



Amazon Web Services (AWS) Cloud 101 for Researchers

What is Cloud Computing?

**the orange asterisks will redirect you to AWS site for more information*

“Cloud” is the on-demand delivery of technology services over the internet to cloud users (e.g. virtual servers, file systems, etc.)*. Users can launch resources anytime for variety of use cases (e.g. offsite backup, DR, HPC, big data analytics, etc.), and only pay for actual usage through pay-as-you-go model or subscription to commitment plans for greater discounts.



And **Amazon Web Services (AWS)**, an Amazon subsidiary, has been providing cloud services for over 14 years.

Reasons why researchers are using AWS cloud:

- (1) **Agility & Speed** in spinning up or down IT resources anytime, unlike procuring physical servers which takes weeks or months. If an experiment fails, users can always de-provision those resources without risk
- (2) **Elasticity** in allowing you to do right-sizing, and provision the right amount of resources to run your application or research workload. Computing capacity be adjusted instantly by manual or auto-scaling, and storage size can be scaled manually (i.e. block volume EBS*) or automatically (i.e. object store S3*)
- (3) **Cost savings** by paying only for resources you provisioned to run your application and store data during your project timeframe, with a much lower variable expense on AWS due to AWS' economies of scale. Unlike procuring physical servers, you stop paying for cloud resources once you terminate/delete them
- (4) **Global deployment in minutes**, with couple of clicks on the console. AWS has physical data centers across 24 geographical regions around the world*, including **Asia Pacific (Hong Kong) Region***. AWS HK Region consists of 3 Availability Zones (groups of data centers) which further reduce latency to local users and help users meet data residency needs. With highest network availability, AWS reports **7x fewer downtime hours** than the next largest cloud provider*.
- (5) **Remote access to IT resources anywhere**, anytime even when you need to work from home

Benefits of using AWS for Research?

Most functionality. AWS offers over 175 fully featured services – **more services and more capabilities than any other cloud provider.** A great cloud adoption entry point is Amazon EC2 (virtual server). With AWS, customers can use and put together services based on actual needs – from compute, storage, databases, to new areas such as serverless, IoT, and machine learning/AI



More machine learning (ML) is built on AWS than anywhere. Twice as many companies use AWS for machine learning than anywhere else. Cater to customers with different capabilities, AWS offers (1) easy-to-use, pre-trained AI services*, (2) fully-managed machine learning service platform*, and (3) ML frameworks* & infrastructure* for custom model. 96% of deep learning today runs on cloud, out of which 81% runs on AWS, driven by breadth of framework and interface choices on AWS, such as TensorFlow (**85% of cloud-based TensorFlow workloads runs on AWS**)



TensorFlow mxnet PYTORCH Caffe2 theano GLUON Keras **and more...**

Most secure. AWS core infrastructure is built to satisfy the security requirements for military, global banks, and high-sensitivity organizations. **On AWS, customers retain content ownership, and can control where your data is stored, how your data is encrypted in-transit and at-rest, and who can access the data.** AWS environment is backed by a deep set of cloud security tools to meet your security objectives and compliance requirements, such as:



HITRUST CSF Certified SCHEMA ISO 9001 ISO 27001 ISO 27017 ISO 27018 AICPA SOC PCI DSS **and more...**

Largest global community of customers and partners. Across the globe, AWS has tens of thousands of APN partners. **AWS Marketplace*** (an online software store that helps users find, buy, and immediately start using software and services that run on AWS) offers more than 4,800 software listings from over 1,400 software vendors, across 39 solution categories

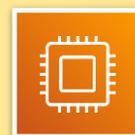


Many customers choose AWS because it not only has a broad and deep service range which derives from customers' needs and feedback, but also provides tutorials* and supports to use these products. You may explore over 175 AWS services by [solution categories here](#) and [free tier](#). As a starting point, let's go over two foundational services which many customers start with on AWS – **Compute** and **Storage**, as well as **cloud management tools***

Compute services on AWS

Compute services in the cloud can be **provisioned in minutes**, compared to on-premises infrastructure where server procurement typically takes 10-12 weeks or longer. AWS offers a comprehensive portfolio of compute services – from [virtual servers](#) to [containers](#) to [serverless](#)

Amazon EC2* is a virtual server service with [wide selection of instance types](#) differing in CPU, GPU, memory, storage, and networking capability for different use cases (e.g. GPU instance for molecular diagnostics workload). Security-wise, a default Virtual Private Cloud* (logically isolated virtual network on AWS) is created for your AWS account automatically the first time you provision Amazon EC2 resources. Users can use the same precautions to secure and protect their applications as they would with on-premises servers, and use the best practices documents and encryption tools provided by AWS to meet different security objectives, such as [Reference Architecture for HIPAA](#)



“Experiments that used to take us weeks now run overnight. This means that our scientists are always engaged and not waiting for results. AWS greatly reduced our turnaround time for scientific inquiry.”



Professor Dragutin Petkovic

[Case study link](#)

Storage services on AWS

Storage is a critical cloud component as it holds information used by applications. AWS storage services range from virtual server's persistent block-level storage volume* to data lake storage platform*, and make collecting, storing, analyzing and sharing data much easier.

“The AWS Cloud enables swift collaboration even with hundreds of TB of data. The ability to have a central area for people to process the data cuts down on bandwidth and the need to buy and maintain vast computational resources.”



Lead Programmer Scientist Narayanan Veeraraghavan

[Case study link](#)

Pricing model & Management tools on AWS

AWS pricing follows **pay-as-you-go model***, so you get charged only for resources launched and consumed. For cost optimization, AWS offers discounted **1-year/3-year commitment plans** for 7 eligible services*, as well as **EC2 Spot Instances*** (up to 90% off on-demand rate) for stateless, fault-tolerant use cases (e.g. HPC)

AWS also offers a range of [management tools*](#) for you to control and operate your cloud environment, such as **AWS Budgets*** to manage AWS service costs and **AWS CloudTrail*** to track AWS account activity

Ready to explore cloud for your research?



“Contact Us” to discuss on your use case and we will support you to:

- **Estimate cost for your research project**
- **Join AWS Research Cloud Program and apply for Cloud Credits**

Contact AWS Hong Kong Public Sector Team for project discussion

Andy Ho | Senior Business Development Manager

hk-aws-jucc-enquiry@amazon.com