



# Infrastructure as a Service (IaaS)

COSC349—Cloud Computing Architecture

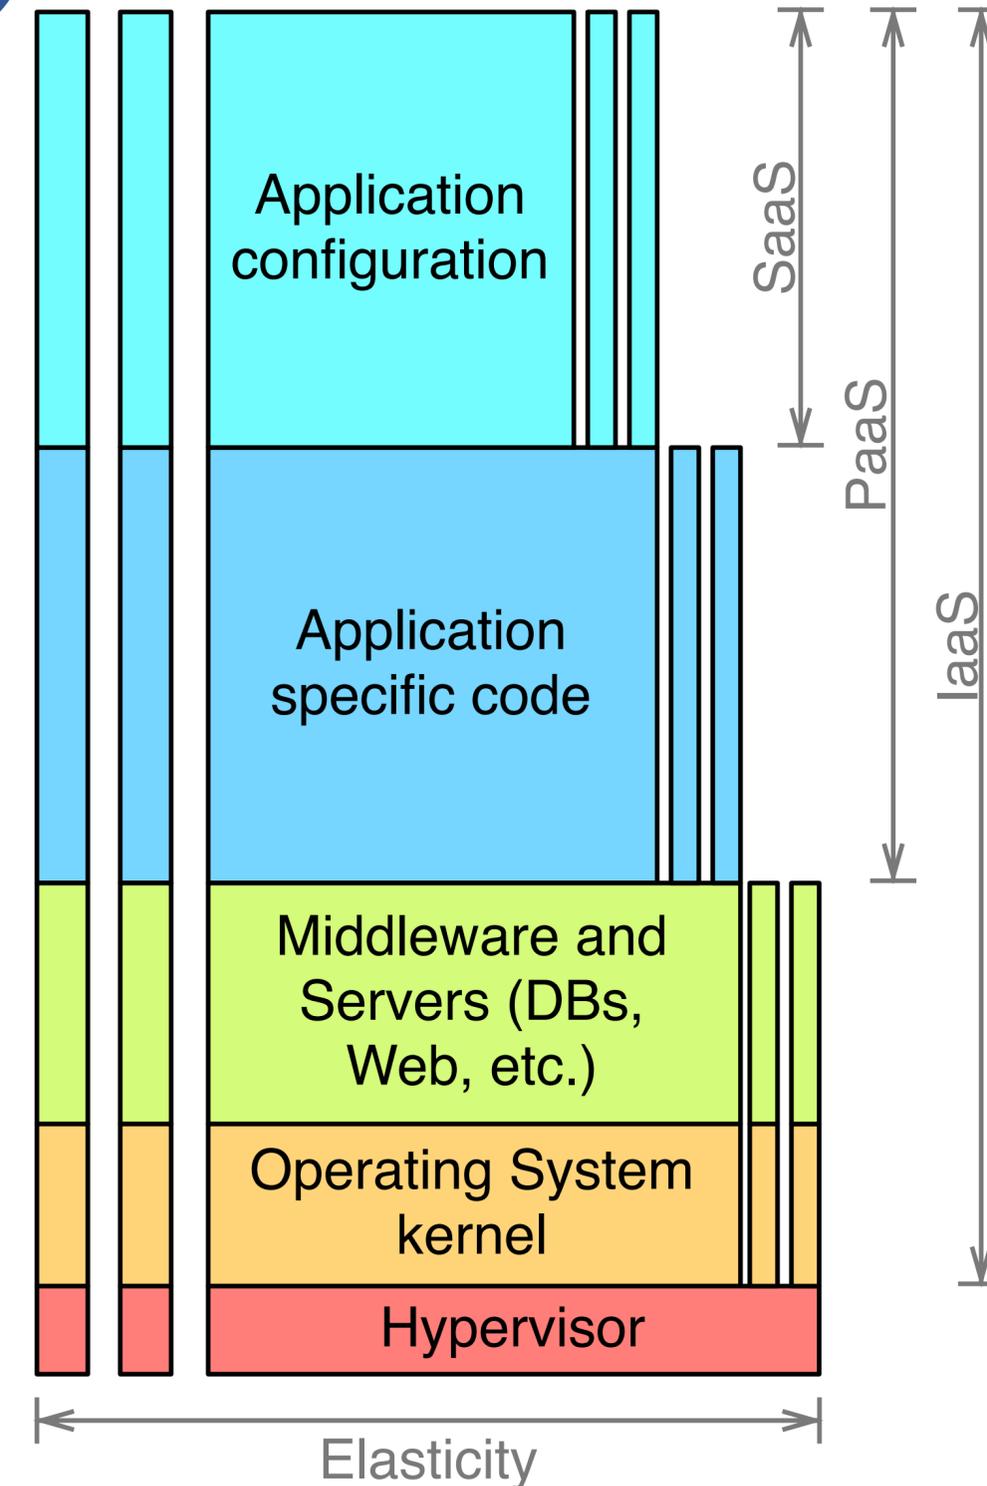
David Eyers

# Learning objectives

- Define **IaaS**
- Give examples of **public IaaS providers**
- Explain benefits and challenges using IaaS
- Describe how IaaS is usually **charged for**
- Sketch how VM **spot pricing** helps clients and providers
- Understand the core components required to be provided by an IaaS platform such as **OpenStack**

# Infrastructure as a Service (IaaS)

- Infrastructure as in **physical hardware**
  - ... but hardware is **virtualised as a service**
- Computing is typical IaaS commodity
  - **Networking** also needed to reach VMs
  - **Storage** typically provided in multiple ways:
    - Install OSs from existing templates
    - Means for VMs' data to be read & written



# Popular IaaS offerings in the public cloud

- Gartner report includes 2022 IaaS market shares:
  - 40.0%—**Amazon** Elastic Compute Cloud (EC2)
  - 21.5%—**Microsoft Azure**
  - 7.7%—**Alibaba**
  - 7.5%—**Google**
  - 4.4%—**Huawei**
  - 18.9%—others
- Top-5 cover over three-quarters of total market share
  - Also, Amazon's lead is significant (but diminishing)

# Advantages of IaaS over other models

- IaaS tenants have the greatest **power and flexibility**
  - Anything that can run virtualised can probably be deployed
- **Low levels of lock-in** to any given provider
  - Typical VM resources are largely interchangeable
    - e.g., VM images can be moved from EC2 to/from local VirtualBox
    - Amazon Linux 2003 is available as a public download
- You can ensure you have high **Quality of Service (\$\$\$)**

# Downsides of IaaS

- Lots of deployment work is necessary before any application code can actually be run—**slow start-up**
  - You need to select the operating system you will use
    - Also, you must keep your operating system updated (*i.e.*, secure!)
  - VMs need to have disks, *etc.*, provisioned
- **All infrastructure management** is up to you
  - You manage scaling up your apps over multiple machines
  - You can build web application load-balancers, *etc.*, but the cloud already has them now, so you will be duplicating effort

# Charging—what do you pay for in IaaS?

- CPU-wise, your use is **opaque to cloud provider**
  - Typically rate for ‘size’ of VM: number of CPU cores; RAM; *etc.*
  - Then multiply by the amount of time VM is ‘powered on’
  - Often **pay by hour units**, but AWS has per-second Linux billing
- Other resources that are likely to be charged for:
  - **Network traffic** to/from Internet & across regions
    - Amazon EC2 doesn’t charge inwards from Internet
  - **Block storage & snapshots**—charged regardless of VM on or off
  - Additional **public IP addresses** for your VMs

# Amazon EC2 instance types

- At launch EC2 had a **one-size-fits-all** approach to VMs
  - 2007-10-16: EC2 gives choice: **small, medium & large**
  - 2020-08-06: At least **298** instance types
  - 2021-06-29: Nearly **400** instance types
- Classifications include
  - General purpose—balance CPU, RAM, storage
  - Compute optimised, Memory optimised, Storage optimised
  - Accelerated—GPGPUs, HPC and FPGAs
  - (Bare metal instance types were not included)

# Amazon's VM spot pricing

- Computers don't wear down proportionally to their use
  - Thus cloud providers have incentive to **ensure high utilisation**
  - However, they need to ensure they don't violate guarantees
    - *i.e.*, claiming CPU allocations are available when they aren't
- Amazon address this (in part) using **spot pricing**
  - Heavily discounted prices on VMs... but these VMs **can be terminated** with just two minutes' notice!
    - You choose: hibernate; stop or terminate your VM
- AWS also discounts long-term use—**reserved instances**

# “Local” IaaS options: Catalyst Cloud

- IaaS platforms exist within New Zealand—ideal when:
  - **data sovereignty** is a concern
  - **fast network response** times for NZ clients is crucial
- Catalyst Cloud—<https://catalystcloud.nz/>
  - Three regions within NZ (~AWS Availability Zones)
  - Likely more expensive than AWS, etc., but price isn't everything
  - Provide free credit for trying out their services
- Spark & others offer IaaS within NZ... soon MS+AWS...

# How do you run an IaaS cloud?

- Hybrid and private cloud models require **local clouds**
- Organisations like telcos have large data centres
  - Increasingly software for those organisations will be virtualised
  - Thus offering cloud services is about **management and billing**
- Existing enterprise virtualisation platforms now do cloud
  - e.g., VMware provides tools for managing fleets of VM hosts
- Numerous open source products such as **OpenStack**
  - Key is the effecting **delegated administration**—*i.e.*, self-service

# OpenStack has many components

- (Taxonomy is mine)
- **Core VM support**
  - Compute: Nova
  - Networking: Neutron
  - Block storage: Cinder
  - Image: Glance
- **Management**
  - Identity: Keystone
  - Dashboard: Horizon
  - Orchestration: Heat
- Workflow: Mistral
- Telemetry: Ceilometer
- Messaging: Zaqar
- DNS: Designate
- Search: Searchlight
- Key manager: Barbican
- Container orchestration: Magnum
- Root Cause Analysis: Vitrage
- Rule-based alarm actions: Aodh
- **Cloud Services**
  - Object storage: Swift
  - Database: Trove
  - Elastic map reduce: Sahara
  - Bare metal: Ironic
  - Shared file system: Manila