

# Cloud Services Models "IaaS"

By

**Asst.L. Mohamed A. Abdul Hamed**

University of Basra

College of Computer Science and Information Technology


Computer Science Department

# Overview

- **Cloud Infrastructure as a Service (IaaS) .**
- **how IaaS works.**
- **IaaS platform and architecture.**
- **Bare-metal-as-a-Service.**
- **Amazon EC2 (Elastic Compute Cloud).**

# Infrastructure as a Service (IaaS)

- Infrastructure as a service (IaaS) is a service model that delivers computer infrastructure on an outsourced basis to support enterprise operations.
- Typically, IaaS provides hardware, storage, servers and data center space or network components; it may also include software.
- Infrastructure as a service (IaaS) is also known as hardware as a service (HaaS).
- An IaaS provider provides policy-based services and is responsible for housing, operating and maintaining the equipment it provides for a client.
- Clients usually **pay on a per-use** or utility computing basis.



# how IaaS works

- IaaS customers access resources and services through a wide area network (WAN), such as the internet, and can use the cloud provider's services to install the remaining elements of an application stack.
- For example, the user can log in to the IaaS platform to create virtual machines (VMs); install operating systems in each VM; deploy middleware, such as databases; create storage buckets for workloads and backups; and install the enterprise workload into that VM.
- Customers can then use the provider's services to track costs, monitor performance, balance network traffic, troubleshoot application issues, manage disaster recovery and more.

# IaaS platform and architecture

- IaaS is made up of a collection of physical and virtualized resources that provide consumers with the basic building blocks needed to run applications and workloads in the cloud.

**1- Physical data centers:** IaaS providers will manage large data centers, typically around the world, that contain the physical machines required to power the various layers of abstraction on top of them and that are made available to end users over the web.

- In most IaaS models, end users do not interact directly with the physical infrastructure, but it is provided as a service to them.

# IaaS platform and architecture

**2- Compute:** IaaS is typically understood as virtualized compute resources, so for the purposes of this article, we will define IaaS compute as a virtual machine.

- Providers manage the hypervisors and end users can then programmatically provision virtual “instances” with desired amounts of compute and memory (and sometimes storage).
- Most providers offer both CPUs and GPUs for different types of workloads.
- Cloud compute also typically comes paired with supporting services like auto scaling and load balancing that provide the scale and performance characteristics that make cloud desirable in the first place.

# IaaS platform and architecture

3- **Network:** Networking in the cloud is a form of Software Defined Networking in which traditional networking hardware, such as routers and switches, are made available programmatically, typically through Application Programming Interface (API).

- More advanced networking use cases involve the construction of multi-zone regions and virtual private clouds, both of which will be discussed in more detail later.

# IaaS platform and architecture

4- **Storage:** The three primary types of cloud storage are block storage, file storage, and object storage.

- Block and file storage are common in traditional data centers but can often struggle with scale, performance and distributed characteristics of cloud.
- Thus, of the three, object storage has thus become the most common mode of storage in the cloud given that it is highly distributed (and thus resilient), it leverages commodity hardware, data can be accessed easily over HTTP, and scale is not only essentially limitless but performance scales linearly as the cluster grows.



# Bare-metal-as-a-Service (BMaaS)

- Bare-metal-as-a-Service (BMaaS) provides an even lower level of control than traditional IaaS.
- In a BMaaS environment, resources are still provisioned on-demand, made available over the internet, and billed on a pay-as-you-go basis (typically in monthly or hourly increments).
- Unlike traditional IaaS, BMaaS does not provide end users with already virtualized compute, network, and storage; instead, it gives direct access to the underlying hardware.


# Bare-metal-as-a-Service (BMaaS)

- Unlike traditional IaaS, BMaaS does not provide end users with already virtualized compute, network, and storage; instead, it gives direct access to the underlying hardware.
- This level of access offers end users almost total control of their hardware specs.
- Given the hardware is neither virtualized nor supporting multiple virtual machines, it also offers end users the greatest amount of potential performance, something of significant value for use cases like HPC and GPU computing, high-performance databases, analytics workloads, and more.

# laaS example



Amazon Machine  
Image (AMI)



Amazon EC2



# Amazon EC2 (Elastic Compute Cloud)

- Amazon Elastic Compute Cloud (Amazon EC2) is a web-based service that allows businesses to run application programs in the Amazon Web Services (AWS) public cloud.
- Amazon EC2 allows a developer to spin up virtual machines (VMs), which provide compute capacity for IT projects and cloud workloads that run with global AWS data centers.
- An AWS user can increase or decrease instance capacity as needed within minutes using the Amazon EC2 web interface or an application programming interface (API).
- A developer can code an application to scale instances automatically with AWS autoscaling.
- A developer can also define an autoscaling policy and group to manage multiple instances at once.

# How EC2 works

- To begin using EC2, developers sign up for an account at Amazon's AWS website.
- They can then use the AWS Management Console, the AWS Command Line Tools (CLI), or AWS Software Developer Kits (SDKs) to manage EC2.
- **Note 1:** The AWS Command Line Interface (AWS CLI) is an Amazon Web Services tool that enables developers to control Amazon public cloud services by typing commands on a specified line.
- **Note 2:** A software development toolkit (SDK) is a set of software tools and programs provided by hardware and software.

# How EC2 works

- A developer then chooses EC2 from the AWS Services dashboard and 'launch instance' in the EC2 console.
- At this point, they select either an Amazon Machine Image (AMI) template or create an AMI containing an operating system, application programs, and configuration settings. The AMI is then uploaded to the Amazon S3 and registered with Amazon EC2, creating an AMI identifier.
- Once this has been done, the subscriber can requisition virtual machines on an as-needed basis.
- Data only remains on an EC2 instance while it is running, but a developer can use an Amazon Elastic Block Store volume for an extra level of durability and Amazon S3 for EC2 data backup.

# How EC2 works

- VM Import/Export allows a developer to import on-premises virtual machine images to Amazon EC2, where they are turned into instances.
- EC2 also offers Amazon **CloudWatch** which monitors Amazon cloud applications and resources, allowing users to set alarms, view graphs, and get statistics for AWS data; and AWS Marketplace, an online store where users can buy and sell software that runs on AWS.



# Amazon EC2 instance types

- Instances allow developers to expand computing capabilities by 'renting' virtual machines rather than purchasing hardware.
- An EC2 instance is used to run applications on the Amazon Web Services infrastructure.
- Amazon EC2 provides different instance types, sizes and pricing structures designed for different computing and budgetary needs.
- In addition to general purpose instances, Amazon EC2 offers an instance type for compute, memory, accelerated computing, and storage-optimized workloads.
- AWS limits how many instances a user can run in a region at a time, depending on the type of instance. Each instance type comes with different size options corresponding to the CPU, memory and storage needs of each enterprise.



# Benefits of EC2

- Getting started with EC2 is easy, and because EC2 is controlled by APIs developers can commission any number of server instances at the same time to quickly increase or decrease capacity.
- EC2 allows for complete control of instances which makes operation as simple as if the machine were in-house.
- The flexibility of multiple instance types, operating systems, and software packages and the fact that EC2 is integrated with most AWS Services -- S3, Relational Database Service (RDS), Virtual Private Cloud (VPC) -- makes it a secure solution for computing, query processing, and cloud storage.

# References

- **Book by Realtime publisher, 'Dan Sullivan', 'The Definitive Guide to Cloud Computing'.**
- **Website, <https://searchaws.techtarget.com/definition/Amazon-Elastic-Compute-Cloud-Amazon-EC2>**
- **Website, [HTTPs://searchcloudcomputing.techtarget.com/definition/Windows-Azure](https://searchcloudcomputing.techtarget.com/definition/Windows-Azure).**