


QUESTIONNAIRE REPORT SHEET


iPQ6AWS CCP.V1.1


iPQ.No	Correct Answer	Points	Explanation
6AWS CCP01	3	15.4	Answer = 3 CloudTrail Insights via AWS Explanation : By continuously examining CloudTrail management events, AWS CloudTrail Insights enables AWS users to recognize and take action in response to suspicious behaviour connected with write API calls. When AWS CloudTrail notices unexpected write management API activity in your account, insights events are logged. When CloudTrail identifies odd activity while Insights is enabled, events are transmitted to the trail's destination S3 bucket. When you see Insights events on the CloudTrail console, you can additionally see the type of insight and the incident time period. Insights events, which differ considerably from other types of events recorded in a CloudTrail trail, are only recorded when CloudTrail notices changes in your account's API usage that are outside of normal usage patterns. By triggering Insights events, AWS CloudTrail Insights can assist you in identifying unexpected API activity in your AWS account. When the volume of API calls deviates from your typical patterns, which CloudTrail Insights refers to as the baseline, it generates Insights events. AWS CloudTrail Insights employs mathematical algorithms to calculate the average amounts of API and service event activity for a given account while continuously monitoring CloudTrail write management events. When behavior deviates from expected patterns, CloudTrail Insights generates insights events and delivers those events to a folder named "/CloudTrail-Insight" in the S3 bucket you've selected as the trail's destination. Insights events for CloudTrail can also be accessed and seen in the AWS Management Console. Utilize AWS CloudTrail Insights to Detect and Address Unusual API Activity, available at: Enabling AWS CloudTrail Insights CloudTrail tracks user activity and API usage. It provides an event history of AWS account activity, including actions taken through the AWS Management Console, AWS SDKs, command line tools, and other AWS services. With the launch of AWS CloudTrail Insights, you can enable machine learning models that detect unusual activity in these logs with just a few clicks. AWS CloudTrail Insights will analyze historical API calls, identifying usage patterns and generating Insight Events for unusual activity.  https://aws.amazon.com/blogs/aws/announcing-cloudtrail-insights-detect-and-address-unusual-api-activity/

iPQ.No	Correct Answer	Points	Explanation
6AWS CCP02	4	15.4	Answer = 4 Simple Queue Service (SQS) by Amazon Explanation : You may decouple and grow serverless applications, distributed systems, and microservices with the help of the fully managed message queuing service known as Amazon Simple Queue Service (SQS). SQS frees developers to concentrate on creating work that stands out from the competition by removing the complexity and overhead involved in managing and operating message-oriented middleware. With SQS, you may communicate between software components at any volume without worrying about message loss or for additional services to be available. Use the AWS console, your preferred command line interface or SDK, and three straightforward commands to get started with SQS in a matter of minutes. The messages in the queue are available until a registered process pulls them to process them in Amazon SQS, which uses a pull mechanism. Since the second application does not have to be present at all times to handle messages from application one, this decouples the design.

iPQ.No	Correct Answer	Points	Explanation
6AWS CCP03	1	15.4	Answer = 1 Application Load Balancing, Amazon CloudFront, Amazon API Gateway, and AWS AppSync Explanation : The AWS Web Application Firewall (AWS WAF) is a web application firewall that aids in defending your web applications or APIs from widespread web exploits that could damage availability, compromise security, or use up excessive resources. By allowing you to build security rules that prohibit typical attack patterns like SQL injection or cross-site scripting and rules that filter out specific traffic patterns you define, AWS WAF offers you control over how traffic to your applications is routed. As part of your CDN solution, you can deploy AWS WAF on Amazon CloudFront, the application load balancer that fronts your EC2 web servers or origin servers for your REST APIs, or AWS AppSync for your GraphQL APIs. The services that AWS customers frequently utilize to distribute content for their websites and apps, including Amazon CloudFront, the Application Load Balancer, Amazon API Gateway, and AWS AppSync, are intimately integrated with AWS WAF. Your rules are applied globally across all AWS Edge Locations that are close to your end customers when you utilize AWS WAF on Amazon CloudFront. This means that performance isn't sacrificed for security. Requests that are blocked are terminated before they reach your web servers. Your rules execute in the region when you use AWS WAF on regional services like Application Load Balancer, Amazon API Gateway, and AWS AppSync, and you may use them to safeguard both internal and internet-facing resources. According to https://aws.amazon.com/waf, how AWS WAF functions


iPQ.No	Correct Answer	Points	Explanation
6AWS CCP04	3	15.4	Answer = 3 to control a workload that displays cyclical load patterns that depend on the day of the week or the hour Explanation : Using previous data from CloudWatch, predictive scaling employs machine learning to forecast capacity requirements. The machine learning system uses the historical data that is currently available, determines the capacity that best fits the historical load pattern, and then iteratively learns from new data to improve projections in the future. Predictive scaling works well in the following circumstances: Cycle traffic, such as heavy resource use during regular business hours and light resource use after hours and on weekends Regular workload patterns that come and go, such as batch processing, testing, or sporadic data analysis During scale-out events, applications that take a long time to initialize have a noticeable latency impact on their performance.

iPQ.No	Correct Answer	Points	Explanation
6AWS CCP05	3	15.4	Answer = 3 Kinesis Data Streams by Amazon Explanation : You can create unique apps that handle or examine streaming data for particular purposes using Amazon Kinesis Data Streams. A constant stream of data from hundreds of thousands of sources, including clickstreams, application logs, and social media, can be added to an Amazon Kinesis data store. Your Amazon Kinesis apps will be able to read and process the data from the stream in a matter of seconds. To quickly move data out of data producers and then process the data constantly, such as to modify the data before sending it to a data store, perform real-time metrics and analytics, or create more complicated data streams for additional processing, Amazon Kinesis Data Streams is helpful. Accelerated log and data feed intake, real-time monitoring and reporting, real-time data analytics, and complicated stream processing are examples of popular use cases for Amazon Kinesis Data Streams. https://aws.amazon.com/kinesis/data-streams/ explains how Amazon Kinesis Data Streams operate.  <pre> graph LR Input[Input Capture and send data to Amazon Kinesis Data Streams] --> KDS[Amazon Kinesis Data Streams Ingests and stores data streams for processing] KDS --> Processing[Processing Services] subgraph Processing [] direction TB KDA[Amazon Kinesis Data Analytics] Spark[Spark on EMR] EC2[Amazon EC2] Lambda[AWS Lambda] Note[Build custom, real-time applications using Amazon Kinesis Data Analytics, stream processing frameworks like Apache Spark, or your code running Amazon EC2 or AWS Lambda] end Processing --> Output[Output Analyze streaming data using your favorite BI tools] </pre>


iPQ.No	Correct Answer	Points	Explanation
6AWS CCP06	3	15.4	Answer = 3 Both A and B Explanation : AWS is in charge of "Security of the Cloud"; specifically, AWS is in charge of safeguarding the infrastructure that powers all of the services provided by the AWS Cloud. The facilities, networking, hardware, and software needed to run AWS cloud services make up this infrastructure. Customers are responsible for patching their guest OS and applications as part of Patch Management, a shared control responsibility of the AWS Shared Responsibility Model. AWS is in charge of patching and addressing infrastructure problems. The customer is in charge of "Security in the Cloud". The AWS cloud services a customer chooses will define the extent of the customer's responsibilities. The quantity of setup work the customer is required to complete as part of their security obligations is determined by this. AWS instructs AWS staff as part of Awareness and Training, a shared control responsibility of the AWS Shared Responsibility Model, but a customer is still responsible for training their own personnel. Amazon Web Services' Shared Responsibility Model is available at https://aws.amazon.com/compliance/shared-responsibility-model/.  The diagram illustrates the AWS Shared Responsibility Model. On the left, a vertical bar is divided into two sections: a blue top section for the CUSTOMER (Responsibility for Security 'IN' the Cloud) and an orange bottom section for AWS (Responsibility for Security 'OF' the Cloud). To the right, a stack of boxes represents the layers of the cloud stack. The top four layers (Customer Data, Platform/Applications/Identity & Access Management, Operating System/Network & Firewall Configuration, and Client-side/Server-side/Network traffic encryption) are blue and fall under the Customer's responsibility. The bottom three layers (Software, Hardware/AWS Global Infrastructure, and Regions/Availability Zones/Edge Locations) are orange and fall under AWS's responsibility. The Software layer is further divided into Compute, Storage, Database, and Networking. The Hardware/AWS Global Infrastructure layer is divided into Regions, Availability Zones, and Edge Locations.


iPQ.No	Correct Answer	Points	Explanation
6AWS CCP07	3	15.4	Answer = 3 Both A and B Explanation : Each group of logical data centers in an AWS region is isolated and physically separate, and each availability zone is referred to as an availability zone (AZ) by AWS. Each AWS region is a physical location where data centers are clustered around the world. AWS customers focused on high availability can design their applications to run in multiple availability zones (AZ) to achieve even greater fault tolerance. AWS infrastructure regions meet the highest levels of security, compliance, and data protection. Each availability zone (AZ) has independent power, cooling, and physical security and is connected via redundant, ultra-low-latency networks. All traffic between availability zones (AZ) in an AWS Region is encrypted. All Availability Zones (AZ) in an AWS Region are interconnected with high-bandwidth, low-latency networking over fully redundant, dedicated metro fiber.
6AWS CCP08	4	15.4	Answer = 4 Security Team Explanation : For your Amazon EC2 instances, a security group serves as a virtual firewall to manage incoming and outgoing traffic. Outbound rules manage the traffic leaving your instance, and inbound rules control the traffic entering it. You have the option to choose one or more security groups when launching an instance. Amazon EC2 utilizes the default security group if you don't provide a security group. Each security group can have rules added to it that permit traffic to or from the instances it is connected to. A security group's rules can be changed at any moment. All instances connected to the security group receive automatic application of new and changed rules. Every rule from every security group connected to an instance is considered by Amazon EC2 when determining whether to let traffic to reach it. AWS and you both share responsibility for security. You must customize security groups, one of AWS's methods for protecting your instances, to fit your security requirements. You can keep your own firewall on any of your instances in addition to employing security groups if your needs can't be fully satisfied by security groups.

iPQ.No	Correct Answer	Points	Explanation
6AWS CCP09	1	15.4	Answer = 1 CodeCommit AWS Explanation : Secure Git-based repositories are hosted by the fully managed source control service AWS CodeCommit. It facilitates team code collaboration in a safe and scalable environment. With CodeCommit, you can stop worrying about maintaining and scaling your own source control system. Source code and binary files may both be safely stored with CodeCommit, and it integrates perfectly with your current Git tools. It is not necessary to host, maintain, back up, and scale up your own source control servers when using AWS CodeCommit. The service automatically scales to match your project's expanding needs. Your files are automatically encrypted in transport and at rest using AWS CodeCommit. Because AWS CodeCommit and AWS Identity and Access Management (AWS IAM) are connected, you may personalize each user's access to your repositories. All Git commands are supported via AWS CodeCommit, which integrates with your current Git tools. With CodeCommit, you can continue to use the graphical clients, continuous integration and delivery systems, and plugins for your favorite development environment.
6AWS CCP10	3	15.4	Answer = 3 Both A and B Explanation : For AWS Elastic Beanstalk, there are no extra fees. To store and execute your application, you establish AWS resources (like EC2 instances or S3 buckets) that you must pay for. There are no minimum payments or up-front commitments; you simply pay for what you use as you use it. The cost of using AWS auto-scaling is free. Only the Amazon CloudWatch monitoring fees and the AWS resources required to execute your applications are your responsibility to pay


iPQ.No	Correct Answer	Points	Explanation
6AWS CCP11	2	15.4	Answer = 2 Config AWS Explanation : You may analyze, audit, and evaluate the configurations of your AWS resources using the service known as AWS Config. Config allows you to automate the comparison of recorded configurations to desired configurations while continuously monitoring and recording your AWS resource configurations. Config allows you to explore in-depth resource configuration histories, examine changes in configurations and interactions between AWS resources, and assess your overall compliance with the configurations outlined in your internal policies. This makes it possible for you to streamline operational troubleshooting, change management, security analysis, and compliance auditing. AWS CloudTrail can assist you determine who made an API call to edit a resource, whereas AWS Config can answer questions like "What did my AWS resource look like at a particular point in time?" An illustration of how AWS Config functions can be seen at https://aws.amazon.com/config/ 


iPQ.No	Correct Answer	Points	Explanation
6AWS CCP12	3	15.4	Answer = 3 Both A and B Explanation : You can easily build and maintain apps in the AWS Cloud with AWS Elastic Beanstalk without needing to understand the technology that powers those applications. The management complexity is lessened by AWS Elastic Beanstalk without limiting choice or control. You only need to upload your application; AWS Elastic Beanstalk will take care of the capacity provisioning, load balancing, scaling, and application health monitoring information automatically. Applications created in Go, Java, .NET, Node.js, PHP, Python, and Ruby are supported by AWS Elastic Beanstalk. AWS Elastic Beanstalk creates the chosen supported platform version and sets up one or more AWS resources, like Amazon EC2 instances, to run your application when you deploy it. For AWS Elastic Beanstalk, there are no extra fees. Only the underlying AWS resources that your application uses are charged for.

iPQ.No	Correct Answer	Points	Explanation
6AWS CCP13	1	15.4	Answer = 1 Wavelength AWS Explanation : An AWS infrastructure product designed for mobile edge computing applications is called AWS Wavelength. Application traffic from 5G devices can reach application servers running in Wavelength Zones without leaving the telecommunications network thanks to AWS infrastructure deployments called Wavelength Zones, which embed AWS compute and storage services within cloud service provider (CSP) data centers at the edge of the 5G network. This prevents the latency that would arise from application traffic having to make numerous Internet hops before reaching its destination, allowing users to fully enjoy the bandwidth and latency advantages provided by current 5G networks. Wavelength can be used to enable low-latency solutions for AWS enterprise clients who develop applications to meet their own use cases, such as IoT, live media creation, and industrial automation. Customers can use Wavelength to perform low-latency operations and processing right where their data is generated, eliminating the need to move sizable amounts of data to be processed in centralized locations. Edge data processing needs include image and video recognition, inference, data aggregation, and responsive analytics. http://aws.amazon.com/wavelength/ describes the operation of AWS Wavelength. How it works Getting started is easy, you simply log-in to the AWS Management Console and enable the Wavelength Zones you want to use for your account. 

iPQ.No	Correct Answer	Points	Explanation
6AWS CCP14	1	15.4	Answer = 1 Kindle Quicksight Explanation : A scalable, embeddable, serverless, machine learning-powered business intelligence (BI) service designed for the cloud is Amazon QuickSight. QuickSight makes it simple to develop and share interactive BI dashboards with machine learning-based insights. QuickSight dashboards are easily integrated into your applications, portals, and websites and accessible from any device. You can easily incorporate interactive dashboards into your applications, websites, and portals with Amazon QuickSight. With the help of QuickSight's comprehensive set of APIs and SDKs, you can quickly alter the dashboards' appearance to better suit your applications. You can control your dashboard versions with Amazon QuickSight, provide dashboard authors access, and share usage statistics with your end users. If customers from several teams or organizations use your application, QuickSight makes sure that their data is always protected and isolated. With its serverless architecture, Amazon QuickSight can scale to tens of thousands of users without your involvement in setting up, configuring, or managing your own servers. Additionally, it guarantees that your users won't experience sluggish dashboards during times when many BI users are attempting to access the same dashboards or datasets. Pay-per-session pricing also makes deployments with lots of users more affordable because you only pay when your customers use the dashboards or reports. Utilizing QuickSight does not require any upfront payments or yearly obligations. Using https://aws.amazon.com/quicksight/, find out how Amazon QuickSight functions.  The diagram illustrates the Amazon QuickSight workflow. It starts with 'Collect and load data' (Clickstreams, sales orders, IoT, financial data, and more), which feeds into 'Data sources' (Seamlessly connect to your data wherever it lives - in the cloud, in 3rd party applications, or on-premises). This data is then processed by 'Amazon QuickSight' (First BI service with Pay-per-Session pricing). The final outputs are 'Interactive dashboards', 'Email reports', and 'Embedded analytics'. Using https://aws.amazon.com/quicksight/, you can connect QuickSight to your data lakes (like Amazon S3).

iPQ.No	Correct Answer	Points	Explanation
6AWS CCP15	1	15.4	Answer = 1 AWS Storage Gateway, Amazon S3 Glacier, and AWS CloudTrail Logs Explanation : All data stored via the AWS Storage Gateway in S3 is by default server-side encrypted using SSE-S3, which is managed by Amazon. Additionally, you may use the Storage Gateway API to configure various gateway types to encrypt stored data using AWS Key Management Service (KMS). The 256-bit Advanced Encryption Standard (AES-256) is automatically server-side encrypted when data is stored in S3 Glacier, and keys are managed by AWS. Before storing data in S3 Glacier, you can alternatively utilize client-side encryption if you'd rather control your own keys. By default, SSE-S3 (Amazon server-side encryption) is used to encrypt the log files that AWS CloudTrail sends to your bucket through S3. Instead, you can utilize server-side encryption with AWS KMS-managed keys (SSE-KMS) for your CloudTrail log files to add a security layer that is directly controlled. You must establish and maintain a KMS key, also known as a customer master key (CMK), in order to use SSE-KMS with CloudTrail.

iPQ.No	Correct Answer	Points	Explanation
6AWS CCP16	3	15.4	Answer = 3 Both A and B Explanation : To build the entrance to a serverless application, Amazon API Gateway can call an AWS Lambda function. REST, HTTP, and WebSocket APIs can be created, published, maintained, monitored, and secured at any scale with Amazon API Gateway, an AWS service. AWS or other web services, as well as data stored in the AWS Cloud, can all be accessed using APIs that are created by API developers. Applications such as workloads operating on Amazon Elastic Compute Cloud (Amazon EC2), code running on AWS Lambda, any web application, or real-time communication applications can access data, business logic, or functionality from your backend services through API Gateway, which functions as a "front door" for applications. Data can be configured to be sent directly to the Amazon Kinesis Data Stream using Amazon API Gateway. In addition to calling HTTP endpoints hosted on AWS Elastic Beanstalk, Amazon EC2, and non-AWS-hosted HTTP-based operations that are reachable over the open Internet, Amazon API Gateway can also establish AWS Step Functions state machines in your account. You can also define a mapping template in API Gateway to create static content that will be returned, which will enable you test your APIs before the backend is prepared. Additionally, you can directly combine API Gateway with other AWS services. As an illustration, you could make an API method available in API Gateway that transmits data to Amazon Kinesis. Using https://aws.amazon.com/api-gateway/, find out how Amazon API Gateway functions. 

iPQ.No	Correct Answer	Points	Explanation
6AWS CCP17	3	15.4	Answer = 3 IoT Core AWS Explanation : You may connect IoT devices to the AWS cloud with AWS IoT Core without having to provision or manage servers. With the ability to manage trillions of messages and billions of devices, AWS IoT Core can consistently and securely process and route messages to AWS endpoints and other devices. Your apps can constantly monitor and communicate with all of your devices, even if they are not connected, thanks to AWS IoT Core. In addition, AWS IoT Core makes it simple to use AWS and Amazon services like AWS Lambda, AWS Kinesis, Amazon S3, Amazon SageMaker, Amazon DynamoDB, AWS CloudWatch, AWS CloudTrail, Amazon QuickSight, and Alexa Voice Service to create IoT applications that gather, process, analyze, and act on data produced by connected devices without having to manage any infrastructure. To connect and manage IoT devices, AWS IoT Core enables you to choose the communication protocol that is most suitable for your use case. AWS IoT Core supports HTTPS (Hypertext Transfer Protocol Secure), MQTT over WSS (WebSockets Secure), LoRaWAN (low-power long-range wide-area network), and MQTT (Message Queuing and Telemetry Transport). In order to prevent data from ever being transmitted between devices and AWS IoT Core without a verified identity, AWS IoT Core offers automated configuration and authentication upon a device's first connection to AWS IoT Core, as well as end-to-end encryption throughout all points of connection. By using policies with precise permissions, you may further secure access to your devices and applications. Capabilities of AWS IoT Core available at https://aws.amazon.com/iot-core/ Publish and subscribe to messages with message broker The Message Broker is a high throughput publish/subscribe (pub/sub) message broker that securely transmits messages to and from all of your IoT devices and applications with low latency. AWS IoT Core supports devices and clients that use the MQTT and the MQTT over WSS protocols to pub/sub to messages, and devices and clients that use the HTTPS protocol to publish messages. 

iPQ.No	Correct Answer	Points	Explanation
6AWS CCP18	3	15.4	Answer = 3 Both A and B Explanation : Follow the general security recommendation to grant only the permissions necessary to complete a task when creating IAM policies. Create regulations that only permit users (and roles) to carry out certain duties after determining what they must do. Give only the bare minimum of permissions at first, and then more as needed. Starting with too lax permissions and attempting to restrict them afterwards is less secure than doing the opposite. Never allow users from one AWS account to access resources in another AWS account by sharing security credentials. Use IAM roles instead. The rights that the IAM users in the other account are granted can be specified in a role that you design. Additionally, you can specify which AWS accounts have IAM users that are authorized to take on the role.

iPQ.No	Correct Answer	Points	Explanation
6AWS CCP19	3	15.4	Answer = 3 Use the AWS-managed master keys that were created in your account automatically for each service Explanation : The main resource in AWS KMS is a KMS key, which may be used to encrypt, decrypt, and re-encrypt data. AWS KMS is replacing the term customer master key (CMK) with AWS KMS key and KMS key. It can also generate data keys that you can utilize outside of AWS KMS. Some AWS services support only an AWS-managed CMK, while others use an AWS-owned CMK or provide you a choice of CMKs. AWS-managed CMKs are CMKs in your account that are produced, managed, and used on your behalf by an AWS service that is connected with the AWS KMS. Only your AWS account can be utilized with AWS-managed CMK. The AWS managed CMKs in your account are visible, together with their essential policies and usage audits in AWS CloudTrail logs. These CMKs cannot be managed, rotated, or had their primary policy altered. Additionally, the service that generates AWS-managed CMKs utilizes them on your behalf; you cannot directly employ them in cryptographic operations. You may find most AWS-managed CMKs by their aliases, which have the format aws/service-name, such as aws/redshift, which are displayed on the AWS-managed keys page of the AWS Management Console for AWS KMS. For CMKs that are maintained by AWS, there is no monthly cost. Some AWS services will pay these costs for you even though they may be subject to fees for use above the free tier.

iPQ.No	Correct Answer	Points	Explanation
6AWS CCP20	3	15.4	Answer = 3 Quantum Ledger Database on Amazon (Amazon QLDB) Explanation : An immutable, transparent, and cryptographically verifiable transaction log is provided by the fully managed ledger database known as the Amazon Quantum Ledger Database (Amazon QLDB), which is controlled by a centrally trusted authority. You can use Amazon QLDB to keep track of all application data changes and to keep a detailed, verifiable history of those changes over time. Ledgers are frequently used to keep track of past financial and economic activities inside a company. For example, tracking the history of credits and debits in banking transactions, confirming the data lineage of an insurance claim, or tracking the movement of an item in a supply chain network are just a few examples of how many organizations build applications with ledger-like functionality. Custom audit tables or audit trails developed in relational databases are frequently used to implement ledger applications. A brand-new category of databases called Amazon QLDB makes it unnecessary to undertake the challenging development process of creating your own ledger-like apps. With QLDB, your data's change history is immutable—it cannot be changed or deleted—and you can quickly confirm using cryptography that there haven't been any unwanted changes to the data used by your application. Every application data update is tracked by QLDB's immutable transactional log, or "journal," which also keeps a complete and verifiable history of modifications over time. Due to its familiar SQL-like API, adaptable document data type, and comprehensive support for transactions, QLDB is simple to use. With the help of QLDB's streaming functionality, you can create event-driven workflows, do real-time analytics, and replicate data to other AWS services in order to facilitate more sophisticated analytical processing. QLDB automatically scales to meet the needs of your application because it is serverless. No servers need to be maintained, and there are no read or write limits to set. You just pay for what you use with QLDB. via https://aws.amazon.com/qldb/ How Amazon Quantum Ledger Database (Amazon QLDB) Operates 

iPQ.No	Correct Answer	Points	Explanation
6AWS CCP21	2	15.4	Answer = 2 Glacier Vault Lock S3 on Amazon Explanation : You may quickly implement and enforce compliance rules for specific S3 Glacier vaults with a vault lock policy using S3 Glacier Vault Lock. In a vault lock policy, you can define controls like "write once, read many" (WORM) and prevent further editing of the policy. The policy cannot be amended after it has been locked. You can lock a vault lock policy to stop alterations in the future, giving your compliance controls robust enforcement. The vault lock policy can be used to implement regulatory and compliance controls, which often ask for stringent restrictions on data access. As an illustration of a vault lock policy, let's say you must keep archives for a year before deleting them. You can achieve this need by setting up a vault lock policy that prevents users from deleting archives until they have been around for a year. This policy can be tested before being locked down. The policy is no longer allowed to be changed after it has been locked.
6AWS CCP22	3	15.4	Answer = 3 Both A and B Explanation : You have the freedom to choose the technological infrastructure and AWS region that are most convenient for your main target market when deploying your apps and workloads to the cloud. You can select the AWS region that is near to the intended location if you have data residency requirements. The location in which your data is physically located remains entirely under your control and ownership, making it simple to satisfy regional compliance and data residency requirements.

iPQ.No	Correct Answer	Points	Explanation
6AWS CCP23	4	15.4	Answer = 4 All of the above Explanation : With the use of AWS CodePipeline, a continuous integration and continuous delivery (CI/CD) service, AWS CodeStar speeds up software deployment. With each commit, an automated pipeline that builds, tests, and deploys your code continually is pre-configured into each project. You can easily change your application code and deploy it to Amazon EC2 and AWS Lambda thanks to AWS CodeStar's integration with AWS CodeDeploy and AWS CloudFormation. AWS CodeStar details can be found at https://aws.amazon.com/codestar/features/. Changes in CodeCommit repositories used as sources for pipelines are discovered by AWS CodePipeline using Amazon CloudWatch Events. There is a corresponding event rule for every source action. In the event that the repository experiences a change, this event rule launches your pipeline. Using https://docs.aws.amazon.com/codepipeline/latest/userguide/integrations-action-type.html, AWS CodePipeline can be integrated with AWS CodeCommit.


iPQ.No	Correct Answer	Points	Explanation
6AWS CCP24	1	15.4	Answer = 1 Management activities Explanation : The record of activity in an AWS account is called an event in AWS CloudTrail. This activity may involve a move made by a user, role, or service that CloudTrail can watch. CloudTrail events offer a history of account activity performed through the AWS Management Console, AWS SDKs, command-line tools, and other AWS services as well as API and non-API account activity. Management events, data events, and AWS CloudTrail Insights events are the three sorts of events that can be recorded in CloudTrail. AWS CloudTrail records all management events by default but excludes data or insights events. Data and insight events carry additional fees. The CloudTrail JSON log format is the same for all event categories. Information regarding management actions taken on resources in your AWS account is provided by management events. Control plane operations is another name for these. Examples include setting up logging, creating rules for data routing, and registering devices.


iPQ.No	Correct Answer	Points	Explanation
6AWS CCP25	3	15.4	Answer = 3 Both A and B Explanation : Adobe document db Because it is an effective and simple data model for developers, data is frequently represented as an object or JSON-like document in application code. By employing the same document model format as their application code, document databases make it simpler for developers to store and query data in a database. Popular document databases like MongoDB and Amazon DocumentDB (which is compatible with MongoDB) offer robust and user-friendly APIs for flexible and iterative development. Kindle Neptune The goal of a graph database is to simplify the development and operation of applications that operate on densely connected datasets. Social networking, recommendation engines, fraud detection, and knowledge graphs are examples of common use cases for a graph database. Fully managed graph database service Amazon Neptune is regarded as a form of NoSQL database.
6AWS CCP26	1	15.4	Answer = 1 Speculate on a website's future traffic patterns Explanation : Machine learning is used by Amazon Forecast, a fully managed service, to produce forecasts that are extremely accurate. Amazon Forecast uses machine learning to mix time series data with extra variables to create forecasts using the same technology as Amazon.com. To get started, Amazon Forecast doesn't require any prior machine learning knowledge. You simply need to supply past data, along with any extra information you think might affect your forecasts. Any time-series data, including retail demand, manufacturing demand, travel demand, revenue, IT capacity, logistics, and web traffic, can be forecast using Amazon Forecast.

iPQ.No	Correct Answer	Points	Explanation
6AWS CCP27	1	15.4	Answer = 1 SaaS (software as a service) Explanation : With software as a service, the service provider gives you a finished product that is operated and managed on your behalf. The majority of the time, when "software as a service" is mentioned, end-user applications are meant. With a SaaS solution, you only need to consider how you will utilize that specific piece of software; you do not need to consider how the service is managed or how the underlying infrastructure is maintained. Web-based email is a typical example of a SaaS application because it allows you to send and receive emails without having to manage feature updates or upkeep the servers and operating systems that the email software is using.


iPQ.No	Correct Answer	Points	Explanation
6AWS CCP28	1	15.4	Answer = 1 Cost Explorer for AWS Explanation : With the use of various filtering dimensions (such as AWS Service, Region, Member Account, etc.), AWS Cost Explorer enables you to examine your AWS costs and consumption both at a high level and in detail. AWS Cost Explorer also enables you to build custom reports from scratch and offers you access to a collection of basic reports to get you started. Using the main graph, the Cost Explorer cost and usage reports, or the Cost Explorer RI report, you can investigate your usage and costs. You may forecast your spending for the upcoming year, check data going back up to the last twelve months, and obtain suggestions for which Reserved Instances to buy. Cost Explorer will help you find areas that require more research and spot trends that can help you better understand your expenditures. The Cost Explorer user interface lets you view your costs and use without paying anything. The Cost Explorer API allows you to programmatically access your data as well. AWS prepares the information about your costs for the current month and the last 12 months when you first sign up for Cost Explorer, and then computes the forecast for the following 12 months. In roughly 24 hours, the data for the current month will be accessible for viewing. It takes a few more days for the rest of your data. Your cost information is updated by Cost Explorer at least once every 24 hours. After you join up, Cost Explorer may show you the current month's prices, the projected costs for the following 12 months, and up to 12 months' worth of historical data (if you have that much).


iPQ.No	Correct Answer	Points	Explanation
6AWS CCP29	1	15.4	Answer = 1 Kindle WorkSpaces Explanation : AWS provides a wide range of worldwide cloud-based services and products, including compute, storage, databases, analytics, networking, machine learning and artificial intelligence, mobile, developer tools, IoT, security, corporate applications, and much more. Some AWS services, such as Amazon Route 53, Amazon Chime, Amazon WorkDocs, Amazon WorkMail, Amazon WorkSpaces, and Amazon WorkLink, are provided internationally rather than locally due to the nature of the service. A controlled, secure desktop-as-a-service (DaaS) option is Amazon WorkSpaces. With the help of Amazon WorkSpaces, you can swiftly scale to give thousands of desktops to workers across the world in either Windows or Linux in just a few short minutes.

iPQ.No	Correct Answer	Points	Explanation
6AWS CCP30	1	15.4	Answer = 1 Kindle Neptune Explanation : Amazon Neptune is a quick, dependable, and fully managed graph database service that makes it simple to create and run programs that operate on densely connected datasets. A purpose-built, high-performance graph database engine that is millisecond latency tuned for storing billions of relationships lies at the heart of Amazon Neptune. For use cases including knowledge graphs, identity graphs, fraud detection, recommendation engines, social networking, life sciences, and more, Amazon Neptune is specifically designed. You may quickly create searches that effectively browse densely connected datasets using Apache TinkerPop Gremlin and SPARQL, which are supported by popular graph models like Property Graph and W3C's RDF. Graph applications like recommendation engines, fraud detection, knowledge graphs, medication discovery, and network security are all powered by Neptune. With read-replicas, point-in-time recovery, continuous backup to Amazon S3, and replication across availability zones, Amazon Neptune is highly available. With support for HTTPS-encrypted client connections and encryption at rest, Neptune is secure. You no longer have to bother about database management activities like hardware provisioning, software patching, setup, configuration, or backups because Neptune is fully handled. Utilize Amazon Neptune to detect fraud: Knowledge Graphs Amazon Neptune helps you build knowledge graph applications. A knowledge graph allows you to store information in a graph model and use graph queries to enable your users to easily navigate highly connected datasets. Neptune supports open source and open standard APIs to allow you to quickly leverage existing information resources to build your knowledge graphs and host them on a fully managed service. For example, if a user is interested in The Mona Lisa, you can also help them discover other works of art by Leonardo da Vinci, or other works of art located in The Louvre. Using a knowledge graph, you can add topical information to product catalogs, build and query complex models of regulatory rules, or model general information, like Wikidata. Learn more about Knowledge Graphs on AWS.  https://aws.amazon.com/neptune


iPQ.No	Correct Answer	Points	Explanation
6AWS CCP31	3	15.4	Answer = 3 Both A and B Explanation : The goal of AWS Control Tower is to offer simple self-service setup and an engaging user interface for continuing governance with guardrails. The AWS Landing Zone solution offers a configurable setup of a landing zone with rich customization options through custom add-ons (like Active Directory, Okta Directory), ongoing modifications through a code deployment and configuration pipeline, and automated creation of new landing zones with pre-configured blueprints (like AWS IAM Identity Center for directory and access). http://aws.amazon.com/controltower/ describes how AWS Control Tower operates.  The diagram illustrates the AWS Control Tower workflow. It starts with the AWS Control Tower icon, which is described as 'The easiest way to set up and govern a secure, compliant multi-account AWS environment'. This leads to a step 'Automate setup of your landing zone based on best-practice blueprints', represented by a grid icon. This is followed by 'Apply guardrails for ongoing governance over your AWS workloads', represented by a shield icon. Next is 'Automate your account provisioning workflow with an account factory', represented by a factory icon. Finally, it leads to 'Get dashboard visibility into your organizational units, accounts, and guardrails', represented by a dashboard icon. Arrows connect these steps in a sequence. You can manage rights in your company by using service control policies (SCPs), a sort of organizational policy. Service control policies (SCPs) provide centralized management of the highest permissible levels for each account in your business. You may make sure your accounts adhere to the access control regulations set out by your company by using service control policies (SCPs). Only organizations with all functionalities enabled have access to SCPs. If your company has just enabled the consolidated billing functionality, SCPs are not accessible.

iPQ.No	Correct Answer	Points	Explanation
6AWS CCP32	2	15.4	Answer = 2 EC2 Instance Savings Plans and Compute Savings Plans Explanation : Savings Plans are a customizable pricing structure that offers savings on your AWS compute use of up to 72%. Regardless of instance family, size, OS, tenancy, or AWS region, this pricing model offers lower charges for using Amazon EC2 instances. It also applies to using AWS Fargate and AWS Lambda. Similar to EC2 Reserved Instances, Savings Plans offer significant savings over On-Demand in exchange for a commitment to use a certain amount of compute power (measured in \$/hour) over the course of one or three years. You can enroll in savings plans for terms of one or three years, and you can manage your plans with ease by utilizing the AWS Cost Explorer's recommendations, performance reports, and budget alerts. Two distinct savings strategies are provided by AWS: The most flexible option is a compute savings plan, which can also lower your expenditures by as much as 66%. These strategies immediately apply to the use of Fargate and Lambda as well as EC2 instances independent of instance family, size, AZ, region, OS, or tenancy. With Compute Savings Plans, for instance, you can switch from C4 to M5 instances at any time and still be charged the Savings Plans rate. You can even migrate a workload from EU (Ireland) to EU (London) or from EC2 to Fargate or Lambda. The cheapest options are EC2 Instance Savings Plans, which promise savings of up to 72% in exchange for a commitment to the use of specific instance families in a location (for example, M5 usage in Northern Virginia). Regardless of AZ, size, OS, or tenancy, this instantly lowers your cost for the chosen instance family in that region. You have the ability to switch between instances within a family in that area thanks to EC2 Instance Savings Plans. Moving from a c5.xlarge running Windows to a c5.2xlarge running Linux, for instance, will instantly qualify you for the Savings Plan pricing. via: https://aws.amazon.com/savingsplans/ How Savings Plans Operate

iPQ.No	Correct Answer	Points	Explanation
6AWS CCP33	4	15.4	Answer = 4 AWS Enterprise On-Ramp, AWS Business Support, and AWS Enterprise Support packages Explanation : You can get real-time advice from an AWS Trusted Advisor to help you deploy your resources according to AWS best practices. The Business and Enterprise Support subscriptions include all AWS Trusted Advisor tests. These tests can aid in streamlining your AWS architecture, boosting security and performance, decreasing overall expenses, and keeping an eye on service caps. With the help of the AWS Support API, you can operateably monitor your Trusted Advisor check requests and status as well as create, manage, and close support issues in the AWS Support Center. The Business, Enterprise On-Ramp, and Enterprise Support subscriptions all include access to the AWS Support API. Utilizing https://aws.amazon.com/premiumsupport/plans/ to compare AWS Support Plans 

iPQ.No	Correct Answer	Points	Explanation
6AWS CCP34	1	15.4	Answer = 1 DynamoDB Accelerator by Amazon Explanation : Scalability and performance are two goals for Amazon DynamoDB. Response times from DynamoDB can typically be calculated in single-digit milliseconds. However, some use cases demand response times measured in microseconds. When accessing eventually consistent data for various use scenarios, Amazon DynamoDB Accelerator (DAX) provides quick response times. With the help of DAX, a DynamoDB-compatible caching service, you can take advantage of quick in-memory performance for demanding applications. Three primary scenarios are covered by DAX: Using DAX as an in-memory cache, eventually consistent read workload response times are decreased by an order of magnitude, from single-digit milliseconds to microseconds. By offering a managed service that is DynamoDB API-compatible, DAX lowers operational and application complexity. Therefore, using it with an existing application just requires minor functional changes. By eliminating the requirement to overprovision read capacity units, DAX increases throughput for workloads that are read-intensive or bursty and may result in operational cost savings. This is particularly useful for programs that repeatedly read individual keys. See https://aws.amazon.com/dynamodb/dax/ for more information on how Amazon DynamoDB Accelerator (DAX) functions.  <pre> graph LR subgraph VPC [Amazon VPC Virtual Private Cloud] subgraph EC2 [Amazon EC2 Instance] App[Your Application] DaxClient[Dax Client] end subgraph DAX [DAX Cluster] C1[CACHE] C2[CACHE] C3[CACHE] end end DaxClient <--> DAX DAX <--> DynamoDB[Amazon DynamoDB] </pre>

iPQ.No	Correct Answer	Points	Explanation
6AWS CCP35	4	15.4	Answer = 4 Resource Groups for AWS Explanation : A resource in AWS is a thing you can interact with. An Amazon S3 bucket, an AWS CloudFormation stack, or an Amazon EC2 instance are a few examples. If you operate with several resources, you might find it more convenient to handle them collectively as opposed to switching between different AWS services for every activity. You probably need to execute bulk actions on a large number of connected resources at once, such as EC2 instances that make up an application layer. Your AWS resources can be organized using AWS Resource Groups. It is simpler to manage and automate tasks on a lot of resources at once when they are organized into resource groups. The account level controls feature permissions for Resource Groups. Users that share your account may access the resource groups you establish as long as they have the appropriate IAM permissions.
6AWS CCP36	4	15.4	Answer = 4 Amazon VPC (Amazon Virtual Private Cloud) Explanation : With the help of the Amazon Virtual Private Cloud (Amazon VPC) service, you can start up AWS services in a custom, logically isolated virtual network. Your virtual networking environment is completely under your control; this includes choosing your own IP address range, setting up subnets, configuring route tables, and setting up network gateways. In order to ensure secure and convenient access to resources and applications, the majority of the resources in your virtual private cloud can be used with both IPv4 and IPv6. The network setup of your VPC can be easily customized using Amazon VPC, one of AWS's core services. Your web servers with internet connection can have a subnet that is open to the public. Additionally, it enables you to install your backend servers, such as database or application servers, on a private subnet that isn't connected to the internet. You can use network access control lists and security groups, among other security measures, to assist manage access to Amazon EC2 instances in each subnet using Amazon VPC.

iPQ.No	Correct Answer	Points	Explanation																																																							
6AWS CCP37	2	15.4	Answer = 2 Snowcone AWS Explanation : The AWS Snowcone, which weighs 4.5 pounds (2.1 kg) and has 8 terabytes of usable storage, is the smallest device in the AWS Snow Family of edge computing, edge storage, and data transmission devices. Snowcone is designed for use outside of a conventional data center and is ruggedized, secure, and safe. Due to its compact size, it is ideal for confined locations or situations where portability is required yet network connectivity is unpredictable. For IoT, vehicle, and drone use cases as well as for first responders, Snowcone may be carried in backpacks. You can run compute applications at the edge, ship the device containing the data to AWS for offline data transfer, or use AWS DataSync to transfer the data online from the edge. Snowcone features numerous encryption and security layers, just like AWS Snowball. Both of these services can be used to gather, process, and transport data to AWS or to perform edge computing tasks. AWS Snowball devices will not fit in locations with space constraints where data migration demands up to 8 terabytes per device are required. Snowcone is built for these situations. According to https://aws.amazon.com/snowcone/, AWS Snowcone functions.  The diagram illustrates the AWS Snowcone workflow: 1. In the AWS Snowball Console, select AWS Snowcone. Create jobs with an Amazon S3 bucket and configure options like Wi-Fi and Amazon EC2 AMIs. 2. AWS prepares and ships the device(s) to you, and you receive your order in approximately 4-6 days. 3. Device arrives. Power it up. Use AWS OpsHub to unlock it. Connect to your network and local machines. 4. Load data with NFS mount and start EC2 instances to process data. Optionally, enable AWS DataSync for online data transfer to S3, Amazon EFS, or Amazon FSx for Windows File Server. 5. When done, shut down and return the device to AWS. The shipping label automatically appears on the e-link screen. 6. When the device arrives at the AWS Region, any stored data is moved to your S3 bucket and verified in about a day. All data is securely erased from the device, and it is sanitized of any customer information. Members of the Snow Family are compared in terms of their features here: Feature comparison <table><tr><th></th><th>AWS Snowcone</th><th>AWS Snowball Edge Storage Optimized</th><th>AWS Snowball Edge Compute Optimized</th><th>AWS Snowmobile</th></tr><tr><td>Usage Scenario</td><td>Edge computing, Data transfer, Edge storage</td><td>Data transfer, Edge storage</td><td>Edge computing, Data transfer</td><td>Data transfer</td></tr><tr><td>Usable HDD Storage</td><td>8 TB</td><td>80 TB</td><td>42 TB</td><td>100 PB</td></tr><tr><td>Usable SSD Storage</td><td>No</td><td>1 TB</td><td>7.68 TB</td><td>No</td></tr><tr><td>Usable vCPUs</td><td>2 vCPUs</td><td>40 vCPUs</td><td>52 vCPUs</td><td>N/A</td></tr><tr><td>Usable Memory</td><td>4 GB</td><td>80 GB</td><td>208 GB</td><td>N/A</td></tr><tr><td>GPU</td><td>No</td><td>No</td><td>nVidia V100 (optional)</td><td>No</td></tr><tr><td>Onboard Computing Options</td><td>AWS IoT Greengrass Amazon EC2 AMIs</td><td>AWS IoT Greengrass Amazon EC2 AMIs</td><td>AWS IoT Greengrass Amazon EC2 AMIs</td><td>N/A</td></tr><tr><td>DataSync</td><td>Yes</td><td>No</td><td>No</td><td>No</td></tr><tr><td>Transfers via NFS</td><td>Yes</td><td>Yes</td><td>Yes</td><td>Yes</td></tr><tr><td>Transfers via S3 API</td><td>No</td><td>Yes</td><td>Yes</td><td>No</td></tr></table>		AWS Snowcone	AWS Snowball Edge Storage Optimized	AWS Snowball Edge Compute Optimized	AWS Snowmobile	Usage Scenario	Edge computing, Data transfer, Edge storage	Data transfer, Edge storage	Edge computing, Data transfer	Data transfer	Usable HDD Storage	8 TB	80 TB	42 TB	100 PB	Usable SSD Storage	No	1 TB	7.68 TB	No	Usable vCPUs	2 vCPUs	40 vCPUs	52 vCPUs	N/A	Usable Memory	4 GB	80 GB	208 GB	N/A	GPU	No	No	nVidia V100 (optional)	No	Onboard Computing Options	AWS IoT Greengrass Amazon EC2 AMIs	AWS IoT Greengrass Amazon EC2 AMIs	AWS IoT Greengrass Amazon EC2 AMIs	N/A	DataSync	Yes	No	No	No	Transfers via NFS	Yes	Yes	Yes	Yes	Transfers via S3 API	No	Yes	Yes	No
	AWS Snowcone	AWS Snowball Edge Storage Optimized	AWS Snowball Edge Compute Optimized	AWS Snowmobile																																																						
Usage Scenario	Edge computing, Data transfer, Edge storage	Data transfer, Edge storage	Edge computing, Data transfer	Data transfer																																																						
Usable HDD Storage	8 TB	80 TB	42 TB	100 PB																																																						
Usable SSD Storage	No	1 TB	7.68 TB	No																																																						
Usable vCPUs	2 vCPUs	40 vCPUs	52 vCPUs	N/A																																																						
Usable Memory	4 GB	80 GB	208 GB	N/A																																																						
GPU	No	No	nVidia V100 (optional)	No																																																						
Onboard Computing Options	AWS IoT Greengrass Amazon EC2 AMIs	AWS IoT Greengrass Amazon EC2 AMIs	AWS IoT Greengrass Amazon EC2 AMIs	N/A																																																						
DataSync	Yes	No	No	No																																																						
Transfers via NFS	Yes	Yes	Yes	Yes																																																						
Transfers via S3 API	No	Yes	Yes	No																																																						

iPQ.No	Correct Answer	Points	Explanation
6AWS CCP38	4	15.4	Answer = 4 OpsHub AWS Explanation : You may manage your AWS Snowball devices using the graphical user interface known as AWS OpsHub, which enables you to quickly install edge computing workloads and make cloud data migration simple. The whole capability of the Snowball devices is available to you in AWS OpsHub with just a few clicks; you can unlock and configure devices, drag and drop data to devices, launch apps, and keep an eye on device metrics. Customers used REST APIs or entered commands via a command-line interface to control Snowball devices in the past. You can now install and manage even huge fleets of Snowball devices more easily with AWS OpsHub, and all while operating without an internet connection. All of the operations currently offered by the Snowball API are combined into AWS OpsHub, which exposes them as a straightforward graphical user interface. Using this interface, you can deploy edge computing apps on Snow Family devices and move data fast and simply to the AWS Cloud. Through the use of AWS Systems Manager, AWS OpsHub offers a unified view of AWS services that are active on Snow Family hardware. Users with various degrees of technical expertise can manage a sizable number of Snow Family devices with the help of AWS OpsHub. You can manage Amazon EC2 instances, unlock devices, transfer files, and view device metrics with just a few clicks. You download, install, and run the AWS OpsHub application on a client computer, such as a laptop, when your Snow device is delivered to your location. After installation, you can unlock the device, begin administering it, and use any local AWS services that are compatible with it. On your smartphone, AWS OpsHub offers a dashboard that provides important information like storage space and running instances. Additionally, it lists some of the AWS services that are compatible with Snow Family devices. You can start transferring files to the device right now.

iPQ.No	Correct Answer	Points	Explanation
6AWS CCP39	2	15.4	Answer = 2 Findings from Amazon GuardDuty, Amazon VPC flow logs, and AWS CloudTrail logs Explanation : With the use of several data sources like Virtual Private Cloud (VPC) Flow Logs, AWS CloudTrail, and Amazon GuardDuty, Amazon Detective can analyze billions of events and automatically produce an interactive, unified view of your resources, users, and their interactions across time. The AWS shared responsibility model, which entails rules and norms for data protection, is adhered to by Amazon Detective. For any accounts where it has been activated, Amazon Detective will begin processing data from AWS CloudTrail logs, VPC flow logs, and Amazon GuardDuty discoveries. Before you enable Amazon Detective on your accounts, you must have Amazon GuardDuty active on those accounts for at least 48 hours. However, you are not limited to looking at the information from GuardDuty using Detective. Your AWS accounts, EC2 instances, AWS users, roles, and IP addresses can all be summarized, analyzed, and visualized in great detail by Amazon Detective. In order to comprehend security concerns or operational account behavior, this information can be very helpful. According to https://aws.amazon.com/detective/

iPQ.No	Correct Answer	Points	Explanation
6AWS CCP40	2	15.4	Answer = 2 Kindle Lightsail Explanation : For developers, small enterprises, students, and other customers that require a solution to build and host their apps on the cloud, Amazon Lightsail is the quickest way to get started with AWS. Developers may deploy and manage websites and web applications in the cloud using Lightsail's compute, storage, and networking capability. For a cheap, dependable monthly cost, Lightsail includes everything you need to swiftly launch your project, such as virtual machines, containers, databases, CDN, load balancers, DNS management, etc. Preconfigured virtual private server (VPS) plans are available that come with everything you need for simple application deployment and management. Projects that need for a few virtual private servers and users who prefer a straightforward management interface should use Amazon Lightsail. Running websites, web apps, blogs, e-commerce sites, basic software, and more are examples of common use cases for Lightsail. A Lightsail package, also known as a bundle, comes with a virtual server that has a set quantity of RAM and vCPUs, SSD-based storage, and a free data transfer quota. Additionally, Amazon Lightsail plans provide DNS control (3 domain zones per account) and static IP addresses (5 per account). Plans with Lightsail are only charged for when you use them because they are hourly, on-demand fees. Numerous preconfigured, one-click-to-launch operating systems, development stacks, and web applications are available through Amazon Lightsail. These include Linux and Windows OS, WordPress, LAMP, CentOS, and more.


iPQ.No	Correct Answer	Points	Explanation
6AWS CCP41	3	15.4	Answer = 3 S3 Intelligent-Tiering (Amazon S3 Tiering) Explanation : The only cloud storage class that automatically reduces costs by shifting objects across four access tiers when access patterns change is Amazon S3 Intelligent-Tiering (S3 Intelligent-Tiering). With no additional operational overhead, the S3 Intelligent-Tiering storage class automatically moves data to the access tier with the lowest cost. This reduces costs. Objects are stored in four access tiers: two low-latency access tiers that are optimized for frequent and infrequent use, and two optional archive access tiers that are built for asynchronous access and are suited for uncommon access. In order for Amazon S3 Intelligent-Tiering to function, objects are stored in four access tiers: two opt-in archive access tiers created for asynchronous access that are optimized for uncommon use, and two low-latency access tiers designed for frequent and infrequent access. The Frequent Access tier is automatically used to store objects that are uploaded or switched to S3 Intelligent Tiering. When an item hasn't been accessed for 30 days in a row, S3 Intelligent-Tiering monitors access patterns and moves it to the Infrequent Access tier. Once one or both of the archive access tiers have been enabled, Amazon S3 Intelligent-Tiering will automatically move items that haven't been accessed for 90 days in a row to the archive access tier and then, after 180 consecutive days, to the deep archive access tier. The items are returned to the Frequent Access tier if they are retrieved subsequently via S3 Intelligent Tiering. When using the S3 Intelligent-Tiering storage class, there are no retrieval fees, and there are no additional tiering fees when objects are moved across access tiers. For data sets with uncertain access patterns, like those from data lakes, or unknown storage access patterns, such those from new applications, it is the perfect storage class.

iPQ.No	Correct Answer	Points	Explanation
6AWS CCP42	3	15.4	Answer = 3 Both A and B Explanation : As the customer plans, deploys, and pro-actively optimizes their AWS solutions, a designated Technical Account Manager (TAM) is their main point of contact and provides assistance, architectural review, and continuing communication to keep them informed and well-prepared. Your TAM acts as your advocate and adviser, focusing on delivering the appropriate resources to support the success and continuing operational health of your AWS infrastructure. It is the cornerstone of the enterprise support strategy. When supporting hundreds of thousands of AWS customers, AWS Trusted Advisor has accumulated operational best practices. The Trusted Advisor examines your AWS setup and notifies you of prospects for cost savings, increased system performance and availability, or assistance in addressing security flaws. The Basic and Developer support plans only include access to the essential Trusted Advisor checks and advice on how to provision your resources in accordance with best practices to boost performance and enhance security.
6AWS CCP43	1	15.4	Answer = 1 Sumerian on Amazon Explanation : Create and execute 3D, augmented reality (AR), and virtual reality (VR) apps with Amazon Sumerian, a managed service. You can create dynamic and immersive scenes that run in your web browser, on mobile devices, and in AR and VR. Getting started with Sumerian is simple, regardless of your level of technical expertise or experience as a web or mobile developer or 3D developer. Because Sumerian is a web-based application, you can easily make connections in your scenes to already-existing AWS services. You can construct scenes directly from your browser. Amazon Sumerian uses the strength of AWS to develop more intelligent and interesting front-end experiences. Use Amazon Lex to quickly integrate conversational interfaces into scenes, and AWS Amplify to include scenes into a web application. Amazon Sumerian uses the most recent WebGL and WebXR standards to develop immersive experiences that can be accessed instantly via a simple URL and are compatible with the majority of hardware platforms for AR and VR. Create your scene once, then use it elsewhere.


iPQ.No	Correct Answer	Points	Explanation
6AWS CCP44	4	15.4	Answer = 4 Efficiency Performance Pillar Explanation : Utilizing IT and computing resources effectively is the subject of the performance efficiency pillar. The correct resource types and sizes should be chosen based on workload requirements, and making informed judgments to preserve efficiency as company demands change is another important aspect. To create and maintain efficient workloads in the cloud, performance efficiency leverages the following design principles: Democratize modern technologies, Go global in minutes. Consider mechanical sympathy, explore more, and use serverless architectures. Visit https://d1.awsstatic.com/whitepapers/architecture/AWS-Performance-Efficiency-Pillar.pdf for further details on the design tenets of the Performance Efficiency Pillar.

iPQ.No	Correct Answer	Points	Explanation
6AWS CCP45	3	15.4	Answer = 3 Both A and B Explanation : The Basic Support Plan, the AWS Developer Support Plan, the AWS Business Support Plan, the AWS Enterprise-On-Ramp Support Plan, and the AWS Enterprise Support Plan are the five support plans offered by AWS Support. Support for account and invoicing inquiries as well as service quota increases are provided under the Basic plan. The other options provide a selection of technical assistance cases with month-to-month billing and no commitment. These Basic Support Plan capabilities are automatically available round-the-clock to all AWS customers: 1. Individualized responses to inquiries about accounts and billing 2. Support boards 3. Checks for service health 4. Technical papers, documentation, and best practice manuals Retrieved from: https://docs.aws.amazon.com/awssupport/latest/user/aws-support-plans.html Features of AWS Support Plans AWS Support offers five support plans: • Basic • Developer • Business • Enterprise On-Ramp • Enterprise Basic Support offers support for account and billing questions and service quota increases. The other plans offer a number of technical support cases with pay-by-the-month pricing and no long-term contracts. All AWS customers automatically have 24x7 access to these features of Basic Support: • One-on-one responses to account and billing questions • Support forums • Service health checks • Documentation, technical papers, and best practice guides Customers with a Developer Support plan have access to these additional features: • Best practice guidance • Client-side diagnostic tools • Building-block architecture support: guidance on how to use AWS products, features, and services together • Supports an unlimited number of support cases that can be opened by one primary contact, which is the AWS account root user. In addition, customers with a Business, Enterprise On-Ramp, or Enterprise Support plan have access to these features: • Use-case guidance – What AWS products, features, and services to use to best support your specific needs. • AWS Trusted Advisor – A feature of AWS Support, which inspects customer environments and identifies opportunities to save money, close security gaps, and improve system reliability and performance. You can access all Trusted Advisor checks. • The AWS Support API to interact with Support Center and Trusted Advisor. You can use the AWS Support API to automate support case management and Trusted Advisor operations. • Third-party software support – Help with Amazon Elastic Compute Cloud (Amazon EC2) instance operating systems and configuration. Also, help with the performance of the most popular third-party software components on AWS. Third-party software support isn't available for customers on Basic or Developer Support plans. • Supports an unlimited number of AWS Identity and Access Management (IAM) users who can open technical support cases. In addition, customers with an Enterprise On-Ramp or Enterprise Support plan have access to these features: • Application architecture guidance – Consultative guidance on how services fit together to meet your specific use case, workload, or application. • Infrastructure event management – Short-term engagement with AWS Support to get a deep understanding of your use case. After analysis, provide architectural and scaling guidance for an event. • Technical account manager – Work with a technical account manager (TAM) for your specific use cases and applications. • White-glove case routing. • Management business reviews.


iPQ.No	Correct Answer	Points	Explanation
6AWS CCP46	2	15.4	Answer =2 Kindle Quicksight Explanation : Automated reference deployments for important AWS Cloud workloads are known as AWS Partner Solutions. Each partner solution uses AWS best practices for security and availability to launch, set up, and operate the compute, network, storage, and other AWS services necessary to deploy a particular workload on AWS. Accelerators known as partner solutions reduce hundreds of laborious processes to just a few simple steps. They are made to order and have a production-oriented design. Among the partner solutions are: 1. An example deployment architecture 2. 3. JSON or YAML scripts or AWS CloudFormation templates that automate and configure the deployment. a deployment guide with detailed explanations of the architecture and implementation and guidelines for adjusting the deployment For customer relationship management (CRM), workforce optimization (WFO), analytics, unified communications (UC), and other use cases, Partner Solutions also include integrations that expand the cloud-based contact center functionality offered by Amazon Connect with essential services and solutions from APN Partners.


iPQ.No	Correct Answer	Points	Explanation
6AWS CCP47	4	15.4	Answer = 4 AWS Guard Duty Explanation : To safeguard your AWS accounts, workloads, and data stored in Amazon Simple Storage Service (Amazon S3), Amazon GuardDuty is a threat detection service that continuously scans for harmful activities and unlawful behavior. While the cloud makes it easier to gather and aggregate account and network activity, it can take a lot of effort for security teams to regularly scan event log data for potential risks. You now have a wise and affordable choice for ongoing threat detection on AWS with Amazon GuardDuty. Potential threats are identified and prioritized by the service using machine learning, anomaly detection, and integrated threat intelligence. Tens of billions of events are analyzed by GuardDuty from a variety of AWS data sources, including AWS CloudTrail event logs, Amazon VPC flow logs, and DNS logs. GuardDuty requires no deployment or upkeep of software or hardware and can be activated with only a few clicks in the AWS Management Console. You can easily establish continuous monitoring of your AWS accounts, workloads, and data kept in Amazon S3 thanks to Amazon GuardDuty. There is no chance that it will have an impact on the performance or availability of your workloads because it runs entirely independently of your resources. With integrated threat intelligence, anomaly detection, and machine learning, it is fully handled. With little effort, existing event management and workflow systems can be easily integrated with Amazon GuardDuty's informative and useful notifications. There are no up-front charges and you just pay for the events that are examined; no additional software deployment or threat intelligence feed subscriptions are necessary. Using https://aws.amazon.com/guardduty/, find out how Amazon GuardDuty functions.  The diagram illustrates the Amazon GuardDuty workflow. It begins with 'Amazon GuardDuty' (a shield icon) which is described as a threat detection service that continuously monitors for malicious or unauthorized behavior to protect AWS accounts, workloads, and data stored in S3. This leads to 'Enable GuardDuty' (a shield with a checkmark icon), where users are instructed to click a few times in the console to monitor all AWS accounts without additional software. The next step is 'Continuously analyze' (a document icon), where GuardDuty automatically analyzes network, account, and data access activity at scale. This is followed by 'Intelligently detect threats' (a target icon), where GuardDuty uses machine learning, anomaly detection, and integrated threat intelligence to identify and prioritize potential threats. The final step is 'Take action' (a computer monitor icon), where users review detailed findings in the console, integrate them into event management or workflow systems, or trigger AWS Lambda for automated remediation or prevention. The data sources for this process are listed as 'CloudTrail Mgmt Events', 'CloudTrail S3 data events', 'VPC Flow Logs', and 'DNS Logs'.


iPQ.No	Correct Answer	Points	Explanation
6AWS CCP48	2	15.4	Answer = 2 Explanation : Make a call to the AWS abuse team You must get in touch with the AWS Abuse team via the Report Amazon AWS Abuse form or by sending an email to abuse@amazonaws.com if you believe that AWS resources are being utilized in an abusive manner. When one of the following abusive behaviors occurs using AWS resources, the AWS Abuse team can help: spam from AWS-owned IP addresses or AWS resources; port scanning; denial-of-service (DoS) or DDoS from AWS-owned IP addresses; intrusion attempts; hosting objectionable or copyrighted content; and malware distribution. Via: https://aws.amazon.com/premiumsupport/knowledge-center/report-aws-abuse/. A list of actions that constitute abusive behavior The AWS Abuse team can assist you when AWS resources are used to engage in the following types of abusive behavior: • Spam: You are receiving unwanted emails from an AWS-owned IP address, or AWS resources are used to spam websites or forums. • Port scanning: Your logs show that one or more AWS-owned IP addresses are sending packets to multiple ports on your server, and you believe this is an attempt to discover unsecured ports. • Denial-of-service (DoS) attacks: Your logs show that one or more AWS-owned IP addresses are used to flood ports on your resources with packets, and you believe that this is an attempt to overwhelm or crash your server or the software running on your server. • Intrusion attempts: Your logs show that one or more AWS-owned IP addresses are used to attempt to log in to your resources. • Hosting objectionable or copyrighted content: You have evidence that AWS resources are used to host or distribute illegal content or distribute copyrighted content without the consent of the copyright holder. • Distributing malware: You have evidence that AWS resources are used to distribute software that was knowingly created to compromise or cause harm to computers or machines on which it is installed.

iPQ.No	Correct Answer	Points	Explanation
6AWS CCP49	1	15.4	Answer = 1 AppStream 2.0 by Amazon Explanation : A fully managed, non-persistent application and desktop streaming service is Amazon AppStream 2.0. On AppStream 2.0, you can safely distribute your desktop programs to any machine while managing them centrally. Without purchasing, provisioning, and maintaining hardware or infrastructure, you may grow to any number of users anywhere in the world. Because AppStream 2.0 is based on AWS, you gain access to a data center and network architecture intended for the most security-conscious businesses. Your apps run on virtual machines designed for certain use cases, and each streaming session dynamically adapts to network conditions to give each end user a smooth and responsive experience. Users can always access the desktop programs they require. Including Chromebooks, Macs, and PCs, AppStream 2.0 broadcasts your applications from AWS to any computer. Your Active Directory, network, cloud storage, and file sharing are all connected through AppStream 2.0. Users use their current login credentials to access applications, and your current security policies control access. AppStream 2.0 is integrated with your IT solutions via extensive APIs. Visit https://aws.amazon.com/appstream2/ to learn how Amazon AppStream 2.0 functions. 

iPQ.No	Correct Answer	Points	Explanation
6AWS CCP50	4	15.4	Answer = 4 Elasticity Explanation : When most people think of cloud computing, they picture how simple it is to get resources when they're needed. This is but one feature of flexibility. When they no longer require resources, the other component is to contract. Scale up and scale down. Scale out and in. When using the cloud, you want to be able to obtain resources as you need them and release them when you're done using them. Amazon Simple Storage Service (Amazon S3), Amazon Simple Queue Service (Amazon SQS), Amazon Simple Notification Service (Amazon SNS), Amazon Simple Email Service (Amazon SES), Amazon Aurora, etc. are among AWS services that perform this as part of their regular operations. Some, like Amazon Relational Database Service (Amazon RDS), demand vertical scaling. Amazon Aurora Serverless and Amazon Athena also qualify as elastic services, as do Amazon EC2, Amazon ECS, AWS Fargate, Amazon EKS, and Amazon DynamoDB.

iPQ.No	Correct Answer	Points	Explanation
6AWS CCP51	4	15.4	Answer = 4 Outposts AWS Explanation : For a completely consistent hybrid experience, AWS Outposts is a fully managed service that provides the same AWS architecture, AWS services, APIs, and tools to almost any data center, co-location space, or on-premises facility. For workloads requiring local data processing, data residency, low-latency access to on-premises systems, and migration of apps with local system requirements, AWS Outposts is the best option. You can access the complete range of AWS services offered in the area to build, operate, and grow your on-premises applications using well-known AWS services and tools. AWS compute, storage, database, and other services run locally on Outposts. Applications that require local data processing or minimal latency can be supported by AWS Outposts. These programs could need to interact with other on-site systems, produce replies to end-user applications in close to real time, or manage equipment that is located on-site. Workloads for automated industrial operations, real-time patient diagnosis or medical imaging, and content and video streaming are a few examples. Outposts can be used to safely store and process customer data that must be kept on-site or in nations without an AWS region. When sending data to the cloud is expensive and unproductive, you can run data-intensive workloads on Outposts and process data locally for improved control over data analysis, backup, and restoration. According to https://aws.amazon.com/outposts, how AWS Outposts operate  <pre> graph LR A[AWS Outposts Run AWS infrastructure and services on premises] --> B[1. Order From the AWS Console select your Outposts configuration consisting of a mix of compute and storage capacity] B --> C[2. Install AWS personnel deliver Outposts to your site, connect Outposts to power, set up network connectivity to AWS Region and your local networks] C --> D[3. Launch Use standard AWS APIs or Management Console to launch EC2 instances on your Outpost] D --> E[4. Build Build and run apps using native AWS services running on Outposts or in the AWS Region] </pre>


iPQ.No	Correct Answer	Points	Explanation
6AWS CCP52	4	15.4	Answer = 4 EC2 Image Builder for Amazon Explanation : Building, testing, and deploying virtual machine and container images for use on AWS or on-premises is made simpler by Amazon EC2 Image Builder. Updates to virtual machine (VM) and container images can be laborious, resource-intensive, and prone to mistakes. At the moment, clients have teams who create automated scripts to maintain images or they manually update and snapshot VMs. Through the use of a straightforward graphical user interface, integrated automation, and security settings from AWS, Amazon EC2 Image Builder considerably minimizes the effort required to maintain images up-to-date and secure. You don't need to create your own automation pipeline or do any manual processes in order to update an image while using Image Builder. Visit https://aws.amazon.com/image-builder/ to learn how to automate server image creation using Amazon EC2 Image Builder. Image Builder provides a one-stop-shop to automate image management processes. Customers can generate an automated pipeline with an intuitive wizard in the AWS console to produce compliant Linux and Windows Server images for use on AWS and on-premises. When software updates become available, Image Builder automatically produces a new image and distributes it to stipulated AWS regions after running tests on it.  Examples of customize software installed on the image includes: 1/ Applications (build environments, business productivity tools, and databases) 2/ OS Updates 3/ Security patches. Examples of secure image with AWS-provided and/or custom templates includes: 1/ Ensure security patches are applied, 2/ Enforce strong passwords, 3/ Turn on full disk encryption, 4/ Close all non-essential open ports, 5/ Enable software firewall, 6/ Enable logging/audit controls. Examples of test image with AWS-provided test and/or your own test includes: 1/ Test that AMI can boot, 2/ Test that sample application can be run, 3/ Test specific patch has been applied, 5/ Test security policy.

iPQ.No	Correct Answer	Points	Explanation
6AWS CCP53	1	15.4	Answer = 1 Direct Connect AWS Explanation : A dedicated network connection from your premises to AWS may be easily established with the help of AWS Direct Connect, a cloud service solution. Your datacenter, office, or colocation environment can be privately connected to AWS via AWS Direct Connect, which in many circumstances can lower network costs, boost bandwidth throughput, and offer a more reliable network experience than Internet-based connections. Your traffic that is in transit is not encrypted via AWS Direct Connect. You can create a dedicated network connection using AWS Direct Connect between your network and one of the AWS Direct Connect locations. This dedicated connection can be divided into numerous virtual interfaces using industry-standard 802.1q VLANs. This preserves network separation between the public and private environments by allowing you to access both public and private resources over the same connection, such as objects stored in Amazon S3 using public IP address space and Amazon EC2 instances running within an Amazon Virtual Private Cloud (VPC) using private IP space. Virtual interfaces can be changed as necessary to accommodate your changing demands. Using https://aws.amazon.com/directconnect/, find out how AWS Direct Connect functions.  The diagram illustrates the five-step process for setting up AWS Direct Connect: AWS Direct Connect: Decide on an AWS Direct Connect location and connection size. Create your connection request(s) on the AWS management console. Download Letter of Authorization (LoA) from the AWS Management console. Provide LoA to an APN Partner and ask them to establish the connection on your behalf. Configure virtual interfaces to establish network connectivity.


iPQ.No	Correct Answer	Points	Explanation
6AWS CCP54	2	15.4	Answer = 2 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. S3 Glacier Deep Archive is an affordable and simple-to-manage replacement to magnetic tape systems, whether they are on-premises libraries or off-premises services. It may also be utilized for backup and disaster recovery use cases. In addition to Amazon S3 Glacier, which is appropriate for archives where data is routinely accessed and some of the data may be needed in only a few minutes, there is also Amazon S3 Glacier Deep Archive. Each item kept in the S3 Glacier Deep Archive is duplicated, spread over at least three different availability zones, guaranteed to be durable for 99.999999999% of the time, and capable of being restored within 12 hours.

iPQ.No	Correct Answer	Points	Explanation
6AWS CCP55	1	15.4	Answer = 1 (Amazon EC2) Amazon Elastic Compute Cloud Explanation : AWS and the customer share responsibility for security and compliance. With AWS running the host operating system, the virtualization layer, and even the physical security of the premises where the service is housed, the operational burden on the user can be lessened thanks to this shared approach. AWS is in charge of "Security of the Cloud"; specifically, AWS is in charge of safeguarding the infrastructure that powers all of the services provided by the AWS Cloud. The facilities, networking, hardware, and software needed to run AWS cloud services make up this infrastructure. The customer is in charge of "Security in the Cloud". The AWS cloud services a customer chooses will define the extent of the customer's responsibilities. The quantity of setup work the customer is required to complete as part of their security obligations is determined by this. For instance, an Infrastructure as a Service (IaaS) offering like Amazon Elastic Compute Cloud (Amazon EC2) necessitates that the customer handle all security configuration and management tasks. Customers are in charge of managing the guest operating system (including updates and security patches), any application software or utilities they install on the instances, and the configuration of the AWS-provided firewall (known as a security group) on each instance when they deploy an Amazon EC2 instance.


iPQ.No	Correct Answer	Points	Explanation
6AWS CCP56	2	15.4	Answer = 2 Elastic Transcoder by Amazon Explanation : You can convert media files you've saved in Amazon Simple Storage Service (Amazon S3) into media files in the formats needed by consumer playback devices using Amazon Elastic Transcoder. For instance, you can convert sizable, high-quality digital media files into formats that viewers can see on connected televisions, smartphones, tablets, and web browsers. The entire media transcoding process is transparently and automatically managed for you by Amazon Elastic Transcoder. There is no need to manage the transcoding infrastructure in any other ways, such as by scaling hardware, managing software, or tuning performance. You only create a transcoding "job," giving the source media file's location and the desired transcoding method. Additionally, Amazon Elastic Transcoder offers transcoding presets for well-liked output formats, so you don't have to rely on your best judgment as to which parameters are most effective for a given device. Through the service API, AWS SDKs, and the AWS Management Console, all of these functions are accessible.

iPQ.No	Correct Answer	Points	Explanation
6AWS CCP57	4	15.4	Answer = 4 (Amazon EFS) Amazon Elastic File System Explanation : For usage with AWS cloud services and on-premises resources, Amazon Elastic File System (Amazon EFS) offers a straightforward, scalable, fully managed elastic NFS file system. 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. With the help of Amazon EFS, your applications can achieve large levels of aggregate throughput and input/output operations per second (IOPS) with consistently low latencies. Amazon EFS is designed to give massively parallel shared access to thousands of Amazon EC2 instances. From home directories to mission-critical applications, Amazon EFS is well suited to handle a wide range of use cases. Customers can move their current corporate applications to the AWS Cloud using EFS. Big data analytics, web serving and content management, application development and testing, processes for the media and entertainment industry, database backups, and container storage are some further use cases. The regular storage class and the rarely access storage class (EFS IA) are the two storage classes that Amazon EFS offers. EFS IA offers cost-optimized price/performance for files that aren't accessed frequently. Files that are not accessed in accordance with the lifespan policy you select will be automatically and transparently migrated into EFS IA by simply enabling EFS lifespan Management on your file system. through https://aws.amazon.com/efs/. How Amazon EFS Works 


iPQ.No	Correct Answer	Points	Explanation
6AWS CCP58	3	15.4	Answer = 3 Disaster Recovery with CloudEndure Explanation : Using block-level replication of the underlying server, CloudEndure Disaster Recovery, a product of the AWS Marketplace, continually replicates server-hosted databases and applications from any source into AWS. You can use AWS Cloud as a disaster recovery region for an on-premises workload and its environment with CloudEndure Disaster Recovery. If an AWS workload consists solely of EC2-hosted databases and apps (and not RDS), it can also be used for disaster recovery. Disaster Recovery with CloudEndure features include: Continuous replication: Block-level replication of your source machines into a staging region is offered by CloudEndure Disaster Recovery continuously and asynchronously. Due to the fact that current applications are always available to be spun up on AWS in the event of a disaster, this enables you to achieve Recovery Point Objectives (RPOs) of under one second. Low-cost staging area: Data is continuously maintained current in your target AWS region's lightweight staging area. Low-cost resources that are automatically supplied and maintained by CloudEndure Disaster Recovery are found in the staging area. Your disaster recovery total cost of ownership (TCO) will go down dramatically as a result of eliminating the requirement for duplicate resources. Automated machine conversion and orchestration: During a disaster or exercise, CloudEndure Disaster Recovery starts a highly automated machine conversion process and a scalable orchestration engine that swiftly spins up hundreds of machines in your target AWS region concurrently. Recovery time goals (RTOs) of minutes are thus made possible. CloudEndure Disaster Recovery replicates whole machines, including the operating system, system state configuration, system disks, databases, programs, and files, in contrast to application-level solutions. Granular point-in-time recovery enables you to restore applications and IT environments that have been harmed by malicious attacks such as ransomware or unintentional system changes. Instead of launching applications in their most recent state in these circumstances, you can

iPQ.No	Correct Answer	Points	Explanation
6AWS CCP59	3	15.4	Answer = 3 AWS Tool with Good Architecture Explanation : You may assess the state of your workloads and evaluate them against the most recent AWS architectural best practices with the help of the AWS Well-Architected Tool. The AWS Well-Architected Framework, on which the tool is based, was created to assist cloud architects in creating safe, effective, high-performing, resilient, and resilient application infrastructure. Simply specify your workload and respond to a series of questions about operational excellence, security, dependability, performance efficiency, and cost optimization to use this free tool, which is accessible in the AWS Management Console. The AWS Well-Architected Tool then offers a strategy for designing for the cloud while adhering to accepted best practices. You have access to the knowledge and best practices employed by AWS architects anytime you need them thanks to the AWS Well-Architected Tool. Following your responses to a series of inquiries regarding your workload, the tool generates an action plan with detailed instructions on how to create better workloads for the cloud. through https://aws.amazon.com/well-architected-tool/. How the AWS Well-Architected Tool functions.  A. Identify the workload to document. Then answer a series of questions about your architecture. B. Review your answers against the five pillars established by the Well-Architected Framework that are visually depicted via the icons in the column above; a) Operational excellence, b) Security, c) Reliability, d) Performance efficiency, & e) Cost optimization. C. You can 1) get videos and documentation 2) generate a report that summarizes your workload review, & 3) see the results of reviews in a single dashboard.

iPQ.No	Correct Answer	Points	Explanation
6AWS CCP60	3	15.4	Answer = 3 Aurora on Amazon Explanation : A relational database engine called Amazon Aurora combines the ease of use and low cost of open-source databases with the performance and dependability of high-end commercial databases. Similar to how Amazon Aurora PostgreSQL provides up to three times the performance of PostgreSQL, Amazon Aurora MySQL provides up to five times the performance of MySQL without requiring any changes to the majority of MySQL applications. Your Amazon Aurora databases are managed by Amazon RDS, which takes care of laborious chores like provisioning, patching, backup, recovery, failure detection, and repair. For each instance of the Amazon Aurora database you use, you just pay a monthly fee. Both up-front fees and long-term commitments are not necessary. A distributed, fault-tolerant, self-healing storage system in Amazon Aurora auto-scales to 128TB for each database instance. With up to 15 low-latency read replicas, point-in-time recovery, continuous backup to Amazon S3, and replication over three Availability Zones (AZs), it offers great performance and availability. The common "mysqldump" and "mysqlimport" utilities can be used to import and export data between MySQL and Amazon Aurora, respectively. The AWS Management Console can be used to convert an RDS MySQL DB Snapshot to Amazon Aurora utilizing the DB Snapshot migration capability of Amazon RDS. Although the length of time varies on the format and quantity of the data set, migration is usually finished in under an hour for clients.

iPQ.No	Correct Answer	Points	Explanation
6AWS CCP61	2	15.2	Answer = 2 CloudTrail by AWS Explanation : Your AWS account can be audited for governance, compliance, operational efficiency, and risk using the service called AWS CloudTrail. You may record, keep track of, and archive account activity linked to operations throughout your AWS infrastructure using CloudTrail. Your AWS account activity, including actions made using the AWS Management Console, AWS SDKs, command-line tools, and other AWS services, is recorded in CloudTrail. The tracking of resource changes, security analysis, and troubleshooting are all made easier by this event history. You may also utilize AWS CloudTrail to find suspicious behavior in your AWS accounts. These tools make operational analysis and troubleshooting easier. AWS CloudTrail records actions performed on your account and provide visibility into user behavior. Every action that is taken is recorded by CloudTrail, including the requester, the services used, the actions taken, the action parameters, and the response elements supplied by the AWS service. You may track changes made to your AWS resources using this information, which also aids in resolving operational problems. It is simpler to guarantee compliance with internal regulations and legal requirements thanks to CloudTrail. through https://aws.amazon.com/cloudtrail/. How AWS CloudTrail Operates 

iPQ.No	Correct Answer	Points	Explanation
6AWS CCP62	3	15.2	Answer = 3 Both A and B Explanation : AWS Elastic Beanstalk monitors resources in your environment and changes their health status to red if they fail to deploy, are incorrectly configured, or become unavailable. This is in addition to Elastic Load Balancing health checks. These verifications show that: 1. The auto-scaling group for the environment is accessible and has a minimum of one instance. 2. The security group for the environment is accessible and set up to accept incoming traffic on port 80. 3. The appropriate load balancer is pointed at by the environment CNAME, which is present. 4. The Amazon Simple Queue Service (Amazon SQS) queue is polled at least once every three minutes in a worker environment. The AWS Elastic Beanstalk service does not publish any metrics to Amazon CloudWatch for basic health reporting. The resources in your environment provide the CloudWatch metrics that are used to generate the graphs on the Monitoring page of the environment interface.
6AWS CCP63	2	15.2	Answer = 2 Optimized for memory instance types Explanation : For workloads that process huge amounts of data in memory, memory-optimized instances are made to execute quickly. For memory-intensive applications, such as in-memory programs, in-memory databases, in-memory analytics tools, high-performance computing (HPC), scientific computing, and other memory-intensive programs, memory optimized instances provide enormous memory sizes. The newest generation of memory-optimized Amazon EC2 instances, powered by Arm-based AWS Graviton2 processors, is known as R6g instances.

iPQ.No	Correct Answer	Points	Explanation
6AWS CCP64	3	15.2	Answer = 3 Backup and restoration approach Explanation : You must compare the advantages of a shorter recovery time goal (RTO) and recovery point objective (RPO) with the expenses of adopting and maintaining a strategy when choosing your DR approach. For the current use case, the backup and restore technique strikes a reasonable balance between advantages and disadvantages. Of all the AWS disaster recovery choices, this is the most affordable. The university can suffer data loss for a few hours, thus backup and restore is the best strategy for the current use case. This method can be used to address the lack of redundancy for workloads deployed to a single availability zone or to mitigate against a regional disaster by duplicating data to other AWS regions. You must redeploy the infrastructure, configuration, and application code in the recovery region in addition to the data. A comparison of various disaster recovery methods is provided at https://aws.amazon.com/blogs/architecture/disaster-recovery-dr-architecture-on-aws-part-iii-pilot-light-and-warm-standby/.  Figure 1. DR strategies

iPQ.No	Correct Answer	Points	Explanation
6AWS CCP65	2	15.2	Answer = 2 CodeArtifact on AWS Explanation : With the help of AWS CodeArtifact, businesses of any size can quickly and easily save, publish, and share the software packages they use in their software development process. This service is fully managed. Developers can access the most recent versions of software by configuring CodeArtifact to automatically fetch software packages and dependencies from open artifact repositories. It is simple to integrate CodeArtifact into current development workflows because it is compatible with widely used package managers and build technologies including Maven, Gradle, npm, yarn, twine, pip, and NuGet. Development teams frequently use both open-source software and programs created by their own company. IT managers must be able to restrict access to and confirm the security of these software programs. Teams require a method for locating current software that has been given the go-ahead by their IT managers. IT managers use central artifact repository systems to store and exchange packages to address these issues. However, many current solutions demand that teams buy licenses for difficult to deploy, scale, and manage software solutions. Pay-as-you-go artifact repository service AWS CodeArtifact grows to meet company demands. There are no servers to maintain or software to update with CodeArtifact. IT managers may quickly set up central repositories that make it simple for development teams to locate and use the software they require. IT managers can also approve software packages and manage their distribution throughout the company, ensuring that only secure software packages are used by development teams. Using https://aws.amazon.com/codeartifact/, find out how AWS CodeArtifact functions.