



# Cloud Computing

Kai Shen

4/20/2011 CSC 258/458 - Spring 2011 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

4/20/2011 CSC 258/458 - Spring 2011 2



## Resource Provisioning Problem

- Provisioning of computing/storage resources for Internet applications
- Challenges:
  - Difficulty to predict load of Internet applications
    - Curse of hugely successful Internet services, DoS attacks
  - Cost of dynamic capacity expansion
  - High upfront cost
    - I want to try some idea and see if people are interested
  - Cost of excess capacity (power)

4/20/2011 CSC 258/458 - Spring 2011 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

4/20/2011 CSC 258/458 - Spring 2011 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, ...

4/20/2011 CSC 258/458 - Spring 2011 5

## Cloud Computing Architecture

- Interface between cloud computing hosts and users
  - Machine interface (VMware/Xen virtual machine)
  - Language interface (Microsoft Common Language Runtime VM, Java)
  - Application interface (Google AppEngine, mapreduce/Hadoop)
- Resource management
  - Executing multiple applications on one machine
  - Migration of applications

4/20/2011 CSC 258/458 - Spring 2011 6

## Cloud Computing Challenges

- Data confidentiality
  - Protection between applications
  - Protection against cloud computing hosts
- Side channel information leaking

4/20/2011 CSC 258/458 - Spring 2011 7

## Stealing RSA key in OpenSSL [HotOS'07, Zhang et al.]

- 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 hardware counters.

4/20/2011 URCS 573 - Spring 2008 8



## 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
- Additional issues:
  - Cache space, multicore resource contention
  - Storage bandwidth resource is non constant
  - Performance variability

4/20/2011 CSC 258/458 - Spring 2011 9



## Cloud Computing Challenges

- Parallel application in a cloud computing environment
  - Processes are shared over a number of machines, shared with other applications
- Problem:
  - When process 1 runs (waiting for pivot value from process 2), process 2 may not even be running

4/20/2011 CSC 258/458 - Spring 2011 10



## Cloud Computing Challenges

- Data transfer bottlenecks
  - Data set up cost
  - Online bandwidth usage

4/20/2011 CSC 258/458 - Spring 2011 11



## 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

4/20/2011 CSC 258/458 - Spring 2011 12



## When not Cloud Computing?

- High cost of data movement
- Strong sensitivity of information leaking
- What I (Google) need alone already has the economy of scale.
- It is really always available?
  - WikiLeaks

4/20/2011 CSC 258/458 - Spring 2011 13



## Distributed Clouds

- EdgeClouds
  - Internet latency is nontrivial
  - Internet bandwidth is a critical resource
  - Internet isn't 100% reliable
- Cloud on your wireless routers (Amal's work)
  - Extreme proximity to clients
  - Energy efficiency
- Cloud collaborations

4/20/2011 CSC 258/458 - Spring 2011 14



## 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.

4/20/2011 CSC 258/458 - Spring 2011 15