

QUESTIONNAIRE REPORT SHEET

iPQ3AWS CCP-CLF.V1.1



AWS CERTIFIED CLOUD PRACTITIONER


CLF-C01 PRACTICE TESTS

iPQ.No	Correct Answer	Points	Explanation
3AWS CCP-CLF01	3	15.4	Answer = 3 Excellence in Operations Explanation : You can learn more about the benefits and drawbacks of decisions you make when developing systems on AWS by using the AWS Well-Architected Framework. You will discover architectural best practices for creating and managing dependable, secure, effective, and affordable cloud systems by utilizing the framework. It gives you a tool to regularly compare your architecture to best practices and pinpoint areas that need improvement. The operational excellence, security, dependability, performance effectiveness, cost optimization, and sustainability are the six pillars on which the AWS Well-Architected Framework is built. The capacity to manage systems to create business value and continuously enhance auxiliary practices and processes is a component of the operational excellence pillar. In the cloud, you can treat your entire environment with the same engineering rigour that you do for application code. Your whole workload (applications, infrastructure) can be defined as code and updated via code. Your operational processes can be codified, and you can trigger them in reaction to events to automate their implementation.

iPQ.No	Correct Answer	Points	Explanation
3AWS CCP-CLF02	3	15.4	Answer = 3 Both A and B Explanation : A network access control list (network ACL) is a set of numbered rules that determines whether to accept traffic by ranking them in order of importance. For your instance, a security group serves as a virtual firewall to manage incoming and outgoing traffic. Rather than at the subnet level, security groups operate at the instance level. Because security groups are stateful, if you submit a request from your instance, any rules for inbound security group traffic are ignored and the response traffic for that request is permitted to enter. Before agreeing to let traffic, a security group considers all rules. Security group basics The following are the basic characteristics of security groups for your VPC: - There are quotas on the number of security groups that you can create per VPC, the number of rules that you can add to each security group, and the number of security groups that you can associate with a network interface. For more information, see Amazon VPC quotas. - You can specify allow rules, but not deny rules. - You can specify separate rules for inbound and outbound traffic. - When you create a security group, it has no inbound rules. Therefore, no inbound traffic originating from another host to your instance is allowed until you add inbound rules to the security group. - By default, a security group includes an outbound rule that allows all outbound traffic. You can remove the rule and add outbound rules that allow specific outbound traffic only. If your security group has no outbound rules, no outbound traffic originating from your instance is allowed. - Security groups are stateful — if you send a request from your instance, the response traffic for that request is allowed to flow in regardless of inbound security group rules. Responses to allowed inbound traffic are allowed to flow out, regardless of outbound rules. Note Some types of traffic are tracked differently from other types. For more information, see Connection tracking in the <i>Amazon EC2 User Guide for Linux Instances</i>. - Instances associated with a security group can't talk to each other unless you add rules allowing the traffic (exception: the default security group has these rules by default). - Security groups are associated with network interfaces. After you launch an instance, you can change the security groups that are associated with the instance, which changes the security groups associated with the primary network interface (eth0). You can also specify or change the security groups associated with any other network interface. By default, when you create a network interface, it's associated with the default security group for the VPC, unless you specify a different security group. For more information about network interfaces, see Elastic network interfaces. - When you create a security group, you must provide it with a name and a description. The following rules apply: - Names and descriptions can be up to 255 characters in length. - Names and descriptions are limited to the following characters: a-z, A-Z, 0-9, spaces, and _-:/()#,@[]+=&;{}!\$*. - A security group name cannot start with sg- as these indicate a default security group. - A security group name must be unique within the VPC. - A security group can only be used in the VPC that you specify when you create the security group. via https://docs.aws.amazon.com/vpc/latest/userguide/VPC_SecurityGroups.html Security group overview An optional layer of protection for your VPC is a network access control list (network ACL), which functions as a firewall for regulating traffic into and out of one or more subnets (i.e., it controls traffic at the subnet level). A

iPQ.No	Correct Answer	Points	Explanation
3AWS CCP-CLF03	3	15.4	Answer = 3 To obtain different invoices for development and production settings, use distinct AWS accounts Explanation : At the end of the month, each AWS account provides its own invoice. By creating separate AWS accounts for development and production environments, you can receive distinct invoices for each environment.
3AWS CCP-CLF04	2	15.4	Answer = 2 ElastiCache by Amazon Explanation : You can easily set up, operate, and grow well-known open-source compatible in-memory data stores in the cloud using Amazon ElastiCache. Build data-intensive applications or improve the performance of your current databases by getting data from in-memory data storage with high throughput and low latency. For real-time use cases like caching, session stores, gaming, geospatial services, real-time analytics, and queuing, Amazon ElastiCache is a popular option. Online analytical processing is not possible with ElastiCache. Using https://aws.amazon.com/elasticache/, find out how Amazon ElastiCache functions.

iPQ.No	Correct Answer	Points	Explanation
3AWSCCP-CLF05	4	15.4	Answer = 4 For both the Amazon Elastic Compute Cloud (Amazon EC2) instance and the on-premises servers, use Amazon CloudWatch logs Explanation : AWS CloudTrail, Route 53, Amazon Elastic Compute Cloud (Amazon EC2) instances, as well as other sources like on-premises servers, can all produce log files that can be monitored, stored, and accessed using Amazon CloudWatch Logs. You can centralize the logs from any system, application, and AWS service you use using Amazon CloudWatch Logs into a single, highly scalable service. Then, it will be simple for you to see them, search for particular error codes or patterns, apply filters based on particular fields, or safely archive them for further examination.
3AWSCCP-CLF06	4	15.4	Answer = 4 To duplicate data between different AWS regions, use S3 cross-region replication (S3 CRR) Explanation : An object storage service called Amazon Simple Storage Service (Amazon S3) provides performance, security, and scalability that are unmatched in the market. Objects can be automatically and asynchronously copied between Amazon S3 buckets thanks to replication. The same AWS account or distinct AWS accounts may possess buckets that are set up for object replication. You can copy things within the same AWS area or between separate AWS regions. Even while Amazon S3 stores your data by default across several geographically distant availability zones, compliance regulations may necessitate that you store data even further away. To meet these requirements, you can replicate data between distant AWS regions using S3 cross-region replication (S3 CRR).

iPQ.No	Correct Answer	Points	Explanation
3AWSCCP-CLF07	1	15.4	Answer = 1 AWS Lambda function versions should be maintained Explanation : According to the AWS Shared Responsibility Model, AWS is in charge of safeguarding the infrastructure that powers all of the AWS Cloud services. The facilities, networking, hardware, and software needed to run AWS cloud services make up this infrastructure. The customer's liability is decided by the AWS cloud services they choose under the AWS Shared liability Model. For services that are abstracted, like Amazon S3 and Amazon DynamoDB, AWS manages the platforms, operating system, and infrastructure layer. Customers access the endpoints to store and retrieve data. Customers are in charge of identifying their assets, managing their data (including encryption settings), and using IAM tools to assign the proper permissions. The customer is in charge of updating an AWS Lambda function's versions for the specified use case. 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 vertical arrows represent the scope of responsibility: a blue arrow pointing up labeled 'CUSTOMER' with 'RESPONSIBILITY FOR SECURITY 'IN' THE CLOUD' and an orange arrow pointing down labeled 'AWS' with 'RESPONSIBILITY FOR SECURITY 'OF' THE CLOUD'. To the right, a stack of boxes represents the layers of the cloud stack. The top three layers (blue) are the customer's responsibility: 'CUSTOMER DATA', 'PLATFORM, APPLICATIONS, IDENTITY & ACCESS MANAGEMENT', and 'OPERATING SYSTEM, NETWORK & FIREWALL CONFIGURATION'. The next layer (dark blue) is split 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 (orange) includes 'COMPUTE', 'STORAGE', 'DATABASE', and 'NETWORKING'. The bottom two layers (orange) are AWS's responsibility: 'HARDWARE/AWS GLOBAL INFRASTRUCTURE' and 'REGIONS', 'AVAILABILITY ZONES', and 'EDGE LOCATIONS'.

iPQ.No	Correct Answer	Points	Explanation
3AWS CCP-CLF08	2	15.4	Answer = 2 Kindle Redshift Explanation : With the help of your existing business intelligence (BI) tools and standard SQL, you can quickly and affordably analyze all of your data with Amazon Redshift, a quick and fully managed cloud data warehouse. It enables you to execute massively parallel queries while leveraging advanced query optimization, columnar storage on high-performance storage, and complicated analytic queries against terabytes to petabytes of structured data.

iPQ.No	Correct Answer	Points	Explanation
3AWS CCP-CLF09	1	15.4	Answer = 1 Effect, Action Explanation : The majority of policies are kept in JSON format on AWS. JSON policy documents that you attach to a user or role are used to define identity-based policies and policies that determine permission limitations. JSON policy documents that you attach to a resource are resource-based policies. These components are found in a JSON policy document: At the top of the document is optional information for the entire policy. Several distinct statements Each statement contains details regarding a single permit. A statement is composed of a number of pieces that together comprise the information. The version of the policy language you want to utilize should be specified. Use the most recent 2012-10-17 version as a best practice. You can add more than one statement in a policy. Use this primary policy element as a container for the following elements. Include an optional statement ID in Sid to distinguish amongst your remarks. Effect: To specify whether the policy permits or prohibits access, use Allow or Deny.

iPQ.No	Correct Answer	Points	Explanation
3AWS CCP-CLF10	3	15.4	Answer = 3 Both A and B Explanation : All users in your account should be subject to multi-factor authentication (MFA), according to AWS. Users have a device that produces a response to an authentication challenge when using multi-factor authentication (MFA). To complete the sign-in procedure, the user's credentials and the device-generated answer are both necessary. Change your credentials frequently. You should frequently change your personal passwords and access keys, as well as those of all IAM users in your account, per AWS's advice. By doing this, you can limit how long a password or access key can be used to access your resources in the event that it is compromised without your awareness. Applying a password policy to your account will force all IAM users to change their passwords on a regular basis. Best practices for AWS IAM security may be found at Security Best Practices in IAM PDF Kindle RSS To help secure your AWS resources, follow these recommendations for the AWS Identity and Access Management (IAM) service. Topics • Lock Away Your AWS Account Root User Access Keys • Create Individual IAM Users • Use Groups to Assign Permissions to IAM Users • Grant Least Privilege • Get Started Using Permissions with AWS Managed Policies • Use Customer Managed Policies Instead of Inline Policies • Use Access Levels to Review IAM Permissions • Configure a Strong Password Policy for Your Users • Enable MFA • Use Roles for Applications That Run on Amazon EC2 Instances • Use Roles to Delegate Permissions • Do Not Share Access Keys • Rotate Credentials Regularly • Remove Unnecessary Credentials • Use Policy Conditions for Extra Security • Monitor Activity in Your AWS Account https://docs.aws.amazon.com/IAM/latest/UserGuide/best


iPQ.No	Correct Answer	Points	Explanation
3AWSCCP-CLF11	3	15.4	Answer = 3 CodeDeploy AWS Explanation : AWS CodeDeploy is a service that automates the deployment of code to any instance, including instances hosted on-premises and in Amazon EC2. You can deliver new features more quickly, avoid downtime during deployment, and handle the complexity of updating your apps with AWS CodeDeploy. AWS CodeDeploy enables deployment automation, doing away with the need for labor-intensive manual tasks. You may quickly deploy to one instance or thousands of instances because the service expands with your infrastructure.
3AWSCCP-CLF12	3	15.4	Answer = 3 Both A and B Explanation : Designed for use with Amazon Elastic Compute Cloud (Amazon EC2) for both throughput and transaction-intensive applications at any scale, Amazon Elastic Block Store (Amazon EBS) is a user-friendly, high-performance block storage solution. On Amazon EBS, a wide variety of workloads are frequently implemented, including relational and non-relational databases, corporate applications, containerized applications, big data analytics engines, file systems, and media workflows. Block-level temporary storage is provided for your EC2 instance by an instance store. Disks that are physically connected to the host computer house this storage. Instance storage is perfect for replicating data across a fleet of instances, such as a load-balanced pool of web servers, or for the temporary storage of information that changes often, such as buffers, caches, scratch data, and other transient material. Data is lost if an instance fails or is terminated since instance storage is only transitory. File sharing between instances is not possible using the Amazon EC2 instance storage.


iPQ.No	Correct Answer	Points	Explanation
3AWS CCP-CLF13	3	15.4	Answer = 372 Explanation : When compared to on-demand instance pricing, reserved instances (RI) provide you significant savings (up to 72%) on your Amazon Elastic Compute Cloud (Amazon EC2) expenses. Reserved Instances (RI) are a billing discount that is applied to the use of on-demand instances in your account; they are not actual physical instances. A reserved instance (RI) can be purchased for a one-year or three-year commitment, with the third year providing a greater savings. Overview of the pricing options for Amazon Elastic Compute Cloud (Amazon EC2) may be found at On-Demand With On-Demand instances, you pay for compute capacity by the hour or the second depending on which instances you run. No longer-term commitments or upfront payments are needed. You can increase or decrease your compute capacity depending on the demands of your application and only pay the specified per hourly rates for the instance you use. On-Demand instances are recommended for: - Users that prefer the low cost and flexibility of Amazon EC2 without any up-front payment or long-term commitment - Applications with short-term, spiky, or unpredictable workloads that cannot be interrupted - Applications being developed or tested on Amazon EC2 for the first time See On-Demand pricing » Spot instances Amazon EC2 Spot instances allow you to request spare Amazon EC2 computing capacity for up to 90% off the On-Demand price. Learn More. Spot instances are recommended for: - Applications that have flexible start and end times - Applications that are only feasible at very low compute prices - Users with urgent computing needs for large amounts of additional capacity See Spot pricing » Savings Plans Savings Plans are a flexible pricing model that offer low prices on EC2 and Fargate usage, in exchange for a commitment to a consistent amount of usage (measured in \$/hour) for a 1 or 3 year term. Dedicated Hosts A Dedicated Host is a physical EC2 server dedicated for your use. Dedicated Hosts can help you reduce costs by allowing you to use your existing server-bound software licenses, including Windows Server, SQL Server, and SUSE Linux Enterprise Server (subject to your license terms), and can also help you meet compliance requirements. Learn more. - Can be purchased On-Demand (hourly). - Can be purchased as a Reservation for up to 70% off the On-Demand price. See Dedicated pricing » Reserved Instances Reserved Instances provide you with a significant discount (up to 75%) compared to On-Demand instance pricing. In addition, when Reserved Instances are assigned to a specific Availability Zone, they provide a capacity reservation, giving you additional confidence in your ability to launch instances when you need them. For applications that have steady state or predictable usage, Reserved Instances can provide significant savings compared to using On-Demand instances. See How to Purchase Reserved Instances for more information. Reserved Instances are recommended for: - Applications with steady state usage - Applications that may require reserved capacity - Customers that can commit to using EC2 over a 1 or 3 year term to reduce their total computing costs https://aws.amazon.com/ec2/pricing

iPQ.No	Correct Answer	Points	Explanation
3AWSCCP-CLF14	4	15.4	Answer = 4 AWS Health Dashboard: The state of your account Explanation : When AWS is having occurrences that could affect you, the AWS Health Dashboard - Your Account Health warns you and offers corrective advice. AWS Health Dashboard - Your account health, notifications are sent out when your AWS resources' health changes, providing you with event visibility and advice to help you immediately identify and fix problems. To find out the most recent status, visit https://health.aws.amazon.com/health/home.
3AWSCCP-CLF15	2	15.4	Answer = 2 North Virginia's US East - us-east-1 Explanation : Using Amazon CloudWatch, you can keep an eye on your anticipated AWS costs. Billing metric data, which represents global charges, is kept in the US East (N. Virginia) region. Along with the expected total of your AWS charges, this information contains the estimated costs for each service in AWS that you utilize.


iPQ.No	Correct Answer	Points	Explanation									
3AWS CCP-CLF16	1	15.4	Answer = 1 Region AWS Explanation : The actual place where AWS develops its data centers is known as an AWS Region. An Availability Zone (AZ) is a set of logical data centers. Each AWS region is made up of a number of geographically distinct, isolated, and physically independent AZs. Regarding AWS regions in the US, please refer to this illustration: Map Key - Regions (Orange circle) - Edge locations (Blue dot) <table><tr><td>US East (Northern Virginia) Region Availability Zones: 6 Launched 2006</td><td>US East (Ohio) Region Availability Zones: 3 Launched 2016</td><td>GovCloud (US-West) Region Availability Zones: 3 Launched 2011</td></tr><tr><td></td><td>US West (Oregon) Region Availability Zones: 4 Launched 2011 Local Zone: 1 Launched 2019</td><td>GovCloud (US-East) Region Availability Zones: 3 Launched 2018</td></tr><tr><td></td><td>US West (Northern California) Region Availability Zones: 3* Launched 2009</td><td>Canada (Central) Region** Availability Zones: 3 Launched 2016</td></tr></table> Learn more at AWS Canada See detailed offerings at all AWS locations >> through - https://aws.amazon.com/about-aws/global-infrastructure/regions_az	US East (Northern Virginia) Region Availability Zones: 6 Launched 2006	US East (Ohio) Region Availability Zones: 3 Launched 2016	GovCloud (US-West) Region Availability Zones: 3 Launched 2011		US West (Oregon) Region Availability Zones: 4 Launched 2011 Local Zone: 1 Launched 2019	GovCloud (US-East) Region Availability Zones: 3 Launched 2018		US West (Northern California) Region Availability Zones: 3* Launched 2009	Canada (Central) Region** Availability Zones: 3 Launched 2016
US East (Northern Virginia) Region Availability Zones: 6 Launched 2006	US East (Ohio) Region Availability Zones: 3 Launched 2016	GovCloud (US-West) Region Availability Zones: 3 Launched 2011										
	US West (Oregon) Region Availability Zones: 4 Launched 2011 Local Zone: 1 Launched 2019	GovCloud (US-East) Region Availability Zones: 3 Launched 2018										
	US West (Northern California) Region Availability Zones: 3* Launched 2009	Canada (Central) Region** Availability Zones: 3 Launched 2016										

iPQ.No	Correct Answer	Points	Explanation
3AWS CCP-CLF17	3	15.4	Answer = 3 Both A and B Explanation : For files accessed less frequently, Amazon Elastic File System (Amazon EFS) - Infrequent Access storage class is cost-optimized. Each time you read from or write to a file on the Amazon Elastic File System (Amazon EFS) - Infrequent Access storage type, you will be charged a price. Your block data is preserved as a point-in-time copy using Amazon EBS snapshots. A complete copy of your data is saved by Amazon EBS to Amazon S3 on the initial snapshot of a volume. Because Amazon EBS Snapshots are stored incrementally, you are only charged for the stored modified blocks.

iPQ.No	Correct Answer	Points	Explanation
3AWS CCP-CLF18	3	15.4	Answer = 3 Both A and B Explanation : Trusted Advisor for AWS 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 Amazon Elastic Compute Cloud (Amazon EC2) instances that have been running at any point in the previous 14 days are checked by AWS Trusted Advisor, which notifies you if the daily CPU utilization is 10% or less and the network I/O is 5 MB or less on four or more days. See https://aws.amazon.com/premiumsupport/technology/trusted-advisor/ for more information on how AWS Trusted Advisor functions.  The diagram illustrates the AWS Trusted Advisor workflow. It starts with a user icon and text: 'As an AWS customer, you want the most value from your investment. Trusted Advisor can help.' An arrow points to a central box labeled 'AWS Trusted Advisor' with a shield icon. Below this box, it says: 'Trusted Advisor scans your AWS infrastructure, compares it to AWS best practices in five categories, and provides recommended actions.' An arrow points from this box to a table of five categories: Cost Optimization, Performance, Security, Fault Tolerance, and Service Limits. Each category has a corresponding icon and a status box. The status boxes are: 'Investigation recommended' (with a warning icon), 'no problem detected' (with a checkmark), 'action recommended' (with a warning icon), 'no problem detected' (with a checkmark), and 'no problem detected' (with a checkmark). A dashed arrow points from the 'Investigation recommended' status box back to the 'AWS Trusted Advisor' box. Below the table, it says 'Recommendations & action links'. Finding low-utilization Amazon EC2 instances with AWS Trusted Advisor is possible by visiting this link: Low utilization Amazon EC2 instances

iPQ.No	Correct Answer	Points	Explanation
3AWSCCP-CLF19	3	15.4	Answer = 3 Both A and B Explanation : AWS and the customer both share responsibilities for security and compliance. With AWS maintaining, managing, and controlling every component, from the host operating system and virtualization layer down to the physical security of the premises in which the service is housed, this shared model might lessen the operational load placed on the customer. Security 'of' the Cloud is the responsibility of AWS. The infrastructure that powers all of the services provided by the AWS Cloud is under the control of AWS. The facilities, networking, hardware, and software used to run AWS Cloud services make up this infrastructure. AWS manages the infrastructure layer, the operating system, and platforms for abstracted services like Amazon S3. For services that are abstracted, like Amazon S3 and Amazon DynamoDB, AWS manages the platforms, operating system, and infrastructure layer. Customers access the endpoints to store and retrieve data. Overview of the AWS Shared Responsibility Model may be found at 

iPQ.No	Correct Answer	Points	Explanation
3AWS CCP-CLF20	3	15.4	Answer = 3 Across numerous Availability Zones (AZ), Regions, and Virtual Private Clouds (VPCs), EC2 instances can access files on an Amazon Elastic File System (Amazon EFS) file system Explanation : For high availability and durability, Amazon EFS is a local service that stores data within and across several Availability Zones (AZs). While on-premises servers can access your file system using AWS Direct Connect or AWS VPN, Amazon EC2 instances can do so across AZs, regions, and VPCs. Overview of Amazon EFS available at: Amazon Elastic File System (Amazon EFS) provides a simple, scalable, fully managed elastic NFS file system for use with AWS Cloud services and on-premises resources. It is built to scale on demand to petabytes without disrupting applications, growing and shrinking automatically as you add and remove files, eliminating the need to provision and manage capacity to accommodate growth. Amazon EFS offers two storage classes: the Standard storage class, and the Infrequent Access storage class (EFS IA). EFS IA provides price/performance that's cost-optimized for files not accessed every day. By simply enabling EFS Lifecycle Management on your file system, files not accessed according to the lifecycle policy you choose will be automatically and transparently moved into EFS IA. The EFS IA storage class costs only \$0.025/GB-month*. While workload patterns vary, customers typically find that 80% of files are infrequently accessed (and suitable for EFS IA), and 20% are actively used (suitable for EFS Standard), resulting in an effective storage cost as low as \$0.08/GB-month*. Amazon EFS transparently serves files from both storage classes in a common file system namespace. Amazon EFS is designed to provide massively parallel shared access to thousands of Amazon EC2 instances, enabling your applications to achieve high levels of aggregate throughput and IOPS with consistent low latencies. Amazon EFS is well suited to support a broad spectrum of use cases from home directories to business-critical applications. Customers can use EFS to lift-and-shift existing enterprise applications to the AWS Cloud. Other use cases include: big data analytics, web serving and content management, application development and testing, media and entertainment workflows, database backups, and container storage. Amazon EFS is a regional service storing data within and across multiple Availability Zones (AZs) for high availability and durability. Amazon EC2 instances can access your file system across AZs, regions, and VPCs, while on-premises servers can access using AWS Direct Connect or AWS VPN. https://aws.amazon.com/efs

iPQ.No	Correct Answer	Points	Explanation
3AWSCCP-CLF21	1	15.4	Answer = 1 Device with virtual multi-factor authentication (MFA) Explanation : A software application that simulates a physical device and runs on a phone or other device creates a six-digit numeric code using a time-synchronized one-time password technique. During sign-in, the user must enter a valid device code on a different webpage. A user must receive a different virtual multi-factor authentication (MFA) device each time. A user cannot authenticate by entering a code from another user's fictitious Multi-Factor Authentication (MFA) device. A virtual Multi-Factor Authentication (MFA) device is an example, such as Google Authenticator: 

iPQ.No	Correct Answer	Points	Explanation
3AWSCCP-CLF22	1	15.4	Answer = 1 (AMI) Amazon Machine Image Explanation : The data needed to launch an instance is provided via an Amazon Machine Image (AMI). When launching an instance, you must supply an Amazon Machine Image (AMI). Among the components of an Amazon Machine Image (AMI) are these: a template for the root disk of the instance (for example, an operating system, an application server, and applications), one or more Amazon Elastic Block Store (Amazon EBS) snapshots, or, for example, store-backed AMIs. Which AWS accounts can launch instances using the Amazon Machine Image (AMI) are controlled by launch permissions. a block device mapping that designates the volumes that should be attached to the instance during launch. The Amazon Machine Image (AMI) lifecycle is depicted in the graphic below: <pre> graph LR create --> EBS[EBS snapshot (or template if instance store backed)] EBS -- register --> AMI1[AMI #1] AMI1 -- launch --> Instance[Instance] AMI1 -- copy --> AMI2[AMI #2] AMI1 -- deregister --> Deregister[] </pre> via - https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/AMIs.html

iPQ.No	Correct Answer	Points	Explanation
3AWS CCP-CLF23	2	15.4	Answer = 2 Make use of Amazon CloudFront to enhance the functionality of your website Explanation : Amazon CloudFront can help your website run more efficiently. Your website's files, including HTML, photos, and video, are made accessible from data centers all over the world (referred to as edge locations) by Amazon CloudFront. Amazon CloudFront automatically sends a file request from a website visitor to a copy of the requested file at the closest edge point. As a result, the visitor receives the content more quickly than if they had requested it from a data center that is further away.

iPQ.No	Correct Answer	Points	Explanation
3AWSACP-CLF24	1	15.4	Answer = 1 Backup routing 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. When you want to configure active-passive failover, you utilize a failover routing policy. When a resource is healthy, traffic can be sent to it. When the initial resource is ill, traffic can be routed to a different resource. The primary and secondary records can direct users to a complex tree of records or an Amazon S3 bucket that has been set up as a website. through – 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. https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/routing-policy.html



iPQ.No	Correct Answer	Points	Explanation
3AWS CCP-CLF25	3	15.4	Answer = 3 The business must simply begin developing new resources in the target AWS Region before migrating the pertinent data and apps there Explanation : The business must first develop resources in the new AWS Region before moving the pertinent data and applications there. The business must first generate resources in the new AWS Region before moving the pertinent data and applications there. The business cannot employ an off-the-shelf product or service to make this move easier.
3AWS CCP-CLF26	4	15.4	Answer = 4 Knowledge Base of AWS Explanation : The answers to the most typical and frequently asked questions are available in the AWS Knowledge Center. When looking for a fix or performing troubleshooting on an AWS service issue, one should start here. https://aws.amazon.com/premiumsupport/knowledge-center/ is the website address for the Knowledge Center.
3AWS CCP-CLF27	3	15.4	Answer = 3 Both A and B Explanation : With the goal of maintaining consistent, predictable performance at the lowest feasible cost, AWS Auto Scaling continuously monitors your applications and automatically modifies the capacity. It's simple to build up application scaling for numerous resources across many services in a matter of minutes using AWS Auto Scaling. You may create scaling plans for resources like Amazon EC2 instances and Spot Fleets, Amazon ECS tasks, Amazon DynamoDB tables and indexes, and Amazon Aurora Replicas using the service's straightforward but effective user interface. Scaling policies allow you to dynamically raise or decrease the number of instances in your group to adapt to changing circumstances. The Auto Scaling group starts or terminates the instances as necessary while the scaling policy is in place, adjusting the group's desired capacity between the minimum and maximum values that you define. Scaling is also possible on a schedule. .

iPQ.No	Correct Answer	Points	Explanation
3AWSCCP-CLF28	3	15.4	Answer = 3 Both A and B Explanation : You can run code with AWS Lambda without setting up or managing servers. 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 handle all the necessary steps to run and grow it with high availability. You only pay for what you really use with AWS Lambda. Your fees are determined by how many times your functions are called as well as how long it takes for your code to run. When AWS Lambda begins operating in response to an event notification or invoke call, including test invokes via the console, it counts each request as one. The execution time of your code is measured, rounded to the nearest 100ms, from the moment it starts until it returns or terminates.

iPQ.No	Correct Answer	Points	Explanation																																								
3AWSCCP-CLF29	3	15.4	Answer = 3 Both A and B Explanation : An AWS resource may have a label called a Cost Allocation Tag applied by you or by AWS. An associated key and value make up each tag. Each tag key for a resource must be distinct, and each tag key can only have one value. Tags can be used to categorize your resources, and cost allocation tags can be used to track your AWS expenses in great detail. User-defined tags and tags generated by AWS are the two different forms of cost allocation tags offered by AWS. You define, create, and apply user-defined tags, while AWS defines, creates, and applies AWS generated tags for you. Before either sort of tag can be seen in Cost Explorer or on a cost allocation report, it must be activated independently. through - https://docs.aws.amazon.com/awsaccountbilling/latest/aboutv2/cost-alloc-tags.html AWS Cost Allocation Tags Overview Using Cost Allocation Tags PDF Kindle RSS A tag is a label that you or AWS assigns to an AWS resource. Each tag consists of a key and a value. For each resource, each tag key must be unique, and each tag key can have only one value. You can use tags to organize your resources, and cost allocation tags to track your AWS costs on a detailed level. After you activate cost allocation tags, AWS uses the cost allocation tags to organize your resource costs on your cost allocation report, to make it easier for you to categorize and track your AWS costs. AWS provides two types of cost allocation tags, an AWS generated tags and user-defined tags. AWS defines, creates, and applies the AWS generated tags for you, and you define, create, and apply user-defined tags. You must activate both types of tags separately before they can appear in Cost Explorer or on a cost allocation report. The following diagram illustrates the concept. In the example, you've assigned and activated tags on two Amazon EC2 instances, one tag called Cost Center and another tag called Stack. Each of the tags has an associated value. You also activated the AWS generated tags, createdBy before creating these resources. The createdBy tag tracks who created a resource. The user-defined tags use the user: prefix, and the AWS generated tag uses the aws: prefix. After you or AWS applies tags to your AWS resources (such as Amazon EC2 instances or Amazon S3 buckets) and you activate the tags in the Billing and Cost Management console, AWS generates a cost allocation report as a comma-separated value (CSV file) with your usage and costs grouped by your active tags. You can apply tags that represent business categories (such as cost centers, application names, or owners) to organize your costs across multiple services. The cost allocation report includes all of your AWS costs for each billing period. The report includes both tagged and untagged resources, so that you can clearly organize the charges for resources. For example, if you tag resources with an application name, you can track the total cost of a single application that runs on those resources. The following screenshot shows a partial report with columns for each tag. <table><thead><tr><th>Total Cost</th><th>user:Owner</th><th>user:Stack</th><th>user:Cost Center</th><th>user:Application</th></tr></thead><tbody><tr><td>0.95</td><td>DBAdmin</td><td>Test</td><td>80432</td><td>Widget2</td></tr><tr><td>0.01</td><td>DBAdmin</td><td>Test</td><td>80432</td><td>Widget2</td></tr><tr><td>3.84</td><td>DBAdmin</td><td>Prod</td><td>80432</td><td>Widget2</td></tr><tr><td>6.00</td><td>DBAdmin</td><td>Test</td><td>78925</td><td>Widget1</td></tr><tr><td>234.63</td><td>SysMgr</td><td>Prod</td><td>78925</td><td>Widget1</td></tr><tr><td>0.73</td><td>DBAdmin</td><td>Test</td><td>78925</td><td>Widget1</td></tr><tr><td>0.00</td><td>DBAdmin</td><td>Prod</td><td>80432</td><td>Portal</td></tr></tbody></table>	Total Cost	user:Owner	user:Stack	user:Cost Center	user:Application	0.95	DBAdmin	Test	80432	Widget2	0.01	DBAdmin	Test	80432	Widget2	3.84	DBAdmin	Prod	80432	Widget2	6.00	DBAdmin	Test	78925	Widget1	234.63	SysMgr	Prod	78925	Widget1	0.73	DBAdmin	Test	78925	Widget1	0.00	DBAdmin	Prod	80432	Portal
Total Cost	user:Owner	user:Stack	user:Cost Center	user:Application																																							
0.95	DBAdmin	Test	80432	Widget2																																							
0.01	DBAdmin	Test	80432	Widget2																																							
3.84	DBAdmin	Prod	80432	Widget2																																							
6.00	DBAdmin	Test	78925	Widget1																																							
234.63	SysMgr	Prod	78925	Widget1																																							
0.73	DBAdmin	Test	78925	Widget1																																							
0.00	DBAdmin	Prod	80432	Portal																																							

iPQ.No	Correct Answer	Points	Explanation
3AWS CCP-CLF30	2	15.4	Answer = 2 Create multiple Availability Zones (AZ) for the Amazon Elastic Compute Cloud (Amazon EC2) instances inside the same AWS Region Explanation : AWS has the idea of a Region, which is a geographical area where AWS groups data centers. Each AWS Region is made up of several (three or more), isolated, and geographically distinct Availability Zones (AZ). Each Availability Zone (AZ) is connected by redundant, ultra-low-latency networks and has independent power, cooling, and physical security. One or more distinct data centers with redundant power, networking, and connectivity make up an Availability Zone (AZ) in an AWS Region. A physical 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 from Availability Zones (AZ) than from a single data center. In order to provide high-throughput, low-latency networking between Availability Zones (AZ), all Availability Zones (AZ) in an AWS Region are interconnected with high-bandwidth, low-latency networking over fully redundant, dedicated metro fiber. Using the following link: 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. https://aws.amazon.com/about-aws/global-infrastructure/regions_az/, AWS Regions and Availability Zones (AZ) are described.


iPQ.No	Correct Answer	Points	Explanation
3AWSCCP-CLF31	3	15.4	Answer = 3 Both A and B Explanation : According to the AWS Shared Responsibility Model, AWS is in charge of safeguarding the infrastructure that powers all of the AWS Cloud services. The facilities, networking, hardware, and software needed to run AWS cloud services make up this infrastructure. The AWS cloud services that the user chooses will define what is their responsibility. 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. upgrades to an Amazon Elastic Compute Cloud (Amazon EC2) instance's operating system The customer is in charge of maintaining security "in" the cloud. As an Infrastructure as a Service (IaaS) offering, Amazon Elastic Compute Cloud (Amazon EC2) requires the customer to carry out all required security configuration and management duties. enabling data encryption for information kept in buckets provided by Amazon Simple Storage Service (Amazon S3) Customers are in charge of managing their data (including encryption options), categorizing their assets, and using IAM tools to assign the proper permissions under the AWS Shared Responsibility Model. through https://aws.amazon.com/compliance/shared-responsibility-model/. AWS Shared Responsibility Model

iPQ.No	Correct Answer	Points	Explanation
3AWS CCP-CLF32	4	15.4	Answer = 4 Recognize the items in a picture Explanation : You can recognize objects, people, text, scenes, and activities in pictures and videos with Amazon Rekognition, and you can also screen out any offensive material. You can detect, analyze, and compare faces using Amazon Rekognition's highly accurate facial analysis and facial search capabilities for a variety of user verification, population counting, and public safety use cases. Use-Cases for Amazon Rekognition: Key features  Labels With Amazon Rekognition, you can identify thousands of objects (such as bike, telephone, building), and scenes (such as parking lot, beach, city). When analyzing video, you can also identify specific activities such as "delivering a package" or "playing soccer". Learn more »  Custom labels With Amazon Rekognition Custom Labels, you can extend the detection capabilities of Amazon Rekognition to extract information from images that is uniquely helpful to your business. For example, you can find your corporate logo in social media, identify your products on store shelves, classify your machine parts in an assembly line, or detect your animated characters in videos. Learn more »  Content moderation Amazon Rekognition helps you identify potentially unsafe or inappropriate content across both image and video assets and provides you with detailed labels that allow you to accurately control what you want to allow based on your needs. Use Amazon A2I to enhance the accuracy of Amazon Rekognition image moderation predictions using human review. Learn more »  Text detection In photos and videos, text appears very differently than neat words on a printed page. Amazon Rekognition can read skewed and distorted text to capture information like store names, forced narratives overlaid on media, street signs, and text on product packaging. Learn more » retrieved from: https://aws.amazon.com/rekognition  Face detection and analysis With Amazon Rekognition, you can easily detect when faces appear in images and videos and get attributes such as gender, age range, eyes open, glasses, facial hair for each. In video, you can also measure how these face attributes change over time, such as constructing a timeline of the emotions expressed by an actor. Learn more »  Face search and verification Amazon Rekognition provides fast and accurate face search, allowing you to identify a person in a photo or video using your private repository of face images. You can also verify identity by analyzing a face image against images you have stored for comparison. Learn more »

iPQ.No	Correct Answer	Points	Explanation
3AWS CCP-CLF33	3	15.4	Answer = 3 The Amazon Machine Image (AMI) you use must be from the same region as the Amazon EC2 instance. The performance of the Amazon EC2 instance is independent of the region of the Amazon Machine Image (AMI) Explanation : The data needed to launch an instance is provided via an Amazon Machine Image (AMI). When launching an instance, you must supply an Amazon Machine Image (AMI). When you require numerous instances with the same configuration, you can launch them all from a single AMI. The Amazon Machine Image (AMI) and the Amazon EC2 instance that will be launched must be in the same region. You can copy the Amazon Machine Image (AMI) to the region where you wish to start the EC2 instance if it already exists in another region. The performance of the Amazon EC2 instance is independent of the region of the Amazon Machine Image (AMI). Overview of the Amazon Machine Images (AMI) service may be found at Amazon Machine Images (AMI) PDF Kindle RSS An Amazon Machine Image (AMI) provides the information required to launch an instance. You must specify an AMI when you launch an instance. You can launch multiple instances from a single AMI when you need multiple instances with the same configuration. You can use different AMIs to launch instances when you need instances with different configurations. An AMI includes the following: - One or more EBS snapshots, or, for instance-store-backed AMIs, a template for the root volume of the instance (for example, an operating system, an application server, and applications). - Launch permissions that control which AWS accounts can use the AMI to launch instances. - A block device mapping that specifies the volumes to attach to the instance when it's launched. Using an AMI The following diagram summarizes the AMI lifecycle. After you create and register an AMI, you can use it to launch new instances. (You can also launch instances from an AMI if the AMI owner grants you launch permissions.) You can copy an AMI within the same Region or to different Regions. When you no longer require an AMI, you can deregister it. <pre> graph LR Create[create] --> EBS[EBS snapshot (or template if instance-store backed)] EBS --> Register[register] Register --> AMI1[AMI #1] AMI1 --> Launch[launch] Launch --> Instance[Instance] AMI1 --> Copy[copy] Copy --> AMI2[AMI #2] AMI1 --> Deregister[deregister] </pre> You can search for an AMI that meets the criteria for your instance. You can search for AMIs provided by AWS or AMIs provided by the community. For more information, see AMI types and Finding a Linux AMI. After you launch an instance from an AMI, you can connect to it. When you are connected to an instance, you can use it just like you use any other server. For information about launching, connecting, and using your instance, see Amazon EC2 instances. https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/AMIs.html.


iPQ.No	Correct Answer	Points	Explanation
3AWS CCP-CLF34	4	15.4	Answer = 4 Agility 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 like computing power, storage, and databases on an as-needed basis from a cloud provider like Amazon Web Services (AWS), rather than purchasing, operating, and maintaining physical data centers and servers. Agility: Agility is the capacity of the cloud to provide you with simple access to a wide range of technologies so that you can develop more quickly and construct almost anything you can think of. From infrastructure services like computation, storage, and databases to the Internet of Things, machine learning, data lakes, analytics, and much more, you can instantly spin up resources as you need them.
3AWS CCP-CLF35	4	15.4	Answer = 4 (Amazon ECS) Amazon Elastic Container Service Explanation : Running, stopping, and managing Docker containers on a cluster is simple using Amazon Elastic Container solution (Amazon ECS), a highly scalable, quick container management solution. You are able to manage the underlying servers yourself because this is not a completely managed service.

iPQ.No	Correct Answer	Points	Explanation
3AWS CCP-CLF36	3	15.4	Answer = 3 Storing documents in a location that is easily accessible to meet audit requirements Explanation : For files that are not accessed frequently, the Amazon EFS Standard-Infrequent Access (EFS Standard-IA) storage class lowers storage costs. It accomplishes this while maintaining Amazon Elastic File System's (Amazon EFS) high availability, robustness, elasticity, and POSIX file system access. If you want to automatically reduce storage costs for files that are visited less frequently while still needing your entire dataset to be easily accessible, AWS suggests using the Amazon EFS Standard-Infrequent Access (EFS Standard-IA) storage class. For instance, doing historical analysis, backup and recovery, or keeping material available to meet audit needs. All Amazon EFS functionalities are compatible with Amazon EFS Standard-Infrequent Access (EFS Standard-IA) storage, which is accessible in all AWS regions where Amazon EFS is offered.

iPQ.No	Correct Answer	Points	Explanation
3AWSCCP-CLF37	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. According to https://aws.amazon.com/premiumsupport/technology/trusted-advisor, Trusted Advisor functions as follows.  Recommendations from AWS Trusted Advisor are available at 

iPQ.No	Correct Answer	Points	Explanation
3AWSCCP-CLF38	4	15.4	Answer = 4 All of the above Explanation : You may schedule your service consumption, service prices, and instance reservations with AWS Budgets. Up to three times each day, AWS updates its budget data. Usually, updates come between eight and twelve hours after the last one. Your unblended costs, subscriptions, refunds, and RIs are tracked by budgets. A cost budget, a use budget, a reservation budget, and a savings plan budget are the four main budget types you can build with AWS Budgets. You can use a cost budget to determine how much you want to spend on a service. Budget for usage: This you in planning how frequently you will use one or more services. Budget for reservations: This aids in keeping track of how your reserved instances (RI) are being used. There are two methods to go about it: Reserved Instance (RI) coverage budgets and Reserved Instance (RI) utilization budgets, which both show you how much of your instance usage is covered by reservations and whether your Reserved Instances (RI) are being used or not.

IPQ.No	Correct Answer	Points	Explanation																								
3AWS CCP-CLF39	2	15.4	Answer = 2 Database scalability is increased through Read Replica Explanation : Setting up, running, and scaling a relational database in the cloud is simple with Amazon Relational Database Service (Amazon RDS). You can build read-only replicas that are synchronized with your main database using replicas. The usage of read replicas enhances reading performance. For better speed, you may also move your read replica to a different AWS region that is closer to your users. An illustration of horizontal resource scaling is replicas. via: https://aws.amazon.com/rds/features/multi-az/ Read 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. Application servers Database server <pre>graph LR subgraph Application_servers [Application servers] A1[] A2[] A3[] end subgraph Database_server [Database server] P[Primary] R[Read replica] end A1 -- Read/write --> P A2 -- Read/write --> P A3 -- Read/write --> P P -- Asynchronous replication --> R R -- Read only --> B[BI/reporting application server]</pre> Exam Warning: Review the distinctions between the Amazon RDS Multi-AZ, Multi-Region, and Read Replica installations for RDS here: Read replicas, Multi-AZ deployments, and multi-region deployments Amazon RDS read replicas complement Multi-AZ deployments. While both features maintain a second copy of your data, there are differences between the two: <table><thead><tr><th>Multi-AZ deployments</th><th>Multi-Region deployments</th><th>Read replicas</th></tr></thead><tbody><tr><td>Main purpose is high availability</td><td>Main purpose is disaster recovery and local performance</td><td>Main purpose is scalability</td></tr><tr><td>Non-Aurora: synchronous replication; Aurora: asynchronous replication</td><td>Asynchronous replication</td><td>Asynchronous replication</td></tr><tr><td>Non-Aurora: only the primary instance is active; Aurora: all instances are active</td><td>All regions are accessible and can be used for reads</td><td>All read replicas are accessible and can be used for readscaling</td></tr><tr><td>Non-Aurora: automated backups are taken from standby; Aurora: automated backups are taken from shared storage layer</td><td>Automated backups can be taken in each region</td><td>No backups configured by default</td></tr><tr><td>Always span at least two Availability Zones within a single region</td><td>Each region can have a Multi-AZ deployment</td><td>Can be within an Availability Zone, Cross-AZ, or Cross-Region</td></tr><tr><td>Non-Aurora: database engine version upgrades happen on primary; Aurora: all instances are updated together</td><td>Non-Aurora: database engine version upgrade is independent in each region; Aurora: all instances are updated together</td><td>Non-Aurora: database engine version upgrade is independent from source instance; Aurora: all instances are updated together</td></tr><tr><td>Automatic failover to standby (non-Aurora) or read replica (Aurora) when a problem is detected</td><td>Aurora allows promotion of a secondary region to be the master</td><td>Can be manually promoted to a standalone database instance (non-Aurora) or to be the primary instance (Aurora)</td></tr></tbody></table>	Multi-AZ deployments	Multi-Region deployments	Read replicas	Main purpose is high availability	Main purpose is disaster recovery and local performance	Main purpose is scalability	Non-Aurora: synchronous replication; Aurora: asynchronous replication	Asynchronous replication	Asynchronous replication	Non-Aurora: only the primary instance is active; Aurora: all instances are active	All regions are accessible and can be used for reads	All read replicas are accessible and can be used for readscaling	Non-Aurora: automated backups are taken from standby; Aurora: automated backups are taken from shared storage layer	Automated backups can be taken in each region	No backups configured by default	Always span at least two Availability Zones within a single region	Each region can have a Multi-AZ deployment	Can be within an Availability Zone, Cross-AZ, or Cross-Region	Non-Aurora: database engine version upgrades happen on primary; Aurora: all instances are updated together	Non-Aurora: database engine version upgrade is independent in each region; Aurora: all instances are updated together	Non-Aurora: database engine version upgrade is independent from source instance; Aurora: all instances are updated together	Automatic failover to standby (non-Aurora) or read replica (Aurora) when a problem is detected	Aurora allows promotion of a secondary region to be the master	Can be manually promoted to a standalone database instance (non-Aurora) or to be the primary instance (Aurora)
Multi-AZ deployments	Multi-Region deployments	Read replicas																									
Main purpose is high availability	Main purpose is disaster recovery and local performance	Main purpose is scalability																									
Non-Aurora: synchronous replication; Aurora: asynchronous replication	Asynchronous replication	Asynchronous replication																									
Non-Aurora: only the primary instance is active; Aurora: all instances are active	All regions are accessible and can be used for reads	All read replicas are accessible and can be used for readscaling																									
Non-Aurora: automated backups are taken from standby; Aurora: automated backups are taken from shared storage layer	Automated backups can be taken in each region	No backups configured by default																									
Always span at least two Availability Zones within a single region	Each region can have a Multi-AZ deployment	Can be within an Availability Zone, Cross-AZ, or Cross-Region																									
Non-Aurora: database engine version upgrades happen on primary; Aurora: all instances are updated together	Non-Aurora: database engine version upgrade is independent in each region; Aurora: all instances are updated together	Non-Aurora: database engine version upgrade is independent from source instance; Aurora: all instances are updated together																									
Automatic failover to standby (non-Aurora) or read replica (Aurora) when a problem is detected	Aurora allows promotion of a secondary region to be the master	Can be manually promoted to a standalone database instance (non-Aurora) or to be the primary instance (Aurora)																									

iPQ.No	Correct Answer	Points	Explanation
3AWSCCP-CLF40	3	15.4	Answer = 3 Both A and B Explanation : You can securely link your on-premises network or branch office location to your Amazon Virtual Private Cloud (Amazon VPC) using AWS Site-to-Site VPN. If your needs are urgent and your bandwidth needs are little to medium, VPN connections are a suitable option. The public internet is used for this connection. The parts of a VPC are a virtual private gateway (VGW), a transit gateway, and a customer gateway. The VPN concentrator on the Amazon side of the AWS site-to-site VPN connection is called a virtual private gateway (VGW). A customer gateway is an AWS resource that offers AWS data regarding your customer gateway device. AWS Site-to-Site VPN components are listed in Components of your Site-to-Site VPN A Site-to-Site VPN connection offers two VPN tunnels between a virtual private gateway or a transit gateway on the AWS side, and a customer gateway on the remote (on-premises) side. A Site-to-Site VPN connection consists of the following components. For more information about Site-to-Site VPN quotas, see Site-to-Site VPN quotas. Contents • Virtual private gateway • Transit gateway • Customer gateway • Customer gateway device Virtual private gateway A <i>virtual private gateway</i> is the VPN concentrator on the Amazon side of the Site-to-Site VPN connection. You create a virtual private gateway and attach it to the VPC from which you want to create the Site-to-Site VPN connection.  https://docs.aws.amazon.com/vpn/latest/s2svpn/how_it_works.html.

iPQ.No	Correct Answer	Points	Explanation
3AWSCCP-CLF41	1	15.4	Answer = 1 Endpoint VPC Explanation : Without the need for an internet gateway, NAT device, VPN connection, or AWS Direct Connect connection, a VPC endpoint enables you to securely connect your VPC to supported AWS services and VPC endpoint services enabled by AWS PrivateLink. Public IP addresses are not necessary for instances in your VPC to communicate with service resources. No traffic leaves the Amazon network to travel between your VPC and the other service. Interface endpoints and gateway endpoints are the two different sorts of VPC endpoints. An interface endpoint is a private IP address on an elastic network interface that is part of your subnet's IP address space and acts as an access point 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. A gateway endpoint is a gateway that you designate in your route table as the destination for a route for traffic going to an approved AWS service. There is support for the following AWS services: (Amazon S3) Amazon Simple Storage Service DynamoDB by Amazon Exam Warning: On the exam, there can be a question centered on this idea.


iPQ.No	Correct Answer	Points	Explanation
3AWSCCP-CLF42	1	15.4	Answer = 1 AWS Shield Advanced provides defense against potential increases in costs brought on by DDoS attacks Explanation : AWS Shield Advanced provides some financial security against any DDoS attack-related increases in your AWS bill. Your Amazon CloudFront distributions, Amazon Route 53 hosted zones, Amazon Elastic Compute Cloud instances, and your AWS Global Accelerator accelerators are all covered by this cost protection. Regardless of the support package, AWS Shield Advanced is a paid service for all clients.
3AWSCCP-CLF43	3	15.4	Answer = 3 Cost and Usage Report for AWS (AWS CUR) Explanation : The most complete set of cost and usage data is found in the AWS Cost and Usage Report (AWS CUR). Your AWS billing reports can be published to an Amazon Simple Storage Service (Amazon S3) bucket that you own using AWS Cost & Usage Report (AWS CUR). You can get reports that show your costs broken down by the hour, month, product, resource, or tags that you define. Once each day, AWS refreshes the CSV-formatted report in your bucket. via: https://docs.aws.amazon.com/cur/latest/userguide/what-is-cur.html AWS Cost & Usage Report (AWS CUR) Overview What are AWS Cost and Usage Reports? PDF RSS The AWS Cost and Usage Reports (AWS CUR) contains the most comprehensive set of cost and usage data available. You can use Cost and Usage Reports to publish your AWS billing reports to an Amazon Simple Storage Service (Amazon S3) bucket that you own. You can receive reports that break down your costs by the hour or month, by product or product resource, or by tags that you define yourself. AWS updates the report in your bucket once a day in comma-separated value (CSV) format. You can view the reports using spreadsheet software such as Microsoft Excel or Apache OpenOffice Calc, or access them from an application using the Amazon S3 API. AWS Cost and Usage Reports tracks your AWS usage and provides estimated charges associated with your account. Each report contains line items for each unique combination of AWS products, usage type, and operation that you use in your AWS account. You can customize the AWS Cost and Usage Reports to aggregate the information either by the hour or by the day. AWS Cost and Usage Reports can do the following: • Deliver report files to your Amazon S3 bucket • Update the report up to three times a day • Create, retrieve, and delete your reports using the AWS CUR API Reference


iPQ.No	Correct Answer	Points	Explanation
3AWSCCP-CLF44	3	15.4	Answer = 3 Testing for Penetration Explanation : For a few popular AWS services, customers can conduct security audits or penetration testing against their AWS infrastructure without prior authorization. Customers are not allowed to do any security audits on the AWS services or infrastructure.
3AWSCCP-CLF45	2	15.4	Answer = 2 Business Explanation : The AWS Cloud Adoption Framework (AWS CAF) makes creative use of AWS and draws on AWS expertise and best practices to assist you in digitally transforming and accelerating your business outcomes. Specific organizational characteristics that support effective cloud conversions are identified by AWS CAF. Six perspectives make up the AWS CAF, which groups its capabilities: business, people, governance, platform, security, and operations.

iPQ.No	Correct Answer	Points	Explanation
3AWS CCP-CLF46	3	15.4	Answer = 3 Amazon Web Services Explanation : Cloud Foundations offers customers a guided approach to assist them in setting up, securing, and deploying new workloads while making sure they are prepared for ongoing cloud operations. By curating AWS services, AWS solutions, partner solutions, and advice, Cloud Foundations assists clients in making the decisions they must.  The diagram illustrates the AWS Cloud Foundations architecture, organized into several key areas: Governance, Risk Management, and Compliance: Includes Log Storage, Data De-identification, Tagging, Service Onboarding, Change Management, and Forensics. Operations: Includes Developer Experience & Tools, Logging & Monitoring, Patching, Rollout/Rollback, Sort/Search for Metadata, Audit & Assessment, and Governance. Security: Includes Identity Management & Access Control, Encryption & Key Management, Secrets Management, Data Isolation, Security Incident Response, Vulnerability & Threat Management, and Application Security. Infrastructure: Includes Network Connectivity, Network Security, Workload Isolation Boundary, and Template Management. Business Continuity: Includes Disaster Recovery, Backups, and Support. Finance: Includes Cloud Financial Management, Records Management, and Resource Inventory. via: https://aws.amazon.com/architecture/cloud-foundations
3AWS CCP-CLF47	2	15.4	Answer = 2 Migration Evaluator for AWS Explanation : AWS Migration Evaluator, formerly TSO Logic, is a free service for building data-driven business cases for planning and migrating to the AWS cloud. Making wise AWS planning and migration decisions is made much easier with the help of the AWS Migration Evaluator. With the help of AWS Migration Evaluator, your company can create a data-driven business case for AWS, receive access to AWS knowledge, see the costs connected to various migration plans, and learn how reusing existing software licensing lowers costs even more.

iPQ.No	Correct Answer	Points	Explanation
3AWSCCP-CLF48	3	15.4	Answer = 3 Both A and B Explanation : (Amazon S3) Amazon Simple Storage Service Objects are automatically encrypted by employing server-side encryption with Amazon S3 managed keys (SSE-S3) in all Amazon S3 buckets, which is configured by default. All the items in your Amazon S3 buckets are subject to this encryption setting. Storage Gateway by AWS You can access nearly infinite cloud storage on-premises with the help of the hybrid cloud storage service known as AWS Storage Gateway. For all three types of gateways—file, volume, and tape—all data that is sent between the gateway and AWS storage is encrypted using SSL.



iPQ.No	Correct Answer	Points	Explanation
3AWS CCP-CLF49	4	15.4	Answer = 4 Kindle Kendra Explanation : Machine learning underlies Amazon Kendra, an intelligent search engine. For your websites and applications, Kendra reimagines corporate search to make it simple for users to locate the content they need, even when it's dispersed across many locations and content repositories within your business. You may stop sifting through mountains of unstructured data by using Amazon Kendra to quickly find the appropriate responses to your queries. Since Amazon Kendra is a fully managed service, there is no need to set up servers or develop, train, or use machine learning models. Unstructured and semi-structured data in HTML, MS Office (.doc,.ppt), PDF, and text formats are supported by Amazon Kendra. No matter where the information is located inside your business, natural language search capabilities may quickly and accurately provide the results you're looking for. Because Amazon Kendra's deep learning models are pre-trained across 14 industry categories, it can extract more accurate results from the start for a variety of commercial use cases. Additionally, you can manually modify the weighting of data sources, authors, freshness, and custom tags to fine-tune search results.

iPQ.No	Correct Answer	Points	Explanation
3AWSCCP-CLF50	3	15.4	Answer = 3 Both A and B Explanation : Snowball AWS AWS Snowball is a data transport solution that uses storage devices that are secure for physical transfer to speed up transporting terabytes to petabytes of data into and out of AWS services. AWS DMS (AWS Database Migration Service) You can easily and securely move databases from on-premises to AWS with the help of the AWS Database Migration Service (AWS DMS). Application downtime is kept to a minimum during migrating because the source database is still fully functional. Your data can be moved to and from the most popular commercial and open-source databases with the help of the AWS Database Migration Service (AWS DMS). AWS Database Migration Service (AWS DMS) is a tool for homogeneous and heterogeneous database migrations and may be accessed  The diagram illustrates two migration paths using AWS Database Migration Service (AWS DMS). In the top path, an Oracle Database (Source) is migrated to Amazon RDS For Oracle Database (Target). In the bottom path, a MySQL Database (Source) is migrated to Amazon Aurora Database (Target). Both paths show the data flowing from the source database through the AWS DMS service to the target database. via https://aws.amazon.com/dms/.

iPQ.No	Correct Answer	Points	Explanation
3AWSCCP-CLF51	1	15.4	Answer = 1 Lambda on AWS Explanation : You can run code with AWS Lambda without setting up or managing servers. Only the computing time that you use is charged. 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. AWS Lambda can be used to start processing data following an upload using Amazon S3. AWS Lambda can be used, for instance, to thumbnail photographs, convert films, index files, analyze logs, validate content, aggregate data in real-time, and filter it. Using https://aws.amazon.com/lambda/, you can see how AWS Lambda runs code in response to a trigger from S3.  <pre> graph LR A[Photograph is taken] --> B[Amazon S3 Photo is uploaded to an S3 Bucket] B --> C[Lambda is triggered] C --> D[AWS Lambda Lambda runs image resizing code] D --> E[Photo is resized into web, mobile, and tablet sizes] </pre>


iPQ.No	Correct Answer	Points	Explanation
3AWSCCP-CLF52	3	15.4	Answer = 3 Fix the bucket policy for Amazon S3 Explanation : 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 static 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. 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 You must modify the bucket's prohibit public access settings if you wish to put up an existing bucket as a static website with public access. You could also need to change your account's settings for barring public access. When implementing the bucket-level and account-level block public access settings, Amazon S3 uses the most stringent combination possible. The following information explains how to change the S3 buckets' public access settings: Block public access (bucket settings) Public access is granted to buckets and objects through access control lists (ACLs), bucket policies, access point policies, or all. In order to ensure that public access to all your S3 buckets and objects is blocked, turn on Block all public access. These settings apply only to this bucket and its access points. AWS recommends that you turn on Block all public access, but before applying any of these settings, ensure that your applications will work correctly without public access. If you require some level of public access to your buckets or objects within, you can customize the individual settings below to suit your specific storage use cases. Learn more  Account settings for Block Public Access are currently turned on Account settings for Block Public Access that are enabled apply even if they are disabled for this bucket.

iPQ.No	Correct Answer	Points	Explanation
3AWS CCP-CLF53	1	15.4	Answer = 1 Spot Example Explanation : An EC2 instance that is not in use but is offered at a discount from the on-demand fee is known as a spot instance. You can dramatically reduce your Amazon EC2 expenditures by requesting unused EC2 instances from Spot Instances at steep discounts (up to 90%). Spot instances are ideal for flexible, interruptible workloads including batch jobs, background processing, and data analysis. These are unsuitable for essential workloads that must operate at a certain period since they can be abruptly terminated. 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, spiky, or unpredictable workloads that cannot be interrupted - Applications being developed or tested on Amazon EC2 for the first time See On-Demand pricing » Spot instances Amazon EC2 Spot instances allow you to request spare Amazon EC2 computing capacity for up to 90% off the On-Demand price. Learn More. Spot instances are recommended for: - Applications that have flexible start and end times - Applications that are only feasible at very low compute prices - Users with urgent computing needs for large amounts of additional capacity See Spot pricing » Savings Plans Savings Plans are a flexible pricing model that offer low prices on EC2 and Fargate usage, in exchange for a commitment to a consistent amount of usage (measured in \$/hour) for a 1 or 3 year term. Dedicated Hosts A Dedicated Host is a physical EC2 server dedicated for your use. Dedicated Hosts can help you reduce costs by allowing you to use your existing server-bound software licenses, including Windows Server, SQL Server, and SUSE Linux Enterprise Server (subject to your license terms), and can also help you meet compliance requirements. Learn more. - Can be purchased On-Demand (hourly). - Can be purchased as a Reservation for up to 70% off the On-Demand price. See Dedicated pricing » Reserved Instances Reserved Instances provide you with a significant discount (up to 75%) compared to On-Demand instance pricing. In addition, when Reserved Instances are assigned to a specific Availability Zone, they provide a capacity reservation, giving you additional confidence in your ability to launch instances when you need them. For applications that have steady state or predictable usage, Reserved Instances can provide significant savings compared to using On-Demand instances. See How to Purchase Reserved Instances for more information. Reserved Instances are recommended for: - Applications with steady state usage - Applications that may require reserved capacity - Customers that can commit to using EC2 over a 1 or 3 year term to reduce their total computing costs https://aws.amazon.com/ec2/pricing

iPQ.No	Correct Answer	Points	Explanation
3AWSCCP-CLF54	1	15.4	Answer = 1 Transit Gateway by AWS Explanation : Through a centralized hub, the AWS Transit Gateway connects on-premises networks and Amazon Virtual Private Clouds (Amazon 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. Without AWS Transit Gateway  Complexity increases with scale. You must maintain routing tables within each VPC and connect to each on-site location using separate network gateways. With AWS Transit Gateway  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.

iPQ.No	Correct Answer	Points	Explanation
3AWSCCP-CLF55	3	15.4	Answer = 3 Both A and B Explanation : The majority of AWS's services are region-specific. However, few services are required to have a worldwide scope by their very nature because of the underlying service they provide. Some of the international services include AWS Web Application Firewall (AWS WAF), Amazon CloudFront, Amazon Route 53, and AWS Identity and Access Management (AWS IAM). Lambda on AWS You can run code using AWS Lambda's computing service without setting up or maintaining servers. AWS Lambda effortlessly scales from a few requests per day to thousands per second and only runs your code when it is required. An area service is AWS Lambda. Rekognition by Amazon You can recognize objects, people, text, scenes, and activities in pictures and videos with Amazon Rekognition, and you can also screen out any offensive material. You can detect, analyze, and compare faces using Amazon Rekognition's highly accurate facial analysis and facial search capabilities for a variety of user verification, population counting, and public safety use cases. An area-specific service is Amazon Recognition.


iPQ.No	Correct Answer	Points	Explanation
3AWS CCP-CLF56	3	15.4	Answer = 3 Both A and B Explanation : DynamoDB by Amazon 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. Solid-state disks (SSDs) are used to store all of your data, and various availability zones (AZ) in an AWS region automatically replicate it to provide data durability and high availability. through https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/Introduction.html Amazon DynamoDB High Availability High Availability and Durability DynamoDB automatically spreads the data and traffic for your tables over a sufficient number of servers to handle your throughput and storage requirements, while maintaining consistent and fast performance. All of your data is stored on solid-state disks (SSDs) and is automatically replicated across multiple Availability Zones in an AWS Region, providing built-in high availability and data durability. You can use global tables to keep DynamoDB tables in sync across AWS Regions. For more information, see Global Tables: Multi-Region Replication with DynamoDB. (Amazon EFS) Amazon Elastic File System 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. It is designed to grow and shrink automatically as you add and remove files, removing the need to provide and manage capacity to meet expansion. It can scale on-demand to petabytes without interfering with applications. For high availability and durability, Amazon EFS is a local service that saves data within and across several availability zones (AZ). Q. What is Amazon Elastic File System?


iPQ.No	Correct Answer	Points	Explanation
3AWS CCP-CLF57	4	15.4	Answer = 4 Config AWS Explanation : The configuration of AWS resources in your AWS account is shown in great detail by AWS Config. You can examine how the configurations and relationships change over time by looking at how the resources are tied to one another and how they were previously configured. The following scenarios are where AWS Config is intended to assist you in managing your application resources: resource administration, auditing and compliance, managing and troubleshooting configuration changes, and security analysis. Visit https://aws.amazon.com/config/ to learn how AWS Config functions.  The diagram illustrates the AWS Config workflow. It starts with a gear icon representing a 'Configuration change occurs in your AWS resources.' An arrow points to a central box labeled 'AWS Config' which contains the text 'AWS Config records and normalizes the changes into a consistent format.' Below this, another box states 'AWS Config automatically evaluates the recorded configurations against the configurations you specify.' An arrow points from the central box to a stack of three service icons: 'AWS Config APIs & Console', 'Amazon SNS', and 'Amazon CloudWatch'. Below these icons, a box mentions 'Amazon S3' and provides details: 'Access change history and compliance results using the console or APIs. CloudWatch Events or SNS alert you when changes occur. Deliver change history and snapshot files to your S3 bucket for analysis.'

iPQ.No	Correct Answer	Points	Explanation																																																																																								
3AWS CCP-CLF58	1	15.4	Answer = 1 Flexible Glacier Retrieval on Amazon S3 Explanation : The best storage class for archiving data that does not need immediate access but needs the flexibility to retrieve substantial amounts of data for free, such as in backup or disaster recovery use cases, is Amazon S3 Glacier Flexible Retrieval (previously S3 Glacier). It is a class of Amazon S3 cloud storage that is safe, dependable, and incredibly affordable for data preservation and long-term backup. It is intended to meet even the most demanding regulatory requirements with its 99.999999999% durability rating and extensive security and compliance features. The use-cases for S3 Glacier Flexible Retrieval can be further reviewed at https://aws.amazon.com/glacier/. MEDIA ASSET WORKFLOWS Media assets such as video and news footage require durable storage and can grow to many petabytes over time. The Amazon S3 Glacier and S3 Glacier Deep Archive storage classes allow you to archive older media content affordably then move it to Amazon S3 for distribution when needed. HEALTHCARE INFORMATION ARCHIVING Hospital systems need to retain petabytes of patient records (LIS, PACS, EHR, etc.) for decades to meet regulatory requirements. The Amazon S3 Glacier and S3 Glacier Deep Archive storage classes help you reliably archive patient record data securely at a very low cost. REGULATORY AND COMPLIANCE ARCHIVING Many enterprises like Financial Services and Healthcare must retain regulatory and compliance archives for extended durations. Amazon S3 Object Lock helps you set compliance controls to meet your objectives, such as SEC Rule 17a-4(f). SCIENTIFIC DATA STORAGE Research organizations generate, analyze, and archive vast amounts of data. With the Amazon S3 Glacier and S3 Glacier Deep Archive storage classes, you avoid the complexities of hardware and facility management and capacity planning. DIGITAL PRESERVATION Libraries and government agencies face data-integrity challenges in their digital preservation efforts. Unlike traditional systems, which can require laborious data verification and manual repair, Amazon S3 performs regular, systematic data integrity checks and is built to be automatically self-healing. MAGNETIC TAPE REPLACEMENT On-premises or offsite tape libraries can lower storage costs but require large upfront investments and specialized maintenance. The Amazon S3 Glacier and S3 Glacier Deep Archive storage classes have no upfront cost and eliminate the cost and burden of maintenance. through https://aws.amazon.com/s3/storage-classes/ S3 Storage Classes Overview 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>		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																																																																																				

iPQ.No	Correct Answer	Points	Explanation																																																																																								
3AWS CCP-CLF59	3	15.4	Answer = 3 S3 One Zone-IA (Infrequent Access) on Amazon Explanation : Data that is viewed less frequently but needs quick access when needed is stored in Amazon S3 One Zone-Infrequent Access (S3 One Zone-IA). Amazon S3 One Zone-Infrequent Access (S3 One Zone-IA) stores data in a single availability zone (AZ), in contrast to other Amazon S3 storage classes, which store data in a minimum of three AZs. For availability of the Amazon S3 storage class, please check this illustration. You only need to keep in mind that Amazon S3 One Zone-Infrequent Access (S3 One Zone-IA) has the lowest availability; more information is available at https://aws.amazon.com/s3/storage-classes/. 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%</td><td>99.9%</td></tr><tr><td>Availability Zones</td><td>≥3</td><td>≥3</td><td>≥3</td><td>1</td><td>≥3</td><td>≥3</td><td>≥3</td></tr><tr><td>Minimum capacity charge per object</td><td>N/A</td><td>N/A</td><td>128 KB</td><td>128 KB</td><td>128 KB</td><td>40 KB</td><td>40 KB</td></tr><tr><td>Minimum storage duration charge</td><td>N/A</td><td>N/A</td><td>30 days</td><td>30 days</td><td>90 days</td><td>90 days</td><td>180 days</td></tr><tr><td>Retrieval charge</td><td>N/A</td><td>N/A</td><td>per GB retrieved</td><td>per GB retrieved</td><td>per GB retrieved</td><td>per GB retrieved</td><td>per GB retrieved</td></tr><tr><td>First byte latency</td><td>milliseconds</td><td>milliseconds</td><td>milliseconds</td><td>milliseconds</td><td>milliseconds</td><td>minutes or hours</td><td>hours</td></tr><tr><td>Storage type</td><td>Object</td><td>Object</td><td>Object</td><td>Object</td><td>Object</td><td>Object</td><td>Object</td></tr><tr><td>Lifecycle transitions</td><td>Yes</td><td>Yes</td><td>Yes</td><td>Yes</td><td>Yes</td><td>Yes</td><td>Yes</td></tr></tbody></table>		S3 Standard	S3 Intelligent-Tiering*	S3 Standard-IA	S3 One Zone-IA†	S3 Glacier Instant Retrieval	S3 Glacier Flexible Retrieval	S3 Glacier Deep Archive	Designed for durability	99.999999999% (11 9's)	99.999999999% (11 9's)	99.999999999% (11 9's)	99.999999999% (11 9's)	99.999999999% (11 9's)	99.999999999% (11 9's)	99.999999999% (11 9's)	Designed for availability	99.99%	99.9%	99.9%	99.5%	99.9%	99.99%	99.99%	Availability SLA	99.9%	99%	99%	99%	99%	99%	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%	99.9%																																																																																				
Availability Zones	≥3	≥3	≥3	1	≥3	≥3	≥3																																																																																				
Minimum capacity charge per object	N/A	N/A	128 KB	128 KB	128 KB	40 KB	40 KB																																																																																				
Minimum storage duration charge	N/A	N/A	30 days	30 days	90 days	90 days	180 days																																																																																				
Retrieval charge	N/A	N/A	per GB retrieved	per GB retrieved	per GB retrieved	per GB retrieved	per GB retrieved																																																																																				
First byte latency	milliseconds	milliseconds	milliseconds	milliseconds	milliseconds	minutes or hours	hours																																																																																				
Storage type	Object	Object	Object	Object	Object	Object	Object																																																																																				
Lifecycle transitions	Yes	Yes	Yes	Yes	Yes	Yes	Yes																																																																																				

iPQ.No	Correct Answer	Points	Explanation
3AWS CCP-CLF60	3	15.4	Answer = 3 Glacier Deep Archive on Amazon S3 Explanation : The cheapest storage class for Amazon S3, Glacier Deep Archive, offers long-term retention and digital preservation for material that might only be accessed occasionally (once or twice a year). It is made for users who need to save data sets for at least seven to ten years to comply with regulatory compliance requirements, especially those in highly regulated fields like finance, healthcare, and government. Backup and recovery from disasters are two other uses for Amazon S3 Glacier Deep Archive. Its first byte latency ranges from 12 to 48 hours during retrieval. via: https://aws.amazon.com/s3/storage-classes/ Amazon S3 Glacier Deep Archive Overview Amazon S3 Glacier Deep Archive (S3 Glacier Deep Archive) S3 Glacier Deep Archive is Amazon S3's lowest-cost storage class and supports long-term retention and digital preservation for data that may be accessed once or twice in a year. It is designed for customers — particularly those in highly-regulated industries, such as the Financial Services, Healthcare, and Public Sectors — that retain data sets for 7-10 years or longer to meet regulatory compliance requirements. S3 Glacier Deep Archive can also be used for backup and disaster recovery use cases, and is a cost-effective and easy-to-manage alternative to magnetic tape systems, whether they are on-premises libraries or off-premises services. S3 Glacier Deep Archive complements Amazon S3 Glacier, which is ideal for archives where data is regularly retrieved and some of the data may be needed in minutes. All objects stored in S3 Glacier Deep Archive are replicated and stored across at least three geographically-dispersed Availability Zones, protected by 99.999999999% of durability, and can be restored within 12 hours. Key Features: • Designed for durability of 99.999999999% of objects across multiple Availability Zones • Lowest cost storage class designed for long-term retention of data that will be retained for 7-10 years • Ideal alternative to magnetic tape libraries • Retrieval time within 12 hours • S3 PUT API for direct uploads to S3 Glacier Deep Archive, and S3 Lifecycle management for automatic migration of objects
3AWS CCP-CLF61	2	15.2	Answer = 2 Snowball AWS Explanation : AWS Snowball is an edge computing and data migration device that is a member of the AWS Snow family. If you need to migrate a lot of data into AWS, offline data transfer with AWS Snowball can help you get around the bandwidth problem without having to pay for more bandwidth. Terabytes of data are moved using AWS Snowball in about a week. It can be used to move items like databases, backups, archives, medical records, analytics datasets, IoT sensor data, and video content, particularly when network conditions make it impossible to realistically estimate the time it will take to transfer significant volumes of data into and out of AWS.

iPQ.No	Correct Answer	Points	Explanation
3AWSCCP-CLF62	4	15.2	Answer = 4 AWS Guard Duty Explanation : To safeguard your AWS account, Amazon GuardDuty is a threat detection service that keeps an eye out for illegal and hostile behavior. AWS CloudTrail (AWS user and API activity in your accounts), Amazon VPC Flow Logs (data on network traffic), and DNS Logs (name query patterns) are just a few of the sources of events that Amazon GuardDuty monitors across all of your AWS accounts. For a period of 90 days, security findings are kept and made accessible via the Amazon GuardDuty dashboard and APIs. The results are discarded 90 days later. You can configure AWS CloudWatch Events to automatically push findings to an Amazon S3 bucket in your account or another data storage for long-term retention if you want to keep them for longer than 90 days. Using https://aws.amazon.com/guardduty/, find out how Amazon GuardDuty functions.  The diagram illustrates the Amazon GuardDuty workflow in five steps: 1. Amazon GuardDuty: A threat detection service that continuously monitors for malicious or unauthorized behavior to protect your AWS accounts and workloads. 2. Enable GuardDuty: With a few clicks in the console, monitor all your AWS accounts without additional security software or infrastructure to deploy or manage. 3. Continuously analyze: Automatically analyze network and account activity at scale, providing broad, continuous monitoring of your AWS accounts. This step is supported by CloudTrail Logs, VPC Flow Logs, and DNS Logs. 4. Intelligently detect threats: GuardDuty combines managed rule-sets, threat intelligence from AWS Security and 3rd party intelligence partners, anomaly detection, and ML to intelligently detect malicious or unauthorized behavior. 5. Take action: Review detailed findings in the console, integrate into event management or workflow systems, or trigger AWS Lambda for automated remediation or prevention.

iPQ.No	Correct Answer	Points	Explanation
3AWS CCP-CLF63	3	15.2	Answer = 3 Store instances Explanation : Your Amazon EC2 instance can access temporary block-level storage through an instance store. Disks that are physically connected to the host computer house this storage. Instance storage is perfect for replicating data across a fleet of instances, such as a load-balanced pool of web servers, or for the temporary storage of information that changes often, such as buffers, caches, scratch data, and other transient material. Data is lost if an instance fails or is terminated since instance storage is only transitory. Overview of the instance store may be found at Amazon EC2 Instance Store PDF Kindle RSS An <i>instance store</i> provides temporary block-level storage for your instance. This storage is located on disks that are physically attached to the host computer. <u>Instance store is ideal for temporary storage of information that changes frequently, such as buffers, caches, scratch data, and other temporary content, or for data that is replicated across a fleet of instances, such as a load-balanced pool of web servers.</u> An instance store consists of one or more instance store volumes exposed as block devices. The size of an instance store as well as the number of devices available varies by instance type. The virtual devices for instance store volumes are ephemeral[0-23]. Instance types that support one instance store volume have ephemeral0. Instance types that support two instance store volumes have ephemeral0 and ephemeral1, and so on.  https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/InstanceStorage.html.

iPQ.No	Correct Answer	Points	Explanation
3AWS CCP-CLF64	2	15.2	Answer = 2 S3 versioning on Amazon Explanation : Multiple versions of the same object can be kept in the same bucket by using versioning. Versioning enables you to save, retrieve, and restore any previous iterations of any object that is kept in your Amazon S3 bucket. Versioning makes it simple to recover from errors in applications as well as from unintentional user actions. You can recover objects from inadvertent deletion or overwriting using versioning-enabled buckets. For instance, when you remove an item, Amazon S3 inserts a delete marker, which becomes the current object version, as opposed to completely destroying the object. via https://docs.aws.amazon.com/AmazonS3/latest/dev/ObjectVersioning.html Amazon S3 Versioning Overview Using versioning PDF Kindle RSS Versioning is a means of keeping multiple variants of an object in the same bucket. You can use versioning to preserve, retrieve, and restore every version of every object stored in your Amazon S3 bucket. With versioning, you can easily recover from both unintended user actions and application failures. When you enable versioning for a bucket, if Amazon S3 receives multiple write requests for the same object simultaneously, it stores all of the objects. If you enable versioning for a bucket, Amazon S3 automatically generates a unique version ID for the object being stored. In one bucket, for example, you can have two objects with the same key, but different version IDs, such as photo.gif (version 111111) and photo.gif (version 121212).  Versioning-enabled buckets enable you to recover objects from accidental deletion or overwrite. For example: • If you delete an object, instead of removing it permanently, Amazon S3 inserts a delete marker, which becomes the current object version. You can always restore the previous version. For more information, see Deleting object versions. • If you overwrite an object, it results in a new object version in the bucket. You can always restore the previous version.

iPQ.No	Correct Answer	Points	Explanation
3AWS CCP-CLF65	2	15.2	Answer = 2 SaaS (software as a service) Explanation : You can get a whole product through software as a service (SaaS), which is run and controlled by the service provider. You don't have to worry about how the service is handled or the underlying infrastructure is maintained while using a Software as a Service (SaaS) solution. You only need to consider how you plan to use that specific software. One example of software as a service (SaaS) is Gmail. Overview of the Different Types of Cloud Computing: Cloud Computing Models There are three main models for cloud computing. Each model represents a different part of the cloud computing stack.  Infrastructure as a Service (IaaS) Infrastructure as a Service, sometimes abbreviated as IaaS, contains the basic building blocks for cloud IT and typically provide access to networking features, computers (virtual or on dedicated hardware), and data storage space. Infrastructure as a Service provides you with the highest level of flexibility and management control over your IT resources and is most similar to existing IT resources that many IT departments and developers are familiar with today.  Platform as a Service (PaaS) Platforms as a service remove the need for organizations to manage the underlying infrastructure (usually hardware and operating systems) and allow you to focus on the deployment and management of your applications. This helps you be more efficient as you don't need to worry about resource procurement, capacity planning, software maintenance, patching, or any of the other undifferentiated heavy lifting involved in running your application.  Software as a Service (SaaS) Software as a Service provides you with a completed product that is run and managed by the service provider. In most cases, people referring to Software as a Service are referring to end-user applications. With a SaaS offering you do not have to think about how the service is maintained or how the underlying infrastructure is managed; you only need to think about how you will use that particular piece of software. A common example of a SaaS application is web-based email where you can send and receive email without having to manage feature additions to the email product or maintaining the servers and operating systems that the email program is running on. https://aws.amazon.com/cloud-computing-types/