon Translate, and you can quickly and effectively translate enormous amounts of text.

iPQ.No	Correct Answer	Points	Explanation
4AWS CCP30	2	15.4	Answer = 2 DynamoDB by Amazon Explanation : At any size, the key-value and document database Amazon DynamoDB offers speed in the single-digit millisecond range. For internet-scale applications, it is a fully managed, multi-Region, multi-master, durable database with built-in security, backup and restore, and in-memory caching. DynamoDB lacks a schema. JSON documents and other structured or semistructured data can be managed with DynamoDB.
4AWS CCP31	1	15.4	Answer = 1 Cost of application software license Explanation : With pay-as-you-go pricing, a cloud services platform delivers computational power, database storage, apps, and other IT resources on demand over the Internet. With cloud computing, there is no need for significant up-front hardware investments or for you to commit a lot of time in administering that infrastructure. As a result, AWS is responsible for covering all expenses related to the data center's physical security, powering computers, and hardware infrastructure. Both the cost of software license and the cost of human resources must be covered by the customer.

iPQ.No	Correct Answer	Points	Explanation
4AWS CCP32	3	15.4	Answer = 3 Both A and B Explanation : An almost real-time stream of system events describing modifications to Amazon Web Services (AWS) resources is provided via Amazon CloudWatch Events. You may match events and rapidly route them to one or more target functions or streams by using straightforward rules. As soon as an operational change takes place, Amazon CloudWatch Events is informed. By delivering messages to respond to the environment, activating functions, making changes, and capturing state data, Amazon CloudWatch Events reacts to these operational changes and takes remedial action as appropriate. You may decouple microservices, distributed systems, and serverless applications using the highly available, reliable, secure, and fully managed pub/sub messaging service known as Amazon Simple Notification Service (Amazon SNS). Additionally, Amazon SNS can be utilized to notify users by email, SMS, and mobile push. through https://aws.amazon.com/sns/. How Amazon SNS Functions  The diagram shows the flow of messages from a Publisher to Amazon SNS, then to an SNS Topic, and finally to Subscribers. The Publisher is labeled 'Publisher' and describes publishing messages from distributed systems, microservices, and other AWS services. Amazon SNS is labeled 'Amazon SNS' and describes a fully-managed pub/sub messaging and event-driven computing service. The SNS Topic is labeled 'SNS Topic' and describes decoupling message publishers from subscribers with topics. The Message Filtering & Fanout step is labeled 'Message Filtering & Fanout' and describes filtering messages according to subscription filter policies and delivering them to subscribers. The Subscribers are labeled 'Subscribers' and describe receiving messages in subscribing serverless functions, queues, microservices, and more. The Subscribers are further categorized into AWS Lambda, Amazon SQS, and HTTP/S. You are able to build an Amazon Simple Notification Service (Amazon SNS) subject to deliver notifications each time the root user of the AWS account logs in. Create an Amazon CloudWatch event rule to keep track on user identity root logins through the AWS Management Console, and when the event occurs, send an email via Amazon SNS.

iPQ.No	Correct Answer	Points	Explanation
4AWS CCP33	3	15.4	Answer = 3 Both A and B Explanation : Sell AWS clients software as a service (SaaS) products. Software that has been combined into specialized Amazon Machine Images (AMIs) by the AWS Marketplace vendors is available for purchase by AWS customers. It is simple to identify, test, purchase, and deploy software that runs on AWS thanks to the AWS Marketplace, a digital catalog with thousands of software listings from independent software suppliers. Qualified partners can promote and sell their applications to AWS customers through the AWS Marketplace. Amazon Machine Image (AMI) and Software as a Service (SaaS) are the two options available on AWS Marketplace for vendors to deploy software to clients. Offering an Amazon Machine Image (AMI) is the recommended method for listing products in the AWS Marketplace. Partners can choose from free or expensive products. Partners may provide compensated goods that are billed by the hour or the month. Customers with existing software licenses can quickly switch to AWS using the Bring-Your-Own-License (BYOL) option, which is also available. Software as a Service (SaaS): The SaaS listing gives our partners a means to promote their software to customers if they provide SaaS solutions running on AWS (and are unable to build their product into an AMI).

iPQ.No	Correct Answer	Points	Explanation
4AWS CCP34	3	15.4	Answer = 3 Both A and B Explanation : AWS Lambda's serverless architecture enables you to run code without setting up or maintaining servers. You may use Lambda to execute code for almost any kind of application or backend service with no administration required. Simply upload your code, and Lambda will take care of all the necessary steps to run and grow it while maintaining high availability. The ability to do real-time data transformations when used with streaming data and serve as a serverless backend for websites and other serverless functions makes it an incredibly useful service. http://aws.amazon.com/lambda/ explains how AWS Lambda functions.  The diagram illustrates the AWS Lambda workflow in four steps: 1. Upload your code to AWS Lambda or write code in Lambda's code editor (represented by a laptop icon). 2. Set up your code to trigger from other AWS services, HTTP endpoints, or in-app activity (represented by a server rack icon). 3. AWS Lambda runs your code only when triggered, using only the compute resources needed (represented by the AWS Lambda logo). 4. Just pay for the compute time you use (represented by a bar chart icon). Code can be run in reaction to events like modifications to DynamoDB tables using AWS Lambda. AWS Lambda can be set up to run code in response to several types of events, including modifications to Amazon S3 buckets, updates to an Amazon DynamoDB table, and bespoke events produced by your software or hardware. Data can be preprocessed using AWS Lambda before being put into Amazon S3 buckets. Data can be filtered, sorted, or transformed using preprocessing scripts that are run using AWS Lambda before being sent to apps or services further down the line.

iPQ.No	Correct Answer	Points	Explanation
4AWS CCP35	3	15.4	Answer = 3 Both A and B Explanation : You establish an AWS account root user identity when you do so, which you use to log in to AWS. Using this root user identification, which consists of the email address and password you used when creating the account, you can access the AWS Management Console. You may also refer to this email address and password combination as your root user credentials. The following are some recommended best practices for AWS account root user security: Use the root user for your AWS account only when it is necessary. For each employee who needs administrator access, you should instead establish a new IAM user. Then add those users to the "Administrators" group, to which the AdministratorAccess managed policy is attached, to make them administrators. Avoid creating an access key for the root user of your AWS account unless absolutely necessary. If your AWS account root user does have an access key, remove it. Never divulge your root user password or access keys for your AWS account to anyone. To protect account-level access to the AWS Management Console, use a strong password. Make your AWS root user account's multi-factor authentication (MFA) enabled. 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. Therefore, protect your root user access key like you would your credit card numbers or any other sensitive secret. Here are some ways to do that: • If you don't already have an access key for your AWS account root user, don't create one unless you absolutely need to. Instead, use your account email address and password to sign in to the AWS Management Console and create an IAM user for yourself that has administrative permissions. • If you do have an access key for your AWS account root user, delete it. If you must keep it, rotate (change) the access key regularly. To delete or rotate your root user access keys, go to the My Security Credentials page in the AWS Management Console and sign in with your account's email address and password. You can manage your access keys in the Access keys section. For more information about rotating access keys, see Rotating Access Keys.

iPQ.No	Correct Answer	Points	Explanation
4AWS CCP36	3	15.4	Answer = 3 Both A and B Explanation : Only if Bob launches his instances in the same availability zone (AZ) as Susan did when she bought her reserved instances will Bob profit financially from Susan's reserved instances (RI). Only if Bob launches his instances in the same availability zone (AZ) as Susan's reserved instances (RI) will Bob receive the cost-benefit on the organization's consolidated bill; for example, if Susan specifies us-west-2a when she purchases her reserved instances, Bob must specify us-west-2a when he launches his instances. However, the actual locations of availability zones (AZs) are independent from one another. Five of the instances are listed by AWS as "reserved instances, and the remaining four instances are regular instances." AWS bills five instances as reserved instances and the other four instances as ordinary instances since Susan owns five reserved instances (RIs).

iPQ.No	Correct Answer	Points	Explanation
4AWS CCP37	3	15.4	Answer = 3 (Amazon S3) Amazon Simple Storage Service Explanation : An object storage service called Amazon Simple Storage Service (Amazon S3) provides performance, security, and scalability that are unmatched in the market. Customers of different sizes and sectors are using Amazon S3's flat, non-hierarchical structure and various management capabilities to arrange their data in ways that are beneficial to their teams and organizations. You set up an Amazon S3 bucket for website hosting and then upload your website's content to the bucket to host a static website there. You must enable website hosting, specify permissions, create and add an index document when configuring a bucket as a static website. Using Amazon S3 to host a static website is described in this document: Hosting a static website on Amazon S3 PDF Kindle RSS You can use Amazon S3 to host a static website. On a <i>static</i> website, individual webpages include static content. They might also contain client-side scripts. By contrast, a <i>dynamic</i> website relies on server-side processing, including server-side scripts such as PHP, JSP, or ASP.NET. Amazon S3 does not support server-side scripting, but AWS has other resources for hosting dynamic websites. To learn more about website hosting on AWS, see Web Hosting. Note You can use the AWS Amplify Console to host a single page web app. The AWS Amplify Console supports single page apps built with single page app frameworks (for example, React JS, Vue JS, Angular JS, and Nuxt) and static site generators (for example, Gatsby JS, React-static, Jekyll, and Hugo). For more information, see Getting Started in the <i>AWS Amplify Console User Guide</i>. To configure your bucket for static website hosting, you can use the AWS Management Console without writing any code. You can also create, update, and delete the website configuration <i>programmatically</i> by using the AWS SDKs. The SDKs provide wrapper classes around the Amazon S3 REST API. If your application requires it, you can send REST API requests directly from your application. To host a static website on Amazon S3, you configure an Amazon S3 bucket for website hosting and then upload your website content to the bucket. When you configure a bucket as a static website, you must enable website hosting, set permissions, and create and add an index document. Depending on your website requirements, you can also configure redirects, web traffic logging, and a custom error document. After you configure your bucket as a static website, you can access the bucket through the AWS Region-specific Amazon S3 website endpoints for your bucket. Website endpoints are different from the endpoints where you send REST API requests. For more information, see Website endpoints. https://docs.aws.amazon.com/AmazonS3/latest/dev/WebsiteHosting.html

iPQ.No	Correct Answer	Points	Explanation
4AWS CCP38	3	15.4	Answer = 3 Both A and B Explanation : Your ability to delegate more of your operational duties to AWS and so increase your agility and innovation is made possible by serverless, which is the cloud's natural architecture. You may create and execute apps and services using serverless technology without giving servers any thought. It does away with infrastructure management duties like capacity provisioning, operating system upkeep, server or cluster provisioning, and patching. http://aws.amazon.com/serverless/ provides an overview of the AWS serverless platform. Compute AWS Lambda lets you run code without provisioning or managing servers. You pay only for the compute time you consume - there is no charge when your code is not running. Lambda@Edge allows you to run Lambda functions at AWS Edge locations in response to Amazon CloudFront events. AWS Fargate is a purpose-built serverless compute engine for containers. Fargate scales and manages the infrastructure required to run your containers. Storage Amazon Simple Storage Service (Amazon S3) provides developers and IT teams with secure, durable, highly-scalable object storage. Amazon S3 is easy to use, with a simple web service interface to store and retrieve any amount of data from anywhere on the web. Amazon Elastic File System (Amazon EFS) provides simple, scalable, elastic file storage. It is built to elastically scale on demand, growing and shrinking automatically as you add and remove files. Data stores Amazon DynamoDB is a fast and flexible NoSQL database service for all applications that need consistent, single-digit millisecond latency at any scale. Amazon Aurora Serverless is an on-demand, auto-scaling configuration for Amazon Aurora (MySQL-compatible edition), where the database will automatically start up, shut down, and scale capacity up or down based on your application's needs. Amazon RDS Proxy is a highly available database proxy that manages thousands of concurrent connections to relational databases, allowing you to build highly scalable, API Proxy Amazon API Gateway is a fully managed service that makes it easy for developers to create, publish, maintain, monitor, and secure APIs at any scale. It offers a comprehensive platform for API management. API Gateway allows you to process hundreds of thousands of concurrent API calls and handles traffic management, authorization and access control, monitoring, and API version management. Application integration Amazon SNS is a fully managed pub/sub messaging service that makes it easy to decouple and scale microservices, distributed systems, and serverless applications. Amazon SQS is a fully managed message queuing service that makes it easy to decouple and scale microservices, distributed systems, and serverless applications. AWS AppSync simplifies application development by letting you create a flexible GraphQL API to securely access, manipulate, and combine data from one or more data sources. Orchestration AWS Step Functions makes it easy to coordinate the components of distributed applications and microservices using visual workflows. Building applications from individual components that each perform a discrete function lets you scale and change applications quickly. Step Functions is a reliable way to coordinate components and step through the functions of your application. Analytics Amazon Kinesis is a platform for streaming data on AWS, offering powerful services to make it easy to load and analyze streaming data, and also providing the ability for you to build custom streaming data applications for specialized needs. Amazon Athena is an interactive query service that makes it easy to analyze data in Amazon S3 using standard SQL. Athena is serverless, so there is no infrastructure to manage, and you pay only for the queries that you run. Developer tooling AWS provides tools and services that aid developers in the serverless application development process. AWS and its partner ecosystem offer tools for continuous integration and delivery, testing, deployments, monitoring and diagnostics, SDKs, frameworks, and integrated development environment (IDE) plugins.

Lambda on AWS

You may execute code for practically any kind of application or backend service using AWS Lambda, all while requiring zero management. Simply upload your code, and Lambda will take care of all the necessary steps to run and grow it while maintaining high availability. You can program your code to run other AWS services automatically or to be called directly from any web or mobile application.


You can run code with AWS Lambda without setting up or managing

iPQ.No	Correct Answer	Points	Explanation
4AWS CCP39	4	15.4	Answer = 4 Fargate AWS Explanation : With both Amazon Elastic Container Service (ECS) and Amazon Elastic Kubernetes Service (EKS), AWS Fargate is a serverless computing engine for containers. You may easily concentrate on creating your applications with Fargate. Fargate eliminates the need to set up and manage servers, enables you to specify and pay for resources per application, and enhances security through application isolation from the start. through https://aws.amazon.com/fargate/. How AWS Fargate Operates

iPQ.No	Correct Answer	Points	Explanation
4AWS CCP40	3	15.4	Answer = 3 Easy route Explanation : A highly scalable and available cloud Domain Name System (DNS) web service is Amazon Route 53. By converting names like www.example.com into the numeric IP addresses that computers use to connect to each other, such as 192.0.2.1, it is intended to give developers and organizations an incredibly dependable and cost-effective way to route end users to Internet services. With no additional Amazon Route 53 routing features, such as weighing or delay, simple routing enables you to configure normal DNS entries. Typically, when using basic routing, traffic is sent to a single resource, such as a web server for your website. through - https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/routing-policy.html Amazon Route 53 Routing Policy Overview 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 example.com 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
4AWS CCP41	2	15.4	Answer = 2 Interface Endpoint for VPC Explanation : An interface endpoint is an elastic network interface (ENI) with a private IP address from your subnet's address space that acts as a point of entry for traffic going to a supported service. AWS PrivateLink, a solution that enables you to access services secretly by using private IP addresses, is the technology that powers interface endpoints. All network traffic between your VPC and services is constrained to the Amazon network by AWS PrivateLink. A Network Address Translation (NAT) device, a virtual private gateway, or an internet gateway are not required. Exam Warning: On the exam, there can be a question centered on this idea. Just keep in mind that VPC gateway endpoints are only supported by Amazon S3 and Amazon DynamoDB. A VPC interface endpoint is used by all other services that enable VPC endpoints (Amazon S3 is one of those services).
4AWS CCP42	3	15.4	Answer = 3 EC2 Instance Connect for Amazon Explanation : Using Secure Shell (SSH), Amazon EC2 Instance Connect offers a quick and safe method of connecting to your instances. With Amazon EC2 Instance Connect, you can govern SSH access to your instances using AWS Identity and Access Management (AWS IAM) policies and principals, doing away with the need to distribute and manage SSH keys. To enable connection request auditing, all connection requests made using Amazon EC2 Instance Connect are recorded in AWS CloudTrail. Your preferred SSH client, the Amazon EC2 Instance Connect CLI, or a browser-based client can all be used to connect to your Linux instances with Amazon EC2 Instance Connect. A Mac, Windows, or Linux-based PC can connect to an EC2 instance via Amazon EC2 Instance Connect.
4AWS CCP43	3	15.4	Answer = 3 DynamoDB by Amazon Explanation : At any size, the key-value and document database Amazon DynamoDB offers speed in the single-digit millisecond range. For internet-scale applications, it is a fully managed, multi-region, multi-master, persistent database with built-in security, backup and restore, and in-memory caching. DynamoDB lacks a schema. JSON documents and other structured or semistructured data can be managed with DynamoDB.

iPQ.No	Correct Answer	Points	Explanation
4AWS CCP44	1	15.4	Answer = 1 Pricing Calculator for AWS Explanation : You can investigate AWS services and generate a cost projection for your use cases on AWS using the AWS Pricing Calculator. You can investigate the price ranges and calculations used to arrive at your estimate, model your solutions before they are built, and discover the instance kinds and contract terms that best suit your requirements. This gives you the information you need to use AWS wisely. You have two options: price out setting up a fresh set of instances and services or plan your AWS prices and use. Visit https://calculator.aws/#/ to use the AWS Pricing Calculator. Additionally, AWS provides a free service called Migration Evaluator (formerly TSO Logic) for building data-driven business reasons for AWS cloud migration and planning.
4AWS CCP45	4	15.4	Answer = 4 Shield Standard AWS Explanation : For applications operating on AWS, AWS Shield is a managed service that guards against distributed denial of service (DDoS) assaults. All AWS clients have access to the AWS Shield Standard at no additional cost. Your web applications running on AWS are automatically guarded against the most typical and frequent DDoS assaults thanks to AWS Shield Standard. By adhering to the recommended DDoS resilience on AWS best practices, you may fully take advantage of the AWS Shield Standard. AWS is responsible for managing the upkeep and configurations for this service because AWS Shield Standard is automatically activated for all AWS customers and does not allow for any changes. Therefore, AWS is responsible for providing this service.
4AWS CCP46	1	15.4	Answer = 1 CodeDeploy AWS Explanation : Application deployments to various compute services, such as Amazon EC2, AWS Fargate, AWS Lambda, and on-premises instances, can be automated with the help of AWS CodeDeploy. Your application deployments are totally automated using CodeDeploy, negating the need for manual processes. Through rolling updates and deployment health monitoring, CodeDeploy guards against application downtime during deployments.


iPQ.No	Correct Answer	Points	Explanation
4AWS CCP47	3	15.4	Answer = 3 EC2 Auto Scaling on Amazon Explanation : You can make sure that you have the appropriate number of Amazon EC2 instances available to handle the load for your application by using Amazon EC2 Auto Scaling. You can define the minimum number of instances in each "auto-scaling group," and Amazon EC2 auto-scaling makes sure that your group never falls below this amount. through https://docs.aws.amazon.com/autoscaling/ec2/userguide/what-is-amazon-ec2-auto-scaling.html Amazon EC2 Auto Scaling Overview What Is Amazon EC2 Auto Scaling? PDF Kindle RSS Amazon EC2 Auto Scaling helps you ensure that you have the correct number of Amazon EC2 instances available to handle the load for your application. You create collections of EC2 instances, called <i>Auto Scaling groups</i>. You can specify the minimum number of instances in each Auto Scaling group, and Amazon EC2 Auto Scaling ensures that your group never goes below this size. You can specify the maximum number of instances in each Auto Scaling group, and Amazon EC2 Auto Scaling ensures that your group never goes above this size. If you specify the desired capacity, either when you create the group or at any time thereafter, Amazon EC2 Auto Scaling ensures that your group has this many instances. If you specify scaling policies, then Amazon EC2 Auto Scaling can launch or terminate instances as demand on your application increases or decreases. For example, the following Auto Scaling group has a minimum size of one instance, a desired capacity of two instances, and a maximum size of four instances. The scaling policies that you define adjust the number of instances, within your minimum and maximum number of instances, based on the criteria that you specify.  The diagram illustrates an Auto Scaling group with four instances represented by orange squares. The first two squares are solid orange, and the last two are dashed orange. Brackets below the squares indicate the following values: Minimum size (1 instance), Scale out as needed (2 instances), Desired capacity (2 instances), and Maximum size (4 instances).


iPQ.No	Correct Answer	Points	Explanation
4AWS CCP48	3	15.4	Answer = 3 Both A and B Explanation : Only if Bob launches his instances in the same availability zone (AZ) as Susan did when she bought her reserved instances will Bob profit financially from Susan's reserved instances (RI). Only if Bob launches his instances in the same availability zone (AZ) as Susan did when she bought her reserved instances will Bob profit financially from Susan's reserved instances (RI). For instance, in order for Bob to receive the cost-benefit on the organization's consolidated bill if Susan selects us-west-2a when purchasing his reserved instances, Bob must also specify us-west-2a when he launches his instances. However, availability zones' (AZs') real locations differ from one account to the next. For instance, Susan's account's US-West-2a Availability Zone (AZ) can be in a different area from Bob's account. Five of the instances are listed by AWS as "reserved instances, and the remaining four instances are regular instances." AWS bills five instances as reserved instances and the other four instances as ordinary instances since Susan owns five reserved instances (RIs).
4AWS CCP49	4	15.4	Answer = 4 The AWS Artifact Explanation : It offers on-demand access to AWS' security and compliance reports, Payment Card Industry (PCI) reports, and certifications from accreditation bodies across geographies and compliance verticals that validate the implementation and operation of AWS' services. AWS Artifact is your go-to, central resource for compliance-related information that matters to your organization.

iPQ.No	Correct Answer	Points	Explanation
4AWS CCP50	3	15.4	Answer = 3 User data for Amazon EC2 instances Explanation : User data for an Amazon EC2 instance is information you entered during instance launch in the form of a bootstrap script or configuration settings. via: https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ec2-instance-metadata.html. Amazon EC2 instance metadata and user information Instance metadata and user data PDF Kindle RSS <i>Instance metadata</i> is data about your instance that you can use to configure or manage the running instance. Instance metadata is divided into categories, for example, host name, events, and security groups. You can also use instance metadata to access <i>user data</i> that you specified when launching your instance. For example, you can specify parameters for configuring your instance, or include a simple script. You can build generic AMIs and use user data to modify the configuration files supplied at launch time. For example, if you run web servers for various small businesses, they can all use the same generic AMI and retrieve their content from the Amazon S3 bucket that you specify in the user data at launch. To add a new customer at any time, create a bucket for the customer, add their content, and launch your AMI with the unique bucket name provided to your code in the user data. If you launch more than one instance at the same time, the user data is available to all instances in that reservation. Each instance that is part of the same reservation has a unique <code>ami-launch-index</code> number, allowing you to write code that controls what to do. For example, the first host might elect itself as an initial master node in a cluster. For a detailed AMI launch example, see Example: AMI launch index value.


iPQ.No	Correct Answer	Points	Explanation
4AWS CCP51	2	15.4	Answer = 2 To make sure that all Amazon S3 resources remain private, use Amazon S3 Block Public Access Explanation : You can control public access to Amazon S3 resources by using the settings for access points, buckets, and accounts provided by the Amazon S3 Block Public Access feature. New buckets, access points, and objects don't provide open access by default. To provide public access, users can change the bucket policies, access point policies, or object permissions. You can restrict public access to these resources by using S3 Block Public Access settings, which take precedence over these rules and restrictions. When Amazon S3 receives a request to access a bucket or an object, it checks to see if the bucket or the account of the bucket owner has a setting in place to prohibit public access. Amazon S3 additionally looks for prohibited public access settings for the access point if the request was made through one. Amazon S3 denies the request if a block public access configuration already in place forbids the requested access. through https://docs.aws.amazon.com/AmazonS3/latest/dev/access-control-block-public-access.html Amazon S3 Block Public Access Overview Using Amazon S3 block public access PDF Kindle RSS The Amazon S3 Block Public Access feature provides settings for access points, buckets, and accounts to help you manage public access to Amazon S3 resources. By default, new buckets, access points, and objects don't allow public access. However, users can modify bucket policies, access point policies, or object permissions to allow public access. S3 Block Public Access settings override these policies and permissions so that you can limit public access to these resources. With S3 Block Public Access, account administrators and bucket owners can easily set up centralized controls to limit public access to their Amazon S3 resources that are enforced regardless of how the resources are created. When Amazon S3 receives a request to access a bucket or an object, it determines whether the bucket or the bucket owner's account has a block public access setting applied. If the request was made through an access point, Amazon S3 also checks for block public access settings for the access point. If there is an existing block public access setting that prohibits the requested access, Amazon S3 rejects the request. Amazon S3 Block Public Access provides four settings. These settings are independent and can be used in any combination. Each setting can be applied to an access point, a bucket, or an entire AWS account. If the block public access settings for the access point, bucket, or account differ, then Amazon S3 applies the most restrictive combination of the access point, bucket, and account settings. When Amazon S3 evaluates whether an operation is prohibited by a block public access setting, it rejects any request that violates an access point, bucket, or account setting. Warning Public access is granted to buckets and objects through access control lists (ACLs), access point policies, bucket policies, or all. To help ensure that all of your Amazon S3 access points, buckets, and objects have their public access blocked, we recommend that you turn on all four settings for block public access for your account. These settings block public access for all current and future buckets and access points. Before applying these settings, verify that your applications will work correctly without public access. If you require some level of public access to your buckets or objects, for example to host a static website as described at Hosting a static website on Amazon S3, you can customize the individual settings to suit your storage use cases.

iPQ.No	Correct Answer	Points	Explanation
4AWS CCP52	3	15.4	Answer = 3 Both A and B Explanation : Your ability to delegate more of your operational duties to AWS and so increase your agility and innovation is made possible by serverless, which is the cloud's natural architecture. You may create and execute apps and services using serverless technology without giving servers any thought. It does away with infrastructure management duties like capacity provisioning, operating system upkeep, server or cluster provisioning, and patching. http://aws.amazon.com/serverless/ provides an overview of the AWS serverless platform. Compute AWS Lambda lets you run code without provisioning or managing servers. You pay only for the compute time you consume - there is no charge when your code is not running. Lambda@Edge allows you to run Lambda functions at AWS Edge locations in response to Amazon CloudFront events. AWS Fargate is a purpose-built serverless compute engine for containers. Fargate scales and manages the infrastructure required to run your containers. Storage Amazon Simple Storage Service (Amazon S3) provides developers and IT teams with secure, durable, highly-scalable object storage. Amazon S3 is easy to use, with a simple web service interface to store and retrieve any amount of data from anywhere on the web. Amazon Elastic File System (Amazon EFS) provides simple, scalable, elastic file storage. It is built to elastically scale on demand, growing and shrinking automatically as you add and remove files. Data stores Amazon DynamoDB is a fast and flexible NoSQL database service for all applications that need consistent, single-digit millisecond latency at any scale. Amazon Aurora Serverless is an on-demand, auto-scaling configuration for Amazon Aurora (MySQL-compatible edition), where the database will automatically start up, shut down, and scale capacity up or down based on your application's needs. Amazon RDS Proxy is a highly available database proxy that manages thousands of concurrent connections to relational databases, allowing you to build highly scalable, API Proxy Amazon API Gateway is a fully managed service that makes it easy for developers to create, publish, maintain, monitor, and secure APIs at any scale. It offers a comprehensive platform for API management. API Gateway allows you to process hundreds of thousands of concurrent API calls and handles traffic management, authorization and access control, monitoring, and API version management. Application integration Amazon SNS is a fully managed pub/sub messaging service that makes it easy to decouple and scale microservices, distributed systems, and serverless applications. Amazon SQS is a fully managed message queuing service that makes it easy to decouple and scale microservices, distributed systems, and serverless applications. AWS AppSync simplifies application development by letting you create a flexible GraphQL API to securely access, manipulate, and combine data from one or more data sources. Orchestration AWS Step Functions makes it easy to coordinate the components of distributed applications and microservices using visual workflows. Building applications from individual components that each perform a discrete function lets you scale and change applications quickly. Step Functions is a reliable way to coordinate components and step through the functions of your application. Analytics Amazon Kinesis is a platform for streaming data on AWS, offering powerful services to make it easy to load and analyze streaming data, and also providing the ability for you to build custom streaming data applications for specialized needs. Amazon Athena is an interactive query service that makes it easy to analyze data in Amazon S3 using standard SQL. Athena is serverless, so there is no infrastructure to manage, and you pay only for the queries that you run. Developer tooling AWS provides tools and services that aid developers in the serverless application development process. AWS and its partner ecosystem offer tools for continuous integration and delivery, testing, deployments, monitoring and diagnostics, SDKs, frameworks, and integrated development environment (IDE) plugins. Lambda on AWS You may execute code for practically any kind of application or backend service using AWS Lambda, all while requiring zero management. Simply upload your code, and Lambda will take care of all the necessary steps to run and grow it while maintaining high availability. You can program your code to run other AWS services automatically or to be called directly from any web or mobile application. You can run code with AWS Lambda without setting up or managing

iPQ.No	Correct Answer	Points	Explanation
4AWS CCP53	3	15.4	Answer = 3 Warm-standby technique Explanation : You must balance the advantages of a lower RTO (recovery time goal) and RPO (recovery point target) against the costs of establishing and maintaining a strategy when choosing your disaster recovery (DR) plan. Both the pilot light and warm standby solutions provide a good ratio of advantages and disadvantages. Data resources in the recovery region, such as Amazon Relational Database Service (Amazon RDS) DB instances or Amazon DynamoDB tables, are used in this technique to duplicate data from the primary region. These data providers are prepared to respond to inquiries. This technique necessitates the creation of a continuous backup in the recovery region in addition to replication. This is because replication will spread the flawed data when "human action"-type events happen and data can be lost or distorted. You need backups to allow you to return to the last known good condition. A working stack is deployed, although at a lower capacity, using the warm standby technique. Requests can be processed by the DR endpoint, but production-level traffic cannot. Even if it is more, it will always be less than the full deployment of production to save money. Warm standby sends a functioning stack to the recovery region, making it simpler to assess region readiness using synthetic transactions. Hot standby, on the other hand, refers to the deployment of the passive stack to the recovery region at full capacity. Disaster recovery (DR) techniques are described in the following article:  Figure 1. DR strategies https://aws.amazon.com/blogs/architecture/disaster-recovery-dr-architecture-on-aws-part-iii-pilot-light-and-warm-standby

iPQ.No	Correct Answer	Points	Explanation
4AWS CCP54	3	15.4	Answer = 3 Fargate AWS Explanation : With both Amazon Elastic Container Service (ECS) and Amazon Elastic Kubernetes Service (EKS), AWS Fargate is a serverless computing engine for containers. You may easily concentrate on creating your applications with Fargate. Fargate eliminates the need to set up and manage servers, enables you to specify and pay for resources per application, and enhances security through application isolation from the start. through https://aws.amazon.com/fargate/. How AWS Fargate Operates 

iPQ.No	Correct Answer	Points	Explanation
4AWS CCP55	1	15.4	Answer = 1 (Amazon S3) Amazon Simple Storage Service Explanation : An object storage service called Amazon Simple Storage Service (Amazon S3) provides performance, security, and scalability that are unmatched in the market. Customers of different sizes and sectors are using Amazon S3's flat, non-hierarchical structure and various management capabilities to arrange their data in ways that are beneficial to their teams and organizations. You set up an Amazon S3 bucket for website hosting and then upload your website's content to the bucket to host a static website there. You must enable website hosting, specify permissions, create and add an index document when configuring a bucket as a static website. Using Amazon S3 to host a static website is described in this document: Hosting a static website on Amazon S3 PDF Kindle RSS You can use Amazon S3 to host a static website. On a <i>static</i> website, individual webpages include static content. They might also contain client-side scripts. By contrast, a <i>dynamic</i> website relies on server-side processing, including server-side scripts such as PHP, JSP, or ASP.NET. Amazon S3 does not support server-side scripting, but AWS has other resources for hosting dynamic websites. To learn more about website hosting on AWS, see Web Hosting. Note You can use the AWS Amplify Console to host a single page web app. The AWS Amplify Console supports single page apps built with single page app frameworks (for example, React JS, Vue JS, Angular JS, and Nuxt) and static site generators (for example, Gatsby JS, React-static, Jekyll, and Hugo). For more information, see Getting Started in the AWS Amplify Console User Guide. To configure your bucket for static website hosting, you can use the AWS Management Console without writing any code. You can also create, update, and delete the website configuration <i>programmatically</i> by using the AWS SDKs. The SDKs provide wrapper classes around the Amazon S3 REST API. If your application requires it, you can send REST API requests directly from your application. To host a static website on Amazon S3, you configure an Amazon S3 bucket for website hosting and then upload your website content to the bucket. When you configure a bucket as a static website, you must enable website hosting, set permissions, and create and add an index document. Depending on your website requirements, you can also configure redirects, web traffic logging, and a custom error document. After you configure your bucket as a static website, you can access the bucket through the AWS Region-specific Amazon S3 website endpoints for your bucket. Website endpoints are different from the endpoints where you send REST API requests. For more information, see Website endpoints. https://docs.aws.amazon.com/AmazonS3/latest/dev/WebsiteHosting.html


iPQ.No	Correct Answer	Points	Explanation
4AWS CCP56	2	15.4	Answer = 2 VPC peering relationship Explanation : A networking link called a VPC peering connection allows you to privately route traffic between two VPCs. Both VPCs allow communication between instances just like if they were on the same network. Between your VPCs, with a VPC in another AWS account, or with a VPC in a different AWS region, you can establish a VPC peering connection. An overview of VPC peering may be found at What is VPC peering? PDF Amazon Virtual Private Cloud (Amazon VPC) enables you to launch AWS resources into a virtual network that you've defined. A VPC peering connection is a networking connection between two VPCs that enables you to route traffic between them using private IPv4 addresses or IPv6 addresses. Instances in either VPC can communicate with each other as if they are within the same network. You can create a VPC peering connection between your own VPCs, or with a VPC in another AWS account. The VPCs can be in different regions (also known as an inter-region VPC peering connection).  AWS uses the existing infrastructure of a VPC to create a VPC peering connection; it is neither a gateway nor a VPN connection, and does not rely on a separate piece of physical hardware. There is no single point of failure for communication or a bandwidth bottleneck. A VPC peering connection helps you to facilitate the transfer of data. For example, if you have more than one AWS account, you can peer the VPCs across those accounts to create a file sharing network. You can also use a VPC peering connection to allow other VPCs to access resources you have in one of your VPCs. https://docs.aws.amazon.com/vpc/latest/peering/what-is-vpc-peering.html.

iPQ.No	Correct Answer	Points	Explanation
4AWS CCP57	2	15.4	Answer = 2 One or more data centers situated nearby Explanation : One or more distinct data centers with redundant power, networking, and connectivity make up an availability zone (AZ) in an AWS region. A geographic area where AWS clusters data centers is referred to as an AWS region. Customers now have the option to run production databases and applications that are more fault-tolerant, highly available, and scalable than would be conceivable in a single data center thanks to AZs. High-bandwidth, low-latency networking over fully redundant, dedicated metro fiber connects every AZ in an AWS region, enabling high-throughput, low-latency networking between AZs. According to https://aws.amazon.com/about-aws/global-infrastructure/regions_az/, AWS Regions and Availability Zones (AZs) are described. Regions AWS has the concept of a Region, which is a physical location around the world where we cluster data centers. We call each group of logical data centers an Availability Zone. Each AWS Region consists of multiple, isolated, and physically separate AZ's within a geographic area. Unlike other cloud providers, who often define a region as a single data center, the multiple AZ design of every AWS Region offers advantages for customers. Each AZ has independent power, cooling, and physical security and is connected via redundant, ultra-low-latency networks. AWS customers focused on high availability can design their applications to run in multiple AZ's to achieve even greater fault-tolerance. AWS infrastructure Regions meet the highest levels of security, compliance, and data protection. AWS provides a more extensive global footprint than any other cloud provider, and to support its global footprint and ensure customers are served across the world, AWS opens new Regions rapidly. AWS maintains multiple geographic Regions, including Regions in North America, South America, Europe, China, Asia Pacific, South Africa, and the Middle East. Availability Zones An Availability Zone (AZ) is one or more discrete data centers with redundant power, networking, and connectivity in an AWS Region. AZ's give customers the ability to operate production applications and databases that are more highly available, fault tolerant, and scalable than would be possible from a single data center. All AZ's in an AWS Region are interconnected with high-bandwidth, low-latency networking, over fully redundant, dedicated metro fiber providing high-throughput, low-latency networking between AZ's. All traffic between AZ's is encrypted. The network performance is sufficient to accomplish synchronous replication between AZ's. AZ's make partitioning applications for high availability easy. If an application is partitioned across AZ's, companies are better isolated and protected from issues such as power outages, lightning strikes, tornadoes, earthquakes, and more. AZ's are physically separated by a meaningful distance, many kilometers, from any other AZ, although all are within 100 km (60 miles) of each other.

iPQ.No	Correct Answer	Points	Explanation
4AWS CCP58	3	15.4	Answer = 3 Enormous scale economies Explanation : The on-demand, pay-as-you-go delivery of IT resources through the Internet is known as cloud computing. You can use technological services, such as computing power, storage, and databases, on an as-needed basis rather than purchasing, owning, and maintaining actual data centers and servers. You can attain a lower variable cost using cloud computing than you could on your own. Providers like AWS are able to reach larger economies of scale thanks to the aggregated usage from hundreds of thousands of clients, which results in cheaper pay-as-you-go costs. Exam Warning: You would undoubtedly be asked about the benefits of cloud computing in comparison to a conventional on-premises setup. Please review the following six advantages of cloud computing. via: Six Advantages of Cloud Computing PDF RSS • Trade capital expense for variable expense – Instead of having to invest heavily in data centers and servers before you know how you're going to use them, you can pay only when you consume computing resources, and pay only for how much you consume. • Benefit from massive economies of scale – By using cloud computing, you can achieve a lower variable cost than you can get on your own. Because usage from hundreds of thousands of customers is aggregated in the cloud, providers such as AWS can achieve higher economies of scale, which translates into lower pay as-you-go prices. • Stop guessing capacity – Eliminate guessing on your infrastructure capacity needs. When you make a capacity decision prior to deploying an application, you often end up either sitting on expensive idle resources or dealing with limited capacity. With cloud computing, these problems go away. You can access as much or as little capacity as you need, and scale up and down as required with only a few minutes' notice. • Increase speed and agility – In a cloud computing environment, new IT resources are only a click away, which means that you reduce the time to make those resources available to your developers from weeks to just minutes. This results in a dramatic increase in agility for the organization, since the cost and time it takes to experiment and develop is significantly lower. • Stop spending money running and maintaining data centers – Focus on projects that differentiate your business, not the infrastructure. Cloud computing lets you focus on your own customers, rather than on the heavy lifting of racking, stacking, and powering servers. • Go global in minutes – Easily deploy your application in multiple regions around the world with just a few clicks. This means you can provide lower latency and a better experience for your customers at minimal cost. https://docs.aws.amazon.com/whitepapers/latest/aws-overview/six-advantages-of-cloud-computing.html

iPQ.No	Correct Answer	Points	Explanation
4AWS CCP59	3	15.4	Answer = 3 Both A and B Explanation : AWS Global Accelerator is a networking solution that aids in enhancing the performance and accessibility of the apps you provide to your international users. By proxying packets at the edge to apps running in one or more AWS regions, Global Accelerator enhances performance for a variety of applications using TCP or UDP. Using https://aws.amazon.com/global-accelerator/, find out how AWS Global Accelerator functions.  For non-HTTP use scenarios, the AWS Global Accelerator is an excellent fit. AWS Global Accelerator is a strong fit for HTTP use cases that expressly demand static IP addresses or deterministic, quick regional failover, as well as for non-HTTP use cases like gaming (UDP), IoT (MQTT), or voice over IP. Static IP addresses are provided by AWS Global Accelerator and serve as a consistent point of entry for your apps. It offers static IP addresses, which give your apps a stable entry point and take the hassle out of configuring unique IP addresses for various AWS regions and availability zones (AZs).
4AWS CCP60	2	15.4	Answer = 2 Amazon Translate Explanation : With its neural machine translation technology, Amazon Translate provides quick, accurate, and cost-effective language translation. You can simply localize content for international consumers, such as websites and applications, and efficiently translate significant amounts of text using Amazon Translate.

iPQ.No	Correct Answer	Points	Explanation
4AWS CCP61	3	15.2	Answer = 3 Both A and B Explanation : An almost real-time stream of system events describing modifications to Amazon Web Services (AWS) resources is provided via Amazon CloudWatch Events. You may match events and rapidly route them to one or more target functions or streams by using straightforward rules. As soon as an operational change takes place, Amazon CloudWatch Events is informed. By delivering messages to respond to the environment, activating functions, making changes, and capturing state data, Amazon CloudWatch Events reacts to these operational changes and takes remedial action as appropriate. You may decouple microservices, distributed systems, and serverless applications using the highly available, reliable, secure, and fully managed pub/sub messaging service known as Amazon Simple Notification Service (Amazon SNS). Additionally, Amazon SNS can be utilized to notify users by email, SMS, and mobile push. through https://aws.amazon.com/sns/. How Amazon SNS Functions <pre> graph LR Publisher[Publisher Publish messages from distributed systems, microservices, and other AWS services] --> SNS[Amazon SNS Fully-managed pub/sub messaging and event-driven computing service] SNS --> Topic[SNS Topic Decouple message publishers from subscribers with topics] Topic --> Filter[Message Filtering & Fanout Filter messages according to subscription filter policies, and deliver them to subscribers] Filter --> Subscribers[Subscribers Receive messages in subscribing serverless functions, queues, microservices, and more] subgraph Subscribers_Box [Subscribers] direction TB Lambda[AWS Lambda] SQS[Amazon SQS] HTTP[HTTP/S] end </pre> You are able to build an Amazon Simple Notification Service (Amazon SNS) subject to deliver notifications each time the root user of the AWS account logs in. Create an Amazon CloudWatch event rule to keep track on user identity root logins through the AWS Management Console, and when the event occurs, send an email via Amazon SNS.

iPQ.No	Correct Answer	Points	Explanation
4AWS CCP62	3	15.2	Answer = 3 Trusted Advisor for AWS Explanation : An online application called AWS Trusted Advisor offers real-time advice to help you supply your resources in accordance with AWS best practices. Regular recommendations from Trusted Advisor help maintain your solutions provisioned at their best, whether for creating new workflows, creating applications, or as part of continuing improvement. AWS Trusted Advisor examines your AWS infrastructure and offers best practice advice in the following five areas: cost reduction, performance improvement, security enhancement, fault tolerance, and service limit adjustments. When volumes seem to be underutilized, AWS Trusted Advisor can analyze Amazon Elastic Block Store (Amazon EBS) volume parameters and issue a warning. When a volume is produced, charges start. A volume is presumably not being utilized if it is detached or has very little write activity (except from boot volumes) for an extended length of time. Using https://aws.amazon.com/premiumsupport/technology/trusted-advisor/, find out how AWS Trusted Advisor functions.  <pre> graph LR A[As an AWS customer, you want the most value from your investment. Trusted Advisor can help.] --> B[AWS Trusted Advisor Trusted Advisor scans your AWS infrastructure, compares it to AWS best practices in five categories, and provides recommended actions] B --> C[Cost Optimization] B --> D[Performance] B --> E[Security] B --> F[Fault Tolerance] B --> G[Service limits] C --> H[Investigation recommended] D --> I[no problem detected] E --> J[action recommended] F --> K[no problem detected] G --> L[no problem detected] H --> M[Recommendations & action links] I --> M J --> M K --> M L --> M </pre>

iPQ.No	Correct Answer	Points	Explanation
4AWS CCP63	4	15.2	Answer = 4 Budgets AWS Explanation : You can create personalized budgets with AWS Budgets, and the system will notify you if your expenses or use go above the predetermined limit. You may set reservation usage or coverage goals using AWS Budgets and get notifications when your utilization falls below the threshold you specify. When the usage of your reserved instances (RI) falls below a specified threshold, you can specify a threshold and receive alerts when that happens. This enables you to determine whether or not your reserved instances (RI) are being used. Amazon EC2, Amazon RDS, Amazon Redshift, Amazon ElastiCache, and Amazon Elasticsearch reservation alerts are all supported. through https://docs.aws.amazon.com/awsaccountbilling/latest/aboutv2/budgets-managing-costs.html AWS Budgets Overview Managing your costs with AWS Budgets PDF Kindle RSS AWS Budgets enable you to plan your service usage, service costs, and instance reservations. Budgets provide you with a way to see the following information: • How close your plan is to your budgeted amount or to the free tier limits • Your usage to date, including how much you have used of your Reserved Instances (RIs) • Your current estimated charges from AWS and how much your predicted usage will incur in charges by the end of the month • How much of your budget has been used AWS Budgets information is updated up to three times a day. Updates typically occur between 8 to 12 hours after the previous update. Budgets track your unblended costs, subscriptions, refunds, and RIs. You can create the following types of budgets: • Cost budgets – Plan how much you want to spend on a service. • Usage budgets – Plan how much you want to use one or more services. • RI utilization budgets – Define a utilization threshold and receive alerts when your RI usage falls below that threshold. This lets you see if your RIs are unused or under-utilized. • RI coverage budgets – Define a coverage threshold and receive alerts when the number of your instance hours that are covered by RIs fall below that threshold. This lets you see how much of your instance usage is covered by a reservation. • Savings Plans utilization budgets – Define a utilization threshold and receive alerts when the usage of your Savings Plans falls below that threshold. This lets you see if your Savings Plans are unused or under-utilized. • Savings Plans coverage budgets – Define a coverage threshold and receive alerts when your Savings Plans eligible usage that is covered by Savings Plans fall below that threshold. This lets you see how much of your instance usage is covered by Savings Plans.

iPQ.No	Correct Answer	Points	Explanation
4AWS CCP64	3	15.2	Answer = 3 Both A and B Explanation : A online service called Amazon Elastic computation Cloud (Amazon EC2) offers safe, scalable computation capacity in the cloud and supports per-second billing. It is the simplest method for accessing the underlying OS and provisioning servers in the AWS cloud. The freedom, control, and significant cost advantages that Amazon EC2 provides developers allow them to use Amazon EC2 instances as their own customized servers supported by the powerful architecture of AWS Cloud, which is a major difference between Amazon EC2 and traditional on-premises hosting solutions. With changing computing needs, Amazon EC2 can scale. Amazon EC2 may be scaled to meet requirements when computing needs unpredictably change. The number of EC2 instances running at any given time can be managed by developers. With Amazon EC2, programmers can start and stop instances whenever they need to. Developers can create, terminate, start, or shut down instances using Amazon EC2 at any moment, and they can entirely alter the configuration of their instances to meet their requirements.

iPQ.No	Correct Answer	Points	Explanation
4AWS CCP65	1	15.2	Answer = 1 A load balancer for applications Explanation : Incoming application traffic is automatically split among numerous targets, including Amazon EC2 instances, containers, IP addresses, and Lambda functions, using elastic load balancing (ELB). It can manage the fluctuating volume of traffic from your application in one availability zone or several. Three different types of load balancers are offered by elastic load balancing (ELB), all of which have the high availability, automatic scaling, and strong security requirements for fault-tolerant systems. obtained from: https://aws.amazon.com/elasticloadbalancing Regions AWS has the concept of a Region, which is a physical location around the world where we cluster data centers. We call each group of logical data centers an Availability Zone. Each AWS Region consists of multiple, isolated, and physically separate AZ's within a geographic area. Unlike other cloud providers, who often define a region as a single data center, the multiple AZ design of every AWS Region offers advantages for customers. Each AZ has independent power, cooling, and physical security and is connected via redundant, ultra-low-latency networks. AWS customers focused on high availability can design their applications to run in multiple AZ's to achieve even greater fault-tolerance. AWS infrastructure Regions meet the highest levels of security, compliance, and data protection. AWS provides a more extensive global footprint than any other cloud provider, and to support its global footprint and ensure customers are served across the world, AWS opens new Regions rapidly. AWS maintains multiple geographic Regions, including Regions in North America, South America, Europe, China, Asia Pacific, South Africa, and the Middle East. Availability Zones An Availability Zone (AZ) is one or more discrete data centers with redundant power, networking, and connectivity in an AWS Region. AZ's give customers the ability to operate production applications and databases that are more highly available, fault tolerant, and scalable than would be possible from a single data center. All AZ's in an AWS Region are interconnected with high-bandwidth, low-latency networking, over fully redundant, dedicated metro fiber providing high-throughput, low-latency networking between AZ's. All traffic between AZ's is encrypted. The network performance is sufficient to accomplish synchronous replication between AZ's. AZ's make partitioning applications for high availability easy. If an application is partitioned across AZ's, companies are better isolated and protected from issues such as power outages, lightning strikes, tornadoes, earthquakes, and more. AZ's are physically separated by a meaningful distance, many kilometers, from any other AZ, although all are within 100 km (60 miles) of each other. The employment of an application load balancer, which offers enhanced request routing aimed at the delivery of current application architectures such as microservices and containers, allows for the load balancing of HTTP and HTTPS traffic.