MODULE 2 · STUDY NOTES

Introduction to Security on AWS

Why cloud security is familiar, the AWS security objectives and their services, the seven design principles, and who secures what under the shared responsibility model

Moving to the cloud does not mean starting security over — it means practicing the same discipline with different tools. Security is the practice of protecting intellectual property from unauthorized access, use, or modification, and the confidentiality-integrity-availability (CIA) triad still defines what you are protecting. What changes in the cloud is that you trade the cost and complexity of guarding physical facilities and hardware for software-based tools, and you inherit a partner. Under the AWS shared responsibility model, AWS secures the cloud — the global infrastructure, hardware, and virtualization layer — while you secure what you put in the cloud. This module builds that picture in three parts: the benefits and security objectives of the AWS Cloud and the services that deliver them, the seven security design principles from the AWS Well-Architected Framework, and the shared responsibility model itself. In the bank business scenario that frames the course, María opens her first meeting with her supervisor John on exactly this foundation — presenting cloud-security fundamentals and the shared responsibility model to show where the bank and AWS each protect the bank's resources with multiple layers of security.

LEARNING OBJECTIVES

- Identify security features and benefits of cloud computing
- Identify the security principles that the AWS Cloud is structured around
- Identify which part of an application the user is responsible to secure in the cloud

2.1 Security in the AWS Cloud

Security is the practice of protecting your intellectual property from unauthorized access, use, or modification. Security in the cloud is similar to security in on-premises data centers — only **without the costs and complexities of protecting facilities and hardware**. Because the cloud has no physical servers or storage devices that you own, you use **software-based security tools** to monitor and protect the flow of information into and out of your AWS resources. Before diving into those tools, it helps to remember *why* organizations move to the cloud in the first place, because several of the benefits directly enable stronger security.

BENEFITS OF MOVING TO THE CLOUD

- **Trade fixed expense for variable expense** — pay only when you consume computing resources, and only for how much you consume, instead of investing heavily in data centers up front.
- **Benefit from massive economies of scale** — aggregated usage from hundreds of thousands of customers translates into lower pay-as-you-go prices.
- **Stop guessing on capacity** — access as much or as little as you need and scale up or down with only a few minutes of notice, avoiding idle or insufficient resources.
- **Increase speed and agility** — new IT resources are a click away, reducing provisioning time from weeks to minutes.
- **Stop spending money to run and maintain data centers** — focus on projects that differentiate your business instead of racking, stacking, and powering servers.
- **Go global in minutes** — deploy across multiple Regions with a few clicks for lower latency and a better customer experience.

Security is familiar. The CIA triad was originally developed to highlight the important aspects of information security within an organization, and it applies just as well in the cloud. The three aspects work together: you limit who can see data, you keep that data trustworthy across its life, and you make sure it is there when it is needed.

KEY TERMS

Confidentiality
Limiting information access and disclosure to authorized users and preventing access by unauthorized people.

Integrity
Maintaining the consistency, accuracy, and trustworthiness of data over its entire lifecycle, preserving data at rest and in transit; this includes source integrity — assurance that the sender of information is who it is supposed to be.

Availability
The readiness of information resources — having access to information when it is needed. An information system that is not available when you need it is almost as useless as no system at all.

At AWS, the goal is to make security as familiar as what you do today: you bring the same security models you already use — providing **visibility, auditability, and controllability** to your resources — and AWS adds the **agility and automation** you need to operate at cloud scale. These five objectives are the organizing idea for the AWS security toolset. Each objective corresponds to questions a security team asks, and AWS offers a service to answer them.

TABLE 2.1 AWS Cloud security objectives and the services that support them

OBJECTIVE	QUESTIONS IT ANSWERS	AWS SERVICE
Controllability	Can I effectively manage users? Provide temporary credentials? Use my own keys?	IAM, STS, CloudHSM
Auditability	Who has access? Who performed what action, when, and from where?	CloudTrail
Visibility	What is in my environment? What changed, and what impact did an action have?	Config
Agility and automation	How do I ensure high availability and apply security checks reproducibly?	CloudFormation

KEY TERMS

AWS IAM
Securely controls access to AWS resources — who can use them (authentication) and what they can do and in what ways (authorization). Define granular permissions for users, groups, and roles; grant different access to different people for different resources; and use federation and multi-factor authentication (MFA).

AWS STS
The Security Token Service creates and provides trusted users with temporary security credentials that control access to AWS resources; they work almost identically to long-term access keys.

AWS CloudHSM
Protects your encryption keys within hardware security modules (HSMs) designed and validated to government standards, so you can generate, store, and manage cryptographic keys where only you have access — without sacrificing application performance.

AWS CloudTrail
Helps answer auditing questions such as what actions a specific user took over a time period, which user acted on a given resource, and the source IP address of an activity.

AWS Config
Discovers your existing AWS resources, exports a complete inventory with all configuration details, and shows how a resource was configured at any point in time — supporting compliance auditing, change tracking, and troubleshooting.

AWS CloudFormation
Deploys environments in a secure and reproducible manner from templates, letting you automate deployment and, for example, pre-provision a forensic clean room rather than building one manually.

> COMMON PITFALL
>
> Cloud security is not about protecting physical hardware. Because the cloud has no servers or storage devices that you own, you protect the *flow of information* with software-based tools. AWS handles the physical facilities, hardware, and networking — that is why you can focus on visibility, auditability, controllability, agility, and automation instead of racking and powering equipment.

2.2 Security design principles

The AWS Cloud is structured around **seven design principles** drawn from the **security pillar of the AWS Well-Architected Framework**. Following these principles helps you strengthen your workload security and can guide your conversations around security and compliance. Each principle names both a mindset and the AWS capabilities that make it practical.

TABLE 2.2 The seven security design principles

PRINCIPLE	WHAT IT MEANS AND HOW AWS SUPPORTS IT
1. Apply least privilege	Grant access only as needed based on job role; start by denying everything. Enforce separation of duties and avoid long-term credentials with IAM roles, SSO/federation, and MFA.
2. Enable traceability	Monitor, alert, and audit actions and changes in near-real time. Implement detective controls with CloudTrail (API calls), CloudWatch (metrics and alarms), Config (configuration history), and GuardDuty (threat detection).
3. Secure all layers	Use a defense-in-depth approach across the network, application, and data-store layers rather than only the outer perimeter, combining several AWS services and hardened, immutable AMIs.
4. Automate security	Automate routine security tasks with a comprehensive set of APIs, and implement infrastructure as code with CloudFormation so your system can detect, review, and respond automatically.
5. Protect data in transit and at rest	Apply encryption and fine-grained access controls, classify data with a tagging schema, and use VPN and TLS connections. AWS builds in server-side encryption (SSE) for S3 and SSL termination on ELB.
6. Prepare for security events	Put incident-response tools and access in place ahead of time and practice with game days. Detailed logs, automated event processing, and a CloudFormation-provisioned clean room enable safe forensics.
7. Minimize the attack surface	Harden operating systems and reduce unused components. Configure security groups and network ACLs in the VPC, and scale to absorb attacks with services such as Auto Scaling and CloudFront.

Least privilege and identity sit at the heart of the principles. Grant access only to those who genuinely need it, by job role, and enforce separation of duties. Reducing reliance on long-term credentials shrinks your attack surface: for **workforce identities**, use AWS Single Sign-On or federation with IAM; for **machine identities** such as EC2 instances or Lambda functions, require **IAM roles** instead of long-term access keys. IAM is the primary service for permissions management — you define principals (accounts, users, roles, services) and build granular policies aligned to them.

Protecting data and shrinking the target round out the set. AWS makes encryption easier to adopt — server-side encryption (SSE) for Amazon S3 and SSL termination on Elastic Load Balancing. To minimize the attack surface, harden and reduce components and configure network controls; services such as **AWS Auto Scaling** and **Amazon CloudFront** help applications *absorb* infrastructure-layer attacks like UDP reflection attacks and SYN floods.

KEY TERMS

Defense in depth
Applying security controls at every layer — network, application, and data store — rather than focusing on a single outer layer.

Immutable (hardened) AMI
A hardened instance configuration persisted to an Amazon Machine Image, so every new instance launched from it — manually or by automatic scaling — inherits the hardened configuration.

Server-side encryption (SSE)
Encryption AWS applies to stored data; for example, SSE for Amazon S3 makes it easier to store your data in encrypted form.

SSL termination
Having Elastic Load Balancing (ELB) handle the entire HTTPS encryption and decryption process on your behalf.

DDoS attack
A distributed denial-of-service attack that aims to make a server unavailable by consuming its resources; a SYN flood is one type, and a UDP reflection attack forges a source address to solicit traffic toward the target.

> COMMON PITFALL
>
> Least privilege does not mean handing broad access to people you trust. The best practice is to **deny access to everything first and grant it as needed** by job role, enforce separation of duties, and lean on temporary credentials and IAM roles rather than long-term access keys. Trust is established through controlled, auditable permissions — not blanket access.

2.3 The shared responsibility model

Security and compliance are shared responsibilities between AWS and the customer. AWS operates, manages, and controls **security *of* the cloud** — securing everything from the host operating system and virtualization layer down to the physical security of the facilities where the services run. This includes protecting the global infrastructure: the hardware, software, networking, and facilities that run all AWS Cloud services. The customer assumes responsibility and management **in the cloud** — the security steps you must take depend on the services you use and the complexity of your system.

AWS – SECURITY OF THE CLOUD	CUSTOMER – SECURITY IN THE CLOUD
the infrastructure that runs all AWS services	*everything you implement and connect*
▪ Physical security of facilities and data centers ▪ AWS Global Infrastructure — Regions, Availability Zones, edge locations ▪ Host operating system and virtualization (hypervisor) layer ▪ Networking hardware and isolation between customers ▪ Foundation services: compute, storage, databases, networking	▪ Customer data, encryption options, and data integrity ▪ Platform, applications, and identity and access management ▪ Operating system, network, and firewall configuration ▪ Client-side encryption and server-side encryption choices ▪ Network traffic protection (encryption, integrity, identity)

Because your responsibility is determined by the services you select, **the model changes depending on the AWS services that you use** — the more managed the service, the more AWS handles for you. Consider a common architecture to see how the line moves.

WORKED EXAMPLE – AN S3, EC2, AND RDS DEPLOYMENT

A company uses Amazon S3 to store data. Its environment also runs EC2 instances and an Amazon RDS instance hosting a MySQL database, all inside a VPC; one EC2 instance runs a web server whose application uses the database. **AWS is responsible** for the global infrastructure — the physical servers, storage hardware, and networking that host the S3 bucket, EC2 instances, and database — and for the hypervisor layer that runs the EC2 instances. **The customer is responsible** for the guest OS on the EC2 instances (Windows or Linux updates and patches), any application software installed, the security groups that control network access to each instance and to the RDS database, and the security of the S3 bucket and its objects — using features such as bucket policies, data encryption, and S3 Block Public Access.

While AWS secures and maintains the cloud infrastructure, **you are responsible for securing everything you put in the cloud.** You maintain complete control over your content and manage critical security requirements around it. The AWS Cloud also gives you greater access to security data and an automated way to respond to security events.

THE CONTENT DECISIONS YOU ALWAYS CONTROL

- **What you choose to store** on AWS.
- **Which AWS services** are used with that content.
- **Which Region** (country) the content is stored in.
- **What format and structure** the content uses — and whether it is masked, anonymized, or encrypted.
- **Who has access** to the content, and how those access rights are granted, managed, and revoked.

TABLE 2.3 Who is responsible? – answers from the module activities

COMPONENT	RESPONSIBLE PARTY
Physical security of the data center	AWS
Virtualization infrastructure	AWS
Guest OS upgrades and patches on the EC2 instance	Customer
EC2 security group settings	Customer
Configuration of applications running on the EC2 instance	Customer
Oracle upgrades/patches when Oracle runs as an Amazon RDS instance	AWS
Oracle upgrades/patches when Oracle runs on an EC2 instance	Customer
S3 bucket access configuration	Customer
Ensuring the AWS Management Console is not hacked	AWS
Configuring the VPC and subnet	Customer
Protecting against network outages in AWS Regions	AWS
Securing the SSH keys	Customer
Network isolation between AWS customers' data	AWS
Enforcing MFA for all user logins	Customer

> COMMON PITFALL
>
> The RDS-versus-EC2 distinction is a favorite exam trap. If a database such as Oracle runs as an **Amazon RDS** instance, AWS performs the backups and patching. If the *same* database runs on an **EC2** instance, the customer owns those upgrades and patches. Choosing a more managed service moves the responsibility line toward AWS — the workload determines the split.

2.4 Managed services organizations (MSO)

One approach to implementing security and governance is to create a **centralized team** responsible for establishing repeatable processes and templates for deploying applications to AWS while maintaining organizational control. Such a team may be internal or an external third party, and is often called a provisioning team, a center of excellence, or a **managed services organization (MSO)**. External vendors are commonly called **managed service providers (MSPs)** — AWS validates AWS Partners under the AWS Managed Service Provider (MSP) Program.

MSOs and MSPs are typically responsible for provisioning accounts, establishing repeatable deployment processes, auditing the deployments of workload owners, and hosting shared services for security, continuous monitoring, connectivity, and authentication. In short, an **MSO essentially creates the guardrails** for security, data protection, and disaster recovery in the company. Adding an MSO lets a workload owner **inherit a significant portion of the security control implementation**, so their own responsibility narrows to the configuration and software unique to their application.

TABLE 2.4 MSO responsibility model – three tiers

TIER	RESPONSIBILITY
Workload owners	Customer data; platform, applications, and IAM. The actual deployment, development, and maintenance of applications — system administrators, developers, and others responsible for specific applications.
MSO / MSP	Operating system, network, and firewall configuration; client-side and server-side encryption; and network traffic protection — the shared guardrails and services the workload owner inherits.
AWS	Software (compute, storage, databases, networking) and hardware plus the AWS Global Infrastructure — Regions, Availability Zones, and edge locations.

ACTIVITIES AN AWS CUSTOMER MSO OFTEN PERFORMS

- **Account provisioning** — establish the initial account, connect it for consolidated billing, and configure basic security before granting the workload owner access.
- **Security oversight** — enable security personnel to monitor applications, for example by establishing an auditor group with cross-account access.
- **Amazon VPC configuration** — manage the VPC, subnets, security group configurations, and network ACLs for tighter control.
- **IAM configuration** — create user groups and assign permissions (internal auditors, an IAM superuser, and function-segregated admin groups).
- **Development and approval of templates** — build preapproved CloudFormation templates that workload owners reuse, inheriting the template's security implementation.
- **AMI creation and management** — maintain a library of common, approved AMIs to enforce hardened, up-to-date machine images.
- **Shared services VPC** — receive continuous monitoring feeds and host shared access management, logging endpoints, and configuration aggregation.

COMMON PITFALL

An MSO does not erase the workload owner's responsibility — it narrows it. The workload owner still owns the configuration and software that are specific to their application; they simply inherit the MSO's implementation of the broader security controls (accounts, VPC, IAM, templates, and AMIs) rather than building each one from scratch.

Module summary

- Security in the cloud practices the same CIA triad — confidentiality, integrity, availability — as on premises, but with software-based tools instead of guarding physical facilities and hardware; the benefits of the cloud (variable expense, economies of scale, elastic capacity, agility, less maintenance, global reach) also enable security at scale.
- AWS structures cloud security around five objectives, each with a supporting service: controllability (IAM for authentication and authorization, STS for temporary credentials, CloudHSM for key protection), auditability (CloudTrail), visibility (Config), and agility and automation (CloudFormation).

- The AWS Well-Architected security pillar defines seven design principles: apply least privilege, enable traceability, secure all layers (defense in depth), automate security, protect data in transit and at rest, prepare for security events, and minimize the attack surface.
- Under the shared responsibility model, AWS is responsible for security OF the cloud (global infrastructure, hardware, networking, host OS, and virtualization), and the customer is responsible for security IN the cloud (guest OS, applications, IAM, security groups, data, and encryption).
- The responsibility line moves with the services you choose — managed services such as Amazon RDS shift patching to AWS, while self-managed workloads (a database on EC2) keep OS and application patching with the customer; either way, customers always control their content (what to store, which services and Region, format, and who has access) and configure VPCs, subnets, security groups, SSH keys, and MFA.
- A managed services organization (MSO) or managed service provider (MSP) centralizes governance — provisioning accounts, approving CloudFormation templates and AMIs, and running shared security services — creating guardrails that workload owners inherit while still owning application-specific configuration.

Review questions

1. What are the three aspects of the CIA triad, and what does each protect?

Answer: Confidentiality limits information access and disclosure to authorized users; integrity maintains the consistency, accuracy, and trustworthiness of data across its lifecycle (at rest and in transit); availability ensures information resources are ready and accessible when needed.

2. Match each AWS Cloud security objective to the service that supports it.

Answer: Controllability → IAM, STS, and CloudHSM; auditability → CloudTrail; visibility → Config; agility and automation → CloudFormation.

3. List the seven security design principles from the AWS Well-Architected security pillar.

Answer: Apply the principle of least privilege, enable traceability, secure all layers, automate security, protect data in transit and at rest, prepare for security events, and minimize the attack surface.

4. A customer runs the same Oracle database two ways — on Amazon RDS and on an EC2 instance. Who patches each, and who owns the guest OS in each case?

Answer: AWS patches Oracle (and manages the underlying OS) when it runs as an Amazon RDS instance, because RDS automates backups and patching; the customer patches Oracle and owns the guest OS when the database runs on an EC2 instance. The responsibility line depends on how managed the service is.

5. Sample exam: who is responsible for configuring security group rules that determine which ports are open to an EC2 Linux instance?

Answer: The customer (choice B). The customer controls network access to EC2 instances by configuring security groups, which filter traffic based on protocols and port numbers; AWS maintains only its own infrastructure devices.

6. What is a managed services organization (MSO), and what does adding one do for a workload owner?

Answer: An MSO (or external MSP) is a centralized team that provisions accounts, approves CloudFormation templates and AMIs, and runs shared security services — creating the guardrails for security, data protection, and disaster recovery. The workload owner inherits much of the security control implementation and narrows their own responsibility to application-specific configuration and software.