AWS BASICS

Beginner-friendly notes for understanding Amazon Web Services

Topic	What you should understand
Cloud Computing	Using computing resources over the internet instead of owning all physical infrastructure.
AWS	Amazon's cloud platform providing compute, storage, databases, networking, security, analytics and more.
Region	A geographic area containing multiple AWS Availability Zones.
Availability Zone	An isolated data-center location within an AWS Region, designed for resilience.
IAM	Identity and Access Management for controlling who can access AWS resources.

1. What is Cloud Computing?

Definition: Cloud computing means renting IT resources—such as servers, storage, databases and networking—on demand over the internet.

Why use it?: You can scale resources quickly, pay for what you use, reduce hardware maintenance, and deploy applications globally.

Common models: IaaS: infrastructure such as virtual machines. PaaS: managed platforms for applications. SaaS: ready-to-use software.

2. Introduction to AWS

AWS: Amazon Web Services is a large cloud platform with services for application hosting, storage, databases, networking, security, monitoring, AI/ML and more.

AWS account: Your account is the boundary in which you create resources, users, permissions, billing and configurations.

Management: Resources can be managed using the AWS Management Console, AWS CLI, SDKs or infrastructure-as-code tools.

3. AWS Global Infrastructure

Regions: Choose a Region based on latency, service availability, compliance requirements and cost.

Availability Zones: A Region normally contains multiple isolated Availability Zones. Deploying across AZs can improve availability.

Edge locations: AWS uses edge locations to serve content closer to users, especially through services such as CloudFront.

4. EC2 — Virtual Servers

What is EC2?: Amazon Elastic Compute Cloud (EC2) provides resizable virtual servers called instances.

Instance basics: Important choices include instance type, AMI, storage, networking, security groups and IAM role.

Instance types: Families are optimized for different workloads: general purpose, compute optimized, memory optimized, storage optimized and accelerated computing.

Example: A small web application might run on a general-purpose instance. Production workloads often use multiple instances behind a load balancer.

5. S3 — Object Storage

What is S3?: Amazon Simple Storage Service stores objects such as images, documents, backups, logs and application files.

Core concepts: Bucket → object → key. A bucket is a container; an object is the stored data plus metadata; the key identifies the object.

Important features: Versioning, lifecycle rules, encryption, access policies, storage classes and replication.

Rule of thumb: Use S3 for durable object/file storage, not as a normal operating-system disk.

6. EBS — Block Storage

What is EBS?: Amazon Elastic Block Store provides persistent block storage volumes commonly attached to EC2 instances.

Use cases: Operating-system disks, application data, databases and workloads that need block storage.

Key idea: EBS volumes are different from S3 objects: EBS behaves like a disk; S3 is object storage.

7. VPC — Networking Basics

What is VPC?: Amazon Virtual Private Cloud gives you a logically isolated network in AWS.

Main components: VPC, subnets, route tables, internet gateway, NAT gateway, security groups and network ACLs.

Public subnet: A subnet with a route that can reach an internet gateway; resources can be publicly reachable only when other required settings are also configured.

Private subnet: Typically used for resources that should not accept direct inbound internet traffic, such as application servers or databases.

8. Security Groups & IAM

Security Group: A stateful virtual firewall associated with resources such as EC2 instances. Rules control allowed inbound and outbound traffic.

IAM user: An identity that can authenticate and receive permissions. For human access, organizations generally prefer centralized identity/federation where possible.

IAM role: An identity with permissions that can be assumed by trusted principals or AWS services. Roles are preferred over storing long-lived access keys on servers.

Least privilege: Give only the permissions required to perform a task.

9. RDS — Managed Databases

What is RDS?: Amazon Relational Database Service makes it easier to run and maintain relational databases.

Engines: Common engines include PostgreSQL, MySQL, MariaDB, Oracle and SQL Server; Amazon Aurora is also available as an AWS-managed relational database.

Benefits: Automated backups, patching options, monitoring, scaling options and high-availability features reduce operational work.

10. Load Balancing & Auto Scaling

Elastic Load Balancing: Distributes incoming traffic across healthy targets such as EC2 instances.

Auto Scaling: Automatically adjusts the number of instances based on policies and demand.

Typical architecture: Users → Load Balancer → multiple EC2 instances → database/storage. This is more resilient than putting everything on one server.

11. CloudWatch & Monitoring

CloudWatch: Collects metrics, logs and events and can trigger alarms.

What to monitor: CPU, memory where available through agents, disk usage, network traffic, application errors, latency and database metrics.

Operational habit: Monitoring should be configured before a production incident happens.

12. AWS Billing & Cost Basics

Pay-as-you-go: Many AWS services charge based on consumption. Some resources continue generating charges while running even if traffic is low.

Common cost sources: EC2 runtime, EBS storage, data transfer, NAT gateways, load balancers, RDS, S3 requests/storage and public IPv4 addresses.

Cost control: Use budgets and billing alerts, right-size resources, delete unused resources, use lifecycle policies and review the Cost Explorer regularly.

13. Beginner AWS Architecture

A common small production-style web application can be organized as:

```
Users → Route 53 / DNS → Load Balancer → EC2 instances in multiple Availability Zones → RDS database

Static files / uploads → S3

Monitoring → CloudWatch

Access control → IAM + Security Groups
```

The exact architecture depends on traffic, availability, security, compliance and cost requirements.

14. Essential AWS CLI Examples

```
Configure the CLI
aws configure

List S3 buckets
aws s3 ls

Create an S3 bucket
aws s3 mb s3://my-example-bucket

Copy a file to S3
aws s3 cp file.txt s3://my-example-bucket/

List EC2 instances
aws ec2 describe-instances
```

15. Beginner Learning Path

• Cloud concepts and shared responsibility model

• AWS Regions, Availability Zones and pricing basics

• IAM and least-privilege access

• VPC, subnets, routes and security groups

• EC2 + EBS

• S3

• RDS

• Load Balancing + Auto Scaling

• CloudWatch and cost management

• Build a small project and document the architecture

16. Quick Interview Questions

Q: What is AWS?
A: A cloud platform offering on-demand infrastructure and managed services.

Q: EC2 vs S3?
A: EC2 provides compute instances; S3 provides object storage.

Q: Region vs Availability Zone?
A: A Region is a geographic area; an AZ is an isolated location within a Region.

Q: Security Group vs IAM?
A: A security group controls network traffic; IAM controls identities and permissions.

Q: Public vs private subnet?
A: The distinction primarily comes from routing to an internet gateway and how resources are exposed.

Q: Why use multiple AZs?
A: To reduce dependence on a single isolated infrastructure location and improve availability.

Study note: AWS services, pricing, limits and best practices change over time. For production decisions, verify the current AWS documentation and pricing for your chosen Region and services.