

Cloud Computing

Kai Shen

4/29/2014

CSC 258/458 - Spring 2014

1

What is Cloud Computing?

"The interesting thing about Cloud Computing is that we've redefined Cloud Computing to include everything that we already do. . . . I don't understand what we would do differently in the light of Cloud Computing other than change the wording of some of our ads."

- Larry Ellison, quoted in the Wall Street Journal, September 26, 2008

"A lot of people are jumping on the [cloud] bandwagon, but I have not heard two people say the same thing about it."

- Andy Isherwood, quoted in ZDnet News, December 11, 2008

4/29/2014

CSC 258/458 - Spring 2014

2

Resource Provisioning Problem

- Provisioning of computing/storage resources for applications
 - Unpredictable load – Internet applications
 - I want to try some idea and see if people are interested
 - Run it once – digitizing old prints of New York Times newspapers
- Challenges:
 - High upfront infrastructure cost
 - Difficulty of dynamic capacity expansion
 - Cost of excess capacity (power)

4/29/2014

CSC 258/458 - Spring 2014

3

Computing As A Utility

- Computing as a utility (like electricity)
 - Illusion of infinite resources on demand
 - No steep upfront cost, pay as needed (very short-term possible)
- Grid computing
 - High-performance computing applications
- Cloud computing
 - Compute-intensive jobs as well as network services, computing as well as storage

4/29/2014

CSC 258/458 - Spring 2014

4



Key Idea

- Pooling resources together, managed centrally (in a data center); many computing jobs and network services share the resource pool
- Benefits:
 - While one application's resource needs may vary dramatically, the sum of many independent applications' resource needs varies much less
 - More resource utilization efficiency through sharing
 - Space sharing as well as time sharing
 - Upfront setup cost is amortized over many applications
 - Cost is proportional to use
 - Commoditizing support for availability, data durability, ...

4/29/2014

CSC 258/458 - Spring 2014

5



Utilizations of Cloud Computing

- Highly available web services with unpredictable load
- Large computing jobs on demand
- Offload piecemeal work to cloud
 - Extension of compute-intensive desktop applications

4/29/2014

CSC 258/458 - Spring 2014

6



Cloud Computing Architecture

- Executing multiple applications on one machine
- Isolation of resources and data through virtualization
- Provide common services to help the application development, further amortizing the costs

4/29/2014

CSC 258/458 - Spring 2014

7



Case Study – Amazon EC2

- Amazon Elastic Compute Cloud (Amazon EC2)
- Interface between cloud computing hosts and users: Virtualization packages applications in
 - VMware ESX and VMware Workstation VMDK images
 - Citrix Xen VHD images
 - Microsoft Hyper-V VHD images
- Additional services
 - Amazon Simple Storage Service (Amazon S3)
 - Amazon Relational Database Service (Amazon RDS)
 - Amazon SimpleDB
 - Amazon Simple Queue Service (Amazon SQS)

4/29/2014

CSC 258/458 - Spring 2014

8



Case Study – Google App Engine

- Interface between cloud computing hosts and users
 - Controlled Java/Python programming environment
- Augmented Java/Python runtime and Linux containers (Linux virtualization) provide isolation and security
- Additional services
 - Hadoop
 - Memcache
 - Task Queues

4/29/2014

CSC 258/458 - Spring 2014

9



Case Study – Microsoft Azure

- I don't know much about – probably use Microsoft Common Language Runtime VM

4/29/2014

CSC 258/458 - Spring 2014

10



Cloud Computing Challenges

- Performance predictability
 - Do we reach required performance if we run the application on these machines?
- Performance factors:
 - CPU numbers, memory size, network/storage bandwidth
 - Resources that can't be strictly partitioned – multicore shared cache, memory bandwidth

4/29/2014

CSC 258/458 - Spring 2014

11



Cloud Computing Challenges

- What if we need to change the machine(s) that run an application?
 - **Migration:** shut down the virtual machine (VM) and re-launch it on a new physical machine
- What if we want to keep it running without much interruption (think of an Internet server serving live users continuously)?
 - **Live migration:** move memory pages while keeping the VM running; some pages are updated after first round ⇒ ...

4/29/2014

CSC 258/458 - Spring 2014

12

Cloud Computing Challenges

- Data confidentiality
 - Protection between applications \Rightarrow virtualization
 - Protection against cloud computing hosts?
- Side channel information leaking
 - Reading processor event counters

4/29/2014


CSC 258/458 - Spring 2014

13

Stealing RSA key in OpenSSL

[Zhang et al. 2007]

- In OpenSSL/RSA, X^d is decomposed into a series of square and multiply operations.
 - e.g., $X^{11} = ((X^2)^2 * X)^2 * X$
- Execution sequence of squares and multiplies can help infer the RSA key d .
- Distinguish square and multiply through processor event counters.



4/29/2014

CSC 258/458 - Spring 2014

14

When not Cloud Computing?

- Strong sensitivity of information leaking
- High cost of data movement
 - Data set up cost, online bandwidth usage
- It is really always available?
 - Some known disruptions of cloud services
 - WikiLeaks

4/29/2014

CSC 258/458 - Spring 2014

15

Bare-metal Cloud

- No virtualization, each application runs on whole machine(s)
 - Get bare-metal speed (without virtualization overhead)
 - Not susceptible for side channel information leaking or other virtualization vulnerabilities
- Resource management between machines in a cloud
 - Resource managers in network switches and intelligent network interface cards

4/29/2014

CSC 258/458 - Spring 2014

16



Distributed Clouds

- EdgeClouds
 - Internet latency is significant, bandwidth is expensive
 - Internet isn't 100% reliable
- Green clouds
 - Distributed clouds using local renewable energy (near solar panel array, wind farm, ...)
 - Energy availability for most renewable energy sources is dynamic and hard to predict ⇒ move work to where green energy is available

4/29/2014

CSC 258/458 - Spring 2014

17



Disclaimer

- Ideas in some of the lecture slides were derived from the paper "Above the Clouds: A Berkeley View of Cloud Computing", by Armbrust et al. The slides are developed for the sole purpose of course instruction at the University of Rochester.

4/29/2014

CSC 258/458 - Spring 2014

18