

Amazon Web Services (AWS) Introduction

AWS Global Infrastructure



- **39** Regions
- **123** Availability Zones
- **750+** Point of Presence
- **245** Countries and territories Served
- **~135** Direct Connect Location

<https://aws.amazon.com/about-aws/global-infrastructure/>

AWS Global Infrastructure Regions & AZs



N AMERICA

Canada Central 3 Ohio 3
Canada West 3 Oregon 4
GovCloud US-East 3 Mexico 3
GovCloud US-West 3
Northern California 3
Northern Virginia 6



S AMERICA

São Paulo 3
Chile



EUROPE

Frankfurt 3 Stockholm 3
Ireland 3 Zurich 3
London 3 Germany
Milan 3
Paris 3
Spain 3



AFRICA

Cape Town 3



MIDDLE EAST

Bahrain 3
Tel Aviv 3
UAE 3
Kingdom of Saudi Arabia



ASIA PACIFIC

*Beijing 3 Osaka 3 Thailand 3
*Ningxia 3 Seoul 4
Hong Kong 3 Singapore 3
Hyderabad 3 Tokyo 4
Jakarta 3 Malaysia 3
Mumbai 3 Taipei 3



AUSTRALIA & NEW ZEALAND

Melbourne 3
Sydney 3
Auckland 3


● Available Region ● Announced # Availability Zone

AWS Region design

AWS Regions are comprised of multiple AZs for **high availability**, **high scalability**, and **high fault tolerance**. Applications and data are replicated in real time and consistent in the different AZs.



Availability zones design for resiliency



100K+ servers at scale

ISOLATED PARTITION

Fully isolated with one or more datacenters

POWER

Highly available, fault tolerant, and scalable

DISTANCE

Physically separated by a meaningful distance – up to 60 miles (100km) of each other

INTERCONNECTION


Datacenters connected via fully redundant and isolated metro fiber

AWS Cloud Infrastructure Designed to Scale with Highest Availability

AWS Investments on Underlying Higher Resilient Infrastructure

Total Incident Hours (Public) by Provider in 2022-2024


Incident hours published = Incident hours published * Regions Impacted * Services Impacted



* Sources: Public Status Pages for Cloud Providers. Incidents Types include = Incidents, Degradation, Disruption, and Outages of Services

Connectivity to AWS infrastructure

Access to AWS from on-premises



Accessing AWS Cloud

AWS Management Console



Command Line Interface (CLI)



Software Development Kit (SDK)



AWS customer di Indonesia

Software & internet



Telecommunication



Education



Travel



Games



Logistics



Media



Mining



Agriculture



Public sector



Financial services



Automotive



Retail



Healthcare



Manufacturing



<https://aws.amazon.com/solutions/case-studies/browse-customer-success-stories>

Breadth and Depth of Services and Functionalities



Analytics

Analytics
Data Exchange
Data Lake
Data Pipelines
Data Warehouse
Elasticsearch

Streaming
ETL
Hadoop / Spark
Interactive SQL Queries
Visualizations



Business Applications

Contact Center
Sharing & Collaboration
Online Meetings & Chat
Voice-Enabled Workplace

Unified Communications
Mobile & Web Apps Without Programming



Blockchain

Blockchain Templates
Ledger Database

Managed Blockchain



Security, Identity, and Compliance

Access Control
Assessment & Reporting
Configuration Compliance
Data Protection
DDOS Protection
Identity Management

Key Management & Storage
Monitoring & Logging
Resource Management
Threat Detection
Web Application Firewall



Storage

Archive Storage
Backup & Restore
Block Storage
Data Transfer
Edge Processing & Computing
File Storage

High-Performance File System
Hybrid Cloud Storage
Object Storage
Unified Communications
Windows File System



Database

Aurora, *a high-performance, relational database Built for the Cloud*
Managed SQL Server
Managed MariaDB
Managed MySQL
Managed Oracle
Managed PostgreSQL

Document Database
Graph Database
In-Memory Caching
Key-Value Store Database
Ledger Database
Time Series Database



Development Tools

Analyze and Debug
Application Lifecycle Management
Authoring
Build & Test
Containers

DevOps Resource Management
One-Click App Development
Patching
Pipeline Orchestration
Resource Templates
Triggers



Compute

Compute
Auto-Scaling
Batch Jobs
Event-Driven Serverless Computing
Instance Types
Managed Virtual Private Servers
Managed Repository for Serverless Apps

Run & Manage Web Apps
Serverless Compute
Virtual Servers
Containers
Container Service
Managed Kubernetes
Store & Retrieve Docker Images



Media Services

Live Video Transport
Media Storage
Transcoding
Video Origination & Packaging

Video Personalization & Monetization
Video Processing & Delivery
Video Streaming Analysis



Hybrid Architecture

AWS Services on Premises
Data Integration
Integrated Devices & Edge Systems
Integrated Identity & Access

Integrated Networking
Integrated Resource & Deployment Management
VMWare Cloud on AWS
Integrated SG



Internet of Things (IoT)

Rules Engine
Device Analytics
Device Gateway
Device SDK
Device Shadows
Event Detection & Response
Local Compute

Local Data Collection
Management & Security
Microcontroller Operating System
Registry
Visual Applications Development



AI and Machine Learning (ML)


ML Frameworks
Deep Learning AMIs & Containers
Hardware Acceleration
ML at the Edge
TensorFlow, PyTorch, MXNet
Sagemaker
Automatic Model Tuning
Data Labeling
Hosted Notebooks
ML Marketplace
Model Hosting
Model Optimization
Model Training
Pre-Built Algorithms
Topic Modeling
Deep Learning Models
Reinforcement Learning
Spot Instances
Batch Predictions
Real-Time Predictions

AI Services
Chatbots
Entity Extraction
Face Analytics
Face Search Forecasting
Image Labeling
Natural Language Processing
Personalization & Recommendation
Sentiment Analysis
Speech Translation
Text & Data Extraction
Text to Speech Translation
Video & Image Analysis
Content Moderation

Breadth and Depth of Services and Functionalities

Source: AWS Functionality Matrix (as of Dec 1st, 2024)

Number of Total Cloud Features by Cloud Provider



Generative AI

<https://aws.amazon.com/ai/generative-ai/>

Over 90% of AI/ML unicorns run on AWS

Unlocking the potential of generative AI




Easiest way to build
with Foundation Models



The most performant
infrastructure for generative AI



Generative AI-powered
applications on AWS




Flexibility to build
your own Foundation Models

AWS for Agentic AI

INTERFACES & PROTOCOLS
(MCP/A2A)

DATA



AWS Security compliance



SNI 27001



SOC 1



SOC 2



SOC 3




CCCS
PIPEDA



CJIS



FERPA



GxP



MPAA



SEC Rule
17a-4(f)



VPAT
Section 508



FISC



G-Cloud



Shared responsibility model




Customers are responsible for their security and compliance **IN** the Cloud



AWS is responsible for the security **OF** the Cloud

Shared Responsibility Model – Services



AWS security services and NIST Cyber Security Framework



NIST Cyber Security Framework

