An Autonomic Computing Engine for Computational Science and Engineering Applications in Cloud Computing and Grid Environments

Moustafa AbdelBaky, Hyunjoo Kim, Ivan Rodero and Manish Parashar

moustafa.a, irodero, parashar@rutgers.edu, hyunjoo.kim@xerox.com

The NSF Cloud and Autonomic Computing Center

Electrical & Computer Engineering Department

Rutgers, The State University of New Jersey, USA

CAC

NSF

RUTGERS

THE STATE UNIVERSITY OF NEW JERSEY

Introduction

Dr. Manish Parashar

parashar@rutgers.edu

CAC

NSF

RUTGERS
THE STATE UNIVERSITY
OF NEW JERSEY

An Introduction to

CLOUD COMPUTING

A Cloudy Weather Forecast

- A Cloudy Outlook
  - About 3.2% of U.S. small businesses, or about 230,000 businesses, use cloud services.
  - Another 3.6%, or 260,000, plan to add cloud services in the next 12 months.
  - Small-business spending on cloud services will increase by 36.2% in 2010 over a year ago, to $2.4 billion from $1.7 billion.
- Source: IDC, 2010

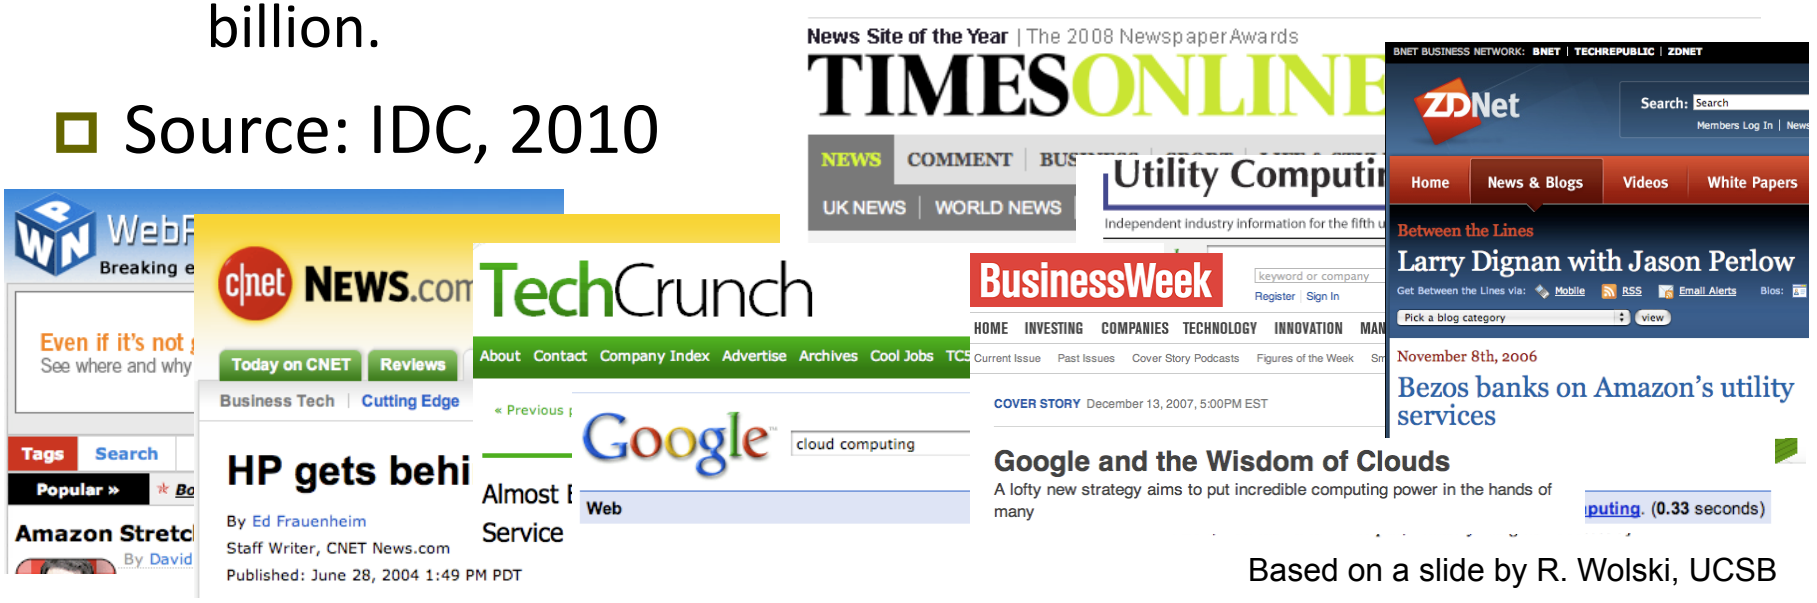

Based on a slide by R. Wolski, UCSB

The Cloud Today

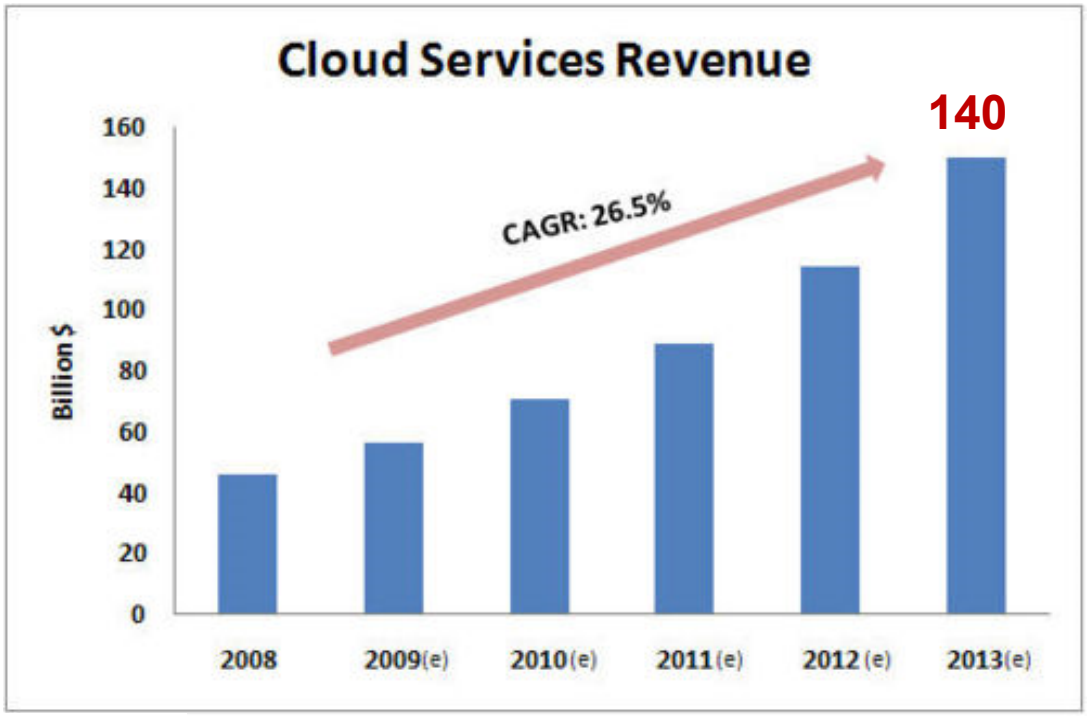

Source: Wikipedia

rackspace cloud

hp

Windows Azure

amazon web services™

YAHOO!

vmware

Time Warner Cable

enomaly elastic computing

Google

CANONICAL

IBM

and more ..

What is a Cloud?

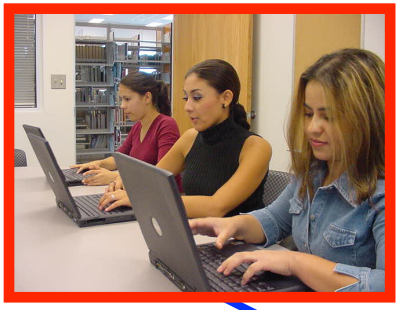
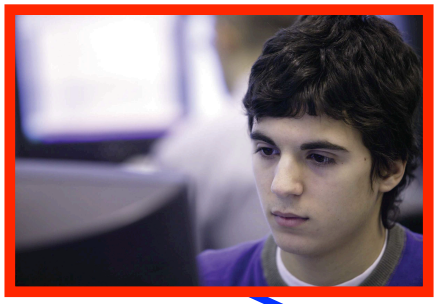

Cloud as a business model

- SLA's
- Web Services
- Virtualization
- SaaS, PaaS, IaaS

Based on sldie by R. Wolski, UCSB

What is a Cloud?


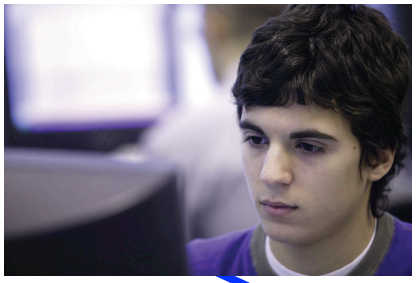

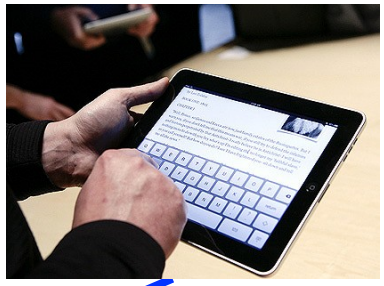

Cloud as an abstraction !

- On demand, elastic access to resources
- Pay as you go
- Ease of use

Based on sldie by
R. Wolski, UCSB

Defining Cloud Computing

- Wikipedia – Cloud computing is Internet-based computing, whereby shared resources, software and information are provided to computers and other devices on-demand like a public utility.
- NIST – A cloud is a computing capability that provides an abstraction between the computing resource and its underlying technical architecture (e.g., servers, storage, networks), enabling convenient, on-demand network access to a shared pool of configurable computing resources that can be rapidly provisioned and released with minimal management effort or service provider interaction

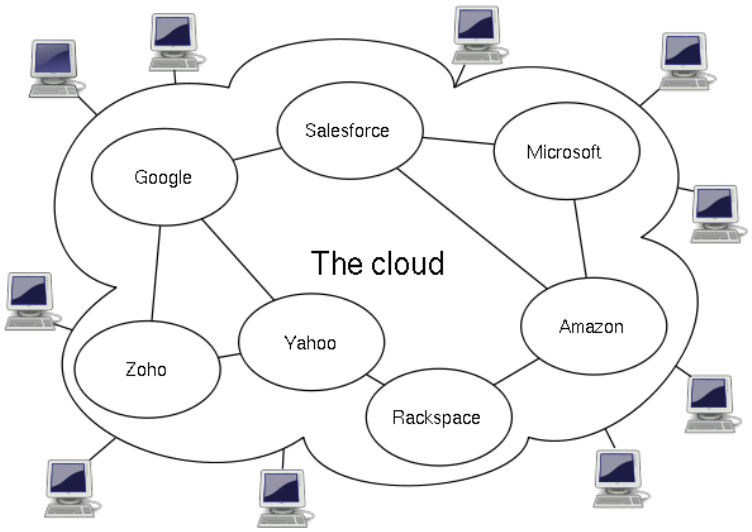

- SLAs
- Web Services
- Virtualization

The Lure ….

- A seductive abstraction – unlimited resources, always on, always accessible!
- Economies of scale
- Multiple entry points
  - *aaS: SaaS, PaaS, IaaS, HaaS
- IT- outsourcing
  - Transform IT from being a capital investment to a utility
  - TCO, capital costs, operation costs
- Potential for on-demand scale-up, scale-down, scale-out
- Pay as you go, for what you use

Why Now – Providers Perspective?

- Datacenters are very expensive and time consuming to build!
  - Cost >$150M and take 24+ months to design and build
  - Increasing size, cost and energy consumption make efficiency and utilization critical
  - *Amortize capital costs across multiple users!*
- Consolidation and virtualization has become essential and possible
- Experience with very large datacenters, Grids
  - Unprecedented economies of scale

Why Now – User Perspective?

- Complex application needs
  - Highly dynamic demands for resources with heterogeneous and dynamic workloads
  - Various and dynamic QoS requirement
    - Throughput, budget, time, etc.
- "Rent" services of clouds: CapEx -> OpEx!
  - Rent required resources from clouds on-demand and pay for what you use
  - Pay by use instead of provisioning for peak
  - Avoid risk of over-provisioning: underutilization
  - Avoid heavy penalty for under-provisioning: Lost revenue, users
- "Cost associativity": 1,000 computers for 1 hour same price as 1 computer for 1,000 hours

The Cloud Stack

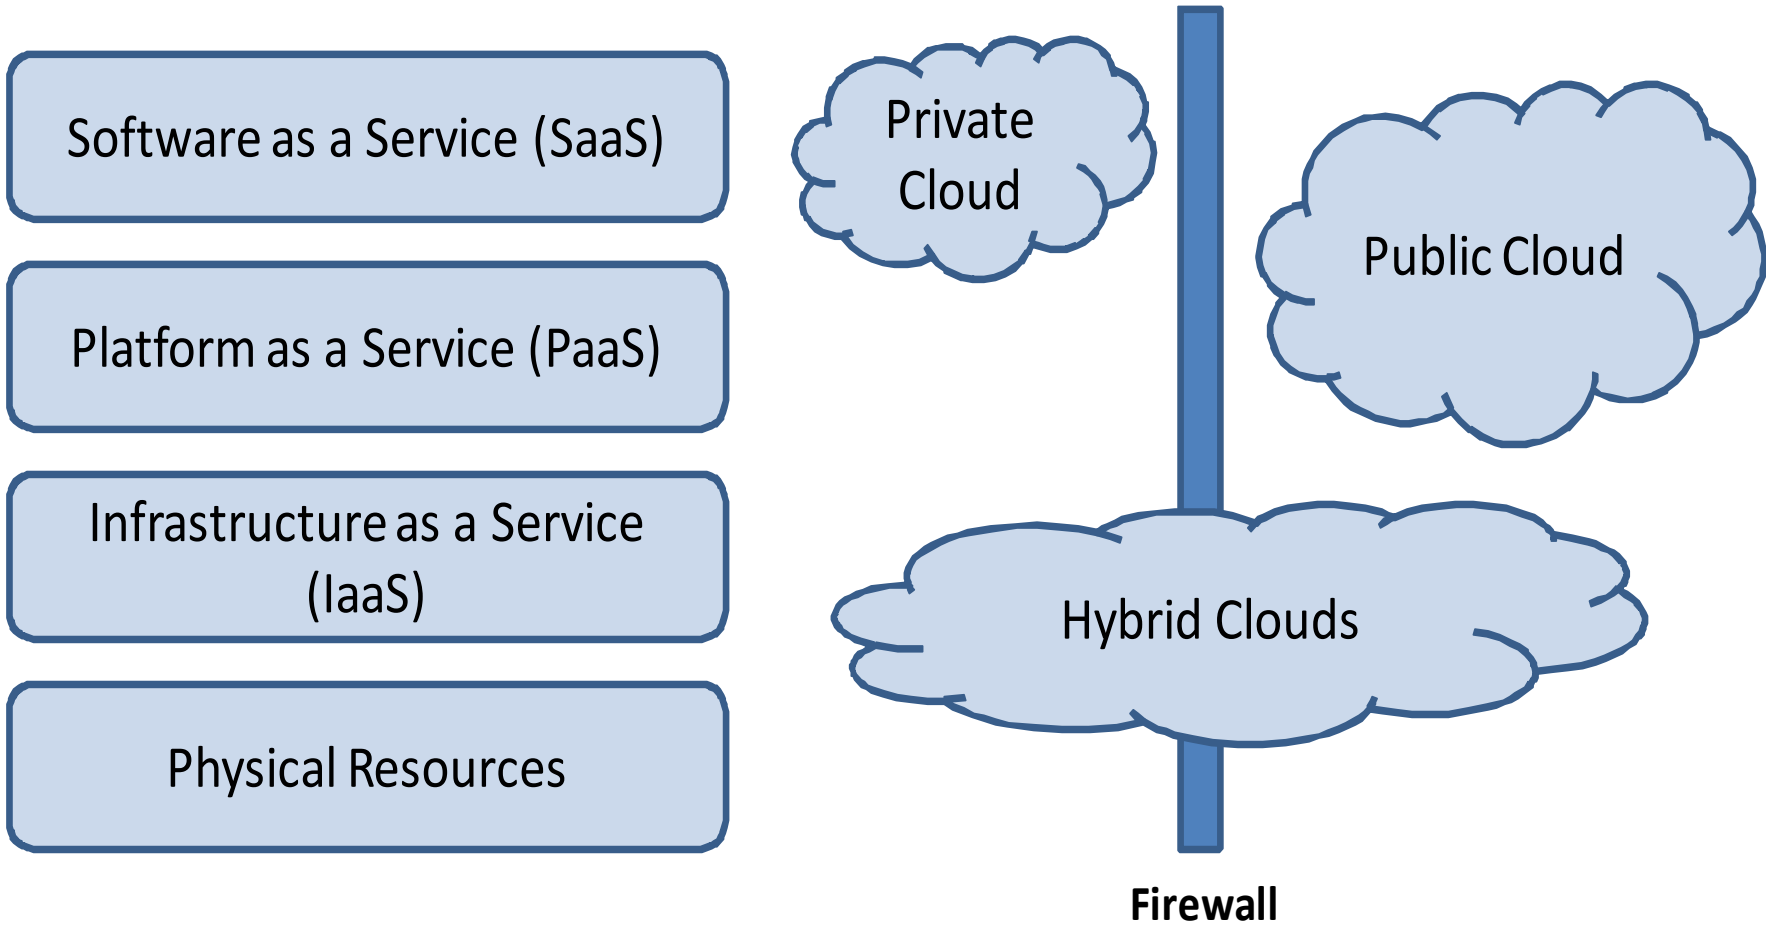

CS&E on the Cloud

- Clouds support different although complementary usage models as compared to more traditional HPC grids

- Some questions
  - Application types and capabilities that can be supported by clouds?
  - Can the addition of clouds enable scientific applications and usage modes that are not possible otherwise?
  - What abstractions and systems are essential to support these advanced applications on different hybrid grid-cloud platforms?

Obvious Candidates

- Nicely parallel
- Minimal synchronization, Modest I/O requirements
- Large messages or very little communication
- Low core counts
- Parallel programming models for data intensive science
  - e.g., BLAST parametric runs
- Customized and controlled environments
  - e.g., Supernova Factory codes have sensitivity to OS/ compiler versions
- Overflow capacity to supplement existing systems
  - e.g., Berkeley Water Center has analysis that far exceeds capacity of desktops

Beyond the Obvious Candidates

- New application formulations
  - Asynchronous, resilient
  - E.g., Asynchronous Replica Exchange Molecular Dynamics, Asynchronous Iterations
- New usage modes
  - Client + Cloud accelerators
  - E.g., Excel + EC2
- New hybrid usage modes
  - Cloud + HPC + Grid

Summary

- The future is *Cloudy*…
  - Cloud becoming a part of production computational environments
  - Clouds will play a role in Science and Engineering
  - Many Cloud Computing Benefits:
    - Shift CapEx to OpEx , Scale OpEx to demand
    - Startups and prototyping, One-off tasks (Wash. Post)
    - Cost associativity
    - Research at scale
- Clouds transforming science
  - New applications, new delivery models, new usage modes, democratization
- A challenging research agenda
  - Not just the Science of Clouds but also ***Science on Clouds***

Questions?

For more info please visit

http://cometcloud.org

RUTGERS
THE STATE UNIVERSITY
OF NEW JERSEY