Cloud and Datacenter Networking

Università degli Studi di Napoli Federico II

Dipartimento di Ingegneria Elettrica e delle Tecnologie dell'Informazione DIETI

Laurea Magistrale in Ingegneria Informatica

Prof. Roberto Canonico

Cloud Computing introduction

Lesson outline

- Cloud Computing: characteristics, service and deployment models
- Cloud Computing and Utility Computing: benefits and risks
- IaaS, PaaS and SaaS solutions and commercial offerings

The Utility Computing vision

- “Computing may someday be organized as a public utility”
  - John McCarthy, MIT Centennial in 1961
- Computing as the fifth utility, after Water, Gas, Electricity and Telephone
- "As of now, computer networks are still in their infancy. But as they grow up and become more sophisticated, we will probably see the spread of ‘computer utilities' which, like present electric and telephone utilities, will service individual homes and offices across the country."
  - Leonard Kleinrock, 1969
- Intuitions behind the Utility Computing vision:
  - Huge computational and storage capabilities made available from large pools
  - Metered billing (pay for what you use)
  - Simple to use interface to access these capabilities (e.g., plugging into an outlet)

Cloud Computing definition by NIST

- Cloud computing is 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
- *This cloud model is composed of:*
  - *five essential characteristics*
  - *four deployment models (public/private/hybrid/community)*
  - *three service models*
- http://www.nist.gov/itl/cloud/
- http://csrc.nist.gov/publications/nistpubs/800-145/SP800-145.pdf

NIST
National Institute of Standards and Technology

Cloud Computing: essential characteristics

1. **On-demand self services** - Cloud Computing services must be provided without requiring human interaction with any service provider
2. **Broad network access** - Cloud capabilities are made available over the Internet and accessed through standard mechanisms that promote use by heterogeneous client platforms
3. **Resource pooling and multi-tenancy** - The provider's computing resources are pooled together to serve multiple consumers using a multiple-tenant model, in which physical and virtual resources are dynamically assigned and reassigned according to consumer demand
4. **Rapid elasticity** - Cloud services must be rapidly and elastically provisioned, in some cases automatically, to quickly scale out and rapidly released to quickly scale in; to the consumer, the capabilities available for provisioning often appear to be unlimited and can be purchased in any quantity at any time
5. **Measured service** - Cloud computing resource usage can be measured, controlled, and reported providing transparency for both the provider and consumer of the service.
   Cloud computing services use a metering capability which enables to control and optimise resource use. This implies that, just like traditional utilities, Cloud services are charged per usage metrics according to a ***pay per use pricing model***

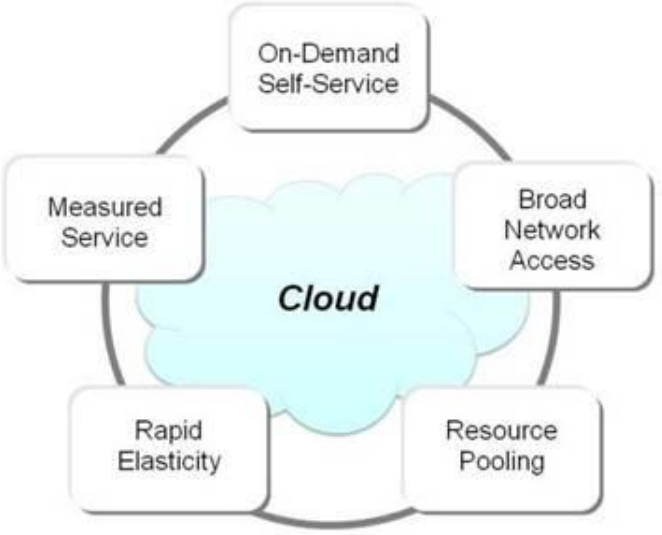

Cloud Computing: deployment models

- **Public cloud**
  - The cloud infrastructure is owned by an organization (the ***Cloud Service Provider***, CSP) whose business is to sell cloud services to the general public
- **Private cloud**
  - The cloud infrastructure is entirely owned by a single organization (e.g. an enterprise) which has internally consolidated its IT infrastructure in a few datacenters that host shared services on virtualized hardware
- **Hybrid cloud**
  - A hybrid cloud environment combines computing resources that are hosted on-site (i.e. within the premises of the organization) and off-site (in a public cloud). By integrating public cloud services, an organization can leverage cloud solutions for specific functions that are too costly to maintain on-premise such as disaster recovery or data backups
- **Community cloud**
  - A community cloud is formed when several organizations with similar requirements share portions of their own IT infrastructures to form a shared cloud infrastructure exclusively built to serve the needs of the organization forming the community

Cloud Computing: basic service models

- IaaS – *Infrastructure as a Service*
  - The cloud infrastructure provides “as a service” the fundamental computing resources (computing, storage, network)
  - Typical user: *system administrator*
- PaaS – *Platform as a Service*
  - The cloud infrastructure provides “as a service” a programmable and configurable execution environment to host packaged applications
  - Typical user: *application developer/administrator*
- SaaS – *Software as a Service*
  - The cloud infrastructure provides “as a service” complete suites of software applications to end users
  - Typical user: *end user*

A variety of services are today offered "on-demand" in a way that is coherent with the five essential characteristics of Cloud Computing → ***Everything as a Service*** **- XaaS**

Service Level Agreements (SLAs)

- Contract between customers and service providers
- An SLA determines the level of service to be provided
- Contains performance metrics (e.g., uptime, throughput, response time)
- Problem management details
- Documented security capabilities
- Contains penalties for non-performance (SLA violations)
- See http://www.sla-zone.co.uk

Cloud Computing: enabling factors

- Virtualization (VMs, containers, etc.)
- Ubiquitous and broadband access to the global Internet
- Standard communication protocols (HTTP, DNS, TCP/UDP, IPv4 and IPv6, etc.)
- Service Oriented Architectures and REST APIs
- Distributed Computing
- Free and Open Source Software
- Service Level Agreements

Traditional Infrastructure Dimensioning Model

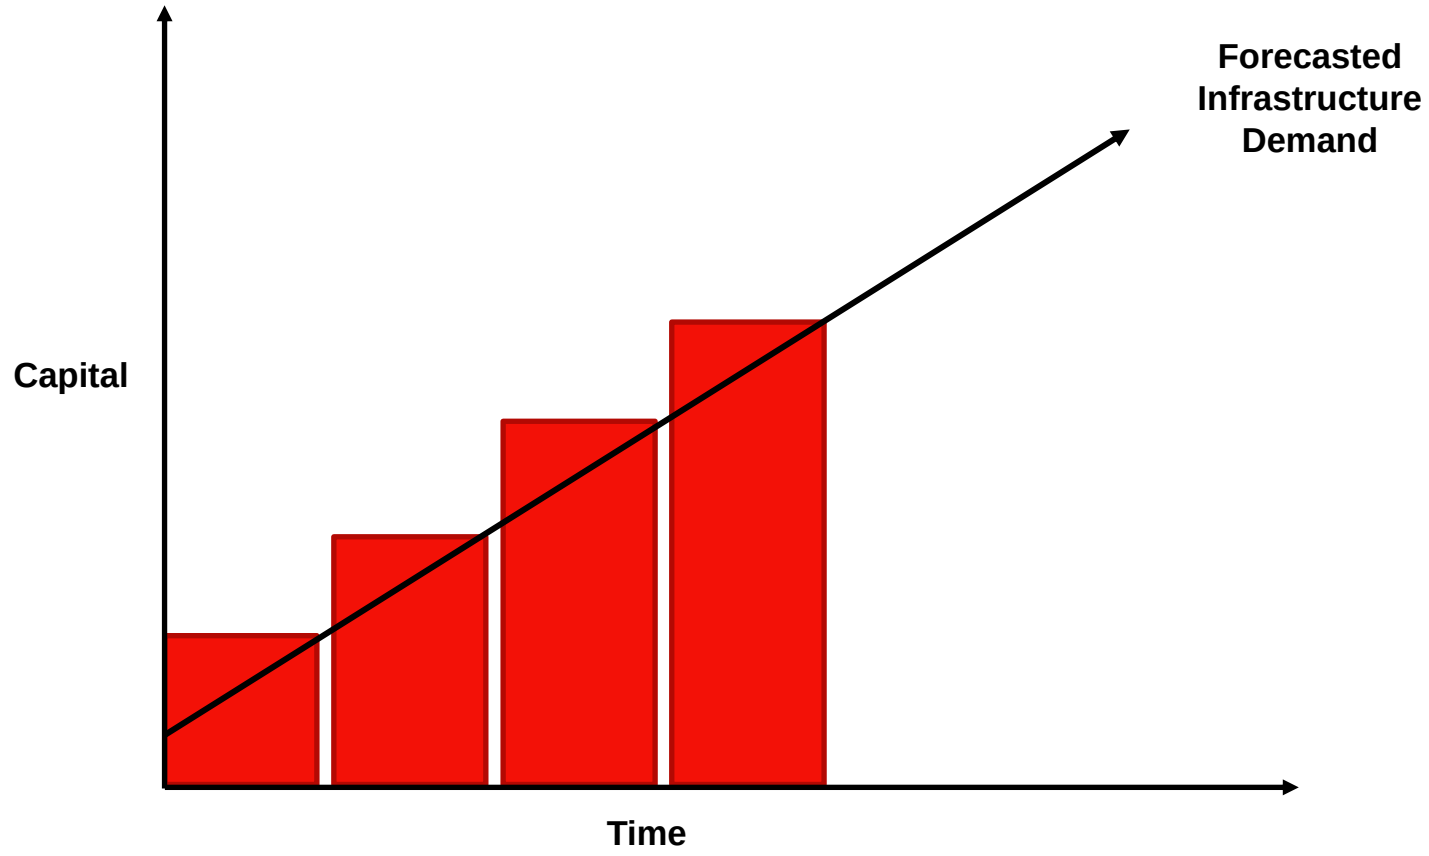

Acceptable Surplus

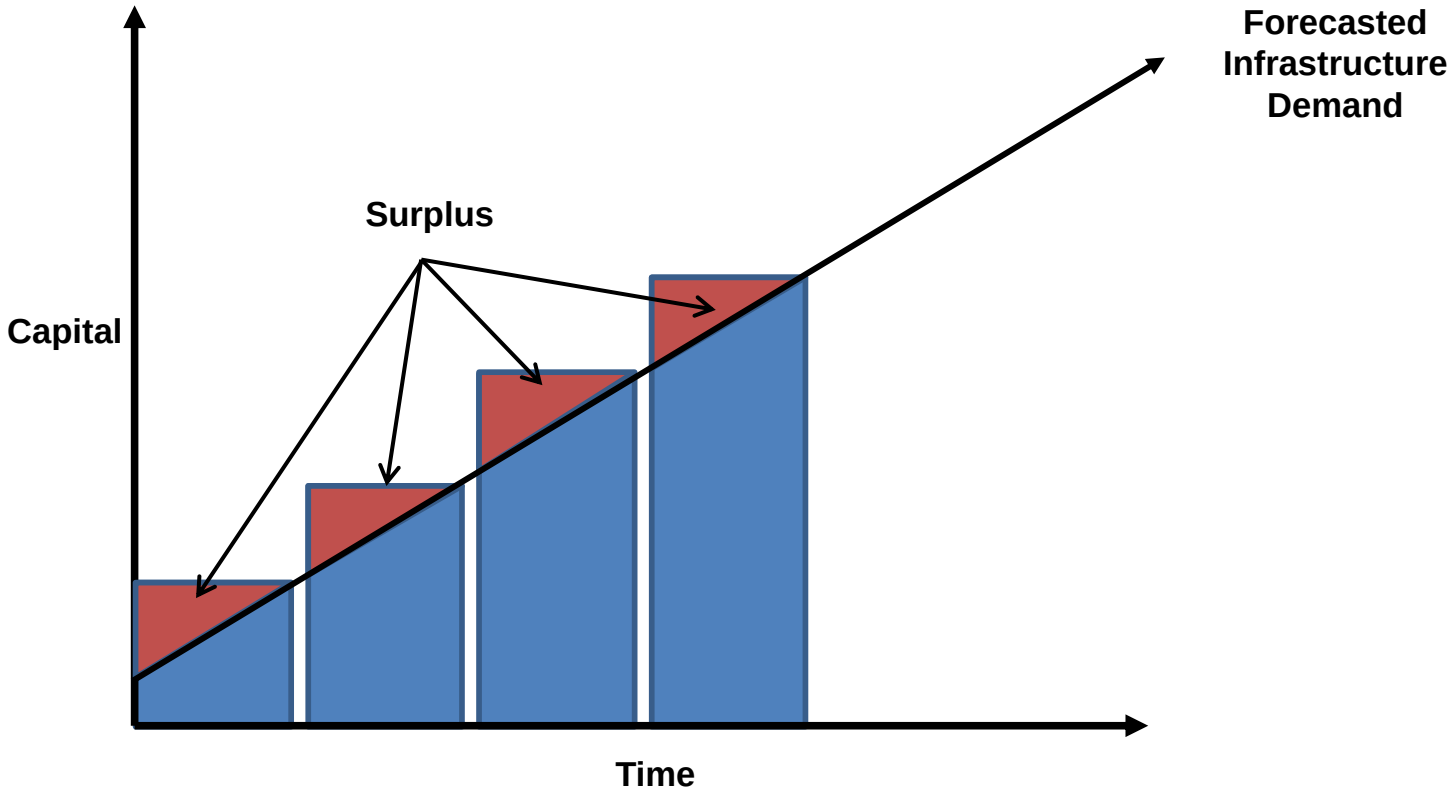

Traditional Dimensioning Model in the real world

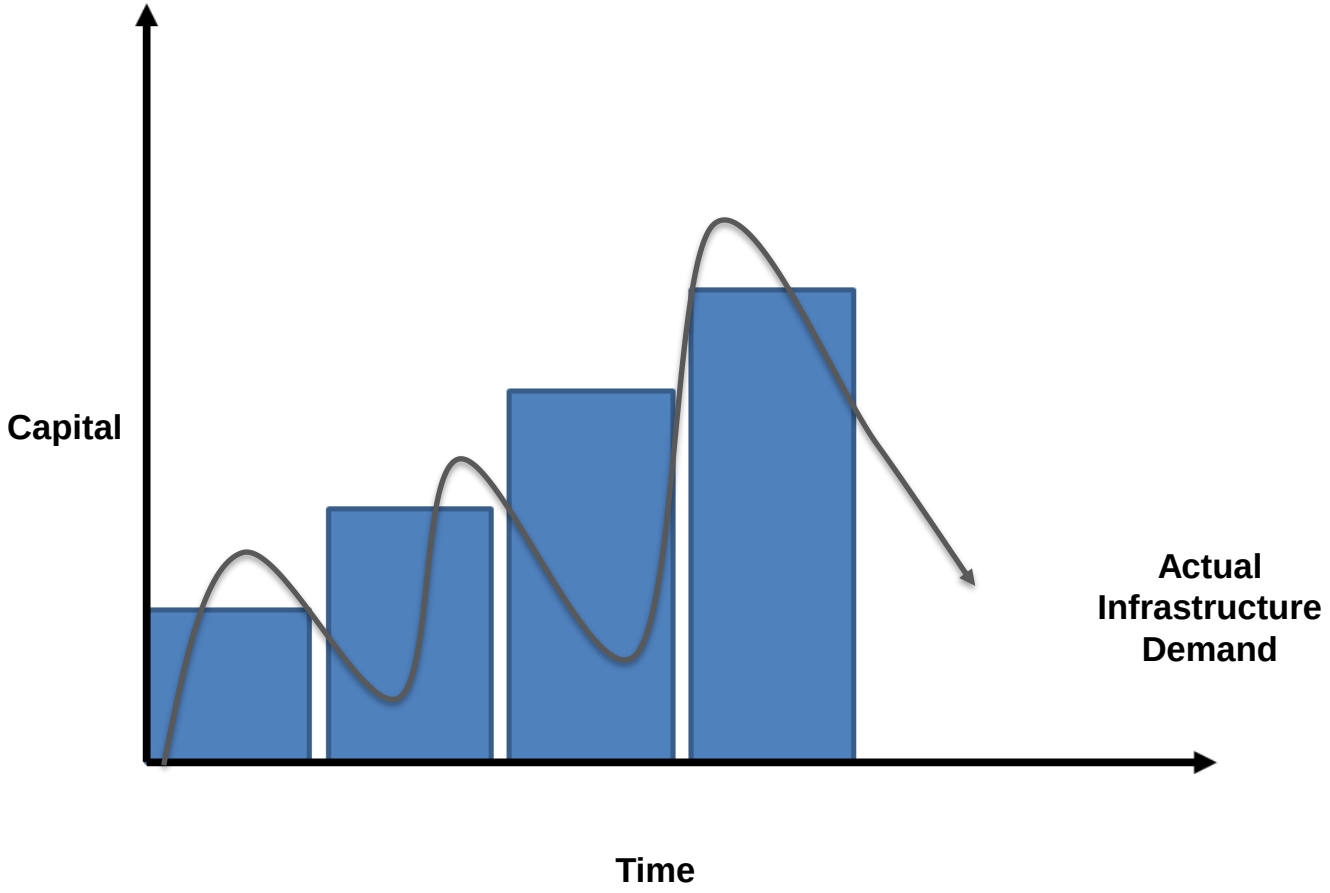

Unacceptable Surplus

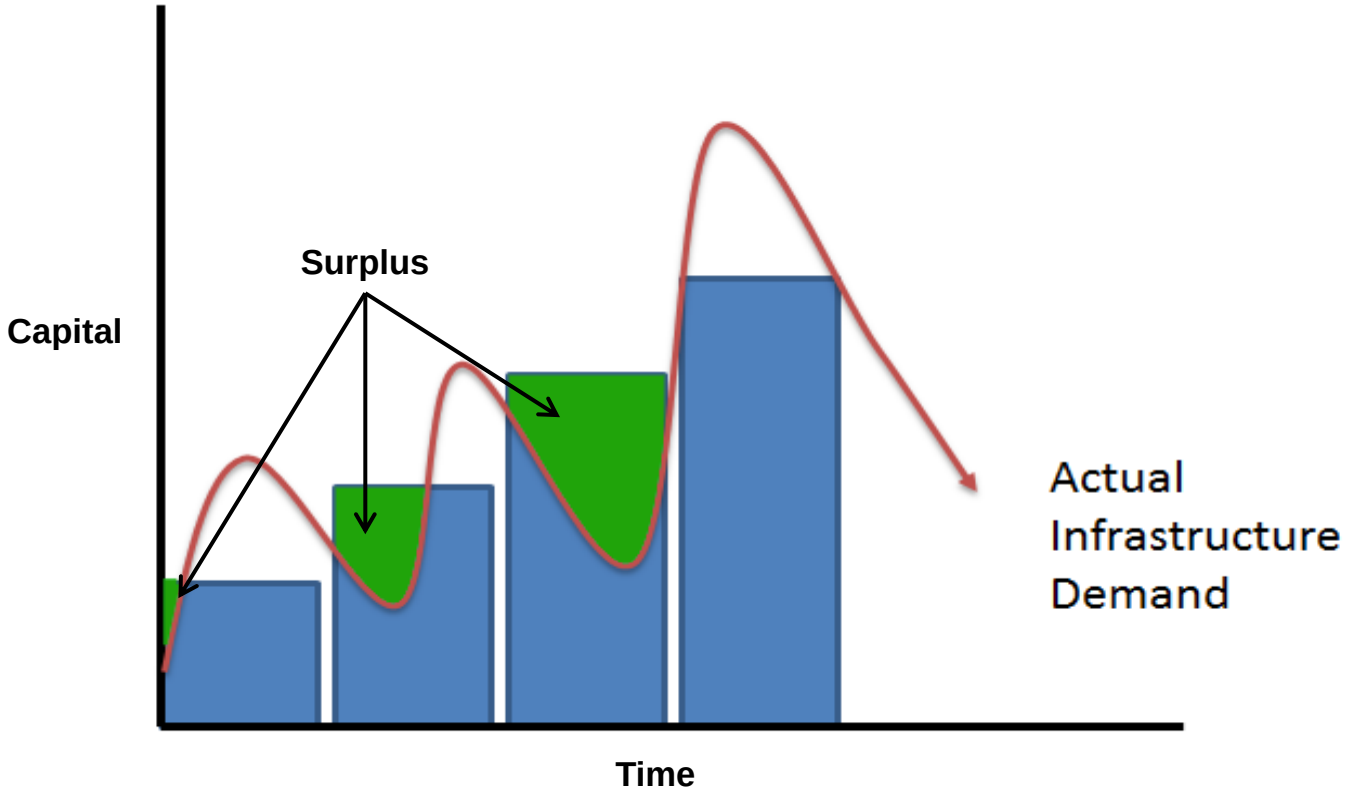

Unacceptable Deficit

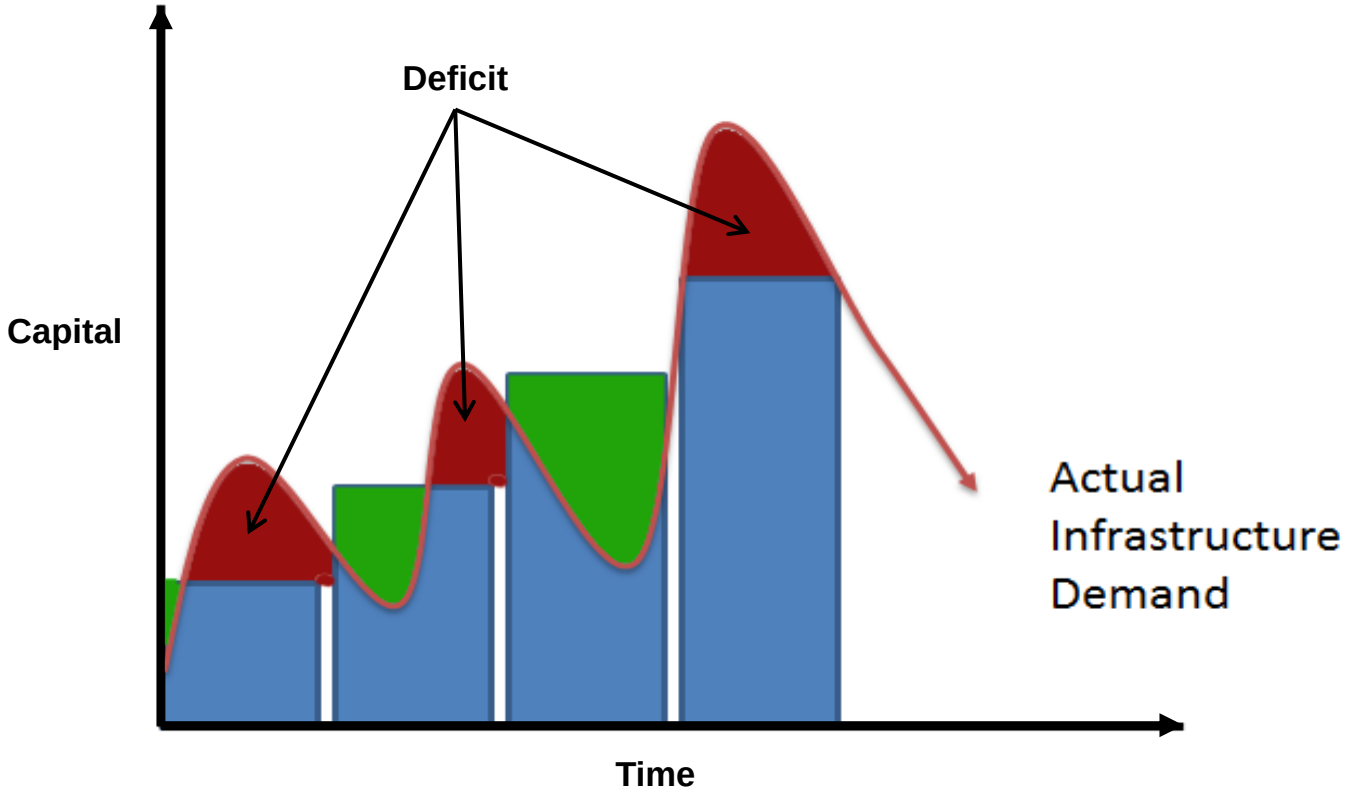

Utility Infrastructure Model

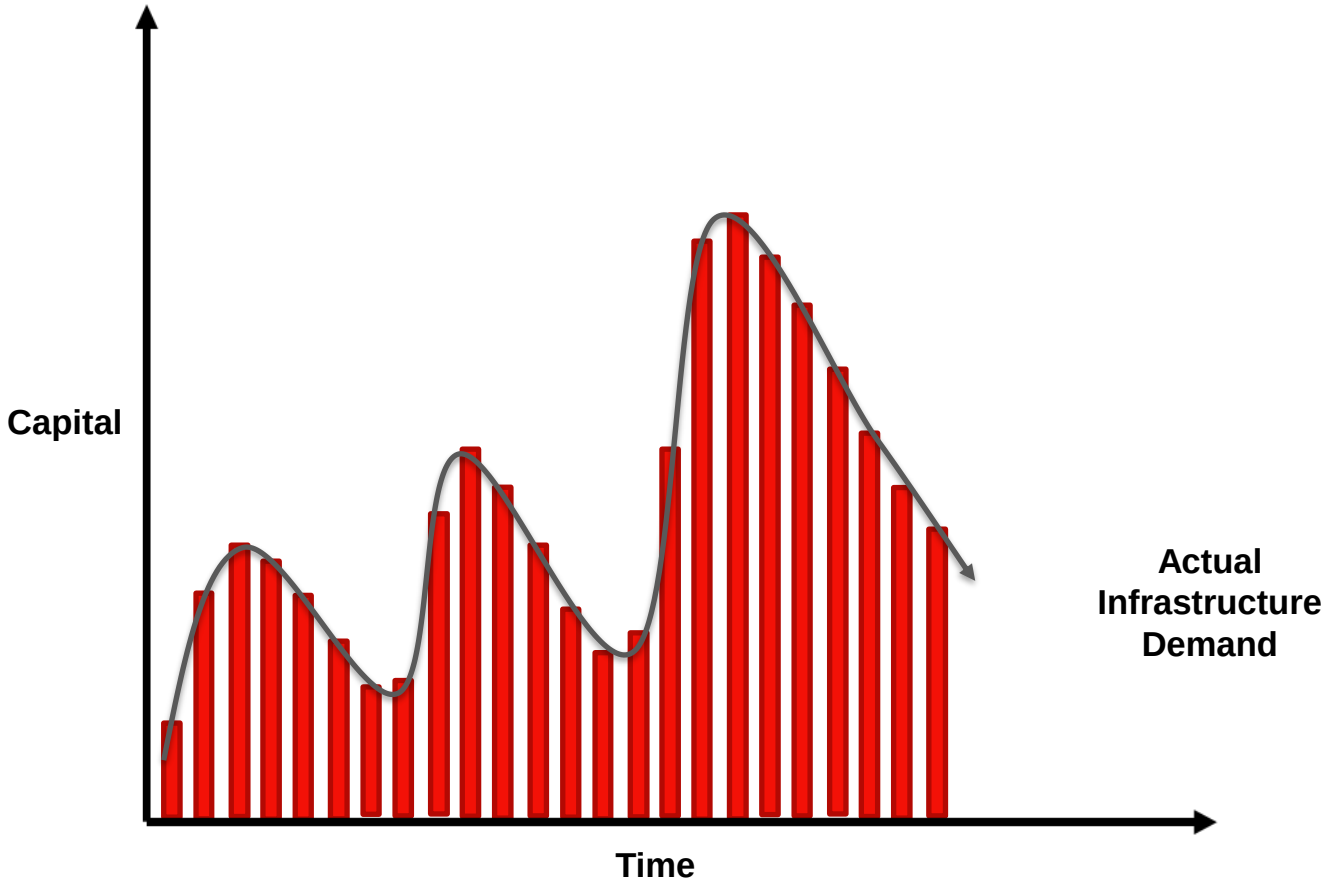

Infrastructure as a Service (IaaS)

- IaaS consists in delivering a technology infrastructure as an on demand scalable service
  - Usually billed based on usage
- IaaS is not managed hosting: traditional managed hosting is a form of web hosting where a user chooses to lease entire server(s) housed in an off-site data center
- IaaS providers rely on large-scale datacenters in which they gather large amounts of IT resources shared among customers (multi-tenancy)
  - Servers
  - Storage systems
  - Network devices
  - Specialized devices: IDS, VPN servers, firewalls, load balancers
- Use of virtualization at all levels (servers, storage, network) allows high levels of utilization of available resources
- Cloud datacenters need very good connectivity to the Internet
- Public IaaS providers have built several datacenters located in several countries all over the five continents for redundancy and to reduce access latencies

IaaS Examples

Platform as a Service (PaaS)

- PaaS provides all of the facilities required to support the complete life cycle of building and delivering web applications and services
- Three kinds of PaaS solutions:
  - Those which are coupled to a specific public IaaS platform
    - Many of these are directly provided by IaaS providers
  - Those which are compatible with several IaaS providers
    - Application developers may choose the IaaS provider
  - Those which are offered as a service that includes the acquisition of IaaS resources
    - The PaaS provider has its own resources or buys them from a 3rd party IaaS provider
    - The application developer does not directly buy IaaS services
- Characteristics of PaaS platforms:
  - Support for the creation of highly scalable and reliable multi-tier sw applications
  - Isolation guarantees in a multi-tenant environment
  - Support for several programming languages
  - Support for several operating systems to execute the hosted applications

PaaS Examples

Software as a Service (SaaS)

- SaaS is a software delivery methodology that provides licensed multi-tenant access to software and its functions remotely as a Web-based service
  - Usually billed based on usage
  - Usually multi tenant environment
  - Highly scalable architecture
- SaaS is not ASP 2.0 !
  - The old-days ASP model concentrated on providing an organization with the ability to move certain application processing duties to leased third-party managed servers
  - ASPs were not necessarily concerned about providing shared services to multiple tenants, but rather hosting a dedicated application on behalf of the customer

SaaS Examples

Software-as-a-Service: Google apps

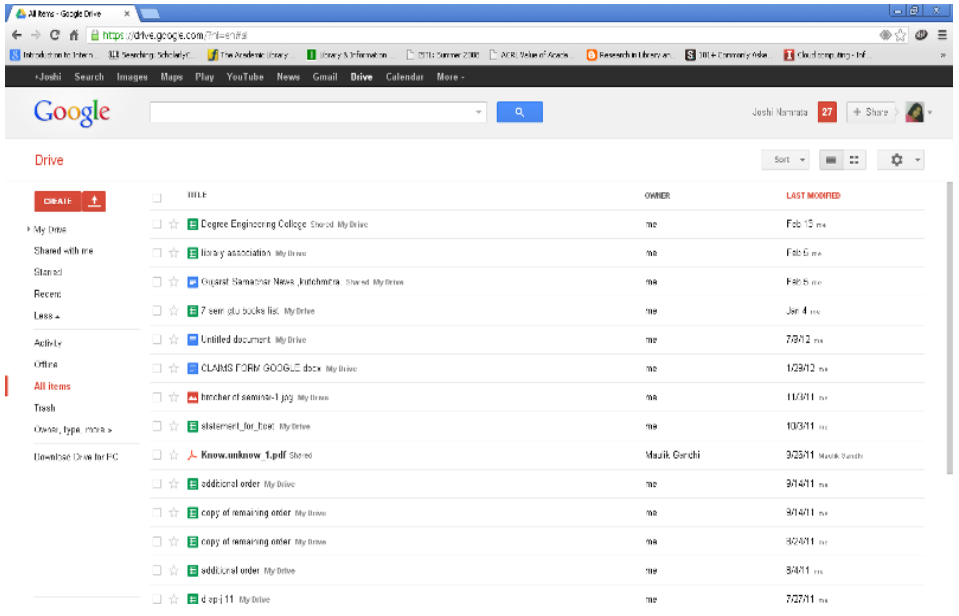

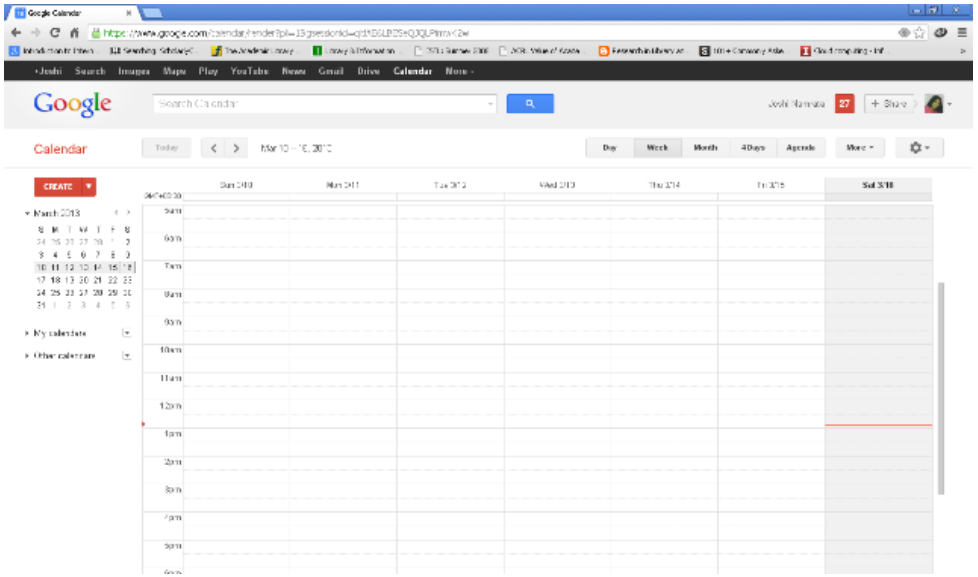

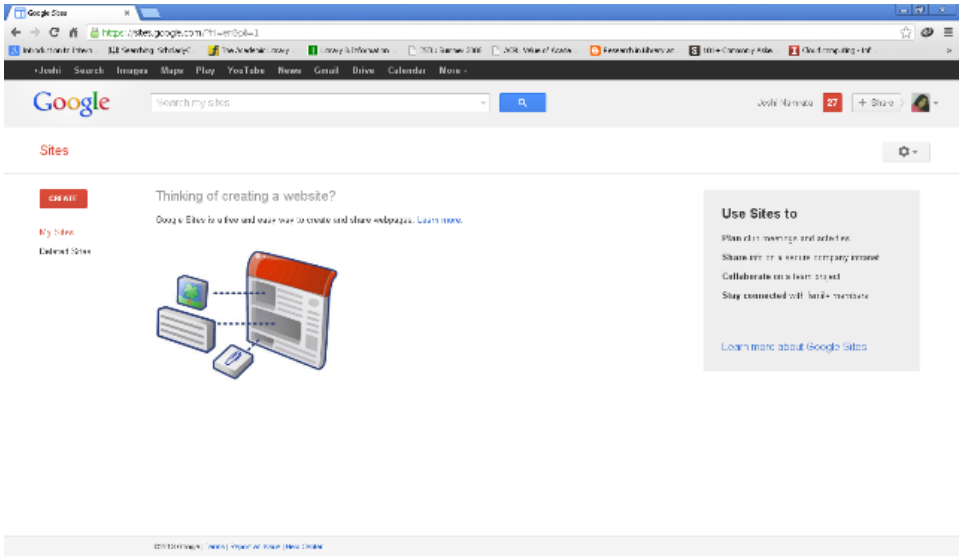

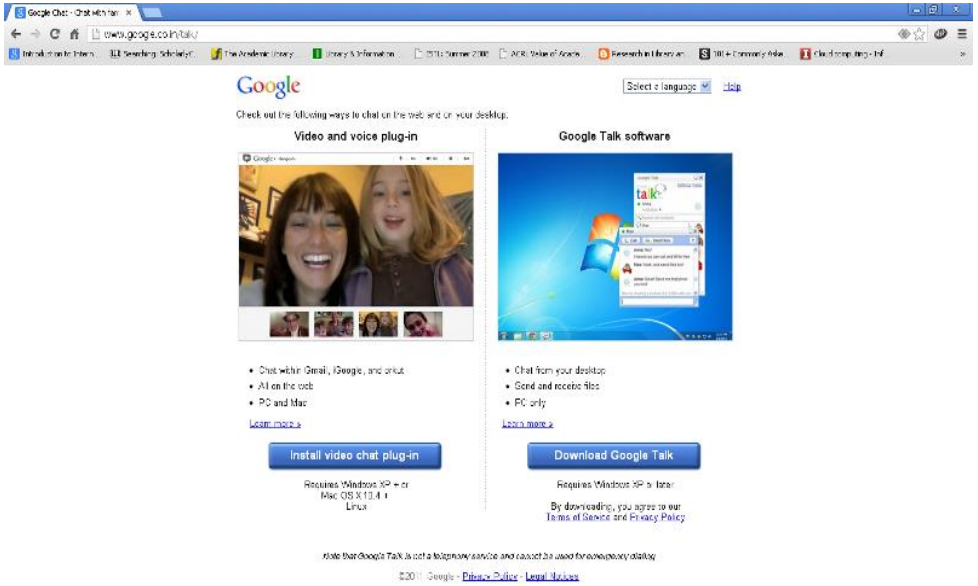

Service Model Architectures

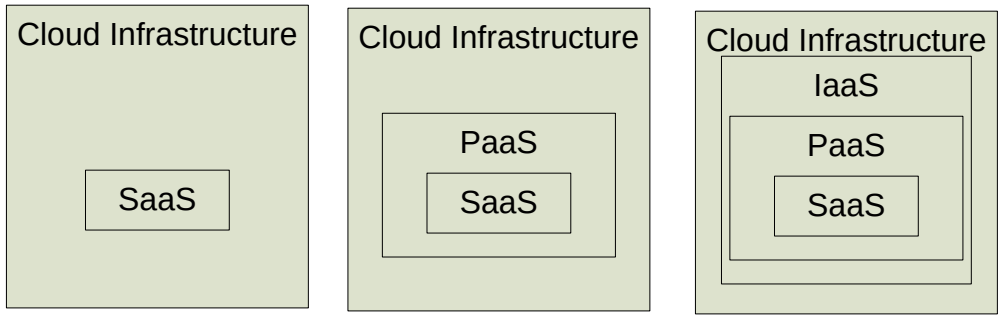

Software as a Service (SaaS) Architectures

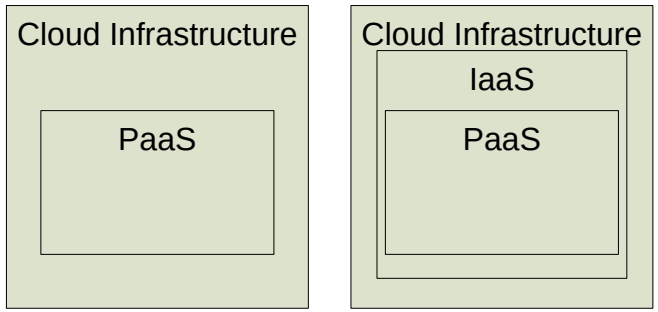

Platform as a Service (PaaS) Architectures

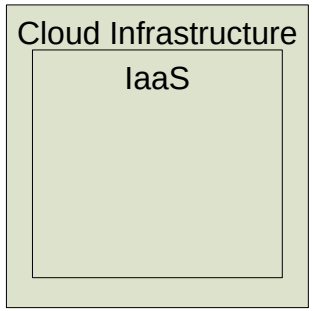

Infrastructure as a Service (IaaS) Architectures

Advantages of CC for service providers & end users

- Not necessary to invest in acquiring, maintaining and update IT infrastructures
- Cost of software licenses included in the service
- Customizability of services
- Cloud providers that have specialization in a particular area can bring advanced services that a single company might not be able to afford or develop
- Scalability, reliability, and efficiency
  - An example of scalability:
    In 2020 "Zoom has grown on AWS to accommodate an increase from 10 million daily meeting participants in December 2019 to more than 300 million a day regularly since April 2020"
    (source ComputerWeekly.com: Zoom signs multi-year preferred cloud provider deal with AWS)

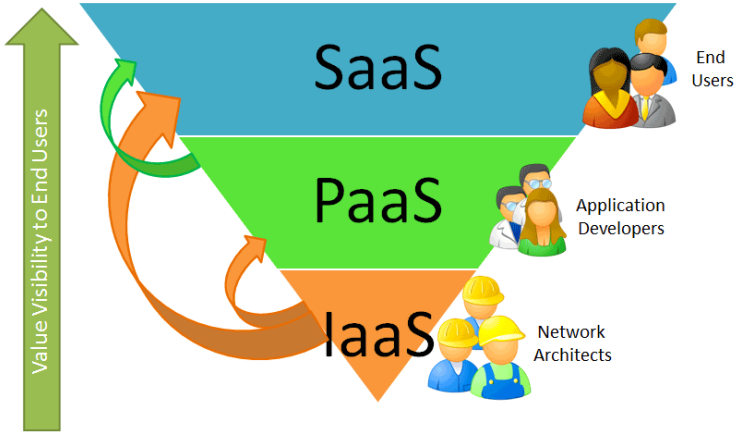

Private Cloud infrastructures

- Managing IT resources “as if they were in the Cloud” may bring the benefits of virtualization and consolidation within an organization’s IT department
- Furthermore, this process paves the way for opening up to Public and Hybrid Cloud adoption
- VMware enriched its portfolio with a Private Cloud management suite
  - VMware vCloud Suite extends the VMware vSphere hypervisor
- In the last few years a few open-source projects have been established to develop software suites to manage IaaS services in Private Cloud environments
  - OpenStack
  - OpenNebula
  - Eucalyptus

- Hardware vendors are also pushing the Private Cloud market
  - HP Elion (HPE) supports both Eucalyptus and OpenStack
  - Dell Red Hat OpenStack cloud
  - Cisco Metapod Private Cloud Solution based on OpenStack
  - Ericsson-Mirantis partnership (OpenStack)

Public vs Private Cloud

- When is Utility Computing preferable to running a Private Cloud?
- The left-hand side of the inequality represents the expected profit from using cloud computing based on virtualized data-centers
- The right-hand side represents the expected profit from using a traditional single-tenant data center
- Both sides perform the same calculation by multiplying profits per user-hour by the total user-hours
- If resources cost the same in both environments, the *true cost* of the traditional data center is greater than or equal to the cost of cloud computing due to the fact that a virtualized datacenter allows the service provider to achieve a resource utilization close to unity (100%) while in a traditional data-center utilization is much less than 1

$$UserHours_{cloud} \times (revenue - Cost_{cloud}) \geq UserHours_{datacenter} \times (revenue - \frac{Cost_{datacenter}}{Utilization})$$

Source: ***Above the Clouds: A Berkeley View of Cloud Computing*** – Armbrust et al., 2009