

# INTRODUCTION TO Cloud Computing

Άγγελος Μιχάλας  
Καθηγητής  
Τμήμα Μηχανικών Πληροφορικής  
ΤΕΙ Δυτικής Μακεδονίας

# Outline

- Introduction
- Cloud Computing
- Examples

# Evolution

- Classical Computing

- Buy & Own
  - Hardware,
  - System Software,
  - Applications often to meet peak needs.
- Install, Configure, Test, Verify
- Manage
- ..
- Finally, use it
- \$\$\$\$.....\$(High CapEx)


- Cloud Computing

- Subscribe
- Use
  
- \$ - pay for what you use, based on **QoS**

**Look to the cloud!**  
Pay for the bandwidth and  
server resources that you  
need. When your work is done  
then turn the whole thing off!



# Evolution of Internet Computing



# Definition (1)

Cloud

Computing

The “Cloud” is the default symbol of the Internet in diagrams.

The broader term of “Computing” encompasses:

- Computation
- Coordination logic
- Storage

Cloud Computing is about moving computing from the single desktop pc/data centers to commercial service providers on the Internet.



# Definition (2)

- Cloud Computing: Internet-based computing, whereby shared resources, software and information are provided to computers and other devices on-demand.
- Numerous previous attempts:
  - on-demand computing, utility computing, ubiquitous computing, autonomic computing, platform computing, edge computing, elastic computing, grid computing, ...

# Resources





# Cloud characteristics

- Uses Internet technologies to offer scalable and elastic services. The term “elastic computing” refers to the ability of dynamically acquiring computing resources and supporting a variable workload.
- The resources used for these services can be metered and the users can be charged only for the resources they used.
- The maintenance and security are ensured by service providers.
- The service providers can operate more efficiently due to specialization and centralization.

# Cloud Service Models

## SaaS

SalesForce.com  
Google Docs



Applications  
& Services



Office Applications, CRM, E-mail,  
Media, Gaming

## PaaS

Heroku  
Microsoft Azure



Development Platform



Development Platforms, Runtime  
Environments for applications, APIs

## IaaS

Amazon EC2, S3  
Joyent, Rightscale



Infrastructure

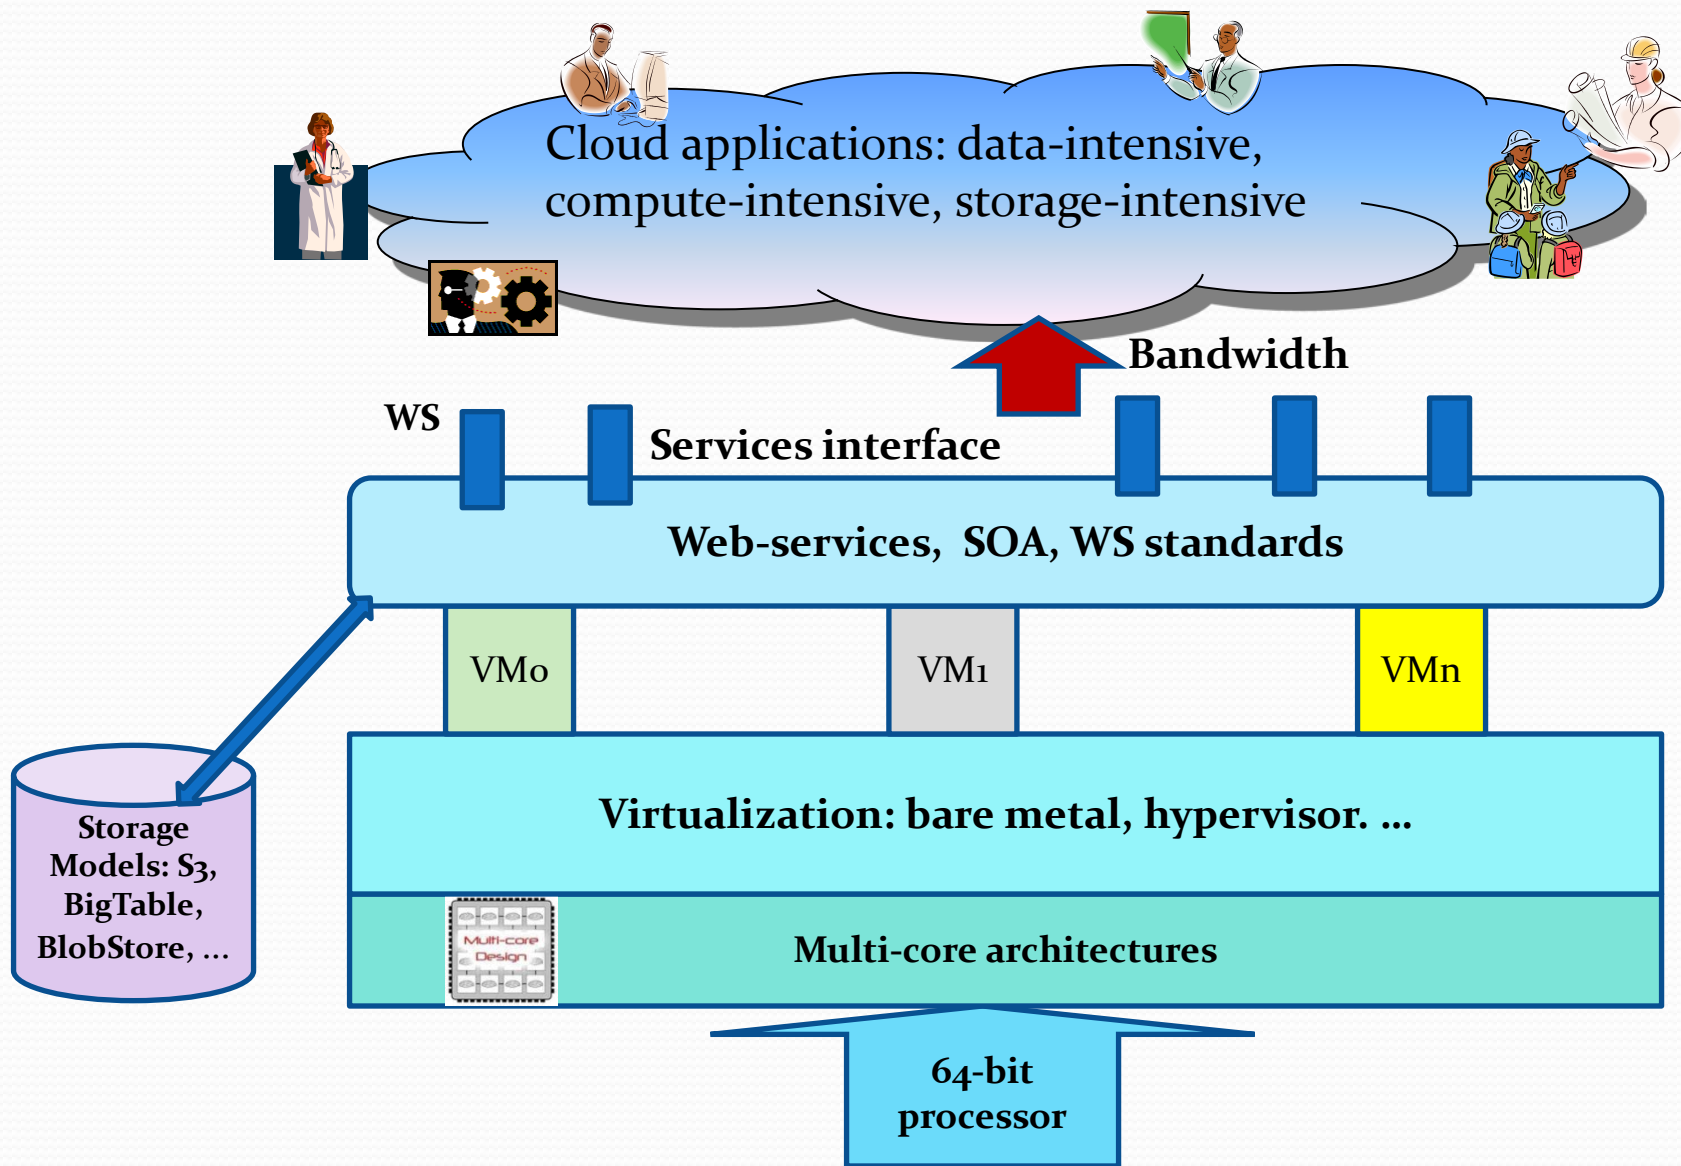


Virtual servers, virtual storage, and  
networking

# Types of clouds

- Public Cloud - the infrastructure is made available to the general public or a large industry group and is owned by the organization selling cloud services.
- Private Cloud – the infrastructure is operated solely for an organization.
- Community Cloud - the infrastructure is shared by several organizations and supports a community that has shared concerns.
- Hybrid Cloud - composition of two or more clouds (public, private, or community) as unique entities but bound by standardized technology that enables data and application portability.

# Cloud supporting technologies



# Virtual machine monitor (VMM / hypervisor)


- Is the software that partitions the resources of computer system into one or more virtual machines (VMs). Allows several operating systems to run concurrently on a single hardware platform.
- A VMM allows
  - Multiple services to share the same platform.
  - Live migration - the movement of a server from one platform to another.
  - Enforces isolation among the systems, thus security.

# Virtual machines (VMs)

- VM - isolated environment that appears to be a whole computer, but actually only has access to a portion of the computer resources.
- Two VM Types:
  - Process VM - a virtual platform created for an individual process and destroyed once the process terminates.
  - System VM - supports an operating system together with many user processes (eg. Java Virtual Machine – JVM).

# Provisioning of virtual resources


- Cloud systems provide virtualized resources as a service
- VM Manager
  - Submission
  - Control
  - Monitoring





# Cloud Applications



- **Scientific/Tech Applications**
- **Business Applications**
- **Customer/Social Applications**



*Science and Technical Applications*



*Business Applications*



*Customer/Social Applications*



# Windows Azure

- Enterprise-level on-demand capacity builder
- Computational cycles and storage available on-request for a cost
- You have to use Azure API to work with the infrastructure offered by Microsoft
- Significant features: web role, worker role , blob storage, table and drive-storage



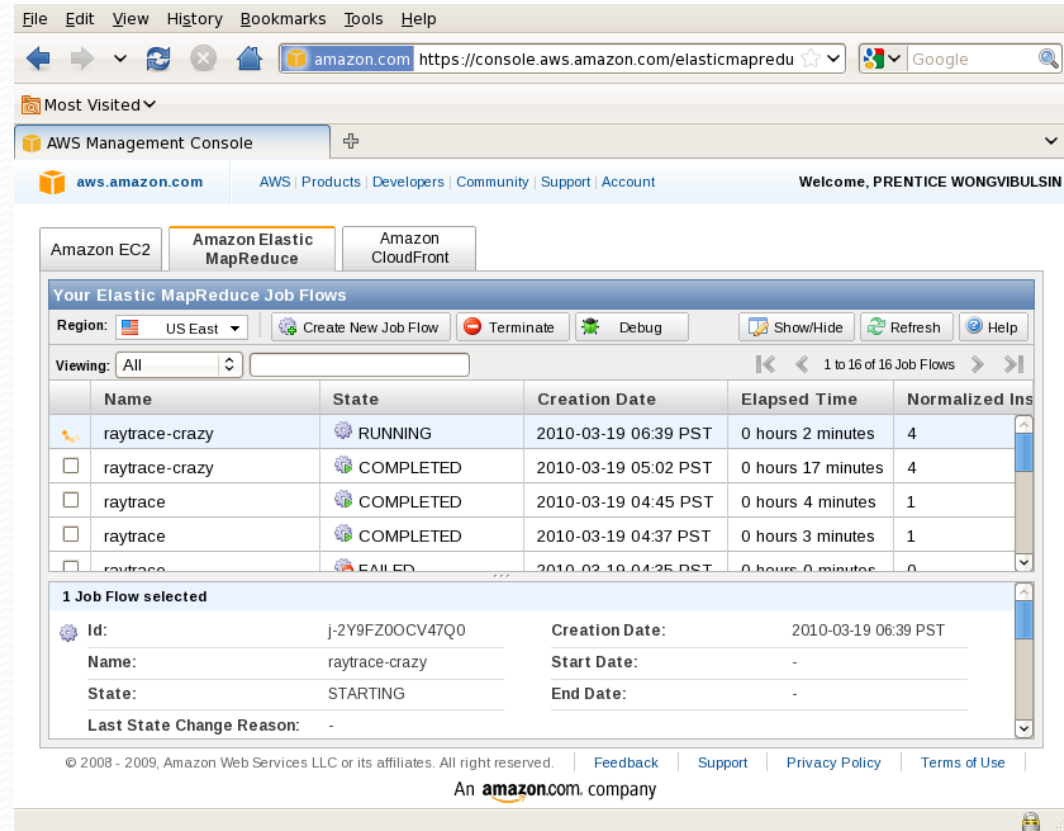
# Amazon web services (AWS)

- Provision of multiple service models:
  - IaaS: EC2, S3, ELB, AutoScaler
  - PaaS: Elastic Beanstalk, EMR
  - SaaS: CloudSearch, Elastic Transcoder



# Amazon Elastic Compute Cloud (EC2)

- A web service for instantiating computing instances with any of the operating systems supported.
- Features: Cloud Management Console, MapReduce Cloud, Cloud monitoring tools



The screenshot displays the AWS Management Console interface for Amazon Elastic MapReduce. The browser address bar shows the URL `https://console.aws.amazon.com/elasticmapredu`. The console header includes navigation links for `aws.amazon.com`, `AWS`, `Products`, `Developers`, `Community`, `Support`, and `Account`, along with a welcome message for `PRENTICE WONGVIBULSIN`.

The main content area is titled "Your Elastic MapReduce Job Flows" and shows a table of job flows in the `US East` region. The table has columns for `Name`, `State`, `Creation Date`, `Elapsed Time`, and `Normalized Ins`. The job flows listed are:

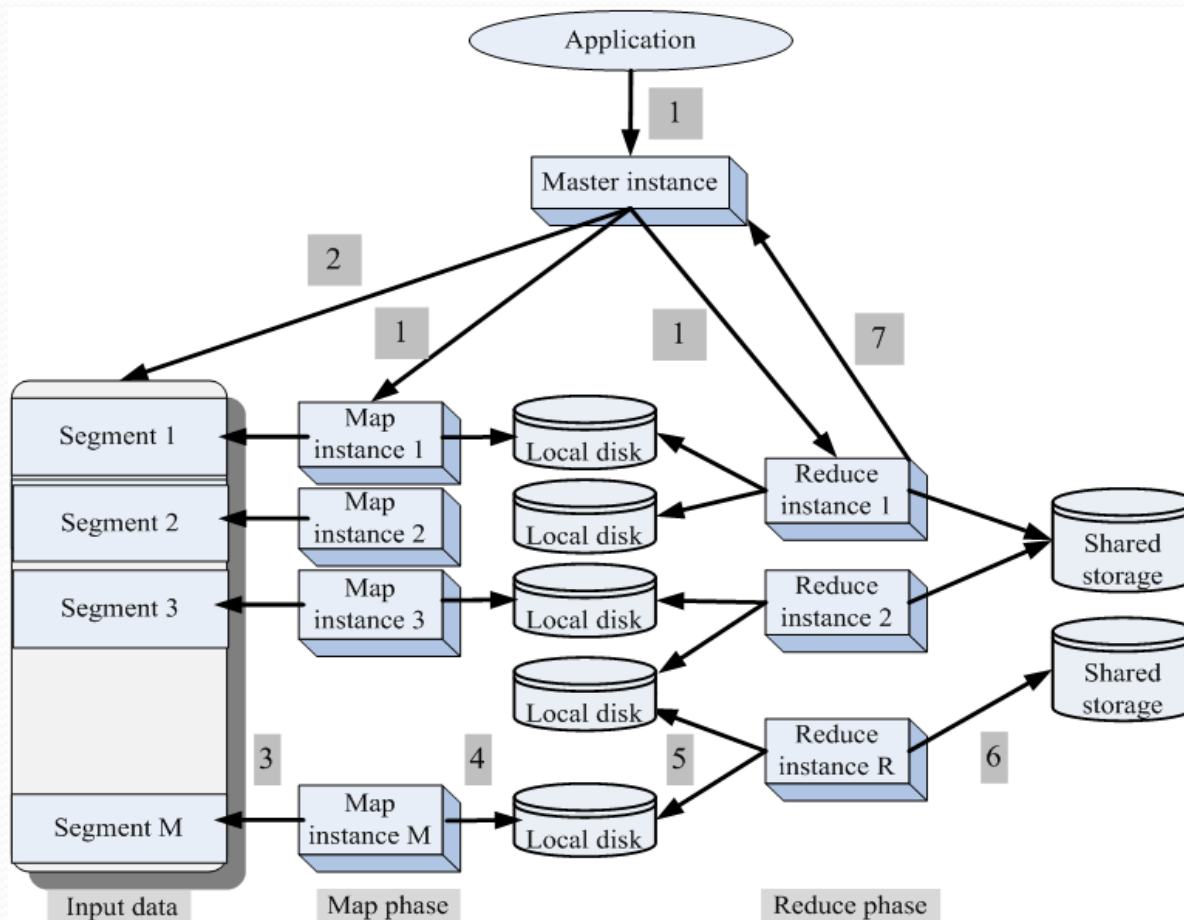
Name	State	Creation Date	Elapsed Time	Normalized Ins
raytrace-crazy	RUNNING	2010-03-19 06:39 PST	0 hours 2 minutes	4
raytrace-crazy	COMPLETED	2010-03-19 05:02 PST	0 hours 17 minutes	4
raytrace	COMPLETED	2010-03-19 04:45 PST	0 hours 4 minutes	1
raytrace	COMPLETED	2010-03-19 04:37 PST	0 hours 3 minutes	1
raytrace	FAILED	2010-03-19 04:35 PST	0 hours 0 minutes	0

Below the table, a detailed view for the selected job flow is shown:

<b>Id:</b>	j-2Y9FZ0OCV47Q0	<b>Creation Date:</b>	2010-03-19 06:39 PST
<b>Name:</b>	raytrace-crazy	<b>Start Date:</b>	-
<b>State:</b>	STARTING	<b>End Date:</b>	-
<b>Last State Change Reason:</b>	-		

The footer of the console includes copyright information: "© 2008 - 2009, Amazon Web Services LLC or its affiliates. All right reserved." and links for `Feedback`, `Support`, `Privacy Policy`, and `Terms of Use`. The Amazon logo and "An amazon.com company" are also present.

# Map - Reduce Philosophy



# MapReduce philosophy

1. An application starts a **master instance**, **M worker instances** for the *Map phase* and later **R worker instances** for the *Reduce phase*.
2. The **master instance** partitions the input data in *M segments*.
3. Each *map instance* reads its input data segment and processes the data.
4. The results of the processing are stored on the local disks of the servers where the map instances run.
5. When all map instances have finished processing their data, the **R reduce instances read the results** of the first phase and merge the partial results.
6. The **final results are written by the reduce instances** to a shared storage server.
7. The **master instance monitors the reduce instances** and when all of them report task completion the application is terminated.

# EC2 pricing

Type	U\$ per hour	Ratio to Small	Compute Units	Virtual Cores	Compute Unit/Core	Memory (GB)	Disk (GB)	Addr. (bits)
Standard Small	0,115	1	1	1	1	1,7	160	32/64
Standard Medium	0,230	2	2	1	2	3,75	410	32/64
Standard Large	0,460	4	4	2	2	7,5	850	64
Standard Extra Large	0,920	8	8	4	2	15	1690	64
High-Memory Extra Large	0,680	5,9	6,5	2	3,25	17,1	420	64
High-Memory Double Extra Large	1,360	14,1	13	4	3,25	34,2	850	64
High-Memory Quadruple Extra Large	2,720	28,2	26	8	3,25	68,4	1690	64
High-CPU Medium	0,230	2	5	2	2,5	1,7	350	32/64
High-CPU Extra Large	0,920	8	20	8	2,5	7	1690	64
Cluster QuadrupleCloud Management Console, MapReduce Cloud, Amazon Machine Image (AMI) Extra Large	---	15,3	33,5	16	2,09	23	1690	64
Eight Extra Large	---	28,2	88	32	2,75	60,5	1690	64

# okeanos



## ► Features:

- Production-quality IaaS cloud
- Everything open source:



synnefo



by



- Persistent VMs
- Commodity Hardware – No SAN, No exotic network hw
- Add plain servers/disks/switches to scale up
- Three clicks to start 1 or 10s of VMs, in 15 secs

# Resources

- An Introduction to Virtualization and Cloud Technologies to Support Grid Computing, Ignacio M. Llorente
- Cloud Computing: Concepts, Technologies and Business Implications, B. Ramamurthy & K. Madurai
- Introduction to Cloud Computing, Dr. Rodrigo Calheiros
- ~OKEANOS presentation at cloudforum 2013, Nectarios Koziris, GRNET
- Cloud Computing Theory and Practice, Dan C. Marinescu



# Questions?

