

QUESTIONNAIRE REPORT SHEET

iPQ1AWS CCP.V1.1

iPQ.No	Correct Answer	Points	Explanation
1AWS CCP01	3	15.4	Answer = 3 Both A and B Explanation : A security group can only contain allow rules. A NAT gateway is managed by Amazon Web Services. A security group controls incoming and outgoing traffic for your instance by acting as a virtual firewall. Instead of acting at the subnet level, security groups act at the instance level. Allow rules can be specified, but not deny rules. You can define distinct rules for incoming and outgoing traffic. Overview of Security Groups: Security group basics The following are the basic characteristics of security groups for your VPC: - There are quotas on the number of security groups that you can create per VPC, the number of rules that you can add to each security group, and the number of security groups that you can associate with a network interface. For more information, see Amazon VPC quotas. - You can specify allow rules, but not deny rules. - You can specify separate rules for inbound and outbound traffic. - When you create a security group, it has no inbound rules. Therefore, no inbound traffic originating from another host to your instance is allowed until you add inbound rules to the security group. - By default, a security group includes an outbound rule that allows all outbound traffic. You can remove the rule and add outbound rules that allow specific outbound traffic only. If your security group has no outbound rules, no outbound traffic originating from your instance is allowed. - Security groups are stateful — if you send a request from your instance, the response traffic for that request is allowed to flow in regardless of inbound security group rules. Responses to allowed inbound traffic are allowed to flow out, regardless of outbound rules. Note Some types of traffic are tracked differently from other types. For more information, see Connection tracking in the Amazon EC2 User Guide for Linux Instances. - Instances associated with a security group can't talk to each other unless you add rules allowing the traffic (exception: the default security group has these rules by default). - Security groups are associated with network interfaces. After you launch an instance, you can change the security groups that are associated with the instance, which changes the security groups associated with the primary network interface (eth0). You can also specify or change the security groups associated with any other network interface. By default, when you create a network interface, it's associated with the default security group for the VPC, unless you specify a different security group. For more information about network interfaces, see Elastic network interfaces. - When you create a security group, you must provide it with a name and a description. The following rules apply: - Names and descriptions can be up to 255 characters in length. - Names and descriptions are limited to the following characters: a-z, A-Z, 0-9, spaces, and _-./!@#%&*~:;(). - A security group name cannot start with sg- as these indicate a default security group. - A security group name must be unique within the VPC. - A security group can only be used in the VPC that you specify when you create the security group. https://docs.aws.com/vpc/latest/userguide/VPC_SecurityGroups.html A Network Access Control List (NACL) is an optional security layer for your VPC that functions as a subnet-level firewall for controlling traffic into and out of one or more subnets. A network ACL has distinct inbound and outbound rules, with each rule having the ability to either permit or deny traffic. NACL Overview: 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 You can use a network address translation (NAT) gateway or a NAT instance to enable instances in a private subnet to connect to the internet or other AWS services, while preventing the internet from connecting to those instances. The NAT Gateway is managed by AWS, while you manage the NAT Instance.

QUESTIONNAIRE REPORT SHEET

iPQ1AWS CCP.V1.1

iPQ.No	Correct Answer	Points	Explanation																					
1AWS CCP02	2	15.4	Answer = 2 Explanation : Dedicated Host Amazon EC2 Dedicated Hosts allow you to use eligible software licenses from vendors such as Microsoft and Oracle on Amazon EC2. An Amazon EC2 Dedicated Host is a physical server fully dedicated for your use, so you can help address corporate compliance requirements. Exam Alert: Please review the differences between dedicated hosts and dedicated instances Differences between Dedicated Hosts and Dedicated Instances Dedicated Hosts and Dedicated Instances can both be used to launch Amazon EC2 instances onto physical servers that are dedicated for your use. There are no performance, security, or physical differences between Dedicated Instances and instances on Dedicated Hosts. However, there are some differences between the two. The following table highlights some of the key differences between Dedicated Hosts and Dedicated Instances: <table><tr><th></th><th>Dedicated Host</th><th>Dedicated Instance</th></tr><tr><td>Billing</td><td>Per-host billing</td><td>Per-instance billing</td></tr><tr><td>Visibility of sockets, cores, and host ID</td><td>Provides visibility of the number of sockets and physical cores</td><td>No visibility</td></tr><tr><td>Host and instance affinity</td><td>Allows you to consistently deploy your instances to the same physical server over time</td><td>Not supported</td></tr><tr><td>Targeted instance placement</td><td>Provides additional visibility and control over how instances are placed on a physical server</td><td>Not supported</td></tr><tr><td>Automatic instance recovery</td><td>Supported. For more information, see Host recovery.</td><td>Supported</td></tr><tr><td>Bring Your Own License (BYOL)</td><td>Supported</td><td>Not supported</td></tr></table> via https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/dedicated-hosts-overview.html.		Dedicated Host	Dedicated Instance	Billing	Per-host billing	Per-instance billing	Visibility of sockets, cores, and host ID	Provides visibility of the number of sockets and physical cores	No visibility	Host and instance affinity	Allows you to consistently deploy your instances to the same physical server over time	Not supported	Targeted instance placement	Provides additional visibility and control over how instances are placed on a physical server	Not supported	Automatic instance recovery	Supported. For more information, see Host recovery .	Supported	Bring Your Own License (BYOL)	Supported	Not supported
	Dedicated Host	Dedicated Instance																						
Billing	Per-host billing	Per-instance billing																						
Visibility of sockets, cores, and host ID	Provides visibility of the number of sockets and physical cores	No visibility																						
Host and instance affinity	Allows you to consistently deploy your instances to the same physical server over time	Not supported																						
Targeted instance placement	Provides additional visibility and control over how instances are placed on a physical server	Not supported																						
Automatic instance recovery	Supported. For more information, see Host recovery .	Supported																						
Bring Your Own License (BYOL)	Supported	Not supported																						

QUESTIONNAIRE REPORT SHEET

iPQ1AWS CCP.V1.1

iPQ.No	Correct Answer	Points	Explanation
1AWS CCP03	3	15.4	Answer = 3 Consulting Partner of APN Explanation : The AWS Partner Network (APN) is a global partner program for technology and consulting firms that use Amazon Web Services to develop customer solutions and services. APN Consulting Partners are professional services firms that assist customers of all sizes in designing, architecting, building, migrating, and managing their workloads and applications on AWS, thereby accelerating their cloud migration. APN Partner Types Overview: APN Partner Types APN Consulting Partners APN Consulting Partners are professional services firms that help customers of all types and sizes design, architect, build, migrate, and manage their workloads and applications on AWS, accelerating their journey to the cloud. APN Consulting Partners often implement Technology Partner solutions in addition to the professional services they offer. APN Consulting Partners include system integrators, strategic consultancies, agencies, managed service providers, and value-added resellers. Learn more » APN Technology Partners APN Technology Partners provide hardware, connectivity services, or software solutions that are either hosted on, or integrated with, the AWS Cloud. Technology Partner products are often delivered as components to broader AWS customer solutions and can be delivered globally by Consulting Partners through AWS Marketplace, bundled solutions, or directly from APN Technology Partners. APN Technology Partners include original equipment manufacturers (OEMs), semiconductor manufacturers, network carriers, SaaS Providers, and independent software vendors (ISVs). Learn more » via: https://aws.amazon.com/partners/

QUESTIONNAIRE REPORT SHEET

iPQ1AWS CCP.V1.1

iPQ.No	Correct Answer	Points	Explanation
1AWS CCP04	3	15.4	Answer = 3 Both A and B Explanation : By default, AWS Shield Standard is activated for all AWS customers. AWS Shield Advanced can be subscribed to for higher levels of protection against attacks. With Shield Advanced, you have exclusive access to advanced, real-time metrics and reports that provide comprehensive visibility into attacks against your AWS resources. AWS Shield Advanced, with the aid of the DRT (DDoS response team), includes intelligent DDoS attack detection and mitigation for not only network layer (layer 3) and transport layer (layer 4) attacks, but also application layer (layer 7) attacks. AWS Shield Advanced offers enhanced DDoS protection for web applications running on Amazon Elastic Compute Cloud, Amazon CloudFront, Amazon Route 53, and AWS Global Accelerator.
1AWS CCP05	3	15.4	Answer = 3 SDK (Software Developer Kit) for AWS Explanation : SDKs reduce the complexity of coding by providing language-specific APIs for AWS services. For instance, the AWS SDK for JavaScript simplifies the use of AWS Services by providing JavaScript developers with a set of consistent and familiar libraries. It supports API lifecycle considerations including credential management, retries, data marshalling, serialization, and deserialization. AWS SDKs are available in multiple programming languages to facilitate the work of developers utilizing various programming and scripting languages. Therefore, the AWS SDK can facilitate the use of language-specific APIs within an application to access AWS services.

QUESTIONNAIRE REPORT SHEET

iPQ1AWS CCP.V1.1

iPQ.No	Correct Answer	Points	Explanation
1AWS CCP06	3	15.4	Answer = 3 Both A and B Explanation : AWS is responsible for "Security of the Cloud" in accordance with the AWS Shared Responsibility Model. This includes protecting the infrastructure that supports all of the AWS cloud services. This infrastructure consists of the hardware, software, networking, and facilities required to operate AWS cloud services. Replacement of faulty Amazon EC2 hardware: Replacement of faulty Amazon EC2 hardware falls under the infrastructure maintenance "of" the cloud. This is AWS's responsibility. AWS is responsible for maintaining Amazon S3 data in multiple availability zones to ensure its durability. AWS is also responsible for ensuring that data on the AWS Cloud is secure, durable, accessible, and reliable. AWS is responsible for safeguarding the data infrastructure against failures. Overview of the Shared Responsibility Model, accessed from: https://aws.amazon.com/compliance/shared-responsibility-model

QUESTIONNAIRE REPORT SHEET

iPQ1AWS CCP.V1.1

iPQ.No	Correct Answer	Points	Explanation
1AWS CCP07	3	15.4	Answer = 3 Both A and B Explanation : Utilize AWS Professional Services to expedite the migration of the infrastructure. The AWS Professional Services organization is a global team of professionals who can assist you in achieving your desired business outcomes when utilizing the AWS Cloud. AWS Professional Services consultants can augment your team with specialized skills and experience, allowing you to achieve results quickly. Therefore, leveraging AWS Professional Services can expedite the infrastructure migration of the startup. Utilize the AWS Partner Network (APN) to develop a custom infrastructure migration solution. The AWS Partner Network (APN) is a global partner program for technology and consulting firms that use Amazon Web Services to develop customer solutions and services. The startup can collaborate with APN's experts to develop a bespoke infrastructure migration solution.

QUESTIONNAIRE REPORT SHEET

iPQ1AWS CCP.V1.1

iPQ.No	Correct Answer	Points	Explanation
1AWS CCP08	3	15.4	Answer = 3 AWS Artifact Explanation : AWS Artifact AWS Artifact serves as your organization's central repository for compliance-related information. It offers on-demand access to AWS's security and compliance reports, as well as select online agreements. Service Organization Control (SOC) reports, Payment Card Industry (PCI) reports, and certifications from accreditation bodies across geographies and compliance verticals validate the implementation and operational effectiveness of AWS security controls. Different types of AWS Artifact Agreements are available to meet the needs of customers subject to specific regulations. For customers who must comply with the Health Insurance Portability and Accountability Act (HIPAA), for instance, the Business Associate Addendum (BAA) is available. It is not a service, but rather a free, self-service portal that provides on-demand access to AWS' compliance reports.


QUESTIONNAIRE REPORT SHEET

iPQ1AWS CCP.V1.1

iPQ.No	Correct Answer	Points	Explanation
1AWS CCP09	3	15.4	Answer = 3 Both A and B Explanation : The company is not required to invest in hardware infrastructure or physical data center security. Therefore, these expense categories would lead to cost savings. The license fee for the SaaS application, developer salary, and project manager salary would remain unchanged. Exam Alert: Please review the following six benefits of cloud computing: You will likely be asked about the benefits of cloud computing versus a traditional on-premises setup. 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. via - https://docs.aws.amazon.com/whitepapers/latest/aws-overview/six-advantages-of-cloud-computing.html


QUESTIONNAIRE REPORT SHEET

iPQ1AWS CCP.V1.1

iPQ.No	Correct Answer	Points	Explanation
1AWS CCP10	1	15.4	Answer = 1 CloudTrail Explanation : CloudTrail can be used to log, monitor, and retain account activity associated with actions across your AWS infrastructure. CloudTrail provides an event log of your AWS account activity, including actions taken via the AWS Management Console, AWS SDKs, and other AWS services. Using CloudTrail:  Via: https://aws.amazon.com/cloudtrail/
1AWS CCP11	1	15.4	Answer = 1 Utilize the Access Key ID and Secret Access Key to programmatically access AWS resources Explanation : Access keys are permanent credentials for an IAM user or the root user of an AWS account. You can use access keys to sign programmatic requests to the AWS CLI or AWS API (directly or through the AWS SDK) with the AWS CLI or AWS API. The components of an access key are an access key ID and a secret access key. You must use both the access key ID and the secret access key to authenticate your requests. Save the access key ID and secret access key in a secure location when creating an access key pair. The secret access key is only accessible at the time of its creation. You must delete your secret access key and generate a new one if you lose it.


QUESTIONNAIRE REPORT SHEET

iPQ1AWS CCP.V1.1

iPQ.No	Correct Answer	Points	Explanation
1AWS CCP12	3	15.4	Answer = 3 Both A and B Explanation : Amazon S3 is an object storage service that provides superior scalability, data availability, security, and performance. Amazon Elastic File System (Amazon EFS) is a simple, scalable, and fully managed elastic NFS file system that can be used with AWS cloud services and on-premises resources. It is designed to scale on-demand to petabytes without disrupting applications, expanding and contracting automatically as files are added or removed, removing the need to provision and manage capacity to accommodate growth.
1AWS CCP13	1	15.4	Answer = 1 Lambda Explanation : AWS Lambda allows you to execute code without managing or provisioning servers. You only pay for the amount of computing time you use. With Lambda, you can execute code for virtually any application or backend service without any administration. Simply upload your code, and Lambda will handle everything necessary to run and scale it with high availability. Lambda's Operation:  via https://aws.amazon.com/lambda/

QUESTIONNAIRE REPORT SHEET

iPQ1AWS CCP.V1.1

iPQ.No	Correct Answer	Points	Explanation
1AWS CCP14	1	15.4	Answer = 1 Configuration Management Explanation : AWS and the customer share responsibility for security and compliance. AWS operates, manages, and controls the components, from the host operating system and virtualization layer to the physical security of the service's operating facilities, in this shared model. Shared controls are controls that apply to both the infrastructure layer and the customer layer, but in completely distinct contexts or perspectives. In a shared control environment, AWS provides the infrastructure requirements and the customer is responsible for implementing control for their use of AWS services. AWS is responsible for maintaining the configuration of its infrastructure devices, while the customer is responsible for configuring their own guest operating systems, databases, and applications. Overview of the Shared Responsibility Model, accessed from:  The diagram illustrates the Shared Responsibility Model as a stack of layers. 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, the layers are detailed as follows: CUSTOMER DATA (Blue block) PLATFORM, APPLICATIONS, IDENTITY & ACCESS MANAGEMENT (Blue block) OPERATING SYSTEM, NETWORK & FIREWALL CONFIGURATION (Blue block) CLIENT-SIDE DATA ENCRYPTION & DATA INTEGRITY AUTHENTICATION (Blue block) SERVER-SIDE ENCRYPTION (FILE SYSTEM AND/OR DATA) (Blue block) NETWORKING TRAFFIC PROTECTION (ENCRYPTION, INTEGRITY, IDENTITY) (Blue block) SOFTWARE (Orange block, containing sub-layers): - COMPUTE - STORAGE - DATABASE - NETWORKING HARDWARE/AWS GLOBAL INFRASTRUCTURE (Orange block, containing sub-layers): - REGIONS - AVAILABILITY ZONES - EDGE LOCATIONS https://aws.amazon.com/compliance/shared-responsibility-model

QUESTIONNAIRE REPORT SHEET

iPQ1AWS CCP.V1.1

iPQ.No	Correct Answer	Points	Explanation				
1AWS CCP15	3	15.4	Answer = 3 Both A and B Explanation : AWS has the concept of a "region," which is a physical location around the globe where AWS clusters its data centers. AWS refers to each group of logical data centers as an Availability Zone. Each AWS region consists of a minimum of three isolated and physically separate AZs within a geographic area. An availability zone (AZ) is one or more data centers with redundant power, networking, and connectivity in an AWS region. Overview of AWS Regions and Availability Zones: <table><tr><th>Regions</th><th>Availability Zones</th></tr><tr><td> AWS has the concept of a Region, which is a physical location around the world where we cluster data centers. We call each group of logical data centers an Availability Zone. Each AWS Region consists of a minimum of three, isolated, and physically separate AZs within a geographic area. Unlike other cloud providers, who often define a region as a single data center, the multiple AZ design of every AWS Region offers advantages for customers. Each AZ has independent power, cooling, and physical security and is connected via redundant, ultra-low-latency networks. AWS customers focused on high availability can design their applications to run in multiple AZs to achieve even greater fault-tolerance. AWS infrastructure Regions meet the highest levels of security, compliance, and data protection. AWS provides a more extensive global footprint than any other cloud provider, and to support its global footprint and ensure customers are served across the world, AWS opens new Regions rapidly. AWS maintains multiple geographic Regions, including Regions in North America, South America, Europe, China, Asia Pacific, South Africa, and the Middle East. </td><td> An Availability Zone (AZ) is one or more discrete data centers with redundant power, networking, and connectivity in an AWS Region. AZs give customers the ability to operate production applications and databases that are more highly available, fault tolerant, and scalable than would be possible from a single data center. All AZs in an AWS Region are interconnected with high-bandwidth, low-latency networking, over fully redundant, dedicated metro fiber providing high-throughput, low-latency networking between AZs. All traffic between AZs is encrypted. The network performance is sufficient to accomplish synchronous replication between AZs. AZs make partitioning applications for high availability easy. If an application is partitioned across AZs, companies are better isolated and protected from issues such as power outages, lightning strikes, tornadoes, earthquakes, and more. AZs are physically separated by a meaningful distance, many kilometers, from any other AZ, although all are within 100 km (60 miles) of each other. Show less </td></tr></table> https://aws.amazon.com/about-aws/global-infrastructure/regions_az	Regions	Availability Zones	AWS has the concept of a Region, which is a physical location around the world where we cluster data centers. We call each group of logical data centers an Availability Zone. Each AWS Region consists of a minimum of three, isolated, and physically separate AZs within a geographic area. Unlike other cloud providers, who often define a region as a single data center, the multiple AZ design of every AWS Region offers advantages for customers. Each AZ has independent power, cooling, and physical security and is connected via redundant, ultra-low-latency networks. AWS customers focused on high availability can design their applications to run in multiple AZs to achieve even greater fault-tolerance. AWS infrastructure Regions meet the highest levels of security, compliance, and data protection. AWS provides a more extensive global footprint than any other cloud provider, and to support its global footprint and ensure customers are served across the world, AWS opens new Regions rapidly. AWS maintains multiple geographic Regions, including Regions in North America, South America, Europe, China, Asia Pacific, South Africa, and the Middle East.	An Availability Zone (AZ) is one or more discrete data centers with redundant power, networking, and connectivity in an AWS Region. AZs give customers the ability to operate production applications and databases that are more highly available, fault tolerant, and scalable than would be possible from a single data center. All AZs in an AWS Region are interconnected with high-bandwidth, low-latency networking, over fully redundant, dedicated metro fiber providing high-throughput, low-latency networking between AZs. All traffic between AZs is encrypted. The network performance is sufficient to accomplish synchronous replication between AZs. AZs make partitioning applications for high availability easy. If an application is partitioned across AZs, companies are better isolated and protected from issues such as power outages, lightning strikes, tornadoes, earthquakes, and more. AZs are physically separated by a meaningful distance, many kilometers, from any other AZ, although all are within 100 km (60 miles) of each other. Show less
Regions	Availability Zones						
AWS has the concept of a Region, which is a physical location around the world where we cluster data centers. We call each group of logical data centers an Availability Zone. Each AWS Region consists of a minimum of three, isolated, and physically separate AZs within a geographic area. Unlike other cloud providers, who often define a region as a single data center, the multiple AZ design of every AWS Region offers advantages for customers. Each AZ has independent power, cooling, and physical security and is connected via redundant, ultra-low-latency networks. AWS customers focused on high availability can design their applications to run in multiple AZs to achieve even greater fault-tolerance. AWS infrastructure Regions meet the highest levels of security, compliance, and data protection. AWS provides a more extensive global footprint than any other cloud provider, and to support its global footprint and ensure customers are served across the world, AWS opens new Regions rapidly. AWS maintains multiple geographic Regions, including Regions in North America, South America, Europe, China, Asia Pacific, South Africa, and the Middle East.	An Availability Zone (AZ) is one or more discrete data centers with redundant power, networking, and connectivity in an AWS Region. AZs give customers the ability to operate production applications and databases that are more highly available, fault tolerant, and scalable than would be possible from a single data center. All AZs in an AWS Region are interconnected with high-bandwidth, low-latency networking, over fully redundant, dedicated metro fiber providing high-throughput, low-latency networking between AZs. All traffic between AZs is encrypted. The network performance is sufficient to accomplish synchronous replication between AZs. AZs make partitioning applications for high availability easy. If an application is partitioned across AZs, companies are better isolated and protected from issues such as power outages, lightning strikes, tornadoes, earthquakes, and more. AZs are physically separated by a meaningful distance, many kilometers, from any other AZ, although all are within 100 km (60 miles) of each other. Show less						

QUESTIONNAIRE REPORT SHEET

iPQ1AWS CCP.V1.1

iPQ.No	Correct Answer	Points	Explanation
1AWS CCP16	1	15.4	Answer = 1 Amazon DynamoDB global table support Explanation : Amazon DynamoDB is a fully managed, serverless, key-value NoSQL database designed to run high-performance applications at any scale. DynamoDB provides integrated security, continuous backups, automated multi-region replication, in-memory caching, and data export tools. DynamoDB global tables replicate data across your choice of AWS regions and automatically scale capacity to accommodate your workloads. With global tables, your globally distributed applications can access data locally in the selected regions for sub-millisecond read and write performance.
1AWS CCP17	2	15.4	Answer = 2 Inspector for Amazon.com Explanation : Amazon Inspector is an automated security assessment service that improves the security and compliance of applications deployed on Amazon EC2 instances. Amazon Inspector evaluates applications for exposure, vulnerabilities, and deviations from best practices.

QUESTIONNAIRE REPORT SHEET

iPQ1AWS CCP.V1.1

iPQ.No	Correct Answer	Points	Explanation
1AWS CCP18	2	15.4	Answer = 2 Weighted route weighting Explanation : Amazon Route 53 is a highly available and scalable cloud Domain Name System (DNS) web service. It is designed to provide developers and businesses with a highly dependable and cost-effective way to route end users to Internet applications by translating domain names such as www.example.com into the numeric IP addresses such as 192.0.2.1 that computers use to communicate. Weighted routing enables you to associate multiple resources with a single domain name (example.com) or subdomain name (acme.example.com) and choose how much traffic is routed to each resource. This is useful for load balancing, testing new software, and other purposes. Overview of the Route 53 Routing Policy: Choosing a routing policy PDF Kindle RSS When you create a record, you choose a routing policy, which determines how Amazon Route 53 responds to queries: • Simple routing policy – Use for a single resource that performs a given function for your domain, for example, a web server that serves content for the example.com website. • Failover routing policy – Use when you want to configure active-passive failover. • Geolocation routing policy – Use when you want to route traffic based on the location of your users. • Geoproximity routing policy – Use when you want to route traffic based on the location of your resources and, optionally, shift traffic from resources in one location to resources in another. • Latency routing policy – Use when you have resources in multiple AWS Regions and you want to route traffic to the region that provides the best latency. • Multivalue answer routing policy – Use when you want Route 53 to respond to DNS queries with up to eight healthy records selected at random. • Weighted routing policy – Use to route traffic to multiple resources in proportions that you specify. https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/routing-policy.html

QUESTIONNAIRE REPORT SHEET

iPQ1AWS CCP.V1.1

iPQ.No	Correct Answer	Points	Explanation																																			
1AWS CCP19	3	15.4	Answer = 3 Both A and B Explanation : Enterprise: AWS Enterprise Support provides customers with concierge-like service focused on helping them achieve their outcomes and find success in the cloud. With Enterprise Support, you get 24x7 technical support from high-quality engineers, tools to automatically manage the health of your environment, consultative architectural guidance delivered in the context of your applications and use-cases, and a designated Technical Account Manager. Business: AWS recommends Business Support if you have production workloads on AWS and want 24x7 phone, email, and chat access to technical support and architectural guidance in the context of your specific use cases. You get full access to AWS Trusted Advisor Best Practice Checks. Exam Alert: Please review the distinctions between the developer, business, and enterprise support plans, as you can anticipate at least a couple questions on the exam: <table><thead><tr><th></th><th>Developer</th><th>Business</th><th>Enterprise On-Ramp</th><th>Enterprise</th></tr></thead><tbody><tr><td></td><td>Recommended if you are experimenting or testing in AWS.</td><td>Minimum recommended tier if you have production workloads in AWS</td><td>Recommended if you have production and/or business critical workloads in AWS.</td><td>Recommended if you have business and/or mission critical workloads in AWS.</td></tr><tr><td>AWS Trusted Advisor Best Practice Checks</td><td>Service Quota and basic Security checks</td><td>Full set of checks</td><td>Full set of checks</td><td>Full set of checks</td></tr><tr><td>AWS Trusted Advisor Priority</td><td></td><td></td><td></td><td>Prioritized recommendations curated by your AWS account team</td></tr><tr><td>Enhanced Technical Support</td><td>Business hours** web access to Cloud Support Associates Unlimited cases with 1 primary contact Prioritized responses on AWS re:Post.</td><td>24/7 phone, web, and chat access to Cloud Support Engineers Unlimited cases and unlimited contacts (IAM supported) Prioritized responses on AWS re:Post Access to AWS Support App in Slack</td><td>24/7 phone, web, and chat access to Cloud Support Engineers Unlimited cases and unlimited contacts (IAM supported) Prioritized responses on AWS re:Post Access to AWS Support App in Slack</td><td>24/7 phone, web, and chat access to Cloud Support Engineers Unlimited cases and unlimited contacts (IAM supported) Prioritized responses on AWS re:Post Access to AWS Support App in Slack</td></tr><tr><td>General guidance:</td><td>< 24 hours**</td><td>< 24 hours</td><td>< 24 hours</td><td>< 24 hours</td></tr><tr><td>System impaired:</td><td>< 12 hours**</td><td>< 12 hours</td><td>< 12 hours</td><td>< 12 hours</td></tr></tbody></table>		Developer	Business	Enterprise On-Ramp	Enterprise		Recommended if you are experimenting or testing in AWS.	Minimum recommended tier if you have production workloads in AWS	Recommended if you have production and/or business critical workloads in AWS.	Recommended if you have business and/or mission critical workloads in AWS.	AWS Trusted Advisor Best Practice Checks	Service Quota and basic Security checks	Full set of checks	Full set of checks	Full set of checks	AWS Trusted Advisor Priority				Prioritized recommendations curated by your AWS account team	Enhanced Technical Support	Business hours** web access to Cloud Support Associates Unlimited cases with 1 primary contact Prioritized responses on AWS re:Post.	24/7 phone, web, and chat access to Cloud Support Engineers Unlimited cases and unlimited contacts (IAM supported) Prioritized responses on AWS re:Post Access to AWS Support App in Slack	24/7 phone, web, and chat access to Cloud Support Engineers Unlimited cases and unlimited contacts (IAM supported) Prioritized responses on AWS re:Post Access to AWS Support App in Slack	24/7 phone, web, and chat access to Cloud Support Engineers Unlimited cases and unlimited contacts (IAM supported) Prioritized responses on AWS re:Post Access to AWS Support App in Slack	General guidance:	< 24 hours**	< 24 hours	< 24 hours	< 24 hours	System impaired:	< 12 hours**	< 12 hours	< 12 hours	< 12 hours
	Developer	Business	Enterprise On-Ramp	Enterprise																																		
	Recommended if you are experimenting or testing in AWS.	Minimum recommended tier if you have production workloads in AWS	Recommended if you have production and/or business critical workloads in AWS.	Recommended if you have business and/or mission critical workloads in AWS.																																		
AWS Trusted Advisor Best Practice Checks	Service Quota and basic Security checks	Full set of checks	Full set of checks	Full set of checks																																		
AWS Trusted Advisor Priority				Prioritized recommendations curated by your AWS account team																																		
Enhanced Technical Support	Business hours** web access to Cloud Support Associates Unlimited cases with 1 primary contact Prioritized responses on AWS re:Post.	24/7 phone, web, and chat access to Cloud Support Engineers Unlimited cases and unlimited contacts (IAM supported) Prioritized responses on AWS re:Post Access to AWS Support App in Slack	24/7 phone, web, and chat access to Cloud Support Engineers Unlimited cases and unlimited contacts (IAM supported) Prioritized responses on AWS re:Post Access to AWS Support App in Slack	24/7 phone, web, and chat access to Cloud Support Engineers Unlimited cases and unlimited contacts (IAM supported) Prioritized responses on AWS re:Post Access to AWS Support App in Slack																																		
General guidance:	< 24 hours**	< 24 hours	< 24 hours	< 24 hours																																		
System impaired:	< 12 hours**	< 12 hours	< 12 hours	< 12 hours																																		

QUESTIONNAIRE REPORT SHEET

iPQ1AWS CCP.V1.1

iPQ.No	Correct Answer	Points	Explanation
1AWS CCP20	3	15.4	Answer = 3 Both A and B Explanation : In a traditional on-premises environment, you must invest heavily in data centers and servers before you know how you'll use them. With cloud computing, you pay for computing resources only when you use them and only for the amount you use. When you make a capacity decision before deploying an application, you often end up sitting on expensive idle resources or dealing with limited capacity. With cloud computing, these problems are eliminated. You can access as much or as little capacity as you need and scale up and down with only a few minutes' notice.
1AWS CCP21	4	15.4	Answer = 4 All of the above Explanation : Exam Alert: Please check out the following six advantages of cloud computing: You would certainly be asked questions on the advantages of cloud computing compared to a traditional on-premises setup: 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. via: https://docs.aws.amazon.com/whitepapers/latest/aws-overview/six-advantages-of-cloud-computing.html


QUESTIONNAIRE REPORT SHEET

iPQ1AWS CCP.V1.1

iPQ.No	Correct Answer	Points	Explanation
1AWS CCP22	3	15.4	Answer = 3 Vertical scaling allows for fault tolerance Explanation : The processes of a "vertically scalable" system are limited to a single computer. In such systems, the only way to increase performance is by adding faster (or more) CPUs, memory, or storage to a single computer. Vertically scalable systems are incapable of fault tolerance because a single instance is susceptible to failure.
1AWS CCP23	3	15.4	Answer = 3 AWS Shield Explanation : AWS Shield is a managed service that protects applications running on AWS from distributed denial of service (DDoS) attacks. AWS Shield offers always-on detection and automatic inline mitigations that minimize application downtime and latency, so AWS Support is not required for DDoS protection. AWS Shield has two levels: standard and advanced. The AWS Shield Standard provides automatic protections to all AWS customers at no additional cost. AWS Shield Standard protects your website or applications from the most common and frequent network and transport layer DDoS attacks. When using AWS Shield Standard in conjunction with Amazon CloudFront and Amazon Route 53, you receive comprehensive protection against all known infrastructure (Layer 3 and 4) attacks. You can subscribe to AWS Shield Advanced for enhanced protection against attacks aimed at your applications running on Amazon Elastic Compute Cloud (EC2), Elastic Load Balancing (ELB), Amazon CloudFront, AWS Global Accelerator, and Amazon Route 53 resources. AWS Shield Advanced provides additional detection and mitigation against large and sophisticated DDoS attacks, visibility into attacks in near real-time, and integration with AWS WAF, a web application firewall.

QUESTIONNAIRE REPORT SHEET

iPQ1AWS CCP.V1.1

iPQ.No	Correct Answer	Points	Explanation
1AWS CCP24	3	15.4	Answer = 3 AWS X-Ray Explanation : AWS X-Ray: AWS X-Ray can be used to analyze and debug serverless and distributed applications, such as those built with a microservices architecture. With X-Ray, you can comprehend how your application and its underlying services are performing in order to identify and troubleshoot performance issues and errors. How X-Rays Work:  The diagram illustrates the AWS X-Ray workflow in four steps: AWS X-Ray: AWS X-Ray helps you analyze and debug modern applications built using microservices and serverless architecture and quantify customer impact. Collect traces: Collect data about the request from each of the underlying application services it passes through. Record traces: X-Ray combines the data gathered from each service into singular units called traces. View service map: View the service map to see trace data such as latencies, HTTP statuses, and metadata for each service. Analyze issues: Drill into the service showing unusual behavior to identify the root issue. via:http://aws.amazon.com/xray/

QUESTIONNAIRE REPORT SHEET

iPQ1AWS CCP.V1.1

iPQ.No	Correct Answer	Points	Explanation
1AWS CCP25	4	15.4	Answer = 4 Inquire with the AWS Abuse Team. Explanation : The AWS Abuse team can assist you when AWS resources are used to engage in abusive behavior. Please see details of the various scenarios that the AWS Abuse team can address: How do I report abuse of AWS resources? Last updated: 2020-04-30 I suspect that Amazon Web Services (AWS) resources are used for abusive or illegal purposes. How do I let AWS know? Resolution The AWS Abuse team can assist you when AWS resources are used to engage in the following types of abusive behavior: • Spam: You are receiving unwanted emails from an AWS-owned IP address, or AWS resources are used to spam websites or forums. • Port scanning: Your logs show that one or more AWS-owned IP addresses are sending packets to multiple ports on your server, and you believe this is an attempt to discover unsecured ports. • Denial-of-service (DoS) attacks: Your logs show that one or more AWS-owned IP addresses are used to flood ports on your resources with packets, and you believe that this is an attempt to overwhelm or crash your server or the software running on your server. • Intrusion attempts: Your logs show that one or more AWS-owned IP addresses are used to attempt to log in to your resources. • Hosting objectionable or copyrighted content: You have evidence that AWS resources are used to host or distribute illegal content or distribute copyrighted content without the consent of the copyright holder. • Distributing malware: You have evidence that AWS resources are used to distribute software that was knowingly created to compromise or cause harm to computers or machines on which it is installed. via:https://aws.amazon.com/premiumsupport/knowledge-center/report-aws-abuse/

QUESTIONNAIRE REPORT SHEET

iPQ1AWS CCP.V1.1

iPQ.No	Correct Answer	Points	Explanation
1AWS CCP26	4	15.4	Answer = 4 Amazon Transcribe can be used to convert speech to text for subsequent analysis. Then, utilize Amazon Polly to speak the text results Explanation : Amazon Transcribe can be used to add speech-to-text functionality to your applications. Amazon Transcribe uses a deep learning technique known as automatic speech recognition (ASR) to accurately convert speech to text quickly. Amazon Transcribe is capable of transcribing customer service calls, automating closed captioning and subtitling, and generating metadata for media assets. Amazon Transcribe Use-Cases: Improving Customer Service By converting audio input into text, Amazon Transcribe helps you build text analytics applications that can search and analyze voice input. Customer contact centers can use Amazon Transcribe to transcribe calls, and mine the data for insights using other AWS services like Amazon Comprehend to extract meaning and intent from conversations. Captioning & Subtitling Workflows Amazon Transcribe can help content producers and media distributors improve reach and accessibility by automatically generating time-stamped subtitles that can be displayed along with the video content. By combining this text with Amazon Translate, you can also easily localize videos. Cataloging Audio Archives You can use Amazon Transcribe to transcribe audio and video assets into fully searchable archives for compliance monitoring and risk management. Convert audio to text and use Amazon Elasticsearch to index and search across your audio/video library. via: https://aws.amazon.com/transcribe/ Amazon Polly can convert text into natural-sounding speech, allowing you to create talking applications. The Text-to-Speech (TTS) service provided by Polly uses advanced deep learning technologies to synthesize human speech that sounds natural. Benefits of Amazon Polly: Natural sounding voices Amazon Polly provides dozens of languages and a wide selection of natural-sounding male and female voices. Amazon Polly's fluid pronunciation of text enables you to deliver high-quality voice output for a global audience. Store & redistribute speech Amazon Polly allows for unlimited replays of generated speech without any additional fees. You can create speech files in standard formats like MP3 and OGG, and serve them from the cloud or locally with apps or devices for offline playback. Real-time streaming Delivering lifelike voices and conversational user experiences requires consistently fast response times. When you send text to Amazon Polly's API, it returns the audio to your application as a stream so you can play the voices immediately. Customize & control speech output Modify Amazon Polly voices to best suit your needs – Amazon Polly supports lexicons and SSML tags which enable you to control aspects of speech, such as pronunciation, volume, pitch, speed rate, etc. Low cost Amazon Polly's pay-as-you-go pricing, low cost per character converted, and unlimited replays make it a cost-effective way to voice your applications.


QUESTIONNAIRE REPORT SHEET

iPQ1AWS CCP.V1.1

iPQ.No	Correct Answer	Points	Explanation
1AWS CCP27	1	15.4	Answer = 1 60 seconds Explanation : The minimum charge for Linux-based EC2 instances is one minute, so this is the correct option.

QUESTIONNAIRE REPORT SHEET

iPQ1AWS CCP.V1.1

iPQ.No	Correct Answer	Points	Explanation
1AWS CCP28	4	15.4	Answer = 4 All of the above Explanation : There are three best practices for cloud-based Reliability: Foundations, Change Management, and Failure Management. Being aware of how change affects a system (change management) enables proactive planning, while monitoring enables rapid identification of trends that could lead to capacity issues or SLA violations. AWS Config is a service for assessing, auditing, and evaluating the configurations of your AWS resources. Config continuously monitors and records the configurations of your AWS resources and enables you to automate the comparison of recorded configurations with desired configurations. Using AWS Config:  https://aws.amazon.com/config/ AWS CloudTrail is a service that facilitates governance, compliance, operational auditing, and risk auditing of your AWS account. CloudTrail allows you to log, continuously monitor, and store account activity related to actions across your AWS infrastructure. CloudTrail provides an event log of your AWS account activity, including actions taken via the AWS

QUESTIONNAIRE REPORT SHEET

iPQ1AWS CCP.V1.1

iPQ.No	Correct Answer	Points	Explanation
1AWS CCP29	3	15.4	Answer = 3 Both A and B Explanation : A VPC endpoint enables you to connect your VPC privately to supported AWS services and VPC endpoint services powered by AWS PrivateLink without requiring an internet gateway, NAT device, VPN connection, or AWS Direct Connect connection. Instances in your VPC do not need public IP addresses in order to communicate with service resources. Between your VPC and the other service, traffic never leaves the Amazon network. Interface endpoints and gateway endpoints are the two kinds of VPC endpoints. An interface endpoint is an elastic network interface with a private IP address from your subnet's IP address range that serves as the entry point for traffic bound for a supported service. AWS PrivateLink, a technology that enables private access to services via private IP addresses, powers interface endpoints. A gateway endpoint is a gateway that you specify as the target of a route in your route table for traffic destined for an AWS service. The subsequent Amazon Web Services are supported: Amazon S3 DynamoDB Exam Alert:


QUESTIONNAIRE REPORT SHEET

iPQ1AWS CCP.V1.1

iPQ.No	Correct Answer	Points	Explanation
1AWS CCP30	4	15.4	Answer = 4 This data transfer would be free of charge for the company Explanation : AWS has three primary cost drivers: computation, storage, and outbound data transfer. In most cases, inbound data transfer and data transfer between AWS services within the same region are free. The outbound data transfer across all services is aggregated and then billed at the outbound data transfer rate. According to AWS pricing, data transfer between S3 and EC2 instances within the same region is free, so moving 500 GB of data from an EC2 instance to an S3 bucket in the same region would incur no data transfer fees.
1AWS CCP31	1	15.4	Answer = 1 CloudTrail Log File Explanation : The AWS CloudTrail service enables governance, compliance, operational auditing, and risk auditing of your AWS account. CloudTrail can be used to record AWS API calls and other account activity and save the information to log files in an Amazon Simple Storage Service (Amazon S3) bucket you specify. By default, CloudTrail encrypts the log files it delivers to your S3 bucket using server-side encryption with Amazon S3-managed encryption keys (SSE-S3).

QUESTIONNAIRE REPORT SHEET

iPQ1AWS CCP.V1.1

iPQ.No	Correct Answer	Points	Explanation
1AWS CCP32	2	15.4	Answer = 2 Physical and ecological controls Explanation : As part of the shared responsibility model, physical and environmental controls are inherited controls, and therefore fall under AWS's purview. Overview of the Shared Responsibility Model, accessed from:  The diagram illustrates the Shared Responsibility Model as a stack of layers. 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, the layers are stacked from top to bottom: 'CUSTOMER DATA' (blue), 'PLATFORM, APPLICATIONS, IDENTITY & ACCESS MANAGEMENT' (blue), 'OPERATING SYSTEM, NETWORK & FIREWALL CONFIGURATION' (blue), a split layer for 'CLIENT-SIDE DATA ENCRYPTION & DATA INTEGRITY AUTHENTICATION' and 'SERVER-SIDE ENCRYPTION (FILE SYSTEM AND/OR DATA)' (blue), 'NETWORKING TRAFFIC PROTECTION (ENCRYPTION, INTEGRITY, IDENTITY)' (blue), 'SOFTWARE' (orange), a split layer for 'COMPUTE', 'STORAGE', 'DATABASE', and 'NETWORKING' (orange), 'HARDWARE/AWS GLOBAL INFRASTRUCTURE' (orange), and a split layer for 'REGIONS', 'AVAILABILITY ZONES', and 'EDGE LOCATIONS' (orange). https://aws.amazon.com/compliance/shared-responsibility-model

QUESTIONNAIRE REPORT SHEET

iPQ1AWS CCP.V1.1

iPQ.No	Correct Answer	Points	Explanation
1AWS CCP33	4	15.4	Answer = 490 Explanation : Amazon EC2 Spot Instances enable you to utilize unused Amazon EC2 cloud capacity. Spot instances are up to 90 percent less expensive than on-demand instances. Spot Instances can be used for a variety of stateless, fault-tolerant, or flexible applications, such as big data, containerized workloads, CI/CD, web servers, and high-performance computing (HPC). 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/

QUESTIONNAIRE REPORT SHEET

iPQ1AWS CCP.V1.1

iPQ.No	Correct Answer	Points	Explanation
1AWS CCP34	3	15.4	Answer = 3 Both A and B Explanation : Credits are applied according to the following sequence: Soonest expiring Fewest number of relevant products oldest account Credit one, which expires in July, is applied to the Amazon EC2 charge for the given use case, leaving you with a \$900 Amazon EC2 charge and a \$500 Amazon S3 charge. The second credit is then applied to the remaining Amazon EC2 usage balance of \$900. Amazon EC2 costs \$850 per month and Amazon S3 costs \$500 per month. All of your credits have been used up.
1AWS CCP35	4	15.4	Answer = 4 Standard-Infrequent Access from Amazon S3 (S3 Standard-IA) Explanation : S3 Standard-IA is designed for data that is accessed infrequently but requires rapid access when necessary. S3 Standard-IA provides the same high durability, high throughput, and low latency as S3 Standard, but at a lower cost per GB for storage and retrieval. This combination of low cost and high performance makes S3 Standard-IA ideally suited for long-term storage, backups, and as a data repository for disaster recovery files.

QUESTIONNAIRE REPORT SHEET

iPQ1AWS CCP.V1.1

iPQ.No	Correct Answer	Points	Explanation
1AWS CCP36	1	15.4	Answer = 1 AWS Budgets Explanation : AWS Budgets enables you to create custom alerts when your costs or usage exceed (or are projected to exceed) your budgeted amount. You can also use AWS Budgets to set reservation utilization or coverage targets and receive alerts when your reservation utilization falls below the defined threshold. Amazon EC2, Amazon RDS, Amazon Redshift, Amazon ElasticCache, and Amazon Elasticsearch reservations support reservation alerts.
1AWS CCP37	4	15.4	Answer = 4 AWS Shield Standard Explanation : AWS Shield Standard protects your website or applications from the most common and frequent network and transport layer DDoS attacks. AWS Shield Standard protects all AWS customers, but Amazon CloudFront and Amazon Route 53 provide enhanced security. The AWS Shield Standard provides automatic protections to all AWS customers at no additional cost.

QUESTIONNAIRE REPORT SHEET

iPQ1AWS CCP.V1.1

iPQ.No	Correct Answer	Points	Explanation
1AWS CCP38	1	15.4	Answer = 1 Utilize AWS Organizations to administer the AWS accounts of all units, and then distribute the reserved EC2 instances to all units Explanation : AWS Organizations enables the centralized management of billing, access, compliance, and security, as well as the sharing of resources across multiple AWS accounts. Using AWS Organizations, you can automate account creation, create groups of accounts that reflect your business requirements, and apply policies for governance to these groups. Additionally, you can streamline billing by configuring a single payment method for all of your AWS accounts. AWS Organizations is available at no additional cost to all AWS customers. Key AWS Organizations Features: CENTRALLY MANAGE POLICIES ACROSS MULTIPLE AWS ACCOUNTS To improve control over your AWS environment, you can use AWS Organizations to create groups of accounts, and then attach policies to a group to ensure the correct policies are applied across the accounts without requiring custom scripts and manual processes. GOVERN ACCESS TO AWS SERVICES, RESOURCES, AND REGIONS AWS Organizations allows you to restrict what services and actions are allowed in your accounts. You can use Service Control Policies (SCPs) to apply permission guardrails on AWS Identity and Access Management (IAM) users and roles. For example, you can apply an SCP that restricts users in accounts in your organization from launching any resources in regions that you do not explicitly allow. AUTOMATE AWS ACCOUNT CREATION AND MANAGEMENT AWS Organizations helps you simplify IT operations by automating AWS account creation and management. The Organizations APIs enable you to create new accounts programmatically, and to add the new accounts to a group. The policies attached to the group are automatically applied to the new account. For example, you can automate the creation of new accounts for workload or application isolation and grant entities in those accounts access only to the necessary AWS services. CONFIGURE AWS SERVICES ACROSS MULTIPLE ACCOUNTS AWS Organizations helps you configure AWS services and share resources across accounts in your organization. For example, Organizations integrates with AWS Single Sign-on to enable you to easily provision access for all of your developers to accounts in your organization from a single place. You can make central changes to access permissions and have them automatically updated on accounts in your organization. CONSOLIDATE BILLING ACROSS MULTIPLE AWS ACCOUNTS You can use AWS Organizations to set up a single payment method for all the AWS accounts in your organization through consolidated billing. With consolidated billing, you can see a combined view of charges incurred by all your accounts, as well as take advantage of pricing benefits from aggregated usage, such as volume discounts for Amazon EC2 and Amazon S3.

QUESTIONNAIRE REPORT SHEET

iPQ1AWS CCP.V1.1

iPQ.No	Correct Answer	Points	Explanation
1AWS CCP39	1	15.4	Answer = 1 Configure the database for automatic failover to the standby in an RDS Multi-AZ deployment Explanation : Amazon RDS automatically creates a primary DB instance and synchronously replicates the data to a standby instance in a different availability zone (AZ) when you provision a multi-AZ DB instance. Amazon RDS performs an automatic failover to the standby (or to a read replica in the case of Amazon Aurora) in the event of an infrastructure failure, allowing you to resume database operations as soon as the failover is complete. Since the endpoint for your DB instance remains unchanged following a failover, your application can resume database operations without requiring administrative intervention.


QUESTIONNAIRE REPORT SHEET

iPQ1AWS CCP.V1.1

iPQ.No	Correct Answer	Points	Explanation
1AWS CCP40	1	15.4	Answer = 1 IaaS is Infrastructure as a Service Explanation : Three types of cloud computing exist: infrastructure as a service (IaaS), platform as a service (PaaS), and software as a service (SaaS). IaaS comprises the fundamental components of cloud IT. Typically, it offers access to networking capabilities, computers (virtual or on dedicated hardware), and data storage space. IaaS offers the greatest degree of IT resource flexibility and management control. EC2 grants you complete management authority over the underlying operating system, virtual network configurations, storage, data, and applications. Consequently, EC2 is an IaaS service. Please review this summary of cloud computing types: 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/.


QUESTIONNAIRE REPORT SHEET

iPQ1AWS CCP.V1.1

iPQ.No	Correct Answer	Points	Explanation
1AWS CCP41	2	15.4	Answer = 2 Enterprise Explanation : AWS offers developer, business, and enterprise support plans to meet the needs of each customer. Every AWS customer receives a basic support plan. AWS Enterprise Support provides clients with concierge-like service with a primary focus on assisting clients in achieving their goals and achieving success in the cloud. With Enterprise Support, you have access to online training with self-paced labs, 24x7 technical support from top-tier engineers, tools and technology to automatically manage the health of your environment, consultative architectural guidance, a Technical Account Manager (TAM) to coordinate access to proactive and preventative programs, and AWS subject matter experts. AWS Enterprise Support Plan Services:  The diagram illustrates the AWS Enterprise Support Plan Services as interconnected gears. The central gears are labeled 'PEOPLE EXPERTISE', 'TOOLS AUTOMATION', and 'STRATEGIC PROGRAMS'. Surrounding these are various support services: TECHNICAL ACCOUNT MANAGER (TAM): Designated technical point of contact to all necessary AWS expertise. TRUSTED ADVISOR (TA): Online resource to help you reduce cost, increase performance, and improve security by optimizing your AWS environment. PERSONAL HEALTH DASHBOARD (PHD): Delivers alerts and remediation guidance when AWS is experiencing events that may impact your environment. SUPPORT API: Programmatic access to AWS Support Center features to create, manage, and close your support cases, and operationally manage your TAM check requests and status. TRAINING: Credits for online self-paced labs provided through an AWS training provider. PEOPLE EXPERTISE: Cloud Support Engineers, Solutions Architects, and product teams are available for guidance. SUPPORT CONCIERGE: Dedicated team of enterprise account specialists to help with billing and account subjects. ABUSE TEAM: Assists you when AWS resources are impacted by things such as Spam, Port scanning, Denial of Service attacks (DDoS), or malware. INFRASTRUCTURE EVENT MANAGEMENT (IEM): Focused planning and support business-critical events (e.g. launches or migrations). WELL-ARCHITECTED REVIEW: Detailed review of your architecture guidance on how to best design your systems. ARCHITECTURE SUPPORT: Consultative reviews of your application architecture and how to align it with AWS. OPERATIONS SUPPORT: Consultative reviews of your cloud operations and advice for optimization. http://aws.amazon.com/premiumsupport/plans/enterprise/

QUESTIONNAIRE REPORT SHEET

iPQ1AWS CCP.V1.1

iPQ.No	Correct Answer	Points	Explanation
1AWS CCP42	3	15.4	Answer = 3 Instance Retail Explanation : An instance store provides temporary storage at the block level for an instance. This data is stored on disks physically connected to the host computer. This is a viable option when you require storage with very low latency but do not require the data to persist after the instance terminates, or when you can take advantage of fault-tolerant architectures. For this use case, the computation application itself has a fault-tolerant architecture, allowing it to handle any failures of instance store volumes automatically. Since the instance store volumes are included in the instance's usage cost, this is the correct answer. Amazon EC2 Instance Store PDF Kindle RSS An <i>instance store</i> provides temporary block-level storage for your instance. This storage is located on disks that are physically attached to the host computer. Instance store is ideal for temporary storage of information that changes frequently, such as buffers, caches, scratch data, and other temporary content, or for data that is replicated across a fleet of instances, such as a load-balanced pool of web servers. An instance store consists of one or more instance store volumes exposed as block devices. The size of an instance store as well as the number of devices available varies by instance type. The virtual devices for instance store volumes are <code>ephemeral[0-23]</code>. Instance types that support one instance store volume have <code>ephemeral0</code>. Instance types that support two instance store volumes have <code>ephemeral0</code> and <code>ephemeral1</code>, and so on.  via: https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/InstanceStorage.html


QUESTIONNAIRE REPORT SHEET

iPQ1AWS CCP.V1.1

iPQ.No	Correct Answer	Points	Explanation
1AWS CCP43	4	15.4	Answer = 4 Business Explanation : AWS recommends Business Support if you have production workloads on AWS and want 24x7 phone, email, and chat access to technical support and architectural guidance for your specific use cases. You have access to AWS Trusted Advisor Best Practice Checks in their entirety. In addition, you have access to infrastructure event management for a fee.

QUESTIONNAIRE REPORT SHEET

iPQ1AWS CCP.V1.1

iPQ.No	Correct Answer	Points	Explanation
1AWS CCP44	3	15.4	Answer = 3 Both A and B Explanation : The basic plan only grants access to the following features: Customer Service and Communities: Access to customer service, documentation, whitepapers, and support forums 24 hours a day, seven days a week AWS Trusted Advisor: Access to the core Trusted Advisor checks and guidance to provision your resources in accordance with best practices to enhance performance and security. AWS Health: Your Account Health Dashboard: A customized view of the health of your AWS services and notifications when your resources are impacted. AWS recommends Developer Support if you are testing or conducting early development on AWS and would like to receive email-based technical support during business hours as well as general architectural guidance as you build and test. This plan provides access to only the core trusted advisor checks and basic security checks from the service quota. Exam Alert: Please review the distinctions between the developer, business, and enterprise support plans, as you can anticipate at least a couple questions on the exam:  The diagram illustrates the progression of AWS support plans. It starts with 'Developer' (recommended for testing), followed by 'Business' (recommended for production workloads), then 'Enterprise On-Ramp' (recommended for production and critical workloads), and finally 'Enterprise' (recommended for mission-critical workloads). A table below shows the 'AWS Trusted Advisor Best Practice Checks' available for each plan: Developer has 'Service Quota and basic Security checks'; Business has 'Full set of checks'; Enterprise On-Ramp has 'Full set of checks'; and Enterprise has 'Full set of checks'. A note at the bottom right states: 'Prioritized recommendations curated by your AWS account team'.

QUESTIONNAIRE REPORT SHEET

iPQ1AWS CCP.V1.1

iPQ.No	Correct Answer	Points	Explanation																																										
1AWS CCP45	4	15.4	Answer = 4 Option for partial upfront payment with a standard three-year term Explanation : You can reserve capacity with Amazon EC2 Reserved Instances and receive a discount on your instance usage in comparison to running on-demand instances. The discounted usage price is reserved for the duration of your contract, allowing you to predict computed costs over the reserved instance's duration. Please review this EC2 Reserved Instances pricing comparison: Standard 1-Year Term <table><tr><th>Payment Option</th><th>Upfront</th><th>Monthly*</th><th>Effective Hourly</th><th>Savings over On-Demand</th><th>On-Demand Hourly</th></tr><tr><td>No Upfront</td><td>\$0</td><td>\$44.53</td><td>\$0.061</td><td>36%</td><td rowspan="3">\$0.096 per Hour</td></tr><tr><td>Partial Upfront</td><td>\$256</td><td>\$21.17</td><td>\$0.058</td><td>39%</td></tr><tr><td>All Upfront</td><td>\$501</td><td>\$0</td><td>\$0.057</td><td>40%</td></tr></table> Standard 3-Year Term <table><tr><th>Payment Option</th><th>Upfront</th><th>Monthly*</th><th>Effective Hourly</th><th>Savings over On-Demand</th></tr><tr><td>No Upfront</td><td>\$0</td><td>\$30.66</td><td>\$0.042</td><td>56%</td></tr><tr><td>Partial Upfront</td><td>\$515</td><td>\$14.60</td><td>\$0.040</td><td>59%</td></tr><tr><td>All Upfront</td><td>\$968</td><td>\$0</td><td>\$0.037</td><td>62%</td></tr></table> https://d0.awsstatic.com/whitepapers/aws_pricing_overview.pdf Therefore, the following are the percentage savings for each option: "There is no option for an upfront payment with the standard 1-year term." - 36% "Option for upfront payment with the standard one-year term" - 40%	Payment Option	Upfront	Monthly*	Effective Hourly	Savings over On-Demand	On-Demand Hourly	No Upfront	\$0	\$44.53	\$0.061	36%	\$0.096 per Hour	Partial Upfront	\$256	\$21.17	\$0.058	39%	All Upfront	\$501	\$0	\$0.057	40%	Payment Option	Upfront	Monthly*	Effective Hourly	Savings over On-Demand	No Upfront	\$0	\$30.66	\$0.042	56%	Partial Upfront	\$515	\$14.60	\$0.040	59%	All Upfront	\$968	\$0	\$0.037	62%
Payment Option	Upfront	Monthly*	Effective Hourly	Savings over On-Demand	On-Demand Hourly																																								
No Upfront	\$0	\$44.53	\$0.061	36%	\$0.096 per Hour																																								
Partial Upfront	\$256	\$21.17	\$0.058	39%																																									
All Upfront	\$501	\$0	\$0.057	40%																																									
Payment Option	Upfront	Monthly*	Effective Hourly	Savings over On-Demand																																									
No Upfront	\$0	\$30.66	\$0.042	56%																																									
Partial Upfront	\$515	\$14.60	\$0.040	59%																																									
All Upfront	\$968	\$0	\$0.037	62%																																									

QUESTIONNAIRE REPORT SHEET

iPQ1AWS CCP.V1.1

iPQ.No	Correct Answer	Points	Explanation
1AWS CCP46	2	15.4	Answer = 2 A customer-managed database instance performs less efficiently than an AWS-managed RDS instance Explanation : Amazon Relational Database Service (Amazon RDS) simplifies the deployment, management, and scalability of cloud-based relational databases. It automates time-consuming administrative tasks such as hardware provisioning, database setup, patching, and backups. It offers scalable and cost-effective capacity. Amazon RDS offers a variety of instance types that are optimized for various relational database use cases. Instance types consist of varying combinations of CPU, memory, storage, and networking capacity, allowing you to select the optimal mix of resources for your database. By selecting the appropriate instance type and size, you can optimize the database for your use case. Due to the fact that RDS instances are optimized for memory, performance, and I/O, the performance of an AWS-managed RDS instance is superior to that of a customer-managed database instance.
1AWS CCP47	3	15.4	Answer = 3 Layer 7 Explanation : AWS WAF is a web application firewall that enables you to monitor HTTP and HTTPS requests forwarded to Amazon API Gateway APIs, Amazon CloudFront, and application load balancers. HTTP and HTTPS requests are part of layer 7, the application layer.

QUESTIONNAIRE REPORT SHEET

iPQ1AWS CCP.V1.1

iPQ.No	Correct Answer	Points	Explanation
1AWS CCP48	1	15.4	Answer = 1 EBS volumes may be attached to a single instance within the same availability zone, whereas EFS file systems may be mounted on instances spanning multiple availability zones Explanation : Amazon Elastic File System (Amazon EFS) is a simple, scalable, and fully managed elastic NFS file system that can be used with AWS cloud services and on-premises resources. It is designed to scale on-demand to petabytes without disrupting applications, expanding and contracting automatically as files are added or removed, removing the need to provision and manage capacity to accommodate growth. The service is designed to be extremely scalable, highly available, and extremely resilient. Amazon EFS file systems distribute data and metadata across multiple availability zones within a single AWS region. The EFS file system can be mounted across multiple availability zones on instances. Amazon Elastic Block Store (EBS) is a high-performance, easy-to-use block storage service designed for use with Amazon Elastic Compute Cloud (EC2) for throughput- and transaction-intensive workloads at any scale. EBS volumes are replicated within an availability zone (AZ) and can easily scale to petabytes of data, making them ideal for mission-critical systems. Attaching an EBS volume to an instance in the same availability zone as the volume is possible.
1AWS CCP49	3	15.4	Answer = 3 Enable client-side encryption using the AWS Encryption Software Development Kit Explanation : Client-side encryption is the process of encrypting data before sending it to Amazon S3. The AWS encryption SDK is a client-side encryption library that is independent of language-specific SDKs.

QUESTIONNAIRE REPORT SHEET

iPQ1AWS CCP.V1.1

iPQ.No	Correct Answer	Points	Explanation																																																																																								
1AWS CCP50	2	15.4	Answer = 2 Business Explanation : AWS recommends Business Support if you have production workloads on AWS and desire 24x7 phone, email, and chat access to technical support and architectural guidance in the context of your specific use cases. You also gain access to infrastructure event management for an additional fee.																																																																																								
1AWS CCP51	4	15.4	Answer = 4 S3 Glacier Deep Archive Explanation : S3 Glacier Deep Archive is the lowest-cost storage class offered by Amazon S3; it supports long-term retention and digital preservation for data that may be accessed once or twice per year. It is intended for clients, particularly those in highly regulated industries such as financial services, healthcare, and the public sector, who are required to retain data sets for 7–10 years or longer to meet regulatory compliance requirements. Additionally, S3 Glacier Deep Archive can be used for backup and disaster recovery. The retrieval time (first byte latency) ranges between 12 and 48 hours. Please review this diagram for retrieval times for S3 storage class data. You do not need to memorize the exact numbers; simply remember that retrieving data from the S3 Glacier Deep Archive takes the longest. Performance across the S3 Storage Classes <table><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><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></table> via: https://aws.amazon.com/s3/storage-classes/		S3 Standard	S3 Intelligent-Tiering*	S3 Standard-IA	S3 One Zone-IA†	S3 Glacier Instant Retrieval	S3 Glacier Flexible Retrieval	S3 Glacier Deep Archive	Designed for durability	99.999999999% (11 9's)	99.999999999% (11 9's)	99.999999999% (11 9's)	99.999999999% (11 9's)	99.999999999% (11 9's)	99.999999999% (11 9's)	99.999999999% (11 9's)	Designed for availability	99.99%	99.9%	99.9%	99.5%	99.9%	99.99%	99.99%	Availability SLA	99.9%	99%	99%	99%	99%	99.9%	99.9%	Availability Zones	≥3	≥3	≥3	1	≥3	≥3	≥3	Minimum capacity charge per object	N/A	N/A	128 KB	128 KB	128 KB	40 KB	40 KB	Minimum storage duration charge	N/A	N/A	30 days	30 days	90 days	90 days	180 days	Retrieval charge	N/A	N/A	per GB retrieved	per GB retrieved	per GB retrieved	per GB retrieved	per GB retrieved	First byte latency	milliseconds	milliseconds	milliseconds	milliseconds	milliseconds	minutes or hours	hours	Storage type	Object	Object	Object	Object	Object	Object	Object	Lifecycle transitions	Yes	Yes	Yes	Yes	Yes	Yes	Yes
	S3 Standard	S3 Intelligent-Tiering*	S3 Standard-IA	S3 One Zone-IA†	S3 Glacier Instant Retrieval	S3 Glacier Flexible Retrieval	S3 Glacier Deep Archive																																																																																				
Designed for durability	99.999999999% (11 9's)	99.999999999% (11 9's)	99.999999999% (11 9's)	99.999999999% (11 9's)	99.999999999% (11 9's)	99.999999999% (11 9's)	99.999999999% (11 9's)																																																																																				
Designed for availability	99.99%	99.9%	99.9%	99.5%	99.9%	99.99%	99.99%																																																																																				
Availability SLA	99.9%	99%	99%	99%	99%	99.9%	99.9%																																																																																				
Availability Zones	≥3	≥3	≥3	1	≥3	≥3	≥3																																																																																				
Minimum capacity charge per object	N/A	N/A	128 KB	128 KB	128 KB	40 KB	40 KB																																																																																				
Minimum storage duration charge	N/A	N/A	30 days	30 days	90 days	90 days	180 days																																																																																				
Retrieval charge	N/A	N/A	per GB retrieved	per GB retrieved	per GB retrieved	per GB retrieved	per GB retrieved																																																																																				
First byte latency	milliseconds	milliseconds	milliseconds	milliseconds	milliseconds	minutes or hours	hours																																																																																				
Storage type	Object	Object	Object	Object	Object	Object	Object																																																																																				
Lifecycle transitions	Yes	Yes	Yes	Yes	Yes	Yes	Yes																																																																																				

QUESTIONNAIRE REPORT SHEET

iPQ1AWS CCP.V1.1

iPQ.No	Correct Answer	Points	Explanation
1AWS CCP52	3	15.4	Answer = 3 AWS Elastic Load Balancing Explanation : Elastic Load Balancing is utilized to distribute incoming application traffic across all running EC2 instances automatically. Elastic load balancing can be used to manage incoming requests by routing traffic optimally so that no single instance is overwhelmed. Your application's load balancer acts as a single point of contact for all incoming web traffic. When a new instance is added, it must register with the load balancer or it will not receive any traffic. When an instance is deleted, it must deregister from the load balancer; otherwise, it will continue to receive traffic.


QUESTIONNAIRE REPORT SHEET

iPQ1AWS CCP.V1.1

iPQ.No	Correct Answer	Points	Explanation
1AWS CCP53	4	15.4	Answer = 4 All of the above Explanation : The following AWS services allow for cost-optimized reservations: Amazon EC2 Reserved Instances: You can reserve capacity and receive a discount on your instance usage with Amazon EC2 Reserved Instances compared to running on-demand instances. Amazon DynamoDB Reserved Capacity: If you can anticipate your demand for Amazon DynamoDB read-and-write throughput, Reserved Capacity offers substantial savings over the standard price of DynamoDB provisioned throughput capacity. Amazon ElastiCache Reserved Nodes: Amazon ElastiCache Reserved Nodes allow you to make a low, one-time payment for each cache node you wish to reserve and, in exchange, receive a substantial discount on the hourly charge for that node. Similar to Amazon EC2 RIs, Amazon RDS RIs may be acquired under No Upfront, Partial Upfront, or All Upfront terms. Database engines Aurora, MySQL, MariaDB, PostgreSQL, Oracle, and SQL Server offer all reserved instance types. Amazon Redshift Reserved Nodes: If you intend to maintain a running Amazon Redshift cluster for an extended period of time, you should consider purchasing reserved-node offerings. These services provide substantial savings compared to on-demand pricing, but they require you to reserve compute nodes and commit to paying for them for either one or three years.


QUESTIONNAIRE REPORT SHEET

iPQ1AWS CCP.V1.1

iPQ.No	Correct Answer	Points	Explanation
1AWS CCP54	4	15.4	Answer = 4 AWS Compute Optimizer Explanation : AWS Compute Optimizer: AWS Compute Optimizer recommends the optimal AWS resources for your workloads to reduce costs and improve performance by analyzing historical utilization metrics with machine learning. Both overprovisioning and underprovisioning resources can result in unnecessary infrastructure costs and poor application performance, respectively. Compute Optimizer helps you choose optimal configurations for three types of AWS resources based on your utilization data: Amazon EC2 instances, Amazon EBS volumes, and AWS Lambda functions. Compute Optimizer recommends up to three options from more than one hundred and forty EC2 instance types, as well as a variety of EBS volume and Lambda function configuration options, to optimally scale your workloads. Compute Optimizer also predicts your workload's CPU utilization, memory utilization, and runtime based on the recommended AWS resource options. Before implementing the recommended options, this helps you understand how your workload would have performed with them. Compute Optimizer's operation is described at  https://aws.amazon.com/compute-optimizer/.


QUESTIONNAIRE REPORT SHEET

iPQ1AWS CCP.V1.1

iPQ.No	Correct Answer	Points	Explanation
1AWS CCP55	4	15.4	Answer = 4 Amazon Macie Explanation : Amazon Macie is a fully managed data security and privacy service that uses machine learning and pattern matching to identify and protect your sensitive AWS data. Macie automatically provides an inventory of Amazon S3 buckets, including a list of unencrypted buckets, buckets that are publicly accessible, and buckets that are shared with AWS accounts other than those defined in AWS Organizations. Then, Macie uses machine learning and pattern matching techniques to identify and alert you of sensitive data, such as personally identifiable information (PII). How Macie Operates:  aws.amazon.com/macie


QUESTIONNAIRE REPORT SHEET

iPQ1AWS CCP.V1.1

iPQ.No	Correct Answer	Points	Explanation
1AWS CCP56	1	15.4	Answer = 1 Direct Connect to AWS Explanation : AWS Direct Connect is a cloud-based service that facilitates the creation of a dedicated network connection between your premises and AWS. Using AWS Direct Connect, you can establish a private virtual interface from your on-premise network directly to your Amazon VPC, giving you a private, high-bandwidth network connection between your network and VPC. This is a private connection that does not traverse the public internet. At least one month is required to establish this physical connection. How Direct Connect Works: via - https://aws.amazon.com/directconnect/ 

QUESTIONNAIRE REPORT SHEET

iPQ1AWS CCP.V1.1

iPQ.No	Correct Answer	Points	Explanation
1AWS CCP57	2	15.4	Answer = 2 EFS Explanation : "EFS": Amazon EFS is a file storage service for use with Amazon EC2. Amazon EFS offers a file system interface, file system access semantics, and concurrently accessible storage for up to a million Amazon EC2 instances. The Amazon EFS protocol is Network File System. Using EFS:  https://aws.amazon.com/efs/
1AWS CCP58	2	15.4	Answer = 2 The AWS account must be capable of operating independently. After that, it can be removed from AWS organizations Explanation : You can only remove a user account from your organization if the account has the necessary information to function independently. You must accept the AWS Customer Agreement, choose a support plan, provide and verify the required contact information, and provide a valid payment method for each account you wish to make independent. AWS charges the payment method for any billable (non-AWS Free Tier) AWS activity that occurs when an account is not associated with an organization.

QUESTIONNAIRE REPORT SHEET

iPQ1AWS CCP.V1.1

iPQ.No	Correct Answer	Points	Explanation
1AWS CCP59	3	15.4	Answer = 3 Both A and B Explanation : Amazon SQS is a fully managed message queuing service that allows you to decouple and scale microservices, distributed systems, and serverless applications. Using SQS, you can send, store, and receive messages between any number of software components without losing messages or requiring the availability of other services. Amazon Simple Notification Service (SNS) is a highly available, durable, secure, and fully managed pub/sub messaging service that enables the decoupling of microservices, distributed systems, and serverless applications. Using Amazon SNS topics, your publisher systems can distribute messages to a large number of subscriber endpoints, such as Amazon SQS queues, AWS Lambda functions, and HTTP/S webhooks, for parallel processing. In addition, SNS can be used to send end-user notifications via mobile push, SMS, and email. Therefore, both SNS and SQS can be utilized to decouple microservices-based application components. Please review this reference architecture for constructing a decoupled order processing system utilizing SNS and SQS:  <pre> graph LR EC2[Amazon EC2] --> Topic[Order Topic] Topic --> SNS[Amazon SNS] SNS --> DevQueue[Developer Queue] SNS --> OrderQueue[Order Dispatcher Queue] SNS --> CustQueue[Customer Order Queue] DevQueue --> Devs[Developers] OrderQueue --> OrderDispatcher[Order Dispatcher] CustQueue --> OrderDispatcher OrderDispatcher --> PriorityQueue[Priority Order Queue] OrderDispatcher --> CustQueue PriorityQueue --> AS1[Auto Scaling] CustQueue --> AS2[Auto Scaling] </pre>

QUESTIONNAIRE REPORT SHEET

iPQ1AWS CCP.V1.1

iPQ.No	Correct Answer	Points	Explanation
1AWS CCP60	1	15.4	Answer = 1 Redshift Explanation : Amazon Redshift Redshift is a fully managed petabyte-scale cloud-based data warehouse product designed for storing and analyzing large-scale data sets.
1AWS CCP61	1	15.4	Answer = 1 Reserved Occurrence Explanation : Reserved Instance: Reserved Instances offer significant Amazon EC2 cost savings (up to 75%) compared to On-Demand Instance pricing. Reserved instances are not actual instances, but rather a discount applied to your account's usage of on-demand instances. You can purchase a reserved instance for either one or three years, with the three-year option offering a greater discount. Reserved instances are impervious to interruption. Therefore, this is the best choice. 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/

QUESTIONNAIRE REPORT SHEET

iPQ1AWS CCP.V1.1

iPQ.No	Correct Answer	Points	Explanation
1AWS CCP62	4	15.4	Answer = 4 Agility Explanation : In the world of cloud computing, "agile" refers to the capacity to rapidly develop, test, and launch business-growth-driving software applications. A further explanation of "agile" is that AWS provides a massive global cloud infrastructure that enables rapid innovation, experimentation, and iteration. Instead of waiting weeks or months for hardware, new applications can be deployed immediately. This quality is known as agility.
1AWS CCP63	1	15.4	Answer = 1 The AWS Product Group routinely Explanation : Amazon Aurora is a relational cloud database compatible with MySQL and PostgreSQL. Amazon Relational Database Service (RDS) fully manages Amazon Aurora, automating time-consuming administrative tasks such as hardware provisioning, database setup, patching, and backups. The AWS Product Team is accountable for applying patches to AWS Aurora's underlying operating system.
1AWS CCP64	1	15.4	Answer = 1 CMK Customer-Managed Explanation : A customer master key (CMK) represents a master key logically. The CMK contains metadata such as the key ID, date of creation, description, and key state. The CMK also contains the encryption and decryption key material. The AWS customer creates and manages these. Using the AWS IAM service, access can be managed to these.

QUESTIONNAIRE REPORT SHEET

iPQ1AWS CCP.V1.1

iPQ.No	Correct Answer	Points	Explanation
1AWS CCP65	2	15.4	Answer = 2 AWS Pricing Calculator Explanation : The AWS Pricing Calculator enables you to investigate AWS services and estimate the price of your use cases on AWS. You can model your solutions prior to constructing them, investigate the price points and calculations underlying your estimate, and locate instance types and contract terms that meet your requirements. This allows you to make informed decisions regarding AWS usage. You can plan your AWS expenses and usage, as well as estimate the cost of deploying a new set of instances and services. Based on the list of AWS services, the AWS Pricing Calculator can provide an estimate of AWS service consumption. <pre> graph LR A["AWS Pricing Calculator Estimate the cost of AWS products and services"] --> B["Add services Search and add AWS services that you need"] B --> C["Configure service Enter the details of your usage to see service costs"] C --> D["View estimate totals See estimated costs per service, service group, and totals"] A --- D </pre> via: https://calculator.aws/#/ The URL for the AWS Pricing Calculator is https://calculator.aws/#/. You should also be aware that AWS is in the process of deprecating the Simple Monthly Calculator, a similar tool. This calculator provides an estimate of AWS service usage fees based on the inputted information. It helps customers and prospects more accurately estimate their monthly AWS bill. This resource is available at https://calculator.s3.amazonaws.com/index.html.