

QUESTIONNAIRE REPORT SHEET

iPQ5AWSCCP-CLF.V1.1





AWS CERTIFIED CLOUD PRACTITIONER


CLF-C01 PRACTICE TESTS


iPQ.No	Correct Answer	Points	Explanation
5AWS CCP01	2	15.4	Answer = 2 Network ACL (network Access Control List) Explanation : An optional layer of protection for your virtual private cloud (VPC) is a network access control list (network ACL), which functions as a firewall for regulating traffic into and out of one or more subnets (i.e., it controls traffic at the subnet level). Each rule in an ACL for a network can either allow or refuse traffic and is divided into inbound and outbound rules. Overview of the Network Access Control List (network ACL) can be found at Network ACL basics The following are the basic things that you need to know about network ACLs: • Your VPC automatically comes with a modifiable default network ACL. By default, it allows all inbound and outbound IPv4 traffic and, if applicable, IPv6 traffic. • You can create a custom network ACL and associate it with a subnet. By default, each custom network ACL denies all inbound and outbound traffic until you add rules. • Each subnet in your VPC must be associated with a network ACL. If you don't explicitly associate a subnet with a network ACL, the subnet is automatically associated with the default network ACL. • You can associate a network ACL with multiple subnets. However, a subnet can be associated with only one network ACL at a time. When you associate a network ACL with a subnet, the previous association is removed. • A network ACL contains a numbered list of rules. We evaluate the rules in order, starting with the lowest numbered rule, to determine whether traffic is allowed in or out of any subnet associated with the network ACL. The highest number that you can use for a rule is 32766. We recommend that you start by creating rules in increments (for example, increments of 10 or 100) so that you can insert new rules where you need to later on. • A network ACL has separate inbound and outbound rules, and each rule can either allow or deny traffic. • Network ACLs are stateless, which means that responses to allowed inbound traffic are subject to the rules for outbound traffic (and vice versa). There are quotas (limits) for the number of network ACLs per VPC, and the number of rules per network ACL. For more information, see Amazon VPC quotas. https://docs.aws.amazon.com/vpc/latest/userguide/vpc-network-acls.html.


iPQ.No	Correct Answer	Points	Explanation
5AWS CCP02	2	15.4	Answer = 2 PaaS (Platform as a Service) Explanation : Infrastructure as a Service (IaaS), Platform as a Service (PaaS), and Software as a Service (SaaS) are the three categories into which cloud computing can be generically categorized. Platform as a Service (PaaS) frees you up to concentrate on the deployment and administration of your apps by removing the requirement to maintain underlying infrastructure (often hardware and operating systems). You don't have to be concerned with any of the other undifferentiated heavy lifting associated with running your application, such as resource acquisition, capacity planning, software maintenance, patching, etc. Please take a look at this summary of cloud computing kinds at Cloud Computing Models There are three main models for cloud computing. Each model represents a different part of the cloud computing stack.  Infrastructure as a Service (IaaS) Infrastructure as a Service, sometimes abbreviated as IaaS, contains the basic building blocks for cloud IT and typically provide access to networking features, computers (virtual or on dedicated hardware), and data storage space. Infrastructure as a Service provides you with the highest level of flexibility and management control over your IT resources and is most similar to existing IT resources that many IT departments and developers are familiar with today.  Platform as a Service (PaaS) Platforms as a service remove the need for organizations to manage the underlying infrastructure (usually hardware and operating systems) and allow you to focus on the deployment and management of your applications. This helps you be more efficient as you don't need to worry about resource procurement, capacity planning, software maintenance, patching, or any of the other undifferentiated heavy lifting involved in running your application.  Software as a Service (SaaS) Software as a Service provides you with a completed product that is run and managed by the service provider. In most cases, people referring to Software as a Service are referring to end-user applications. With a SaaS offering you do not have to think about how the service is maintained or how the underlying infrastructure is managed; you only need to think about how you will use that particular piece of software. A common example of a SaaS application is web-based email where you can send and receive email without having to manage feature additions to the email product or maintaining the servers and operating systems that the email program is running on. https://aws.amazon.com/types-of-cloud-computing/.
5AWS CCP03	4	15.4	Answer = 4 Kindle Cognito Explanation : You can quickly and easily add user registration, sign-in, and access control to your online and mobile apps with Amazon Cognito. You can utilize SAML identity solutions, your own identity system, or social identity providers like Facebook, Twitter, or Amazon with Amazon Cognito to authenticate users.

iPQ.No	Correct Answer	Points	Explanation
5AWS CCP04	4	15.4	Answer = 4 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 data center, 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. Using https://aws.amazon.com/directconnect/, find out how AWS Direct Connect functions.  <pre> graph LR A["AWS Direct Connect Decide on an AWS Direct Connect location and connection size"] --> B["Create your connection request(s) on the AWS management console"] B --> C["Download Letter of Authorization (LoA) from the AWS Management console"] C --> D["Provide LoA to an APN Partner and ask them to establish the connection on your behalf"] D --> E["Configure virtual interfaces to establish network connectivity"] </pre>

iPQ.No	Correct Answer	Points	Explanation
5AWS CCP05	3	15.4	Answer = 3 Configuration of the Amazon Elastic Compute Cloud's (Amazon EC2) firewall and networking Explanation : "Security in the cloud" is the responsibility of the customers. It includes setting up the operating system, network, and application firewall. Exam Warning: Please thoroughly review the AWS Shared Responsibility Model because there will be many questions on it on the test (found at: https://aws.amazon.com/compliance/shared-responsibility-model/).  The diagram illustrates the AWS Shared Responsibility Model. On the left, two large chevron-shaped boxes represent the primary responsibilities: a blue box for the CUSTOMER (Responsibility for security 'in' the cloud) and an orange box for AWS (Responsibility for security 'of' the cloud). To the right, a stack of boxes details the layers of responsibility. The top four layers (dark blue) are the Customer's responsibility: CUSTOMER DATA, PLATFORM, APPLICATIONS, IDENTITY & ACCESS MANAGEMENT, OPERATING SYSTEM, NETWORK & FIREWALL CONFIGURATION, and a split layer for CLIENT-SIDE DATA ENCRYPTION & DATA INTEGRITY AUTHENTICATION, SERVER-SIDE ENCRYPTION (FILE SYSTEM AND/OR DATA), and NETWORKING TRAFFIC PROTECTION (ENCRYPTION, INTEGRITY, IDENTITY). The next layer (orange) is SOFTWARE, which includes COMPUTE, STORAGE, DATABASE, and NETWORKING. The bottom two layers (orange) are HARDWARE/AWS GLOBAL INFRASTRUCTURE, which includes REGIONS, AVAILABILITY ZONES, and EDGE LOCATIONS.

iPQ.No	Correct Answer	Points	Explanation
5AWS CCP06	4	15.4	Answer = 4 Systems Manager for AWS Explanation : You may automate processes across your AWS resources and centralize operational data from several AWS services using AWS Systems Manager. You can arrange resources into logical categories, such as apps, various application stack layers, or development and production environments. You can choose a resource group in AWS Systems Manager to observe its most recent API activity, resource configuration changes, related notifications, operational alerts, software inventory, and patch compliance status. Depending on your operational needs, you can also take action on each group of resources. To ensure comprehensive visibility and control over your operations, AWS Systems Manager offers a central location to see and manage your AWS resources. http://aws.amazon.com/systems-manager/ describes the operation of AWS Systems Manager.  The diagram illustrates the AWS Systems Manager workflow. It starts with a box labeled 'AWS Systems Manager' which states 'Systems Manager helps you safely manage and operate your resources at scale'. An arrow points to a 'Group resources' step, which says 'Create groups of resources across different AWS services, such as applications or different layers of an application stack'. Another arrow points to a 'Visualize data' step, which says 'View aggregated operational data by resource group'. A final arrow points to a 'Take action' step, which says 'Respond to insights and automate operational actions across resource groups'. The entire process is enclosed in a light blue rounded rectangle.
5AWS CCP07	3	15.4	Answer = 3 Both A and B Explanation : Incoming application traffic is automatically split among numerous targets, such as Amazon EC2 instances, containers, and IP addresses, through elastic load balancing (ELB). It can manage the fluctuating volume of traffic from your application in a single availability zone (AZ) or across several AZs. The three load balancer types provided by Elastic Load Balancing (ELB) are Application Load Balancer (best suited for HTTP and HTTPS traffic), Network Load Balancer (best suited for TCP traffic), and Classic Load Balancer. All three load balancer types have high availability, automatic scaling, and robust security.


iPQ.No	Correct Answer	Points	Explanation
5AWS CCP08	4	15.4	Answer = 4 AWS CDK, or AWS Cloud Development Kit Explanation : The AWS Cloud Development Kit (AWS CDK) is an open-source platform for creating software that allows you to describe the resources for your cloud applications using well-known programming languages. The AWS Cloud Development Kit (AWS CDK) models your apps using the comfort and expressiveness of programming languages. It gives you access to high-level building blocks known as constructs that preconfigure cloud resources with tested defaults, enabling you to create cloud applications without specialized knowledge. Through AWS CloudFormation, AWS CDK creates your resources in a secure, repeatable manner. Additionally, it enables you to create and distribute your own original bespoke structures that take into account the needs of your business, which will speed up the beginning of new projects. To put it simply, you create AWS CDK projects using the AWS CDK framework, which are then performed to produce AWS CloudFormation templates. AWS CDK's functionality is explained at  https://aws.amazon.com/cdk/.

iPQ.No	Correct Answer	Points	Explanation
5AWS CCP09	4	15.4	Answer = 4 Monitor changes to the configuration Explanation : AWS Config has a capability for tracking configuration changes. 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. Visit https://aws.amazon.com/config/ to learn how AWS Config functions. 

iPQ.No	Correct Answer	Points	Explanation
5AWS CCP10	2	15.4	Answer = 2 CloudFormation by AWS Explanation : Developers and system administrators can easily establish and manage a group of linked AWS resources with AWS CloudFormation, provisioning and updating them in a systematic and predictable way. To specify your AWS resources, any associated dependencies, and any runtime settings necessary to run your application, you can use the AWS CloudFormation sample templates or construct your own templates. As a result, you can standardize the infrastructure components utilized throughout your business, enabling configuration compliance and quicker troubleshooting. This gives you a single source of truth for all of your resources. You may calculate the cost of your resources using AWS CloudFormation templates. Visit https://aws.amazon.com/cloudformation/ to learn more about AWS CloudFormation.  <pre> graph LR A[Code infrastructure Code your infrastructure from scratch with the CloudFormation template language, in either YAML or JSON format, or start from many available sample templates] --> B[Amazon S3 Check out your template code locally, or upload it into an S3 bucket] B --> C[AWS CloudFormation Use AWS CloudFormation via the browser console, command line tools or APIs to create a stack based on your template code] C --> D[Output AWS CloudFormation provisions and configures the stacks and resources you specified on your template] </pre>

iPQ.No	Correct Answer	Points	Explanation
5AWS CCP11	2	15.4	Answer = 2 IAM Identity Center for AWS Explanation : AWS Single Sign-On (AWS SSO) is replaced by AWS IAM Identity Center. For the purpose of streamlining access management to several AWS accounts, AWS applications, and other SAML-enabled cloud services, it is developed on top of AWS Identity and Access Management (IAM). You can create or link your workforce users for use across AWS in IAM Identity Center. You have the option of controlling access to both your cloud applications and AWS accounts separately or jointly. Users can be created in IAM Identity Center natively or imported from an existing employee directory. You may define, adapt, and assign fine-grained access using IAM Identity Center's unified management interface. Your workforce users can access their assigned AWS accounts or cloud applications using a user portal. IAM Identity Center enables you to quickly and easily manage employee access to numerous AWS accounts, SAML-enabled cloud services (including Salesforce, Microsoft 365, and Box), and specially created in-house applications from a single location. Using https://aws.amazon.com/iam/identity-center/, learn how AWS IAM Identity Center functions. 

iPQ.No	Correct Answer	Points	Explanation																																																												
5AWS CCP12	3	15.4	Answer = 3 Business Support by AWS Explanation : If you use AWS for production workloads and need access to technical support and architectural advice on a 24/7 basis through phone, email, and chat for your unique use cases, AWS recommends AWS Business Support. AWS Trusted Advisor Best Practice Checks are fully accessible to you. Additionally, it is the least expensive support package for offering cloud support experts better technical support. Exam Warning: As you can anticipate at least a few questions on the exam, please review the distinctions between the AWS Developer Support, AWS Business Support, and AWS Enterprise Support plans: <table><thead><tr><th></th><th>Developer</th><th>Business</th><th>Enterprise On-Ramp</th><th>Enterprise</th></tr><tr><td></td><td><i>Recommended if you are experimenting or testing in AWS.</i></td><td><i>Minimum recommended tier if you have production workloads in AWS</i></td><td><i>Recommended if you have production and/or business critical workloads in AWS.</i></td><td><i>Recommended if you have business and/or mission critical workloads in AWS.</i></td></tr></thead><tbody><tr><td>AWS Trusted Advisor Best Practice Checks</td><td>Service Quota and basic Security checks</td><td>Full set of checks</td><td>Full set of checks</td><td>Full set of checks</td></tr><tr><td>AWS Trusted Advisor Priority</td><td></td><td></td><td></td><td>Prioritized recommendations curated by your AWS account team</td></tr><tr><td>Enhanced Technical Support</td><td>Business hours** web access to Cloud Support Associates Unlimited cases with 1 primary contact Prioritized responses on AWS re:Post</td><td>24/7 phone, web, and chat access to Cloud Support Engineers Unlimited cases and unlimited contacts (IAM supported) Prioritized responses on AWS re:Post Access to AWS Support App in Slack</td><td>24/7 phone, web, and chat access to Cloud Support Engineers Unlimited cases and unlimited contacts (IAM supported) Prioritized responses on AWS re:Post Access to AWS Support App in Slack</td><td>24/7 phone, web, and chat access to Cloud Support Engineers Unlimited cases and unlimited contacts (IAM supported) Prioritized responses on AWS re:Post Access to AWS Support App in Slack</td></tr><tr><td>Case Severity / Response Times*</td><td>General guidance: < 24 hours** System impaired: < 12 hours**</td><td>General guidance: < 24 hours System impaired: < 12 hours Production system impaired: < 4 hours Production system down: < 1 hour</td><td>General guidance: < 24 hours System impaired: < 12 hours Production system impaired: < 4 hours Production system down: < 1 hour Business-critical system down: < 30 minutes</td><td>General guidance: < 24 hours System impaired: < 12 hours Production system impaired: < 4 hours Production system down: < 1 hour Business/Mission-critical system down: < 15 minutes</td></tr><tr><td>Architectural Guidance</td><td>General</td><td>Contextual to your use-cases</td><td>Consultative review and guidance based on your applications (one-per-year)</td><td>Consultative review and guidance based on your applications</td></tr><tr><td>Programmatic Case Management</td><td></td><td>AWS Support API</td><td>AWS Support API</td><td>AWS Support API</td></tr></tbody></table> <table><thead><tr><th></th><th>Developer</th><th>Business</th><th>Enterprise On-Ramp</th><th>Enterprise</th></tr><tr><td></td><td><i>Recommended if you are experimenting or testing in AWS.</i></td><td><i>Minimum recommended tier if you have production workloads in AWS</i></td><td><i>Recommended if you have production and/or business critical workloads in AWS.</i></td><td><i>Recommended if you have business and/or mission critical workloads in AWS.</i></td></tr></thead><tbody><tr><td>Third-Party Software Support</td><td></td><td>Interoperability and configuration guidance and troubleshooting</td><td>Interoperability and configuration guidance and troubleshooting</td><td>Interoperability and configuration guidance and troubleshooting</td></tr><tr><td></td><td>Access to Support Automation Workflows with prefixes AWSsupport</td><td>Access to Infrastructure Event Management for additional fee</td><td>Infrastructure Event Management (one-per-year)</td><td>Infrastructure Event Management Access to proactive reviews, workshops, and more</td></tr></tbody></table>		Developer	Business	Enterprise On-Ramp	Enterprise		<i>Recommended if you are experimenting or testing in AWS.</i>	<i>Minimum recommended tier if you have production workloads in AWS</i>	<i>Recommended if you have production and/or business critical workloads in AWS.</i>	<i>Recommended if you have business and/or mission critical workloads in AWS.</i>	AWS Trusted Advisor Best Practice Checks	Service Quota and basic Security checks	Full set of checks	Full set of checks	Full set of checks	AWS Trusted Advisor Priority				Prioritized recommendations curated by your AWS account team	Enhanced Technical Support	Business hours** web access to Cloud Support Associates Unlimited cases with 1 primary contact Prioritized responses on AWS re:Post	24/7 phone, web, and chat access to Cloud Support Engineers Unlimited cases and unlimited contacts (IAM supported) Prioritized responses on AWS re:Post Access to AWS Support App in Slack	24/7 phone, web, and chat access to Cloud Support Engineers Unlimited cases and unlimited contacts (IAM supported) Prioritized responses on AWS re:Post Access to AWS Support App in Slack	24/7 phone, web, and chat access to Cloud Support Engineers Unlimited cases and unlimited contacts (IAM supported) Prioritized responses on AWS re:Post Access to AWS Support App in Slack	Case Severity / Response Times*	General guidance: < 24 hours** System impaired: < 12 hours**	General guidance: < 24 hours System impaired: < 12 hours Production system impaired: < 4 hours Production system down: < 1 hour	General guidance: < 24 hours System impaired: < 12 hours Production system impaired: < 4 hours Production system down: < 1 hour Business-critical system down: < 30 minutes	General guidance: < 24 hours System impaired: < 12 hours Production system impaired: < 4 hours Production system down: < 1 hour Business/Mission-critical system down: < 15 minutes	Architectural Guidance	General	Contextual to your use-cases	Consultative review and guidance based on your applications (one-per-year)	Consultative review and guidance based on your applications	Programmatic Case Management		AWS Support API	AWS Support API	AWS Support API		Developer	Business	Enterprise On-Ramp	Enterprise		<i>Recommended if you are experimenting or testing in AWS.</i>	<i>Minimum recommended tier if you have production workloads in AWS</i>	<i>Recommended if you have production and/or business critical workloads in AWS.</i>	<i>Recommended if you have business and/or mission critical workloads in AWS.</i>	Third-Party Software Support		Interoperability and configuration guidance and troubleshooting	Interoperability and configuration guidance and troubleshooting	Interoperability and configuration guidance and troubleshooting		Access to Support Automation Workflows with prefixes AWSsupport	Access to Infrastructure Event Management for additional fee	Infrastructure Event Management (one-per-year)	Infrastructure Event Management Access to proactive reviews, workshops, and more
	Developer	Business	Enterprise On-Ramp	Enterprise																																																											
	<i>Recommended if you are experimenting or testing in AWS.</i>	<i>Minimum recommended tier if you have production workloads in AWS</i>	<i>Recommended if you have production and/or business critical workloads in AWS.</i>	<i>Recommended if you have business and/or mission critical workloads in AWS.</i>																																																											
AWS Trusted Advisor Best Practice Checks	Service Quota and basic Security checks	Full set of checks	Full set of checks	Full set of checks																																																											
AWS Trusted Advisor Priority				Prioritized recommendations curated by your AWS account team																																																											
Enhanced Technical Support	Business hours** web access to Cloud Support Associates Unlimited cases with 1 primary contact Prioritized responses on AWS re:Post	24/7 phone, web, and chat access to Cloud Support Engineers Unlimited cases and unlimited contacts (IAM supported) Prioritized responses on AWS re:Post Access to AWS Support App in Slack	24/7 phone, web, and chat access to Cloud Support Engineers Unlimited cases and unlimited contacts (IAM supported) Prioritized responses on AWS re:Post Access to AWS Support App in Slack	24/7 phone, web, and chat access to Cloud Support Engineers Unlimited cases and unlimited contacts (IAM supported) Prioritized responses on AWS re:Post Access to AWS Support App in Slack																																																											
Case Severity / Response Times*	General guidance: < 24 hours** System impaired: < 12 hours**	General guidance: < 24 hours System impaired: < 12 hours Production system impaired: < 4 hours Production system down: < 1 hour	General guidance: < 24 hours System impaired: < 12 hours Production system impaired: < 4 hours Production system down: < 1 hour Business-critical system down: < 30 minutes	General guidance: < 24 hours System impaired: < 12 hours Production system impaired: < 4 hours Production system down: < 1 hour Business/Mission-critical system down: < 15 minutes																																																											
Architectural Guidance	General	Contextual to your use-cases	Consultative review and guidance based on your applications (one-per-year)	Consultative review and guidance based on your applications																																																											
Programmatic Case Management		AWS Support API	AWS Support API	AWS Support API																																																											
	Developer	Business	Enterprise On-Ramp	Enterprise																																																											
	<i>Recommended if you are experimenting or testing in AWS.</i>	<i>Minimum recommended tier if you have production workloads in AWS</i>	<i>Recommended if you have production and/or business critical workloads in AWS.</i>	<i>Recommended if you have business and/or mission critical workloads in AWS.</i>																																																											
Third-Party Software Support		Interoperability and configuration guidance and troubleshooting	Interoperability and configuration guidance and troubleshooting	Interoperability and configuration guidance and troubleshooting																																																											
	Access to Support Automation Workflows with prefixes AWSsupport	Access to Infrastructure Event Management for additional fee	Infrastructure Event Management (one-per-year)	Infrastructure Event Management Access to proactive reviews, workshops, and more																																																											

iPQ.No	Correct Answer	Points	Explanation
5AWS CCP13	3	15.4	Answer = 3 Both A and B Explanation : A group of Amazon EC2 instances that are automatically scaled and managed as a logical unit is called an auto-scaling group. Its size can be modified to satisfy demand either manually or automatically. When using AWS services, AWS Auto Scaling can help you maximize usage and cost efficiency so that you only pay for the resources you really use. Using https://aws.amazon.com/autoscaling/, find out how AWS Auto Scaling functions.  The diagram illustrates the AWS Auto Scaling workflow. It starts with a box labeled 'AWS Auto Scaling' with the subtitle 'Unified scaling for your cloud applications'. An arrow points to 'Explore your applications', which includes icons for a magnifying glass, a database, and a mail icon. Another arrow points to 'Discover what you can scale', showing a database icon and a scaling icon. This is followed by 'Choose what to optimize', with icons for 'COST' and 'PERFORMANCE'. The final step is 'Track scaling as it happens', represented by a monitor icon showing a line graph. Arrows connect these steps in a sequence, with a feedback loop from the final step back to the initial 'AWS Auto Scaling' box. When deployed in a particular availability zone (AZ), Amazon EC2 Reserved Instances (RI) offer a large discount (up to 72%) compared to on-demand pricing and guarantee capacity. Overview of EC2 Pricing Options: On-Demand With On-Demand instances, you pay for compute capacity by the hour or the second depending on which instances you run. No longer-term commitments or upfront payments are needed. You can increase or decrease your compute capacity depending on the demands of your application and only pay the specified per hourly rates for the instance you use. On-Demand instances are recommended for: • Users that prefer the low cost and flexibility of Amazon EC2 without any up-front payment or long-term commitment • Applications with short-term, spiky, or unpredictable workloads that cannot be interrupted • Applications being developed or tested on Amazon EC2 for the first time See On-Demand pricing » Spot instances Amazon EC2 Spot instances allow you to request spare Amazon EC2 computing capacity for up to 90% off the On-Demand price. Learn More. Spot instances are recommended for: • Applications that have flexible start and end times • Applications that are only feasible at very low compute prices • Users with urgent computing needs for large amounts of additional capacity See Spot pricing » Savings Plans Savings Plans are a flexible pricing model that offer low prices on EC2 and Fargate usage, in exchange for a commitment to a consistent amount of usage (measured in \$/hour) for a 1 or 3 year term. Reserved Instances Reserved Instances provide you with a significant discount (up to 75%) compared to On-Demand instance pricing. In addition, when Reserved Instances are assigned to a specific Availability Zone, they provide a capacity reservation, giving you additional confidence in your ability to launch instances when you need them. For applications that have steady state or predictable usage, Reserved Instances can provide significant savings compared to using On-Demand instances. See How to Purchase Dedicated Hosts A Dedicated Host is a physical EC2 server dedicated for your use. Dedicated Hosts can help


iPQ.No	Correct Answer	Points	Explanation
5AWS CCP14	2	15.4	Answer = 2 Service Health on the AWS Health Dashboard Explanation : The only place to find information regarding the functionality and availability of AWS services is the AWS Health Dashboard (Service Health). You can sign in to read personalized communications regarding your specific AWS account or business, as well as view the general status of AWS services. To find out the current status, visit this page at https://health.aws.amazon.com/health/status. You can subscribe to an RSS feed in the AWS Health Dashboard (Service Health) to receive notifications when a service is interrupted.
5AWS CCP15	4	15.4	Answer = 4 Access Logs for Amazon Simple Storage Service (Amazon S3) Explanation : For each request made to an Amazon S3 bucket, server access logging keeps a thorough record. Many applications can benefit from server access logs. Information from access logs, for instance, may be helpful during security and access audits. Additionally, it may assist you in comprehending your Amazon S3 bill and learning more about your clientele.


iPQ.No	Correct Answer	Points	Explanation
5AWS CCP16	1	15.4	Answer = 1 All availability zones (AZs) in an AWS region are covered by a VPC Explanation : A logically isolated area of the AWS Cloud where you can launch AWS resources in a virtual network that you define can be created using the Amazon Virtual Private Cloud (Amazon VPC). Your virtual networking environment is completely under your control, and this includes choosing your own IP address range, setting up subnets, configuring route tables, and setting up network gateways. All Availability Zones (AZs) in a region are covered by an Amazon VPC.
5AWS CCP17	3	15.4	Answer = 3 Hybrid Cloud Explanation : A hybrid deployment is a means to link current resources that are not in the cloud with existing infrastructure and applications. The most popular hybrid deployment strategy involves extending and growing an organization's on-premises infrastructure into the cloud while linking cloud services to the internal system. through - https://aws.amazon.com/types-of-cloud-computing/ An Overview of Cloud Computing Deployment Models Cloud Computing Deployment Models  Cloud A cloud-based application is fully deployed in the cloud and all parts of the application run in the cloud. Applications in the cloud have either been created in the cloud or have been migrated from an existing infrastructure to take advantage of the benefits of cloud computing. Cloud-based applications can be built on low-level infrastructure pieces or can use higher level services that provide abstraction from the management, architecting, and scaling requirements of core infrastructure.  Hybrid A hybrid deployment is a way to connect infrastructure and applications between cloud-based resources and existing resources that are not located in the cloud. The most common method of hybrid deployment is between the cloud and existing on-premises infrastructure to extend, and grow, an organization's infrastructure into the cloud while connecting cloud resources to internal system. For more information on how AWS can help you with your hybrid deployment, please visit our hybrid page.  On-premises Deploying resources on-premises, using virtualization and resource management tools, is sometimes called "private cloud". On-premises deployment does not provide many of the benefits of cloud computing but is sometimes sought for its ability to provide dedicated resources. In most cases this deployment model is the same as legacy IT infrastructure while using application management and virtualization technologies to try and increase resource utilization.

iPQ.No	Correct Answer	Points	Explanation
5AWS CCP18	4	15.4	Answer = 4 Several Availability Zones (AZs) can be used to bind Amazon EBS Elastic Volumes Explanation : You can attach a robust, block-level storage device called an Amazon EBS Elastic Volume to your instances. You can utilize a volume as you would a physical hard drive after attaching it to an instance. The volume and the instance must be in the same Availability Zone (AZ) when utilizing Amazon EBS Elastic Volumes.
5AWS CCP19	1	15.4	Answer = 1 (AMI) Amazon Machine Image Explanation : The data needed to launch an instance is provided via an Amazon Machine Image (AMI). When launching an instance, you must supply an Amazon Machine Image (AMI). When you require numerous instances with the same configuration, you can launch them all from a single Amazon Machine Image (AMI). Using an Amazon Machine Image (AMI) is explained in this user guide: <pre> graph LR create --> EBS[EBS snapshot (or template if instance store backed)] EBS -- register --> AMI1[AMI #1] AMI1 -- launch --> Instance[Instance] AMI1 -- copy --> AMI2[AMI #2] AMI1 -- deregister --> deregister[deregister] </pre> https://docs.aws.amazon.com/AWSEC2/latest/AMIs.html


iPQ.No	Correct Answer	Points	Explanation
5AWS CCP20	1	15.4	Answer = 1 WSA Glue Explanation : Customers can easily prepare and load their data for analytics with AWS Glue, a fully managed extract, transform, and load (ETL) service. The purpose of an AWS Glue task is batch ETL data processing. How AWS Glue functions may be found at: https://aws.amazon.com/glue/
5AWS CCP21	3	15.4	Answer = 3 Both A and B Explanation : Exam Warning: Please consider the following six benefits of cloud computing: You will undoubtedly be asked about how cloud computing differs from a conventional on-premises configuration. via: Whitepapers.aws.amazon.com/latest/aws-overview/six-advantages-of-cloud-computing.html Six Advantages of Cloud Computing PDF RSS • Trade capital expense for variable expense – Instead of having to invest heavily in data centers and servers before you know how you're going to use them, you can pay only when you consume computing resources, and pay only for how much you consume. • Benefit from massive economies of scale – By using cloud computing, you can achieve a lower variable cost than you can get on your own. Because usage from hundreds of thousands of customers is aggregated in the cloud, providers such as AWS can achieve higher economies of scale, which translates into lower pay as-you-go prices. • Stop guessing capacity – Eliminate guessing on your infrastructure capacity needs. When you make a capacity decision prior to deploying an application, you often end up either sitting on expensive idle resources or dealing with limited capacity. With cloud computing, these problems go away. You can access as much or as little capacity as you need, and scale up and down as required with only a few minutes' notice. • Increase speed and agility – In a cloud computing environment, new IT resources are only a click away, which means that you reduce the time to make those resources available to your developers from weeks to just minutes. This results in a dramatic increase in agility for the organization, since the cost and time it takes to experiment and develop is significantly lower. • Stop spending money running and maintaining data centers – Focus on projects that differentiate your business, not the infrastructure. Cloud computing lets you focus on your own customers, rather than on the heavy lifting of racking, stacking, and powering servers. • Go global in minutes – Easily deploy your application in multiple regions around the world with just a few clicks. This means you can provide lower latency and a better experience for your customers at minimal cost.


iPQ.No	Correct Answer	Points	Explanation
5AWS CCP22	4	15.4	Answer = 4 Amazon SQS (Amazon Simple Queue Service) Explanation : The serverless applications, distributed systems, and microservices that you may decouple and scale using Amazon Simple Queue Service (Amazon SQS), a fully managed message queuing service. Developers can now concentrate on creating work that stands out from the crowd thanks to Amazon SQS, which removes the complexity and overhead involved in managing and operating message-oriented middleware. You can send, store, and receive messages across software components using Amazon Simple Queue Service (Amazon SQS) at any volume without losing messages or needing other services to be available.
5AWS CCP23	3	15.4	Answer = 3 Both A and B Explanation : Persistent block storage volumes are offered by Amazon Elastic Block Store (Amazon EBS) for usage with Amazon EC2 instances in the AWS Cloud. In order to provide high availability and durability and to safeguard you against component failure, each Amazon EBS volume is automatically replicated inside its Availability Zone (AZ). The reliable and low-latency performance required to operate your applications is provided by Amazon EBS volumes. With Amazon EBS, you can quickly scale up or decrease your usage while only paying for what you actually provision. The primary fees for EBS volumes include the volume type (depending on performance), the provisioned storage capacity in GB per month, the provisioned IOPS per month, the storage used for snapshots, and outbound data transmission.


iPQ.No	Correct Answer	Points	Explanation
5AWS CCP24	3	15.4	Answer = 3 Snowmobile AWS Explanation : Extremely huge amounts of data are transferred to AWS using the exabyte-scale data transfer service known as AWS Snowmobile. Per snowmobile, a 45-foot-long ruggedized shipping container drawn by a semi-trailer truck, you can transfer up to 100 PB. Massive volumes of data, such as video libraries, image repositories, or even a whole data center transfer, may be moved to the cloud with ease with AWS Snowmobile. Snowmobiling is a faster, more efficient, and more secure way to transfer data.
5AWS CCP25	3	15.4	Answer = 3 Kindle Macie Explanation : Amazon Macie is a fully managed data security and privacy solution that finds and safeguards sensitive data in AWS using machine learning and pattern matching. Amazon Macie finds sensitive data at scale in an economical manner using machine learning and pattern matching. A sizable and expanding range of sensitive data categories, including personally identifiable information (PII) like names, addresses, and credit card numbers, are automatically detected by Amazon Macie. Additionally, it provides you with continuous insight into the safety and protection of your data stored in Amazon S3. Using https://aws.amazon.com/macie/, find out how Amazon Macie functions. 

iPQ.No	Correct Answer	Points	Explanation
5AWS CCP26	3	15.4	Answer = 3 Both A and B Explanation : Responsibility of AWS: "Security of the Cloud": AWS is in charge of safeguarding the hardware 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. Exam Warning: Please thoroughly review the AWS Shared Responsibility Model because there will be many questions on it on the test (found at: https://aws.amazon.com/compliance/shared-responsibility-model/).  The diagram illustrates the AWS Shared Responsibility Model. On the left, two large vertical blocks represent the primary responsibilities: 'CUSTOMER' (Responsibility for Security 'IN' the Cloud) in blue and 'AWS' (Responsibility for Security 'OF' the Cloud) in orange. To the right, a stack of layers shows the components of the cloud environment. The top three layers are under the Customer's responsibility: 'CUSTOMER DATA', 'PLATFORM, APPLICATIONS, IDENTITY & ACCESS MANAGEMENT', and 'OPERATING SYSTEM, NETWORK & FIREWALL CONFIGURATION'. Below these are three boxes for security measures: 'CLIENT-SIDE DATA ENCRYPTION & DATA INTEGRITY AUTHENTICATION', 'SERVER-SIDE ENCRYPTION (FILE SYSTEM AND/OR DATA)', and 'NETWORKING TRAFFIC PROTECTION (ENCRYPTION, INTEGRITY, IDENTITY)'. The next layer is 'SOFTWARE', which includes 'COMPUTE', 'STORAGE', 'DATABASE', and 'NETWORKING'. Below that is 'HARDWARE/AWS GLOBAL INFRASTRUCTURE', which includes 'REGIONS', 'AVAILABILITY ZONES', and 'EDGE LOCATIONS'. The bottom two layers, 'SOFTWARE' and 'HARDWARE/AWS GLOBAL INFRASTRUCTURE', are under AWS's responsibility.

iPQ.No	Correct Answer	Points	Explanation
5AWS CCP27	4	15.4	Answer = 4 Quick Starts was formerly known as AWS Partner Solutions Explanation : AWS Partners and Amazon Web Services (AWS) solution architects work together to create automated reference deployments known as AWS Partner Solutions. Partner Solutions assists you with AWS best practices deployment of well-known technologies. You can use your surroundings in minutes after condensing hundreds of manual processes into a few actions. 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.
5AWS CCP28	2	15.4	Answer = 2 Cost Explorer for AWS Explanation : With the help of many filtering dimensions (such as AWS Service, AWS Region, Linked Account, etc.), AWS Cost Explorer enables you to investigate 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. In addition to the current AWS organization-level suggestions in AWS Cost Explorer, customers can also obtain Savings Plan recommendations at the member (linked) account level.

iPQ.No	Correct Answer	Points	Explanation
5AWS CCP29	1	15.4	Answer = 1 Trusted Advisor for AWS Explanation : By optimizing your AWS setup, AWS Trusted Advisor is an online tool that can help you lower costs, boost performance, and enhance security. AWS Trusted Advisor offers in-the-moment advice to assist you in provisioning your resources in accordance with AWS best practices. See https://aws.amazon.com/premiumsupport/technology/trusted-advisor/ for more information on how AWS Trusted Advisor functions.  <pre> graph LR A[As an AWS customer, you want the most value from your investment. Trusted Advisor can help.] --> B[AWS Trusted Advisor Trusted Advisor scans your AWS infrastructure, compares it to AWS best practices in five categories, and provides recommended actions] B --> C[Cost Optimization] B --> D[Performance] B --> E[Security] B --> F[Fault Tolerance] B --> G[Service limits] C --> H[Investigation recommended] D --> I[no problem detected] E --> J[action recommended] F --> K[no problem detected] G --> L[no problem detected] H --> M[Recommendations & action links] I --> M J --> M K --> M L --> M </pre>


iPQ.No	Correct Answer	Points	Explanation
5AWS CCP30	2	15.4	Answer = 2 ElastiCache by Amazon Explanation : An in-memory cache in the cloud may be quickly deployed, run, and scaled thanks to Amazon ElastiCache, a web service. The service enhances web application performance by enabling you to access data from quick, managed in-memory caches rather than solely relying on slower disk-based databases. ElastiCache can store some values to reduce the strain on a database while Amazon EC2 instances are heavily reading data from it. Using https://aws.amazon.com/elasticache/, find out how Amazon ElastiCache functions. 

iPQ.No	Correct Answer	Points	Explanation
5AWS CCP31	2	15.4	Answer = 2 AWS WAF (AWS Web Application Firewall) Explanation : AWS Web Application Firewall (AWS WAF) is a web application firewall that enables you to set up rules that permit, deny, or monitor (count) web requests based on criteria that you specify. This helps secure web applications from assaults. IP addresses, HTTP headers, HTTP bodies, URI strings, SQL injection, and cross-site scripting (XSS) are some examples of these circumstances. AWS WAF is a web application firewall that enables you to manage access to your content and track HTTP and HTTPS requests sent to Amazon CloudFront. Your rules are applied across all AWS Edge Locations, which are places all over the world near your end customers, when you utilize the AWS web application firewall (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. According to https://aws.amazon.com/waf, how AWS WAF functions  <pre> graph LR A[AWS Firewall Manager Manage multiple AWS WAF deployments] --> B[AWS WAF Protect your web applications from common web exploits] B --> C[Amazon CloudFront Application Load Balancer Amazon API Gateway] C --> D[Create a policy Build your own rules using the visual rule builder, code in JSON, or deploy managed rules maintained by AWS and/or sellers from AWS Marketplace] D --> E[Block & Filter Protect against exploits and vulnerabilities such as SQLi/XSS attacks; filter out unwanted traffic by defining specific patterns or by IP address] E --> F[Monitor Use Amazon CloudWatch for incoming traffic metrics & Amazon Kinesis Firehose for request details; then tune rules based on metrics & log data] </pre>

iPQ.No	Correct Answer	Points	Explanation
5AWS CCP32	3	15.4	Answer = 3 Both A and B Explanation : As your workloads on AWS expand and scale, AWS Organizations enables you to centrally manage your environment. Whether you run a huge corporation or a developing start-up, AWS organizations enable you to share resources among your AWS accounts and manage invoicing, access, compliance, and security from a single location. You can automatically generate accounts, group accounts to meet your business requirements, and apply policies to these groups for governance using AWS Organizations. Setting up a single payment method for all of your AWS accounts can also streamline billing. You can use AWS Organizations to define centralized setups and resource sharing across accounts in your business through interactions with other AWS services. All AWS users have access to AWS Organizations at no extra cost. For improved resource isolation and to have distinct per-account service limits, you should create accounts per department based on regulatory requirements (using Service Control Policies (SCP)). You can limit the services and activities that are permitted in your accounts using AWS Organizations. To apply permission boundaries to AWS Identity and Access Management (IAM) users and roles, use the Service Control Policies (SCP).

iPQ.No	Correct Answer	Points	Explanation
5AWS CCP33	3	15.4	Answer = 3 Both A and B Explanation : You can investigate AWS services and generate a cost projection for your use cases on AWS using the AWS Pricing Calculator. You can investigate the price ranges and calculations used to arrive at your estimate, model your solutions before they are built, and discover the instance kinds and contract terms that best suit your requirements. This gives you the information you need to use AWS wisely. You may budget your AWS expenses and use or estimate them by launching a fresh batch of instances and services. Visit https://calculator.aws/#/ to use the AWS Pricing Calculator. Using server, storage, network, and IT labor costs from an on-premises or conventional hosting environment, the AWS Pricing Calculator compares the costs of your applications against those of AWS. As a result, you must provide every component that is pertinent to these points of comparison. The costs of IT labor also include server administration. The costs associated with servers, storage, and networks include power and cooling.
5AWS CCP34	1	15.4	Answer = 1 Automation Explanation : Implementing automation will increase the efficiency of your business and the stability of your system. AWS Elastic Beanstalk, Auto Scaling, AWS Lambda, and other services can automate application architecture to increase performance, scalability, and resilience.


iPQ.No	Correct Answer	Points	Explanation
5AWS CCP35	2	15.4	Answer = 2 Kindle LightSail Explanation : The simplest approach to start and maintain a virtual private server (VPS) with AWS is to use Amazon Lightsail. For a low, dependable cost, Amazon Lightsail plans give everything you require to kick out your project, including a virtual machine, SSD-based storage, data transfer, Domain Name System (DNS) management, and a static IP address. It's excellent that a well-known IT solution can be swiftly launched and made available for usage even by those with little prior cloud computing experience.
5AWS CCP36	3	15.4	Answer = 3 OpsWorks by AWS Explanation : Chef and Puppet instances are controlled via the configuration management service AWS OpsWorks. You can use code to automate the configuration of your servers using automation technologies like Chef and Puppet. You can utilize AWS OpsWorks to leverage Chef and Puppet to automate the configuration, deployment, and management of servers across your on-premises compute infrastructures or Amazon Elastic Compute Cloud (Amazon EC2) instances.

iPQ.No	Correct Answer	Points	Explanation
5AWS CCP37	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. Your AWS resource configurations are continuously monitored and recorded by AWS Config, which also enables you to automatically compare the recorded configurations to the desired configurations. With AWS Config, you may investigate intricate resource configuration histories, examine changes in links across AWS resources, and assess your overall compliance with the configurations outlined in your internal rules. This makes it possible for you to streamline operational troubleshooting, change management, security analysis, and compliance auditing. Visit https://aws.amazon.com/config/ to learn how AWS Config functions.  The diagram illustrates the AWS Config workflow. It starts with a gear icon representing a 'Configuration change occurs in your AWS resources.' An arrow points to a central box labeled 'AWS Config' which contains the text 'AWS Config records and normalizes the changes into a consistent format.' Below this box is another box stating 'AWS Config automatically evaluates the recorded configurations against the configurations you specify.' An arrow points from the central box to a vertical stack of services: 'AWS Config APIs & Console', 'Amazon SNS', 'Amazon CloudWatch', and 'Amazon S3'. A final text block at the bottom right explains: 'Access change history and compliance results using the console or APIs. CloudWatch Events or SNS alert you when changes occur. Deliver change history and snapshot files to your S3 bucket for analysis.'

iPQ.No	Correct Answer	Points	Explanation																																													
5AWS CCP38	4	15.4	Answer = 4 For each department, make tags Explanation : Your AWS resources can receive tags as metadata that you can apply. Each tag is a label made up of a key and value that the user defines. You may manage, identify, organize, find, and filter materials with the use of tags. Resources can be categorized using tags according to their owner, environment, purpose, or other factors. Typically, you link AWS expenses with conventional cost-allocation dimensions using business identifiers like cost center/business unit, client, or project. But any tag can be included in a cost allocation report. This enables you to link expenditures to technical or security aspects like particular apps, environments, or compliance initiatives. By way of https://docs.aws.amazon.com/general/latest/gr/aws_tagging.html, an illustration of tagging for cost optimization <table><tr><th>Total Cost</th><th>user:Owner</th><th>user:Stack</th><th>user:Cost Center</th><th>user:Application</th></tr><tr><td>0.95</td><td>DbAdmin</td><td>Test</td><td>80432</td><td>Widget2</td></tr><tr><td>0.01</td><td>DbAdmin</td><td>Test</td><td>80432</td><td>Widget2</td></tr><tr><td>3.84</td><td>DbAdmin</td><td>Prod</td><td>80432</td><td>Widget2</td></tr><tr><td>6.00</td><td>DbAdmin</td><td>Test</td><td>78925</td><td>Widget1</td></tr><tr><td>234.63</td><td>SysEng</td><td>Prod</td><td>78925</td><td>Widget1</td></tr><tr><td>0.73</td><td>DbAdmin</td><td>Test</td><td>78925</td><td>Widget1</td></tr><tr><td>0.00</td><td>DbAdmin</td><td>Prod</td><td>80432</td><td>Portal</td></tr><tr><td>2.47</td><td>DbAdmin</td><td>Prod</td><td>78925</td><td>Portal</td></tr></table>	Total Cost	user:Owner	user:Stack	user:Cost Center	user:Application	0.95	DbAdmin	Test	80432	Widget2	0.01	DbAdmin	Test	80432	Widget2	3.84	DbAdmin	Prod	80432	Widget2	6.00	DbAdmin	Test	78925	Widget1	234.63	SysEng	Prod	78925	Widget1	0.73	DbAdmin	Test	78925	Widget1	0.00	DbAdmin	Prod	80432	Portal	2.47	DbAdmin	Prod	78925	Portal
Total Cost	user:Owner	user:Stack	user:Cost Center	user:Application																																												
0.95	DbAdmin	Test	80432	Widget2																																												
0.01	DbAdmin	Test	80432	Widget2																																												
3.84	DbAdmin	Prod	80432	Widget2																																												
6.00	DbAdmin	Test	78925	Widget1																																												
234.63	SysEng	Prod	78925	Widget1																																												
0.73	DbAdmin	Test	78925	Widget1																																												
0.00	DbAdmin	Prod	80432	Portal																																												
2.47	DbAdmin	Prod	78925	Portal																																												

iPQ.No	Correct Answer	Points	Explanation
5AWS CCP39	4	15.4	Answer = 4 Configuration of the lifespan of Amazon Simple Storage Service (Amazon S3) Explanation : Configure the Amazon S3 Lifecycle of your objects to manage them so that they are cost-effectively stored over the course of their lives. An S3 Lifecycle configuration is a set of guidelines that specify the activities Amazon S3 should take in relation to a collection of items. There are two different kinds of activities: expiration actions and transition actions, which specify when items move from one storage class to another and when they are deleted by Amazon S3 on your behalf. In this specific use scenario, a transition action would be used.

iPQ.No	Correct Answer	Points	Explanation
5AWS CCP40	3	15.4	Answer = 3 Both A and B Explanation : The capacity to manage systems to create business value and continuously enhance auxiliary practices and processes is a component of the operational excellence pillar. To find probable sources of failure and remove or lessen them, do "pre-mortem" activities. Test your failure scenarios and confirm that you comprehend their effects. Test your response protocols to make sure they work and that teams are comfortable using them. To evaluate workloads and team reactions to simulated situations, schedule regular game days. Create workloads that allow for frequent component updates. When possible, implement minor, reversible adjustments that can be undone if they don't work. The AWS Well-Architected Framework aids in your comprehension of the advantages and disadvantages of the choices you make while developing systems on AWS. Using the Framework will teach you architectural best practices for creating and managing dependable, secure, successful, and affordable cloud systems. It gives you a tool to regularly compare your architectures to best practices and pinpoint areas that need improvement. Operational Excellence, Security, Reliability, Performance Efficiency, Cost Optimization, and Sustainability are the six pillars on which the AWS Well-Architected Framework is built. Overview of the AWS Well-Architected Framework's six pillars:

iPQ.No	Correct Answer	Points	Explanation
5AWS CCP41	3	15.4	Answer = 3 A client gateway Explanation : A customer gateway device on your end of a site-to-site VPN connection is a hardware or software appliance. The device must be configured by you or your network administrator in order to use a site-to-site VPN connection. By setting up an AWS Site-to-Site VPN connection and configuring routing to pass traffic over it, you can enable access to your remote network from your VPC. via - https://docs.aws.amazon.com/vpn/latest/s2svpn/your-cgw.html - more information on customer gateway devices 


iPQ.No	Correct Answer	Points	Explanation
5AWS CCP42	1	15.4	Answer = 1 Cloud Data Sync Explanation : AWS DataSync is a secure online data transfer service that simplifies, automates, and accelerates moving terabytes of data to and from AWS storage services. Easily migrate or replicate large data sets without having to build custom solutions or oversee repetitive tasks. In addition to self-managed object storage, AWS Snowcone, Amazon Simple Storage Service (Amazon S3) buckets, Amazon Elastic File System (Amazon EFS) file systems, and Amazon FSx for Windows File Server file systems, DataSync can copy data between Network File System (NFS) shares or Server Message Block (SMB) shares. AWS DataSync can be used to continuously transfer data between on-premises systems and AWS for processing. In industries where moving active files into AWS quickly is necessary, DataSync can help speed up your crucial hybrid cloud storage workflows. This includes machine learning in life sciences, video production in media and entertainment, and big data analytics in financial services. To prevent delays in dependent processes, AWS DataSync offers timely delivery. You can define exclude filters, include filters, or both to select which files, folders, or objects get transferred each time your job runs. To speed up data transfer, AWS DataSync uses a transfer protocol that it designed that is independent of the storage protocol. The protocol performs optimizations on how, when, and what data is sent over the network. Network optimizations performed by DataSync include incremental transfers, in-line compression, and sparse file detection, as well as in-line data validation and encryption. Data Transfer between on-premises and AWS using AWS DataSync: via https://aws.amazon.com/datasync/ 

iPQ.No	Correct Answer	Points	Explanation
5AWS CCP43	3	15.4	Answer = 3 Both A and B Explanation : Shared Controls: Controls that apply to both the infrastructure layer and the customer layer but in completely separate contexts or perspectives In a shared control environment, AWS provides the requirements for the infrastructure, and the customer must provide their own control implementation within their use of AWS services. AWS manages configuration for its infrastructure devices, but customers are in charge of setting up their own guest operating systems, databases, and applications. Configuration of the operating system (OS) "Security in the cloud" is the responsibility of the customers. It comprises client data as well as the guest operating system settings. Operating system configuration as a whole is an AWS shared responsibility, but be careful: the host operating system configuration is the responsibility of AWS, and the guest operating system configuration is the responsibility of the customer. Exam Warning: Please thoroughly review the AWS Shared Responsibility Model because there will be many questions on it on the test (found at: https://aws.amazon.com/compliance/shared-responsibility-model/).  The diagram illustrates the AWS Shared Responsibility Model. On the left, a blue arrow points upwards, labeled 'CUSTOMER' and 'RESPONSIBILITY FOR SECURITY 'IN' THE CLOUD'. On the right, three stacked blue boxes represent AWS's responsibilities: 'CUSTOMER DATA', 'PLATFORM, APPLICATIONS, IDENTITY & ACCESS MANAGEMENT', and 'OPERATING SYSTEM, NETWORK & FIREWALL CONFIGURATION'.

iPQ.No	Correct Answer	Points	Explanation
5AWS CCP44	4	15.4	Answer = 4 Amazon Lex Explanation : Amazon Lex is a service for creating voice- and text-based conversational user interfaces. The same conversational engine that powers Amazon Alexa also powers Amazon Lex, which offers excellent speech recognition and language understanding capabilities. This paves the way for the integration of sophisticated, natural language "chatbots" into both new and existing applications. Use Cases for Amazon Lex  via: https://aws.amazon.com/lex/
5AWS CCP45	2	15.4	Answer = 2 AWS Identity and Access Management (AWS IAM) is a global service, thus there is nothing to do Explanation : AWS Identity and Access Management (IAM) enables you to manage access to AWS services and resources securely. Using IAM, you may establish and manage IAM users and IAM user groups and utilize permissions to grant and prohibit their access to AWS services. Global service provided by AWS IAM. Users created within IAM can access their accounts all across the world and deploy resources in every location.

iPQ.No	Correct Answer	Points	Explanation
5AWS CCP46	3	15.4	Answer = 3 It is a platform as a service (PaaS) that allows you to deploy and scale web applications and services Explanation : Developers now have it even simpler to quickly deploy and manage applications in the AWS Cloud thanks to AWS Elastic Beanstalk. Developers simply upload their applications, and AWS Elastic Beanstalk automatically handles the deployment details of capacity provisioning, load balancing, auto-scaling, and application health monitoring. You only control the applications and the data, making it a platform as a service (PaaS). Review https://aws.amazon.com/types-of-cloud-computing/ for a summary of the different types of cloud computing. Cloud Computing Models There are three main models for cloud computing. Each model represents a different part of the cloud computing stack.  Infrastructure as a Service (IaaS) Infrastructure as a Service, sometimes abbreviated as IaaS, contains the basic building blocks for cloud IT and typically provide access to networking features, computers (virtual or on dedicated hardware), and data storage space. Infrastructure as a Service provides you with the highest level of flexibility and management control over your IT resources and is most similar to existing IT resources that many IT departments and developers are familiar with today.  Platform as a Service (PaaS) Platforms as a service remove the need for organizations to manage the underlying infrastructure (usually hardware and operating systems) and allow you to focus on the deployment and management of your applications. This helps you be more efficient as you don't need to worry about resource procurement, capacity planning, software maintenance, patching, or any of the other undifferentiated heavy lifting involved in running your application.  Software as a Service (SaaS) Software as a Service provides you with a completed product that is run and managed by the service provider. In most cases, people referring to Software as a Service are referring to end-user applications. With a SaaS offering you do not have to think about how the service is maintained or how the underlying infrastructure is managed; you only need to think about how you will use that particular piece of software. A common example of a SaaS application is web-based email where you can send and receive email without having to manage feature additions to the email product or maintaining the servers and operating systems that the email program is running on.

iPQ.No	Correct Answer	Points	Explanation
5AWS CCP47	3	15.4	Answer = 3 Both A and B Explanation : You receive a sizable discount (up to 75%) on Reserved Instances (RI) compared to the cost of an on-demand instance. Besides, when reserved instances (RI) are allocated to a certain availability zone (AZ), they give a capacity reservation, providing you additional assurance in your ability to run instances when you need them. A 1-year or 3-year term is available for the purchase of standard and convertible reserved instances. Overview of EC2 Pricing Options: On-Demand With On-Demand instances, you pay for compute capacity by the hour or the second depending on which instances you run. No longer-term commitments or upfront payments are needed. You can increase or decrease your compute capacity depending on the demands of your application and only pay the specified per hourly rates for the instance you use. On-Demand instances are recommended for: - Users that prefer the low cost and flexibility of Amazon EC2 without any up-front payment or long-term commitment - Applications with short-term, spiky, or unpredictable workloads that cannot be interrupted - Applications being developed or tested on Amazon EC2 for the first time See On-Demand pricing » Spot instances Amazon EC2 Spot instances allow you to request spare Amazon EC2 computing capacity for up to 90% off the On-Demand price. Learn More. Spot instances are recommended for: - Applications that have flexible start and end times - Applications that are only feasible at very low compute prices - Users with urgent computing needs for large amounts of additional capacity See Spot pricing » Savings Plans Savings Plans are a flexible pricing model that offer low prices on EC2 and Fargate usage, in exchange for a commitment to a consistent amount of usage (measured in \$/hour) for a 1 or 3 year term. Dedicated Hosts A Dedicated Host is a physical EC2 server dedicated for your use. Dedicated Hosts can help you reduce costs by allowing you to use your existing server-bound software licenses, including Windows Server, SQL Server, and SUSE Linux Enterprise Server (subject to your license terms), and can also help you meet compliance requirements. Learn more. - Can be purchased On-Demand (hourly). - Can be purchased as a Reservation for up to 70% off the On-Demand price. See Dedicated pricing » Reserved Instances Reserved Instances provide you with a significant discount (up to 75%) compared to On-Demand instance pricing. In addition, when Reserved Instances are assigned to a specific Availability Zone, they provide a capacity reservation, giving you additional confidence in your ability to launch instances when you need them. For applications that have steady state or predictable usage, Reserved Instances can provide significant savings compared to using On-Demand instances. See How to Purchase Reserved Instances for more information. Reserved Instances are recommended for: - Applications with steady state usage - Applications that may require reserved capacity - Customers that can commit to using EC2 over a 1 or 3 year term to reduce their total computing costs via: https://aws.amazon.com/ec2/pricing

iPQ.No	Correct Answer	Points	Explanation
5AWS CCP48	2	15.4	Answer = 2 Apple Batch Explanation : Developers, scientists, and engineers can run hundreds of thousands of batch computing jobs on AWS quickly and effectively thanks to AWS Batch. You can use AWS Batch to plan, schedule, and execute your batch computing workloads across the full range of AWS compute services. Based on the volume and particular resource requirements of the batch jobs submitted, AWS Batch dynamically allocates the best amount and kind of compute resources (such as memory-optimized instances or CPUs). AWS Batch Provisions computes resources and optimizes the job distribution based on the volume and resource requirements of the submitted batch jobs. Please study the typical AWS Batch use cases: Use cases Financial services: Post-trade analytics Automate the analysis of the day's transaction costs, execution reporting, and market performance.  Life sciences: Drug screening for biopharma Rapidly search libraries of small molecules for drug discovery.  Digital media: Visual effects rendering Automate content rendering workloads and reduce the need for human intervention due to execution dependencies or resource scheduling. 


iPQ.No	Correct Answer	Points	Explanation
5AWS CCP49	3	15.4	Answer = 3 AWS Security Token Service (AWS STS) Explanation : The AWS Security Token Service (AWS STS) is a web service that enables you to request temporary, limited-privilege credentials for AWS Identity and Access Management (AWS IAM) users or for users that you authenticate (federated users). You can create and distribute trusted users with temporary security credentials that can restrict access to your AWS resources using the AWS Security Token Service (AWS STS). Temporary security credentials work almost identically to the long-term access key credentials that your IAM users can use, with the following differences: Temporary security credentials are short-term, as the name implies. They can be adjusted to last ranging from a few minutes to several hours. After the credentials expire, AWS no longer recognizes them or allows any kind of access to API requests made with them. Temporary security credentials are not saved with the user but are produced dynamically and delivered to the user as required. The user can ask for new security credentials when (or even before) the current ones run out as long as they are still authorized to do so. Temporary security credentials are generated by the AWS Security Token Service (AWS STS). AWS STS is a universal service by default, and its endpoint is https://sts.amazonaws.com. However, you can also opt to make AWS STS API calls to endpoints in any other supported region.

iPQ.No	Correct Answer	Points	Explanation
5AWS CCP50	1	15.4	Answer = 1 Reserved instance (RI) convertible Explanation : Purchase a convertible reserved instance (RI) if you need greater flexibility, such as the option to use alternative instance families, operating systems, or tenancies over the reserved instance (RI) term. You can acquire a converted reserved instance (RI) for a 1-year or 3-year period and receive a large discount (up to 54%) compared to an on-demand instance. When workloads are likely to alter, a convertible reserved instance (RI) may be helpful. In this example, a convertible reserved instance (RI) enables you to adjust as needs evolve while still enjoying discounts and capacity reservations. Amazon EC2 Pricing Options Overview On-Demand With On-Demand instances, you pay for compute capacity by the hour or the second depending on which instances you run. No longer-term commitments or upfront payments are needed. You can increase or decrease your compute capacity depending on the demands of your application and only pay the specified per hourly rates for the instance you use. On-Demand instances are recommended for: - Users that prefer the low cost and flexibility of Amazon EC2 without any up-front payment or long-term commitment - Applications with short-term, spiky, or unpredictable workloads that cannot be interrupted - Applications being developed or tested on Amazon EC2 for the first time See On-Demand pricing » Spot instances Amazon EC2 Spot instances allow you to request spare Amazon EC2 computing capacity for up to 90% off the On-Demand price. Learn More. Spot instances are recommended for: - Applications that have flexible start and end times - Applications that are only feasible at very low compute prices - Users with urgent computing needs for large amounts of additional capacity See Spot pricing » Savings Plans Savings Plans are a flexible pricing model that offer low prices on EC2 and Fargate usage, in exchange for a commitment to a consistent amount of usage (measured in \$/hour) for a 1 or 3 year term. Dedicated Hosts A Dedicated Host is a physical EC2 server dedicated for your use. Dedicated Hosts can help you reduce costs by allowing you to use your existing server-bound software licenses, including Windows Server, SQL Server, and SUSE Linux Enterprise Server (subject to your license terms), and can also help you meet compliance requirements. Learn more. - Can be purchased On-Demand (hourly). - Can be purchased as a Reservation for up to 70% off the On-Demand price. See Dedicated pricing » Reserved Instances Reserved Instances provide you with a significant discount (up to 75%) compared to On-Demand instance pricing. In addition, when Reserved Instances are assigned to a specific Availability Zone, they provide a capacity reservation, giving you additional confidence in your ability to launch instances when you need them. For applications that have steady state or predictable usage, Reserved Instances can provide significant savings compared to using On-Demand instances. See How to Purchase Reserved Instances for more information. Reserved Instances are recommended for: - Applications with steady state usage - Applications that may require reserved capacity - Customers that can commit to using EC2 over a 1 or 3 year term to reduce their total computing costs : via: https://aws.amazon.com/ec2/pricing/

iPQ.No	Correct Answer	Points	Explanation
5AWS CCP51	1	15.4	Answer = 1 SageMaker on Amazon Explanation : Amazon SageMaker is a fully managed platform that enables developers and data scientists to rapidly and effectively construct, train, and deploy machine learning models at any scale. All the obstacles that often prevent developers from using machine learning are eliminated by Amazon SageMaker.

iPQ.No	Correct Answer	Points	Explanation																																																																																								
5AWS CCP52	1	15.4	Answer = 1 Amazon S3 One Zone-Infrequent Access (S3 One Zone-IA) Explanation : Data that is accessed less frequently but needs quick access when needed is stored in Amazon S3 One Zone-Infrequent Access (S3 One Zone-IA). Unlike other S3 storage classes, which store data in a minimum of three availability zones (AZs), S3 One Zone-IA stores data in a single availability zone (AZ) and costs 20% less than S3 Standard-IA. S3 One Zone-IA is ideal for customers who want a lower-cost option for infrequently accessed data but do not require the availability and resilience of S3 Standard or S3 Standard-IA. It's a good choice for storing secondary backup copies of on-premises data or easily re-creatable data. You can also use it as cost-effective storage for data that is replicated from another AWS area using S3 cross-region replication (S3 CRR). Exam Warning: If you have any questions about S3's storage classes, please review this comprehensive comparison here: Performance across the S3 Storage Classes <table><thead><tr><th></th><th>S3 Standard</th><th>S3 Intelligent-Tiering*</th><th>S3 Standard-IA</th><th>S3 One Zone-IA†</th><th>S3 Glacier Instant Retrieval</th><th>S3 Glacier Flexible Retrieval</th><th>S3 Glacier Deep Archive</th></tr></thead><tbody><tr><td>Designed for durability</td><td>99.999999999% (11 9's)</td><td>99.999999999% (11 9's)</td><td>99.999999999% (11 9's)</td><td>99.999999999% (11 9's)</td><td>99.999999999% (11 9's)</td><td>99.999999999% (11 9's)</td><td>99.999999999% (11 9's)</td></tr><tr><td>Designed for availability</td><td>99.99%</td><td>99.9%</td><td>99.9%</td><td>99.5%</td><td>99.9%</td><td>99.99%</td><td>99.99%</td></tr><tr><td>Availability SLA</td><td>99.9%</td><td>99%</td><td>99%</td><td>99%</td><td>99%</td><td>99.9%</td><td>99.9%</td></tr><tr><td>Availability Zones</td><td>≥3</td><td>≥3</td><td>≥3</td><td>1</td><td>≥3</td><td>≥3</td><td>≥3</td></tr><tr><td>Minimum capacity charge per object</td><td>N/A</td><td>N/A</td><td>128 KB</td><td>128 KB</td><td>128 KB</td><td>40 KB</td><td>40 KB</td></tr><tr><td>Minimum storage duration charge</td><td>N/A</td><td>N/A</td><td>30 days</td><td>30 days</td><td>90 days</td><td>90 days</td><td>180 days</td></tr><tr><td>Retrieval charge</td><td>N/A</td><td>N/A</td><td>per GB retrieved</td><td>per GB retrieved</td><td>per GB retrieved</td><td>per GB retrieved</td><td>per GB retrieved</td></tr><tr><td>First byte latency</td><td>milliseconds</td><td>milliseconds</td><td>milliseconds</td><td>milliseconds</td><td>milliseconds</td><td>minutes or hours</td><td>hours</td></tr><tr><td>Storage type</td><td>Object</td><td>Object</td><td>Object</td><td>Object</td><td>Object</td><td>Object</td><td>Object</td></tr><tr><td>Lifecycle transitions</td><td>Yes</td><td>Yes</td><td>Yes</td><td>Yes</td><td>Yes</td><td>Yes</td><td>Yes</td></tr></tbody></table> https://aws.amazon.com/s3/storage-classes/		S3 Standard	S3 Intelligent-Tiering*	S3 Standard-IA	S3 One Zone-IA†	S3 Glacier Instant Retrieval	S3 Glacier Flexible Retrieval	S3 Glacier Deep Archive	Designed for durability	99.999999999% (11 9's)	99.999999999% (11 9's)	99.999999999% (11 9's)	99.999999999% (11 9's)	99.999999999% (11 9's)	99.999999999% (11 9's)	99.999999999% (11 9's)	Designed for availability	99.99%	99.9%	99.9%	99.5%	99.9%	99.99%	99.99%	Availability SLA	99.9%	99%	99%	99%	99%	99.9%	99.9%	Availability Zones	≥3	≥3	≥3	1	≥3	≥3	≥3	Minimum capacity charge per object	N/A	N/A	128 KB	128 KB	128 KB	40 KB	40 KB	Minimum storage duration charge	N/A	N/A	30 days	30 days	90 days	90 days	180 days	Retrieval charge	N/A	N/A	per GB retrieved	per GB retrieved	per GB retrieved	per GB retrieved	per GB retrieved	First byte latency	milliseconds	milliseconds	milliseconds	milliseconds	milliseconds	minutes or hours	hours	Storage type	Object	Object	Object	Object	Object	Object	Object	Lifecycle transitions	Yes	Yes	Yes	Yes	Yes	Yes	Yes
	S3 Standard	S3 Intelligent-Tiering*	S3 Standard-IA	S3 One Zone-IA†	S3 Glacier Instant Retrieval	S3 Glacier Flexible Retrieval	S3 Glacier Deep Archive																																																																																				
Designed for durability	99.999999999% (11 9's)	99.999999999% (11 9's)	99.999999999% (11 9's)	99.999999999% (11 9's)	99.999999999% (11 9's)	99.999999999% (11 9's)	99.999999999% (11 9's)																																																																																				
Designed for availability	99.99%	99.9%	99.9%	99.5%	99.9%	99.99%	99.99%																																																																																				
Availability SLA	99.9%	99%	99%	99%	99%	99.9%	99.9%																																																																																				
Availability Zones	≥3	≥3	≥3	1	≥3	≥3	≥3																																																																																				
Minimum capacity charge per object	N/A	N/A	128 KB	128 KB	128 KB	40 KB	40 KB																																																																																				
Minimum storage duration charge	N/A	N/A	30 days	30 days	90 days	90 days	180 days																																																																																				
Retrieval charge	N/A	N/A	per GB retrieved	per GB retrieved	per GB retrieved	per GB retrieved	per GB retrieved																																																																																				
First byte latency	milliseconds	milliseconds	milliseconds	milliseconds	milliseconds	minutes or hours	hours																																																																																				
Storage type	Object	Object	Object	Object	Object	Object	Object																																																																																				
Lifecycle transitions	Yes	Yes	Yes	Yes	Yes	Yes	Yes																																																																																				

iPQ.No	Correct Answer	Points	Explanation
5AWS CCP53	1	15.4	Answer = 1 Access advisor for AWS Identity and Access Management (IAM) Explanation : The IAM Access Advisor reveals the service permissions granted to a user and when those services were last accessed. You can use this information to revise your policies. In conclusion, you may spot unused rights and adjust your IAM policies accordingly.

iPQ.No	Correct Answer	Points	Explanation
5AWS CCP54	4	15.4	Answer = 4 Amazon CodeGuru Explanation : With the help of Amazon CodeGuru, developers may find the most expensive lines of code in an application and enhance the quality of their code. Use Amazon CodeGuru to automate code reviews during application development, continuously monitor application performance in production, offer suggestions and visual cues for enhancing code quality and application performance, and cut costs. When developing applications, Amazon CodeGuru Reviewer uses machine learning and automated reasoning to find critical flaws, security holes, and difficult-to-find bugs. It then makes suggestions to improve the quality of the code. By assisting developers in comprehending the behavior of their applications during runtime, locating and eliminating code inefficiencies, enhancing performance, and drastically lowering compute costs, Amazon CodeGuru Profiler helps identify an application's most expensive lines of code. https://aws.amazon.com/codeguru/ explains how Amazon CodeGuru functions.  The diagram illustrates the Amazon CodeGuru workflow. It starts with a box labeled 'Amazon CodeGuru' with a lightbulb icon and the text 'Find your most expensive lines of code'. An arrow points to 'Write & Review Code' (with a document icon), which is associated with 'CodeGuru Reviewer' and the text 'Built-in code reviews with actionable recommendations'. Another arrow points to 'Build & Test' (with a document and checkmark icon), associated with 'CodeGuru Profiler' and the text 'Detect and optimize the expensive lines of code'. This is followed by 'Deploy' (with a rocket icon). An arrow then points to 'Run' (with a server icon), associated with 'CodeGuru Profiler' and the text 'Continuously detect anomalies and most expensive lines of code in production'. Finally, an arrow points to 'Improve' (with a lightbulb and bar chart icon), associated with 'CodeGuru Profiler' and the text 'Fix performance issues and reduce cost'. A feedback loop arrow connects 'Improve' back to 'Write & Review Code'.

iPQ.No	Correct Answer	Points	Explanation
5AWS CCP55	3	15.4	Answer = 3 Changing out bad Amazon EC2 instances Explanation : With the aid of Amazon EC2 Auto Scaling, you may automatically add or remove Amazon EC2 instances in accordance with the criteria you specify while still maintaining the availability of your applications. To keep your fleet healthy and available, you can make use of Amazon EC2 Auto Scaling's fleet management features. To add or remove EC2 instances, you can also use Amazon EC2 Auto Scaling's dynamic and predictive scaling features. When an instance becomes unhealthy, Amazon EC2 Auto Scaling can identify it, terminate it, and then start a new one.
5AWS CCP56	1	15.4	Answer = 1 (Amazon S3) Amazon Simple Storage Service Explanation : An Amazon EBS Elastic Volume can have a point-in-time snapshot created of it that you can use as a starting point for new volumes or data backups. When you periodically take a snapshot of a volume, the snapshots are incremental; only the blocks that have changed since the previous snapshot are saved in the new snapshot. You can back up the data on your Amazon EBS Elastic Volumes to Amazon Simple Storage Service (Amazon S3) by taking point-in-time snapshots.

iPQ.No	Correct Answer	Points	Explanation
5AWS CCP57	1	15.4	Answer = 1 Use AWS Key Management Service (AWS KMS) to encrypt data Explanation : The security pillar includes the ability to protect information, systems, and assets while delivering business value through risk assessments and mitigation strategies. Encrypting data is part of the design concept "Protect data in transit and at rest." Classify your data into sensitivity levels and employ technologies such as encryption, tokenization, and access control when applicable. AWS Key Management Service (AWS KMS) makes it easy for you to produce and control keys needed for encryption. It is an essential Security Pillar service. The AWS Well-Architected Framework aids in your comprehension of the advantages and disadvantages of the choices you make while developing systems on AWS. By using the framework, you will learn architectural best practices for designing and operating reliable, secure, efficient, and cost-effective systems in the cloud. It gives you a tool to regularly compare your architecture to best practices and pinpoint areas that need improvement. The AWS Well-Architected Framework is centered on six pillars: operational excellence, security, reliability, performance efficiency, cost optimization, and sustainability. Overview of the AWS Well-Architected Framework's six pillars: AWS Well-Architected and the Six Pillars Framework Overview The AWS Well-Architected Framework describes key concepts, design principles, and architectural best practices for designing and running workloads in the cloud. By answering a few foundational questions, learn how well your architecture aligns with cloud best practices and gain guidance for making improvements. 


iPQ.No	Correct Answer	Points	Explanation
5AWS CCP58	4	15.4	Answer = 4 AWS CloudHSM Explanation : Utilizing dedicated Hardware Security Module (HSM) instances within the AWS cloud, the CloudHSM service from Amazon Web Services (AWS) assists you in adhering to contractual, regulatory, and corporate compliance requirements for data security. You can securely create, store, and manage the cryptographic keys required for data encryption using AWS CloudHSM in a way that only you have access to the keys. How AWS CloudHSM works: via https://aws.amazon.com/cloudhsm/

iPQ.No	Correct Answer	Points	Explanation
5AWS CCP59	4	15.4	Answer = 4 Cost and Usage Report for AWS (AWS CUR) Explanation : The most complete set of cost and usage data is found in the AWS Cost & Usage Report (AWS CUR). Your Amazon Simple Storage Service (Amazon S3) bucket can be used to publish your AWS billing reports using Cost and Usage Reports. You can get reports that show your costs broken down by the hour, month, product, resource, or tags that you define. What are AWS Cost and Usage Reports? PDF RSS The AWS Cost and Usage Reports (AWS CUR) contains <u>the most comprehensive set of cost and usage data available</u>. You can use Cost and Usage Reports to publish your AWS billing reports to an Amazon Simple Storage Service (Amazon S3) bucket that you own. <u>You can receive reports that break down your costs by the hour or month, by product or product resource, or by tags that you define yourself.</u> AWS updates the report in your bucket once a day in comma-separated value (CSV) format. You can view the reports using spreadsheet software such as Microsoft Excel or Apache OpenOffice Calc, or access them from an application using the Amazon S3 API. AWS Cost and Usage Reports tracks your AWS usage and provides estimated charges associated with your account. Each report contains line items for each unique combination of AWS products, usage type, and operation that you use in your AWS account. You can customize the AWS Cost and Usage Reports to aggregate the information either by the hour or by the day. AWS Cost and Usage Reports can do the following: • Deliver report files to your Amazon S3 bucket • Update the report up to three times a day • Create, retrieve, and delete your reports using the AWS CUR API Reference via: https://docs.aws.amazon.com/cur/latest/userguide/what-is-cur.html AWS Cost & Usage Report (AWS CUR) Overview
5AWS CCP60	1	15.4	Answer = 1 Per second payment Explanation : With on-demand instances, you only pay for the Amazon EC2 instances you use. Utilizing on-demand instances relieves you of the expenses and hassles associated with organizing, acquiring, and maintaining hardware and converts what are typically high fixed costs into considerably lower variable costs. Pay-per-second pricing is used whenever a Windows EC2 on-demand instance is running.

iPQ.No	Correct Answer	Points	Explanation
5AWS CCP61	1	15.2	Answer = 1 The AWS Fault Injection Simulator Explanation : To do fault injection experiments on AWS, use the AWS Fault Injection Simulator (AWS FIS), a fully managed tool that makes it simpler to increase an application's performance, observability, and resilience. Chaos engineering, which is the technique of straining an application in testing or production environments by generating disruptive events, such as a sharp rise in CPU or memory use, watching how the system reacts, and applying improvements, uses fault injection experiments to do this. The real-world conditions needed to find hidden bugs, monitor blind spots, and track performance bottlenecks that are challenging to find in distributed systems are created by teams using fault injection experiments. Teams can gain confidence in the behavior of their applications by using the AWS Fault Injection Simulator (AWS FIS), which makes it easier to set up and run controlled fault injection experiments across a variety of AWS services. Teams may quickly set up experiments using pre-built templates that produce the required disruptions using AWS Fault Injection Simulator (AWS FIS). Teams can execute experiments in production with the help of the AWS Fault Injection Simulator (AWS FIS), which offers controls and safeguards including automatically rolling back or ending the experiment if certain criteria are satisfied. Teams can run complex scenarios with typical distributed system failures occurring in parallel or building sequentially over time with just a few clicks in the console, giving them the ability to simulate real-world conditions to find hidden weaknesses.

iPQ.No	Correct Answer	Points	Explanation
5AWS CCP62	3	15.2	Answer = 3 Both A and B Explanation : A highly scalable and available cloud Domain Name System (DNS) web service is Amazon Route 53. By converting names like www.example.com into the numeric IP addresses that computers use to connect to each other, such 192.0.2.1, it is intended to give developers and organizations an incredibly dependable and cost-effective way to route end users to Internet services. You can search for and register available domain names using Amazon Route 53's domain name registration services, or you can transfer already-registered domain names to Route 53 for management. Your application, web servers, and other resources can all be kept under observation by Amazon Route 53.

iPQ.No	Correct Answer	Points	Explanation
5AWS CCP63	4	15.2	Answer = 4 Amazon Customize Explanation : With the help of Amazon Personalize, programmers can create applications using the same machine learning (ML) technology that powers real-time personalized recommendations on Amazon.com. Any digital channel can have its end-user experience customized using Amazon Personalize. Examples include recommending products for e-commerce, news stories and content for publishing, media, and social networks, hotels for travel websites, credit cards for banks, and matches for dating services. Websites, mobile apps, emails, and messaging services can all be used to deliver these suggestions and personalized experiences. When a user interacts with a product or service through a physical channel, such as a meal delivery service, Amazon Personalize can be used to personalise the user experience. The major use cases that Amazon Personalize provides are as follows: - individual recommendations - comparable items - Reranking a user's list of items according to their preferences - Personalized promotions and notifications

iPQ.No	Correct Answer	Points	Explanation
5AWS CCP64	3	15.2	Answer = 3 Both A and B Explanation : Amazon CloudWatch is a monitoring service for AWS cloud resources and the applications you run on AWS. You can use Amazon CloudWatch to gather and track metrics, collect and monitor log files, and create alarms. Amazon CloudWatch can track custom metrics produced by your applications and services, any log files your applications produce, and AWS resources like Amazon Elastic Compute Cloud (Amazon EC2) instances, Amazon DynamoDB tables, and Amazon Relational Database Service (Amazon RDS) DB instances. Gaining system-wide visibility into resource utilization, application performance, and operational health is possible with Amazon CloudWatch. These insights can help you react and maintain the efficiency of your application. Using https://aws.amazon.com/cloudwatch/, find out how Amazon CloudWatch functions.  The diagram illustrates the Amazon CloudWatch workflow. It starts with 'Amazon CloudWatch' (Complete visibility into your cloud resources and applications). This leads to 'Collect' (Metrics and logs from all your AWS resources, applications, and services that run on AWS and on-premises servers). 'Collect' leads to 'Monitor' (Visualize applications and infrastructure with CloudWatch dashboards; correlate logs and metrics side by side to troubleshoot and set alerts with CloudWatch Alarms). 'Monitor' leads to 'Act' (Automate response to operational changes with CloudWatch Events and Auto Scaling). 'Act' leads to 'Analyze' (Up to 1-second metrics, extended data retention (15 months), and real-time analysis with CloudWatch Metric Math). Finally, 'Analyze' leads to a box containing four outcomes: Application Monitoring, System-wide Visibility, Resource Optimization, and Unified Operational Health.
5AWS CCP65	1	15.2	Answer = 1 Horizontal scaling Explanation : The only way to increase performance in a vertically scalable system, which can only run its processes on a single computer, is to add more resources to that computer, such as faster (or more) CPUs, memory, or storage.