Name: Tan Li Sin

Matrix Number: A21EC0231

Section: 05 (TIS)

Title: AWS Academy Cloud Foundations Badges

1. Define the AWS Cloud. (Module 1 - Cloud Concepts Overview)

- AWS is a secure cloud platform that offers a broad set of global cloud-based products.
- AWS provides you with on-demand access to the compute, storage, network, database, and other IT resources and management tools.
- AWS offers flexibility.
- You pay only for the individual services you need, for as long as you use them.
- AWS services work together like building blocks.

2. Explain the AWS pricing philosophy. (Module 2 - Cloud Economics and Billing)

While the number and types of services offered by AWS have increased dramatically, our philosophy on pricing has not changed. At the end of each month, you pay for what you use. You can start or stop using a product at any time. No long-term contracts are required.
This utility-style pricing model:

- Pay for what you use
- Pay less when you reserve
- Pay less when you use more
- Pay even less as AWS grows

3. Identify the global infrastructure components of AWS. (Module 3 - AWS Global Infrastructure Overview)

The AWS Global Infrastructure consists of Regions and Availability Zones. Choice of a Region is typically based on compliance requirements or to reduce latency. Each Availability Zone is physically separate from other Availability Zones and has

redundant power, networking, and connectivity. Edge locations and regional edge caches improve performance by caching content closer to users.

AWS Region: physical geographical location with one or more Availability Zones.
Availability Zone: fully isolated partition of the AWS infrastructure.
AWS Data Centers: designed for security.
AWS Global Infrastructure features: elastic and scalable, fault-tolerant, high availability.

4. Describe security and compliance measures of the AWS Cloud including AWS Identity and Access Management (IAM). (Module 4 - AWS Cloud Security)

Security and compliance are a shared responsibility between **AWS and the customer**.
AWS is responsible for protecting the infrastructure that runs all the services that are offered in the AWS Cloud. This infrastructure is composed of the hardware, software, networking, and facilities that run the AWS Cloud services.
The customer is responsible for the encryption of data at rest and data in transit. The customer should also ensure that the network is configured for security and that security credentials and logins are managed safely
AWS Identity and Access Management (IAM) allows you to control access to compute, storage, database, and application services in the AWS Cloud.

- Use IAM to manage access to AES resources
- Example-Control who can terminate Amazon EC2 instances
- Define fine-grained access rights
- IAM is no cost AWS account feature

5. Create an AWS Virtual Private Cloud (Amazon VPC). (Module 5 - Networking and Content Delivery)

- Use the VPC Wizard to create a VPC an Internet Gateway and two subnets in a single Availability Zone. An Internet gateway (IGW) is a VPC component that allows communication between instances in your VPC and the Internet.

- After creating a VPC, you can add subnets. Each subnet resides entirely within one Availability Zone and cannot span zones. If a subnet's traffic is routed to an Internet Gateway, the subnet is known as a public subnet. If a subnet does not have a route to the Internet gateway, the subnet is known as a private subnet.
- The wizard will also create a NAT Gateway, which is used to provide internet connectivity to EC2 instances in the private subnets.

6. Demonstrate when to use Amazon Elastic Compute Cloud (EC2), AWS Lambda, and AWS Elastic Beanstalk. (Module 6 – Compute)

AWS offers many computing services because different use cases benefit from different computing environments. The optimal compute service or services that you use will depend on your use case.

Services	Key Concepts	Characteristics	Ease of Use
Amazon EC2	• Infrastructure as a service (IaaS) • Instance-based • Virtual machines	Provision virtual machine that you can manage as you choose	A familiar concept to many IT professionals
AWS Lambda	• Serverless computing • Function-based • Low cost	Write and deploy code that runs on a schedule or that can be triggered by events Use when possible (architect for the cloud)	A relatively new concept for many IT staff members, but easy to use after you learn how
AWS Elastic Beanstalk	• Platform as a service (PaaS) • For web applications	Focus on your code (building your application) Can easily tie into other services-database, Domain Name System (DNS), etc.	Fast and easy to get started.

7. Differentiate between Amazon S3, Amazon EBS, Amazon EFS, and Amazon S3 Glacier. (Module 7 – Storage)

Storage Options	Amazon S3	Amazon EBS	Amazon EFS	Amazon S3 Glacier
Type of Storage	Object storage (photos, videos, documents, etc.)	Block storage for an EC2 instance	File system storage for multiple EC2 instances	Object storage (photos, videos, documents, etc.)
Defining Features	Can be accessible to any service or person	High performance for workloads of a single EC2 instance	Strong consistency, concurrent accessibility, and file locking features	store data at an extremely low cost (even in comparison to Amazon S3), but cannot retrieve data immediately when want it
Use Cases	Web applications, content management, photos, videos, backups, big data	Boot volumes, transactional and NoSQL databases, data warehousing & ETL	Home directories, database backups, developer tools, container storage, big data analytics	Media asset archiving, Healthcare information archiving, Regulatory and compliance archiving, Scientific data archiving, Digital preservation, Magnetic tape replacement
Max Storage Style	Unlimited	One volume: 16 TB or 64 TB	Unlimited	Unlimited
Max File Size	One object: 5 TB	Max file size = max volume size	Single file: 47.9 TB	One object: 40 TB
Latency	ms	Lower latency than EFS and S3	Low, uses Max I/O mode for higher performance	Minutes/hours

8. Demonstrate when to use AWS Database services including Amazon Relational Database Service (RDS), Amazon DynamoDB, Amazon Redshift, and Amazon Aurora. (Module 8 – Databases)

Amazon RDS:

- Web and mobile applications—High throughout, Massive storage scalability, High availability
- Ecommerce applications—Low-cost database, Data security, Fully managed solution
- Mobile and online games—Rapidly grow capacity, Automatic scaling, Database monitoring

Amazon DynamoDB:

- Mobile, web, gaming, ad tech, and Internet of Things (IoT) applications—Provide consistent, virtually unlimited storage, items can have different attributes, single-digit millisecond latency, no limits on table size or throughput

Amazon Redshift:

- Enterprise data warehouse (EDW)—Migrate at a pace that customers are comfortable with, Experiment without large upfront cost or commitment, Respond faster to business needs
- The big data—Low price point for small customers, Managed service for ease of deployment and maintenance, Focus more on data and less on database management
- Software as a service (SaaS)—Scale the data warehouse capacity as demand grows, Add analytic functionality to applications, Reduce hardware and software costs

Amazon Aurora:

- Fully managed by Amazon RDS
- Aurora automates database management tasks (hardware provisioning, software patching, setup, configuration, or backups)—High performance and scalability, High availability and durability, Multiple levels of security, Reduce database cost

9. Explain AWS Cloud architectural principles. (Module 9 - Cloud Architecture)

The AWS Well-Architected Framework is organized into five pillars each pillar includes a set of design principles.

Operational Excellence pillar:	Security pillar:	Reliability pillar:
>Perform operations as code >Make frequent, small, reversible changes >Refine operations procedures frequently >Anticipate failure >Learn from all operational failures	>Implement a strong identity foundation >Enable traceability >Apply security at all layers >Automate security best practices >Protect data in transit and at rest >Keep people away from data >Prepare for security events	>Automatically recover from failure >Test recovery procedures >Scale horizontally to increase aggregate workload availability >Stop guessing capacity >Manage change in automation

Performance Efficiency pillar:	Cost Optimization pillar:
>Democratize advanced technologies >Go global in minutes >Use serverless architectures >Experiment more often >Consider mechanical sympathy	>Implement Cloud Financial Management >Adopt a consumption model >Measure overall efficiency >Stop spending money on undifferentiated heavy lifting >Analyze and attribute expenditure

10. Explore key concepts related to Elastic Load Balancing (ELB), Amazon CloudWatch, and Auto Scaling. (Module 10 - Auto Scaling and Monitoring)

Elastic Load Balancing:

- Distributes incoming application or network traffic across multiple targets (such as Amazon EC2 instances, containers, IP addresses, and Lambda functions) in one or more Availability Zones.
- Elastic Load Balancing supports three types of load balancers:
  - Application Load Balancer
  - Network Load Balancer
  - Classic Load Balancer
- Elastic Load Balancing offers several monitoring tools for continuous monitoring and logging for auditing and analytics.

Amazon CloudWatch:

- Helps you monitor your AWS resources—and the applications that you run on AWS—in real-time.
- CloudWatch enables:
  - Collect and track standard and custom metrics
  - Set alarms to automatically send notifications to SNS topics or perform Amazon EC2 Auto Scaling or Amazon EC2 actions based on the value of the metric or expression relative to a threshold over a number of time periods.
  - Define rules that match changes in your AWS environment and route these events to targets for processing

  Auto Scaling:

  - Scaling enables you to respond quickly to changes in resource needs.
  - Amazon EC2 Auto Scaling helps you maintain application availability, and enables you to automatically add or remove EC2 instances according to your workloads
  - An Auto Scaling group is a collection of EC2 instances
  - A launch configuration is an instance configuration template
  - Implement dynamic scaling with Amazon EC2 Auto Scaling, Amazon CloudWatch, and Elastic Load Balancing