

Cloud Computing

Nikos Parlavantzas

Course overview

- Objectives
 - Understand the main cloud computing concepts, architectures and technologies
 - Gain practical experience with modern cloud platforms

Course contents

- Lectures
 - Introduction to cloud computing
 - IaaS
 - PaaS
- Practicals
 - AWS

Today's lecture

- Overview of cloud computing
- Types of clouds
- Virtualisation

Collaboration service demand skyrocketed amid the pandemic

2,900%

Growth of daily participants between December and April 2020

zoom

2.5 million

New connected users in 1 week, up 25%

 **slack**



12 million

New daily active users in 1 week, up 37.5%

Microsoft Teams

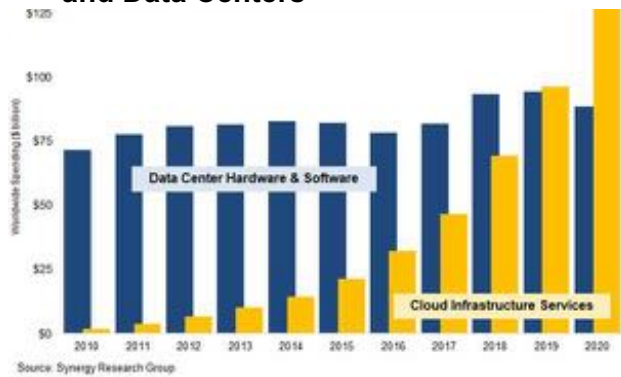
15.8 million

New subscribers between January and March


NETFLIX

All made possible by the cloud


Worldwide Enterprise Spending on Cloud and Data Centers



Total EU expenditure on Public Cloud



http://ec.europa.eu/newsroom/document.cfm?doc_id=41184



<https://www.grandviewresearch.com/industry-analysis/cloud-computing-industry>

What is the cloud?



An old dream

"If computers of the kind I have advocated become the computers of the future, then computing may someday be organized as a **public utility** just as the telephone system is a public utility...The computer utility could become the basis of a new and important industry"




*John McCarthy
MIT Centennial (1961)*

Power grid analogy



- *Private power generators*



- *Electricity as utility*

What is the cloud?



Delivering computing resources as a service



What is the cloud?

- *“A model for enabling ubiquitous, convenient, on-demand network access to a shared pool of configurable computing resources (e.g. networks, servers, storage, applications, and services) that can be rapidly provisioned and released with minimal management effort or service provider interaction”*
 - National Institute of Standards and Technology (NIST)

Essential characteristics

- On-demand self-service
- Broad network access
- Resource pooling
- Elasticity
- Metered service

On-demand self-service

- Consumers can easily manage resource capacity as needed automatically
 - No human interaction with cloud provider
 - Based on APIs and user interfaces



Broad network access

- Consumers are physically separated from the computing capabilities
- Those capabilities are available over a network and accessed through standard mechanisms and devices



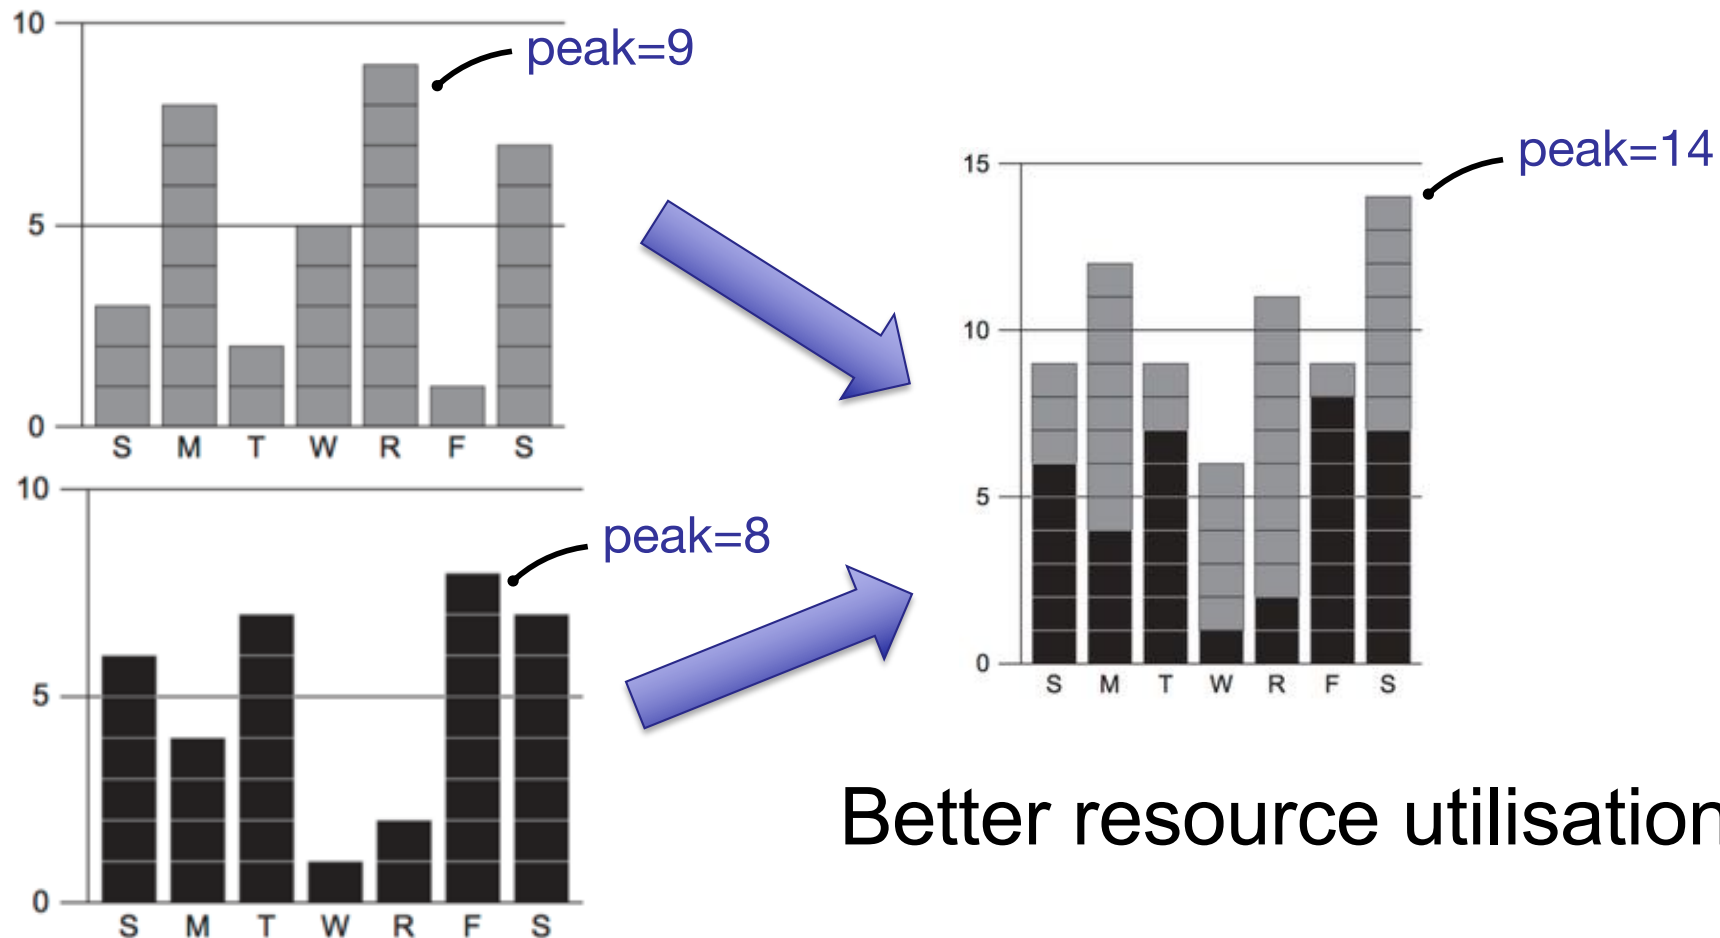
Resource pooling

- Computing resources are put into a common pool and shared among multiple consumers
 - Dynamic resource allocation based on consumer demand



Resource pooling


- Aggregate demand is smoother than individual customer demands



Better resource utilisation


Elasticity

- Ability to rapidly add and remove resources in response to demand
 - Enables matching resources to demand
 - Vertical scaling: bigger servers
 - Horizontal scaling: more servers



Facing dynamic demand


- Over-provisioning




Resource waste

Facing dynamic demand

- Under-provisioning




Lost revenue



Lost users

Facing dynamic demand

- Elastic provisioning



Meeting demand with
minimum cost

Metered service

- The resource usage of consumers is continuously measured
 - Forms basis for billing and ensuring fair use

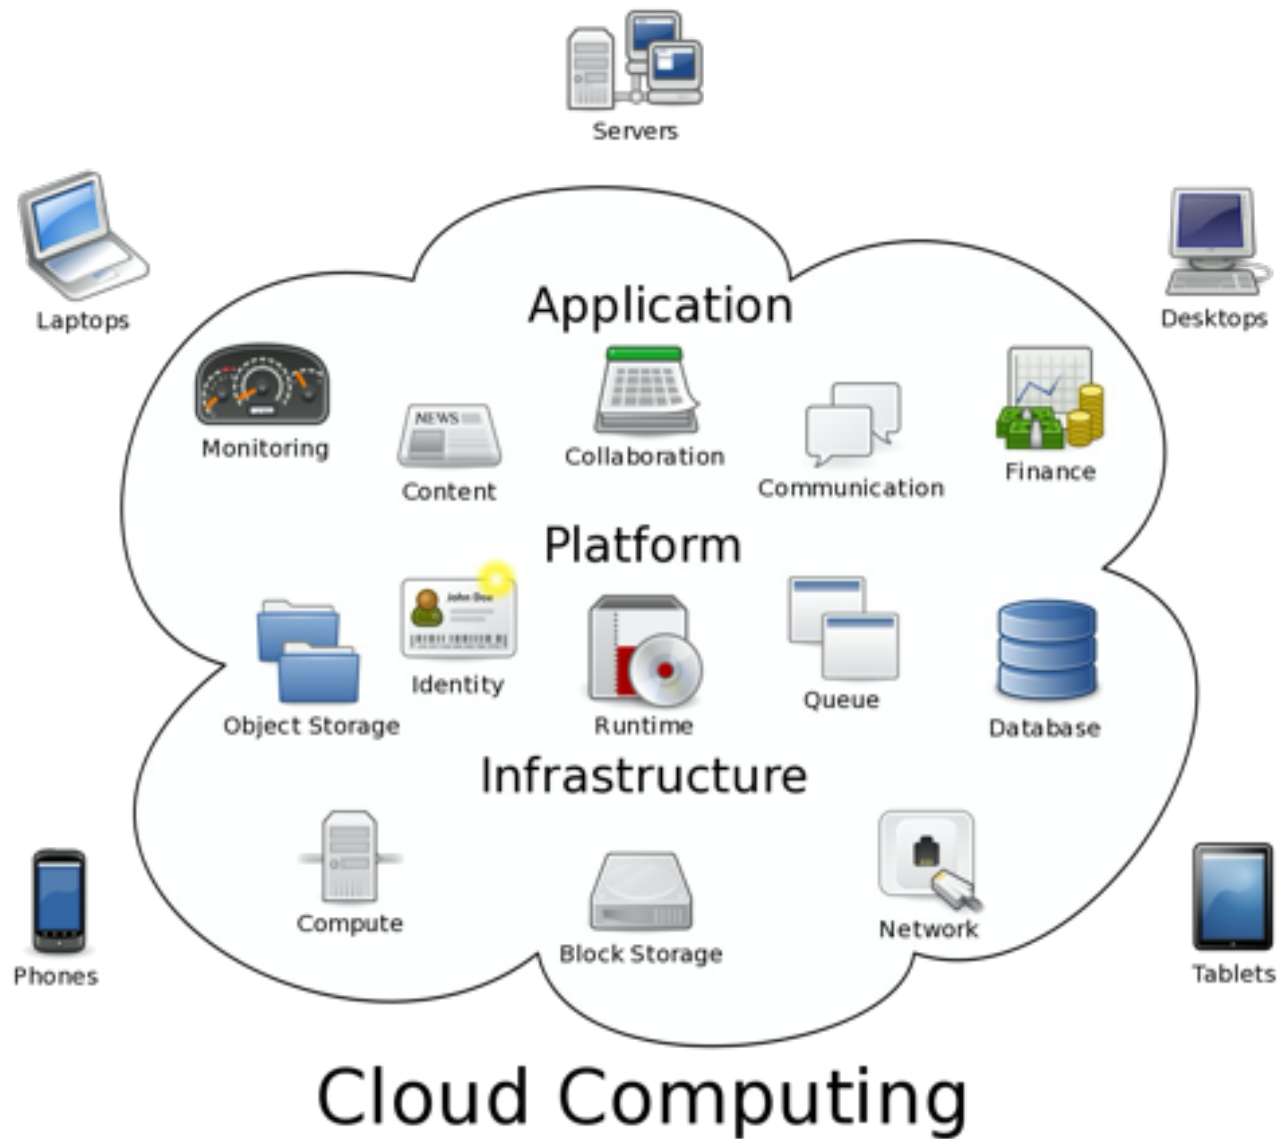


Cloud pricing

- Usage-based pricing (pay-per-use)
 - Most common model
 - E.g., per server hour, per GB transferred
- Subscription Pricing
 - Fixed fee per time period
- Dynamic Pricing
 - E.g., AWS Spot Instances
- Many variations
 - E.g., AWS Reserved Instances


Types of clouds

Everything as a Service




Service models

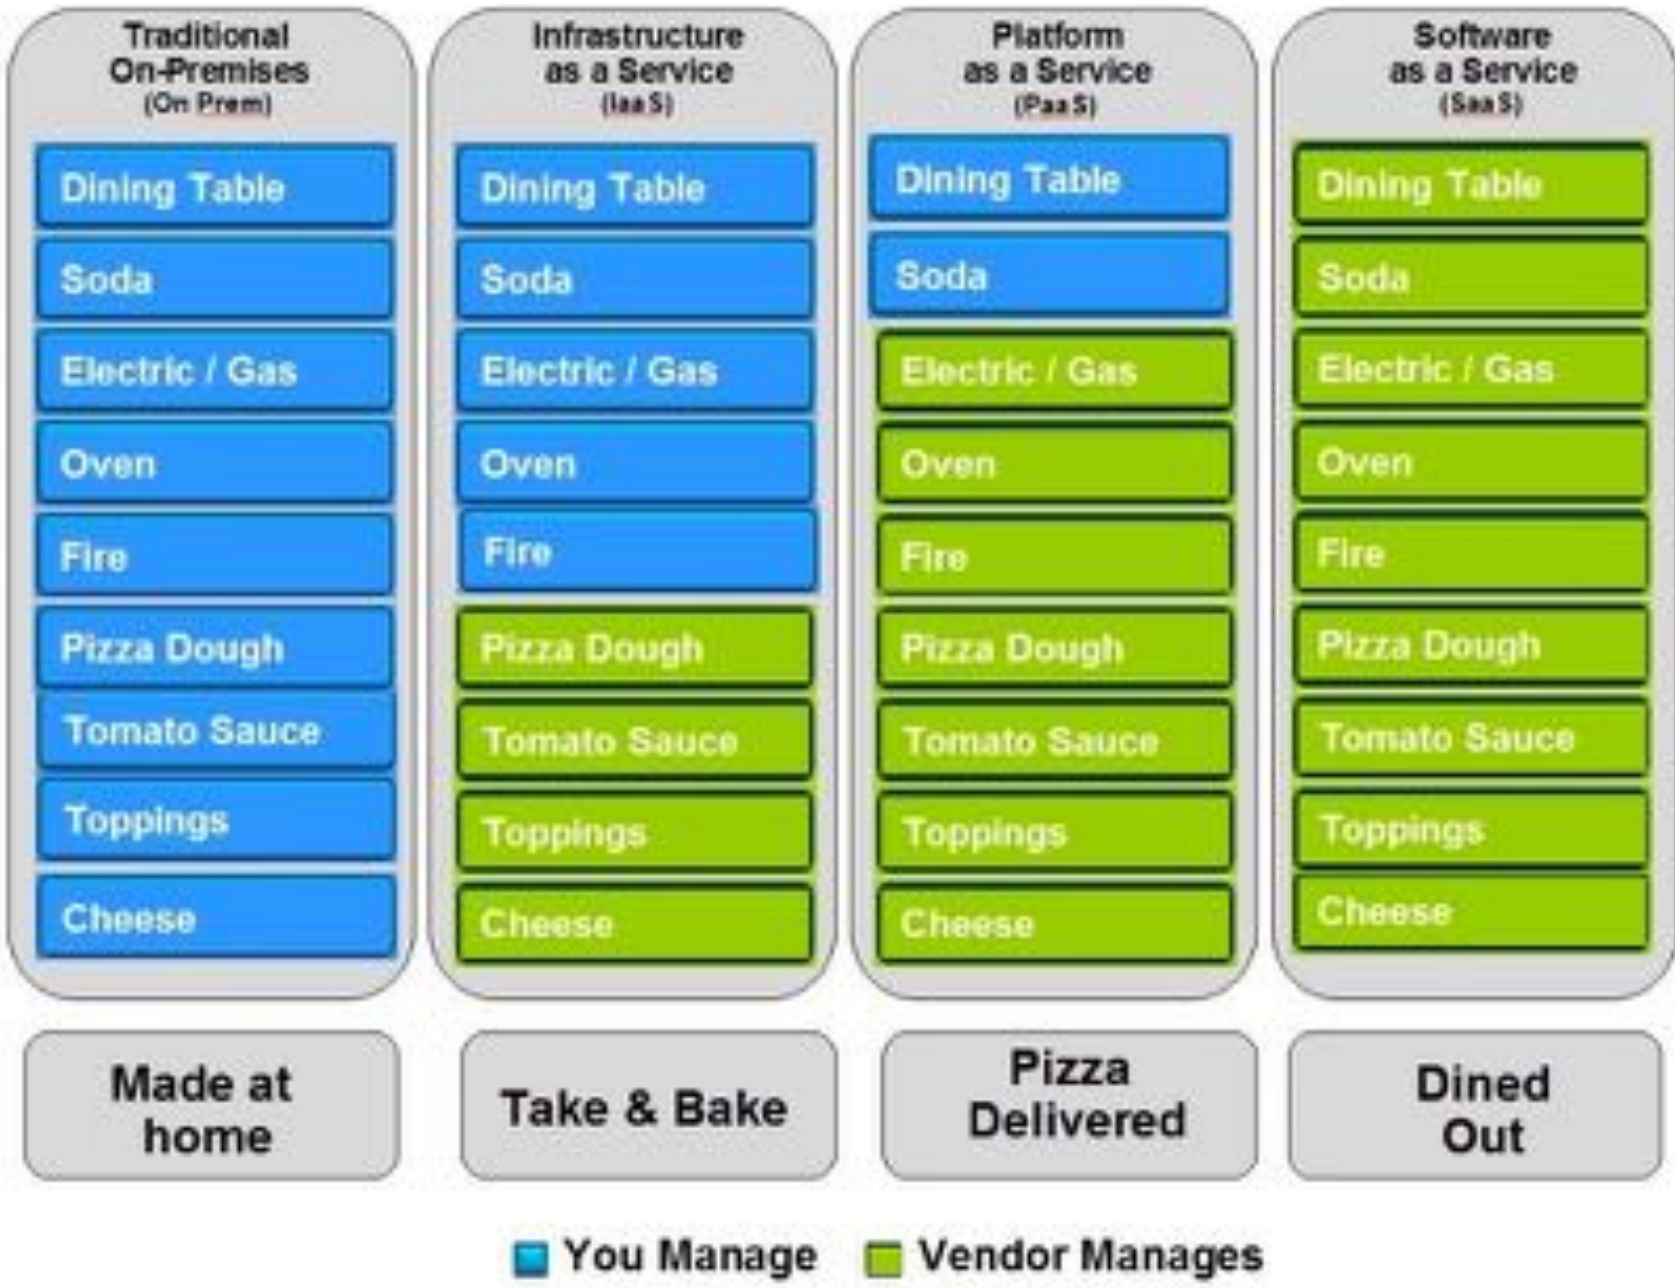
- **Infrastructure as a Service (IaaS)**
 - offers storage and compute resources
- **Platform as a Service (PaaS)**
 - offers development environments used to create applications
- **Software as a Service (SaaS)**
 - offers purpose-built business applications



Service models



Pizza as a service



Infrastructure as a Service

- The cloud provider delivers raw computing resources (typically virtualised)
 - servers, storage, networking, ...
- Consumers use these resources to deploy and run arbitrary software, including operating systems and applications



Platform as a Service

- The cloud provider delivers a complete application development and hosting environment
 - APIs, IDE plug-ins, services, tools, ...
- Consumers write, deploy and manage their applications using this environment



Software as a Service

- The cloud provider delivers running applications
- Consumers access the applications using thin clients
- Typically priced using a subscription fee

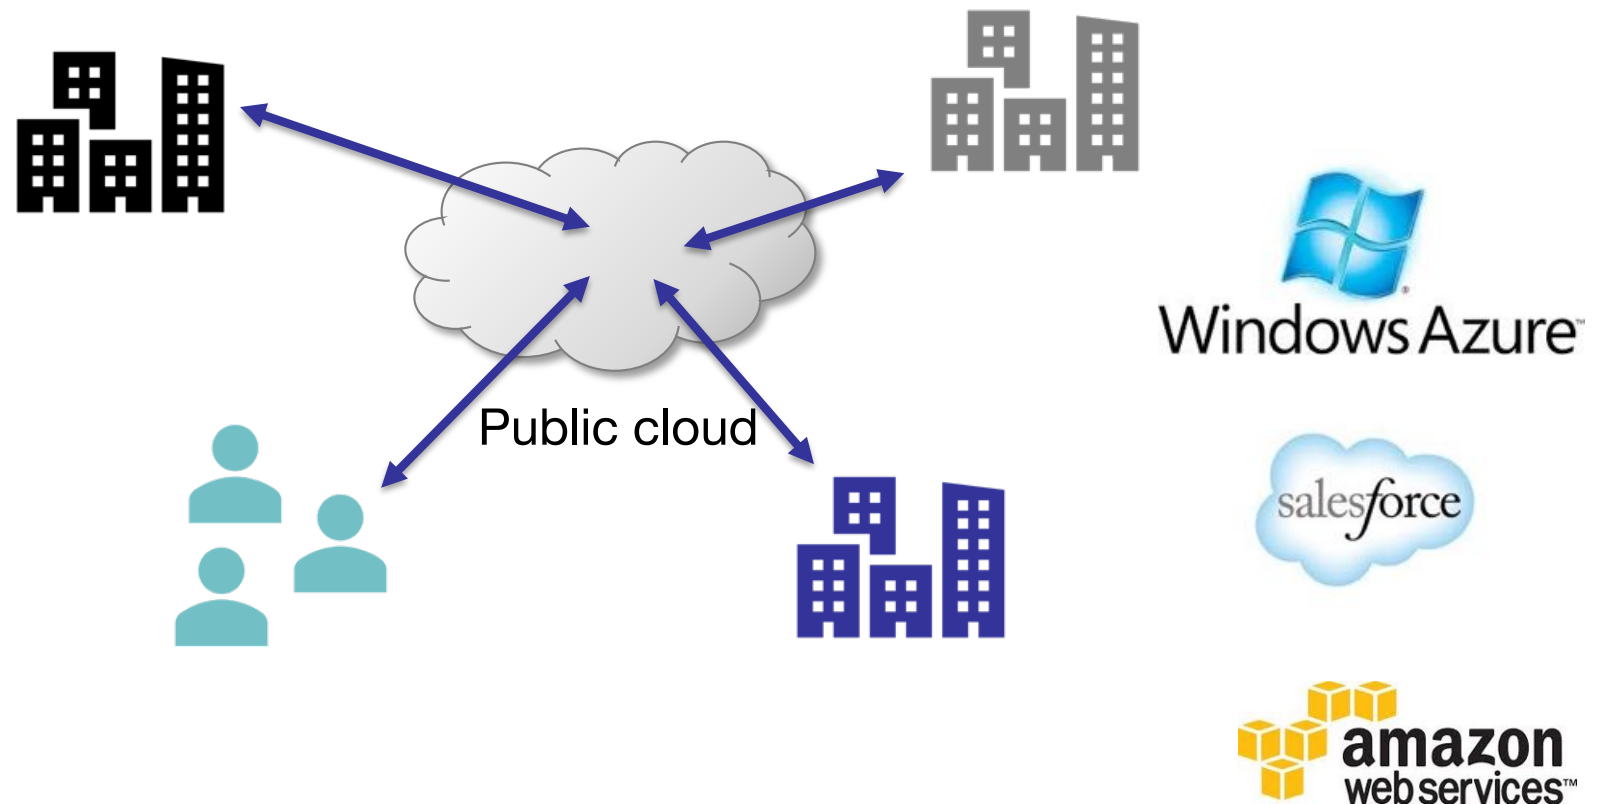


Deployment models

- Public cloud
- Private cloud
- Hybrid cloud

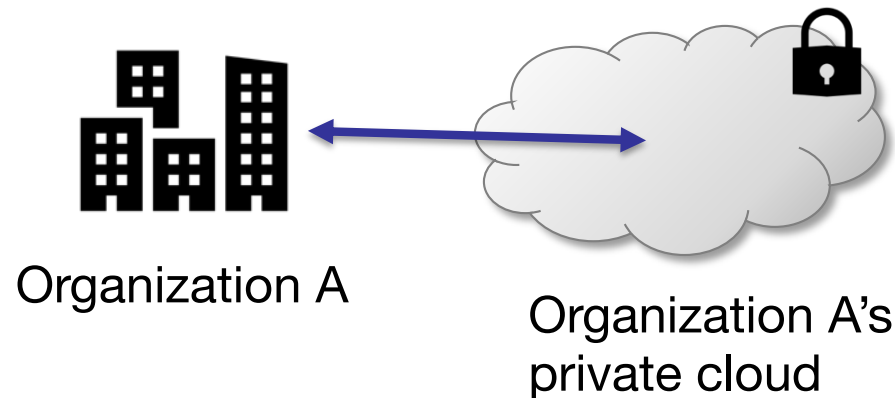
Public cloud

- Open to use to the general public; owned by an organisation selling cloud services



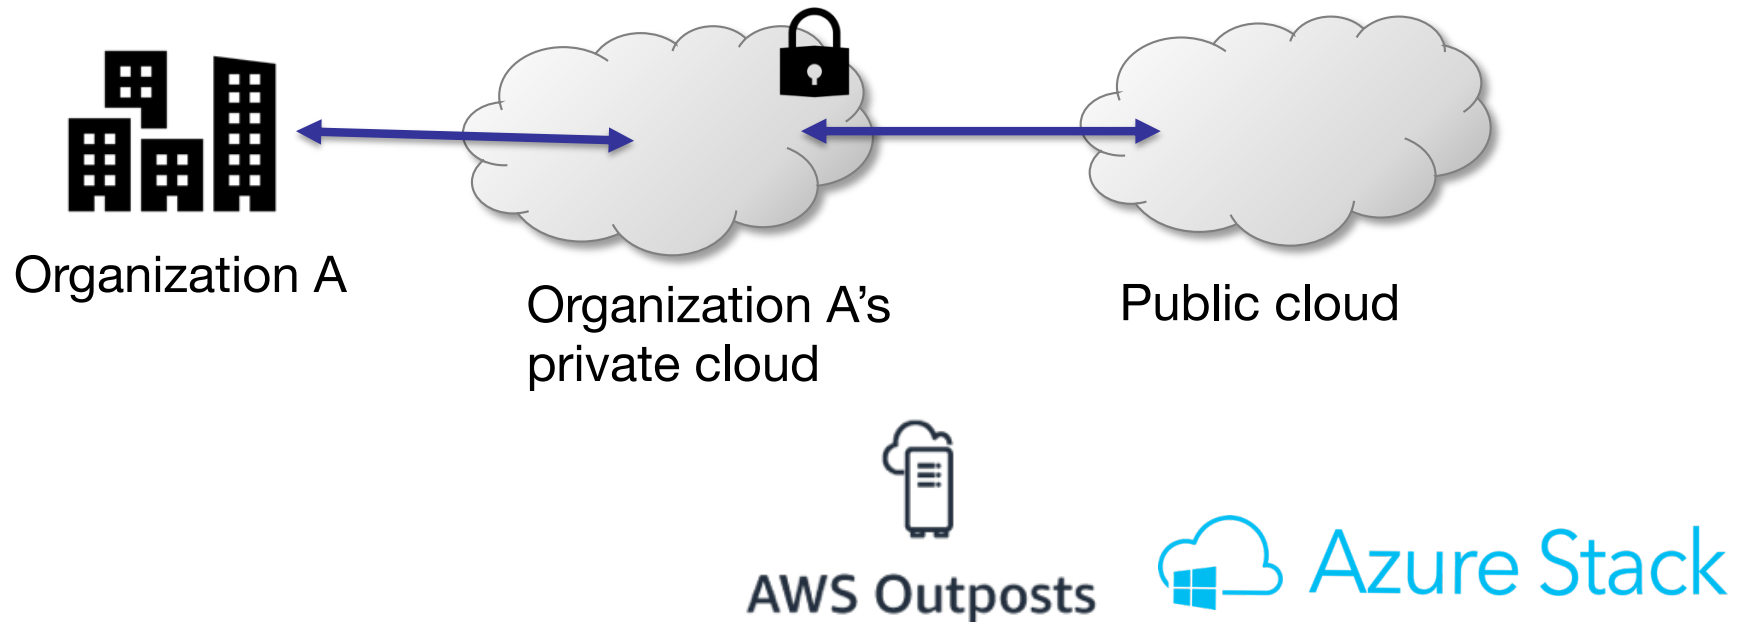
Private cloud

- Available for exclusive use by a single organisation; it may be managed by the organisation or a third party and may exist on premise or off premise



Hybrid cloud

- A combination of two or more clouds enabling data and application portability
e.g., “*cloud bursting*” for handling demand spikes



Cloud benefits

- Economic benefits
 - Lower cost of using resources because of elasticity and usage-based pricing
 - Even if clouds are more expensive on a unit-cost basis, the cloud can be cheaper overall
 - Lower cost of provisioning resources because of pooling
 - Change from capital to operational expenses



Cloud benefits

- Agility benefits
 - Reduced time to market
 - No need for months of planning, purchasing, provisioning, and configuring.
- Security benefits
 - Running secure data centres is *not* the core competence of typical enterprises



Cloud challenges

- Data confidentiality
 - Where is my data?
 - Who can have access ?
- Performance and availability
 - How long do I wait?
 - Is the data transfer speed high enough?
 - What if there is an outage in the cloud?
- Vendor lock-in
 - Can I move from one cloud to another?



Enabling technology: Virtualisation

Virtualisation

- Technology that allows multiple virtual machines to run on a single physical machine



Virtualisation features

- *Isolation*: each virtual machine (VM) is isolated from the host physical machine and the other VMs
 - E.g., a VM may crash or be compromised without affecting other VMs
- *Encapsulation*: the complete state of a VM can be manipulated as a unit
 - E.g., VMs can be migrated across physical machines, suspended, or resumed
 - E.g., one can dynamically modify the resources allocated to a VM

Virtualisation types

- Hosted virtualisation
 - VirtualBox, VMware Workstation, ...
- Bare-metal (native) virtualisation
 - Xen, KVM, vSphere, Microsoft Hyper-V, ...



Hosted Architecture

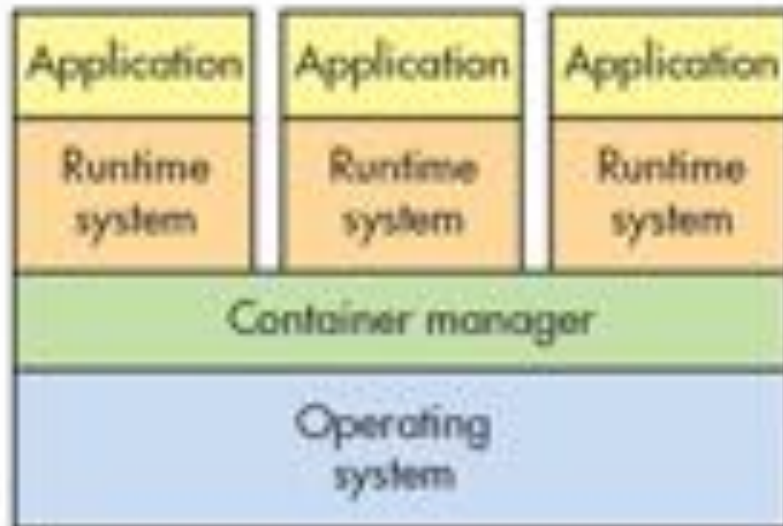


Bare-Metal Architecture

Virtualisation types

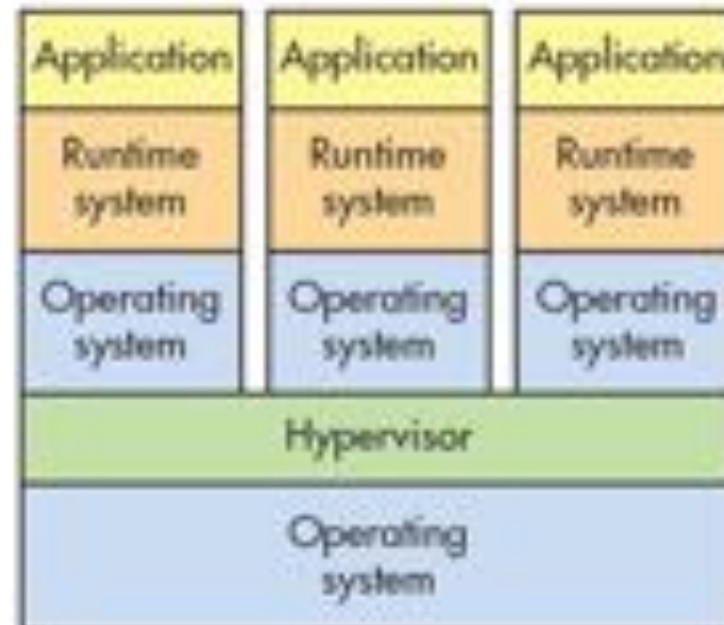
- OS-level virtualisation (container-based virtualisation)
 - Docker, LXC, ...

Containers



vs.

VMs



Virtualisation benefits

- Faster application development, testing, and deployment
- Cost reduction through server consolidation
- Enhanced availability
 - E.g., migration to accommodate planned maintenance; restoring from snapshots



Virtualisation challenges

- Performance degradation
 - Compared to running directly on physical machine
- Performance isolation
 - Interference between VMs/containers
- Security risks
 - Complexity and potential vulnerabilities introduced by virtualisation layer



Summary

- Cloud computing is about offering resources as a service; its main features include elasticity, metered service and on-demand self-service
- Cloud offerings can be categorised as IaaS, PaaS, or SaaS as well as public, private, or hybrid
- Virtualisation is a critical enabling technology for the cloud

References

- *The Cloud at Your Service*, Jothy Rosenberg and Arthur Mateos, Manning Publications Co., 2010
- *A view of cloud computing*, Michael Armbrust, et al., Commun. ACM 53, 4, 2010
- *Cloudonomics: The Business Value of Cloud Computing*, Joe Weinman, Wiley, 2012
- *Google App Engine*,
<https://developers.google.com/appengine>