Introduction to Cloud Computing

@ 北京航天航空大学软件学院

By 邓侃 Ph.D
dengkan@smartclouder.com
02/28/2012

Mobile Internet and Big Data A Revolutionary Era

1 Billion US$

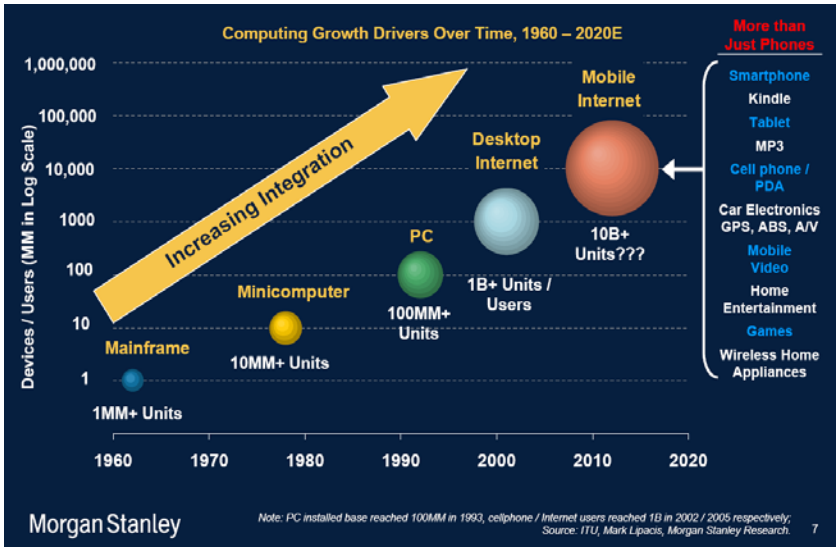

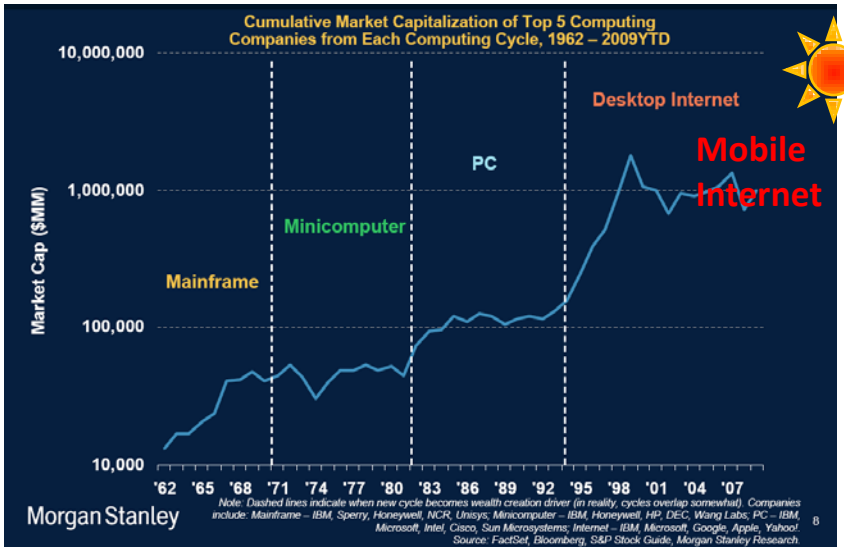

- 5 phases of computing growth, since 1960's.
  1. Main-frame, 2. Minicomputer, 3. PC, 4. Internet, 5. Mobile Internet.
- Every phase, the total amount of user-time, increased 10 times.
  The sum of the top 5 companies' market value increased 10 times every phase.
- With mobile internet, the big amount of user-times, induces big data。
  The technical challenge is how to deal with big data.
- The solution to the big data challenge, is cloud computing.

Intel Pentium4 CPU's power is 10,000 MIPS

MIPS: Million Instructions Per Second.

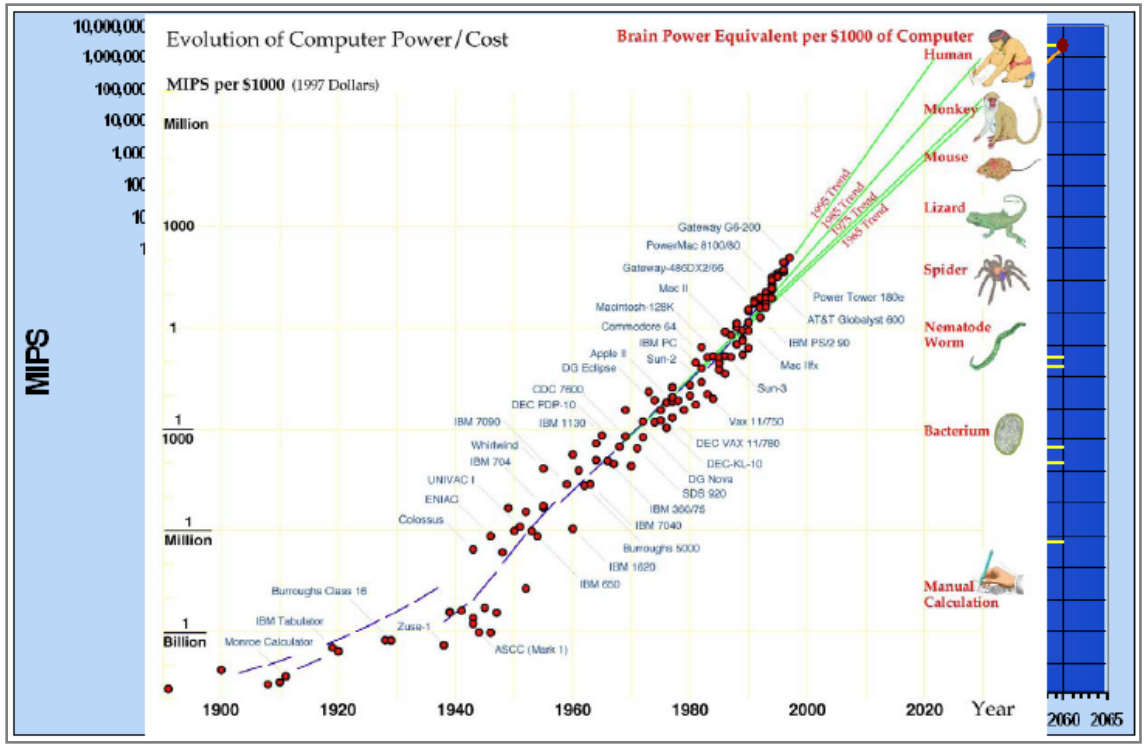

- 1965, Moore's Law:

  The number of transistors in IC doubles every 2 years, or even 18 months.

- Still, the power of a single CPU, cannot beat the human brain power.

  Solution: use many computers.

- Challenge, to orchestrate many computers working together.

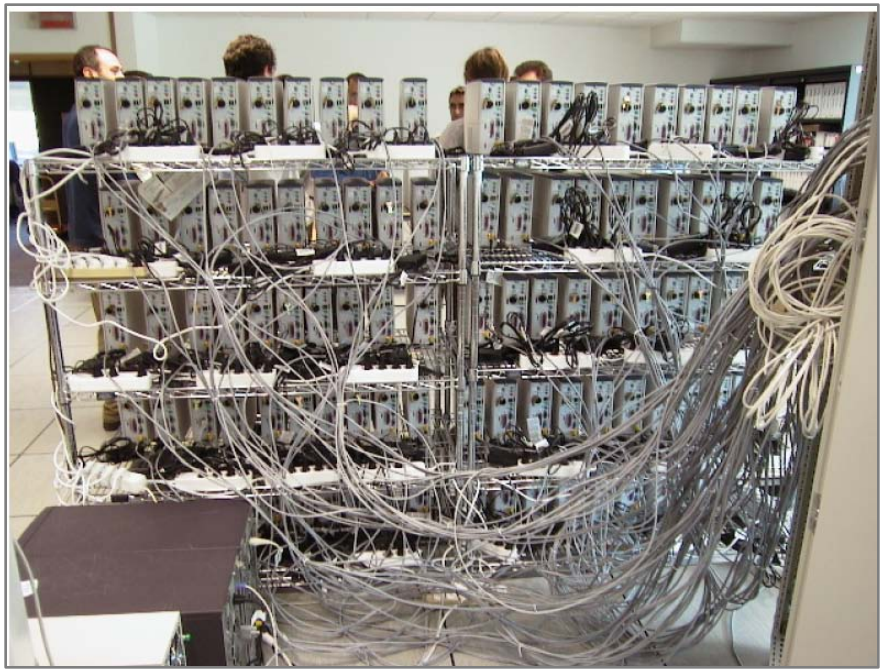

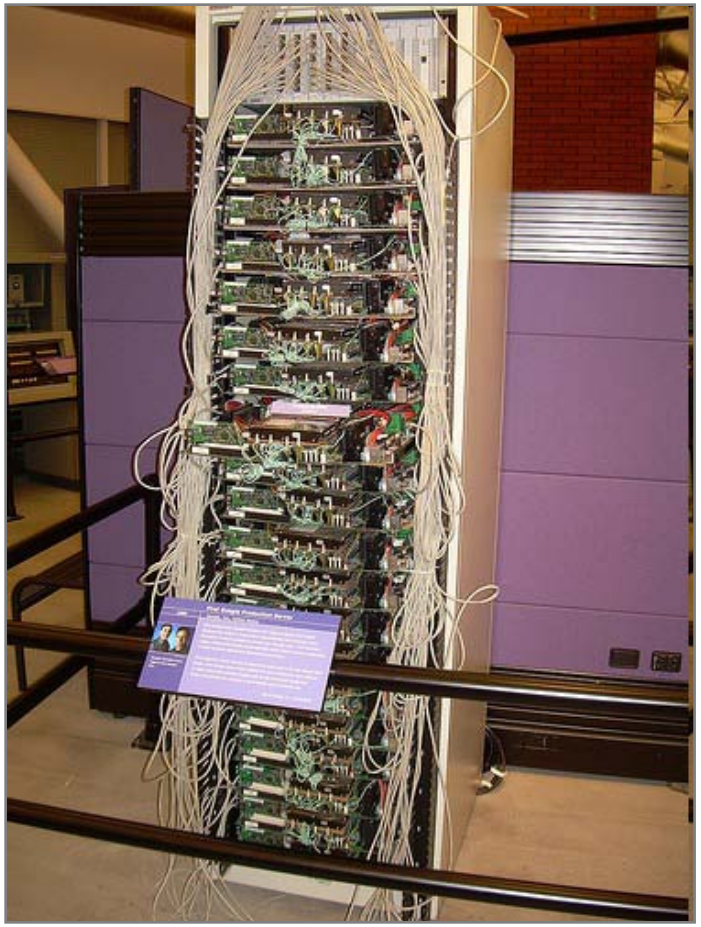

Google’s initial cloud

- Cloud computing can be built with commodity PC servers.
- The most successful cloud so far, was by two graduate students.
  Larry Page from University of Maryland, (北航 in the US).
  Sergey Brin from UIUC, (北邮 in the US).

Sergey Brin & Larry Page

Andy Bechtolsheim

- Sergey and Larry wanted to build a search engine.
  Need the power of super-computer,
  to store every webpage, of every website, globally, every historic version.
  And to process the big data, to build search index.
- Raised fund from Andy Bechtolsheim, in 1997.
  Andy, CMU alumni, cofounder of Sun Microsystems, very rich.
- But Andy only gave them 100K US$.
  The most successful investment, but also the most stupid one.

- Why was Andy not positive on Google?

  4 technical difficulties.

  The two boys might not have the skillset.

Andy Bechtolsheim

- Scalability:

  Big storage space for big data, Googol (10^100) scale!

  Big paralleled computing to process them.

  Never succeeded in human's history.

  And the data is increased every second.

- Reliability:

  Using commodity machines,

  One single machine's failure should not break down the entire system.

- Elasticity:

  The load fluctuation on different modules are different,

  Schedule the same machines, to work for different modules at different time.

- Security:

  Dynamically separate the machines into clusters, mutually inaccessible.

- Sergey and Larry's answer was,
  "O, yah, our company's name is Google!
  We deal with big data."
- Google runs the world's largest cloud,
  for 15 years continuously, reliably.

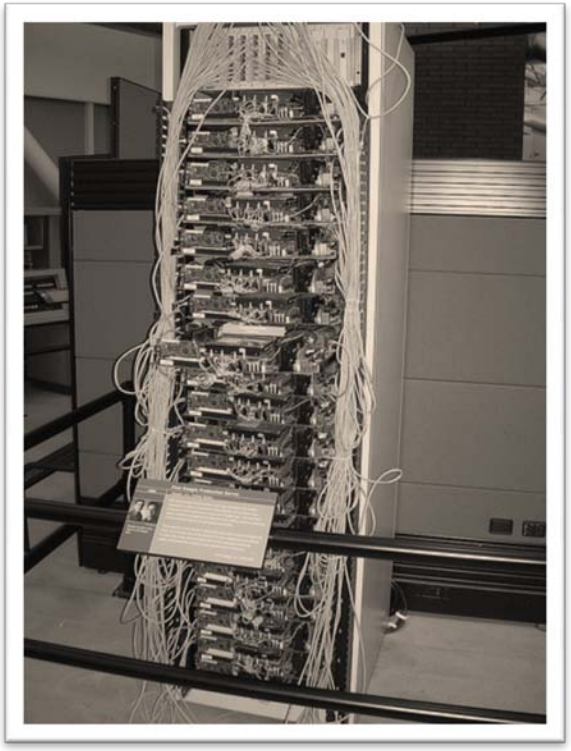

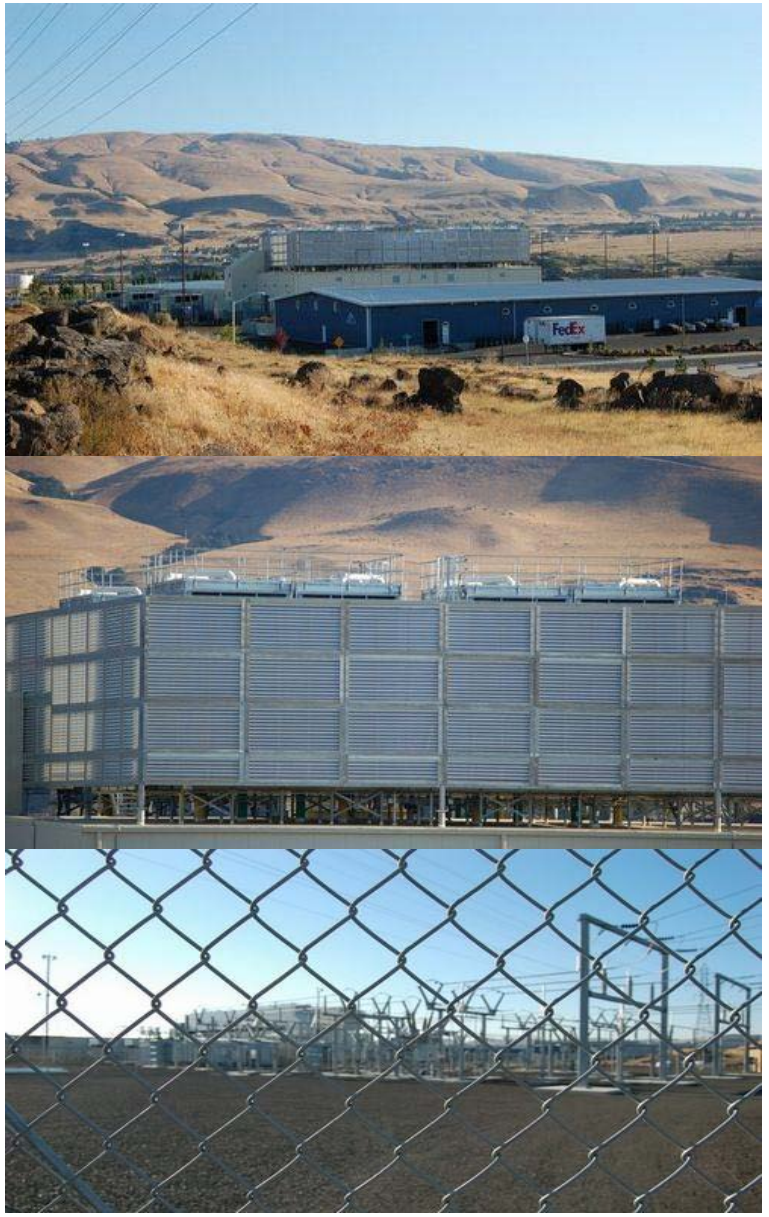

Cloud Computing Problems to Solve

- Scalability: add more machines, without modify the current system.
- Twitter was launched in May 2006.

  Dec 2007, Twitter users increased to 66K.

  Dec 2008, Twitter users grew to 5 millions.

  April 2009, over 100 million.
- Weibo was launched in Sept 2009.

  Nov 2009, Weibo users increased to 1 million.

  April 2010, over 10 million.

  Aug 2010, over 30 million.

  Oct 2010, over 50 million.
- China's population

  makes itself the best test-bed

  of cloud computing technology.

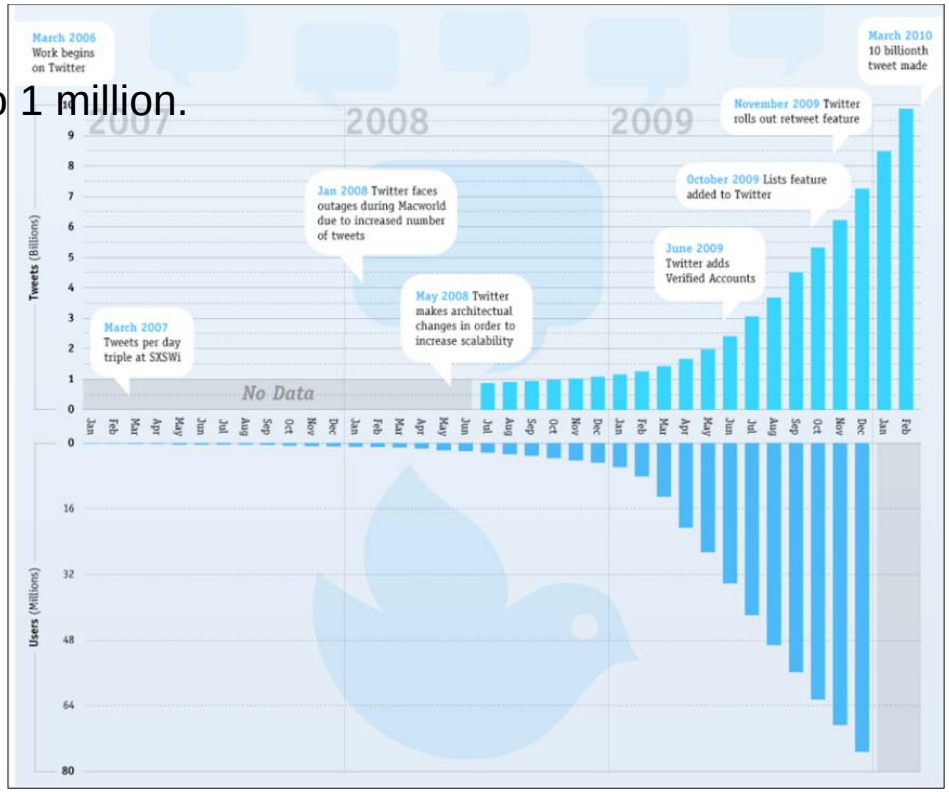

- Reliability: one single machine's failure, don't break down the entire system.
- Oct 29, 2009, T-mall kicked-off 50% discount.
- Half hour after the event started,
  支付宝 slowed down significantly.
  Another half hour later, the service shut down.
  One hour later, the service recovered.
- During the one hour that service was down,
  billion yuan's business was lost.

- Elasticity: use the same machines, for different business, at different time.
- Does 支付宝 need to keep the huge amount of machines,
  only to prepare for the annual sales? NO!
- Superbowl is the most popular sport event in the US.
  During the game, Twitter's load is 40% higher than the usual one.
  During the exciting moment,
  Twitter's load is 150% higher than usual .
- But unlike 支付宝,
  Twitter doesn't keep a lot of machines.
  Twitter borrowed machines
  temporarily from a third-party.
- A lesson learned from Twitter,
  to dynamically allocate machines,
  among different business,
  automatically,
  in real-time.

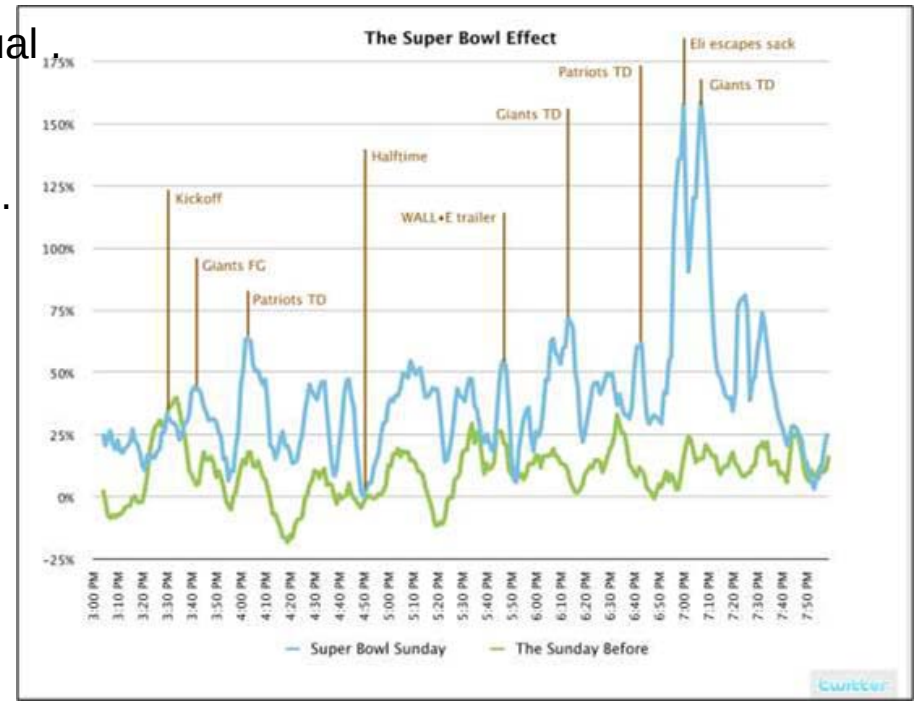

- Security: prevent data leak.
- Cloud can contain multiple business.
- Each business runs in its own LAN.
  Mutually inaccessible.

Cloud Computing
Technical Components

- Data flow and control:
  push the cloud to run faster.
- Anatomy of Twitter.
- Cache for fast read.
- Queue for async tasks.
- Pub/Sub for messaging.

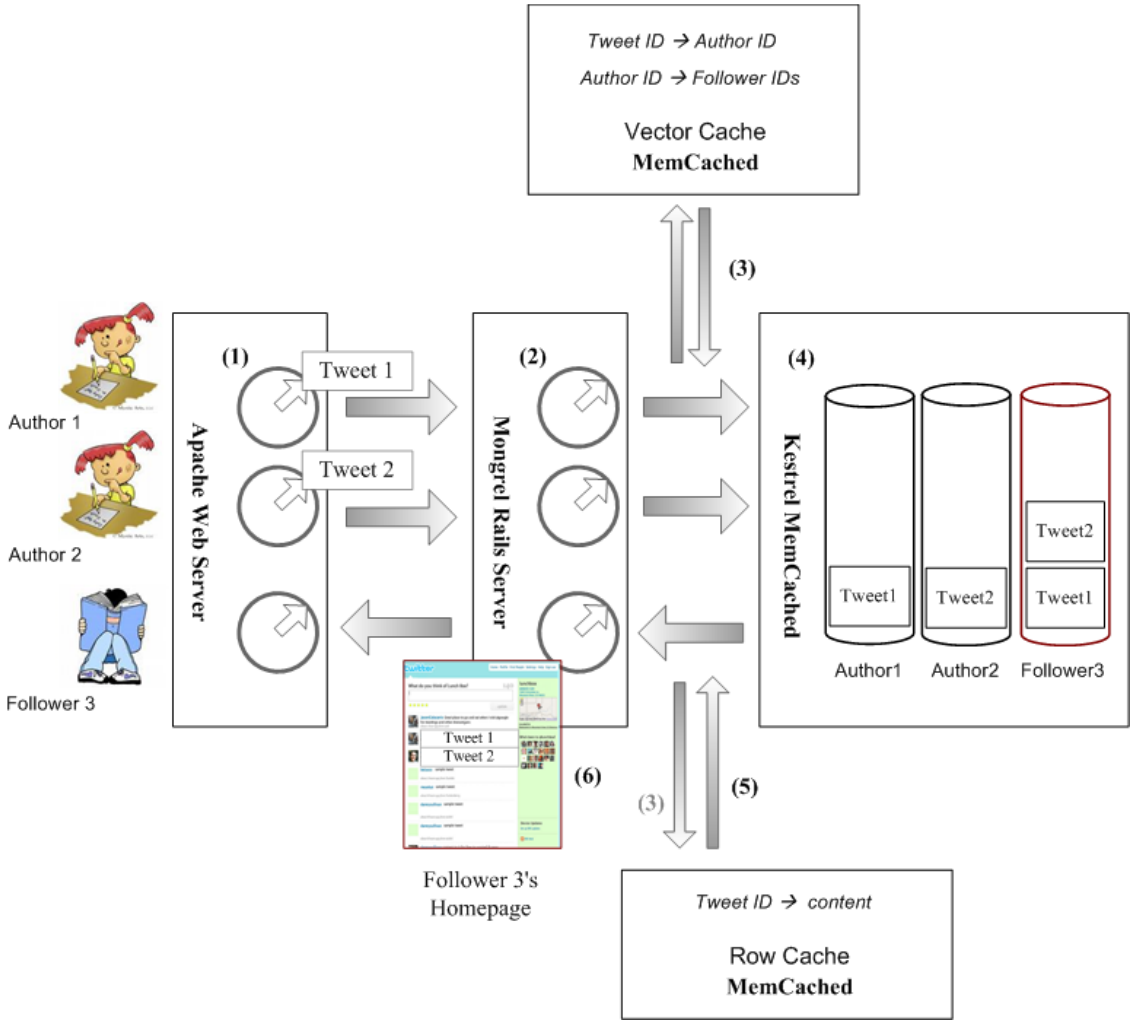

- Distributed File System:
  Scalable file storage.
- Google File System.
  (Hadoop HDFS)
- Master and Namespace.
- Chunk vs. File
- Replica vs. Fragmentations

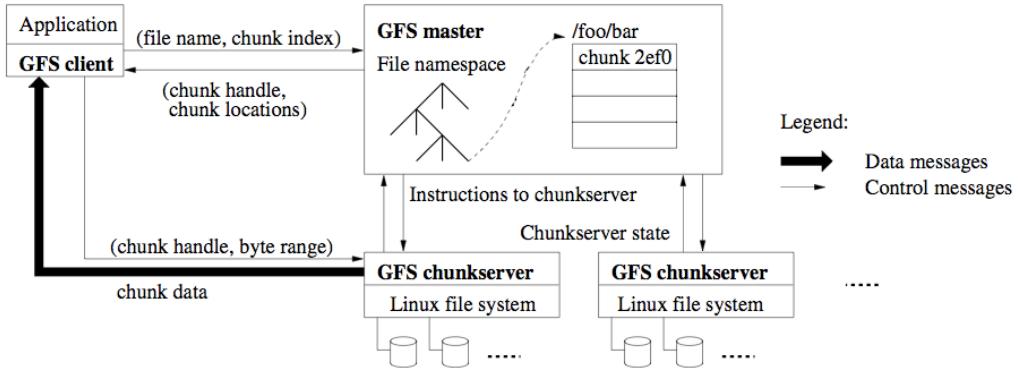

Figure 1: GFS Architecture

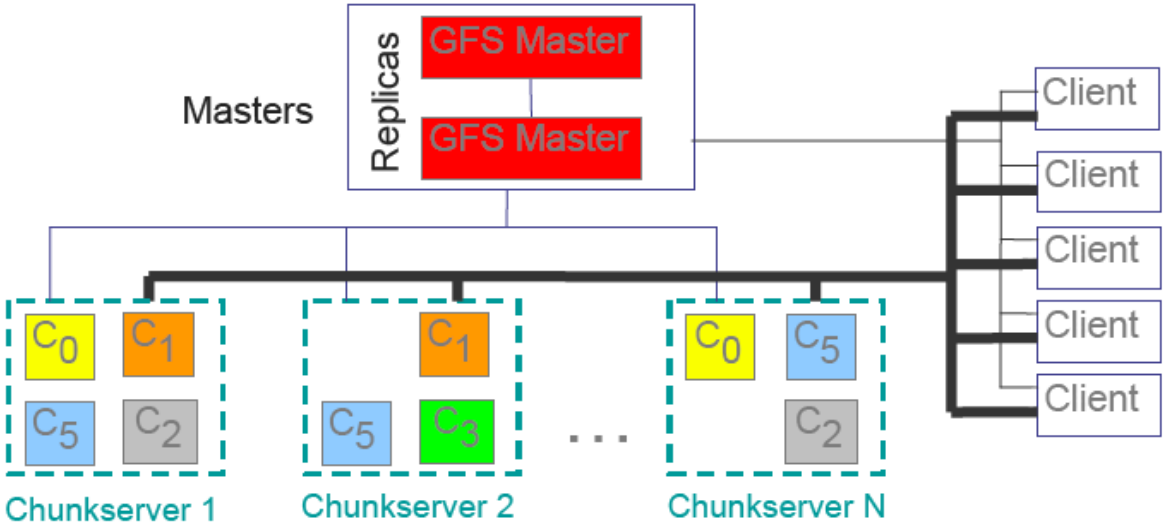

- Distributed Database:
  Scalable database.
- Google Bigtable.
  (Hadoop HBase)
- Distributed Index.
- Distributed ACID Transaction.
- Distributed lock.

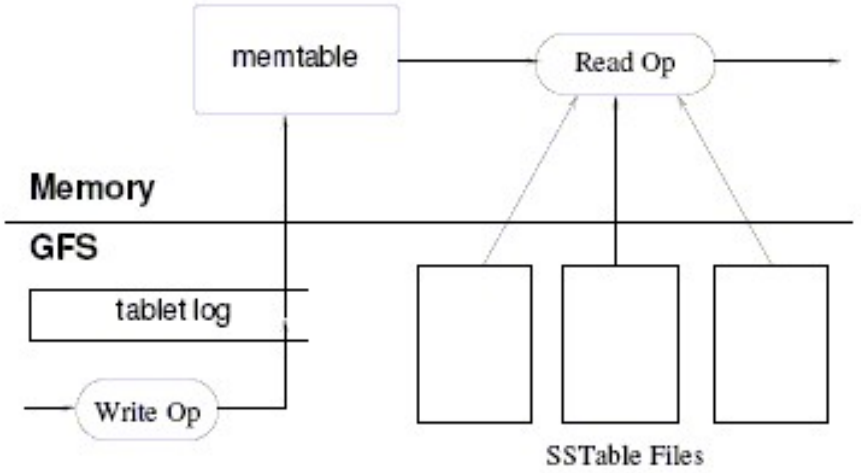

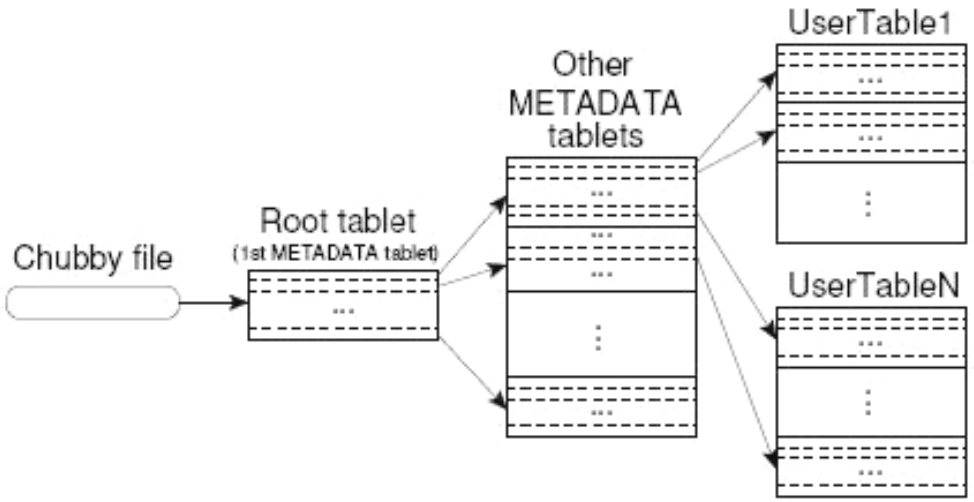

- Distributed Lock:
  Guarantee multiple read single write in distributed system.
- With replica,
  each data one lock or plural.
- How to deal with inconsistency?
- How to raise master,
  by Paxos protocol?

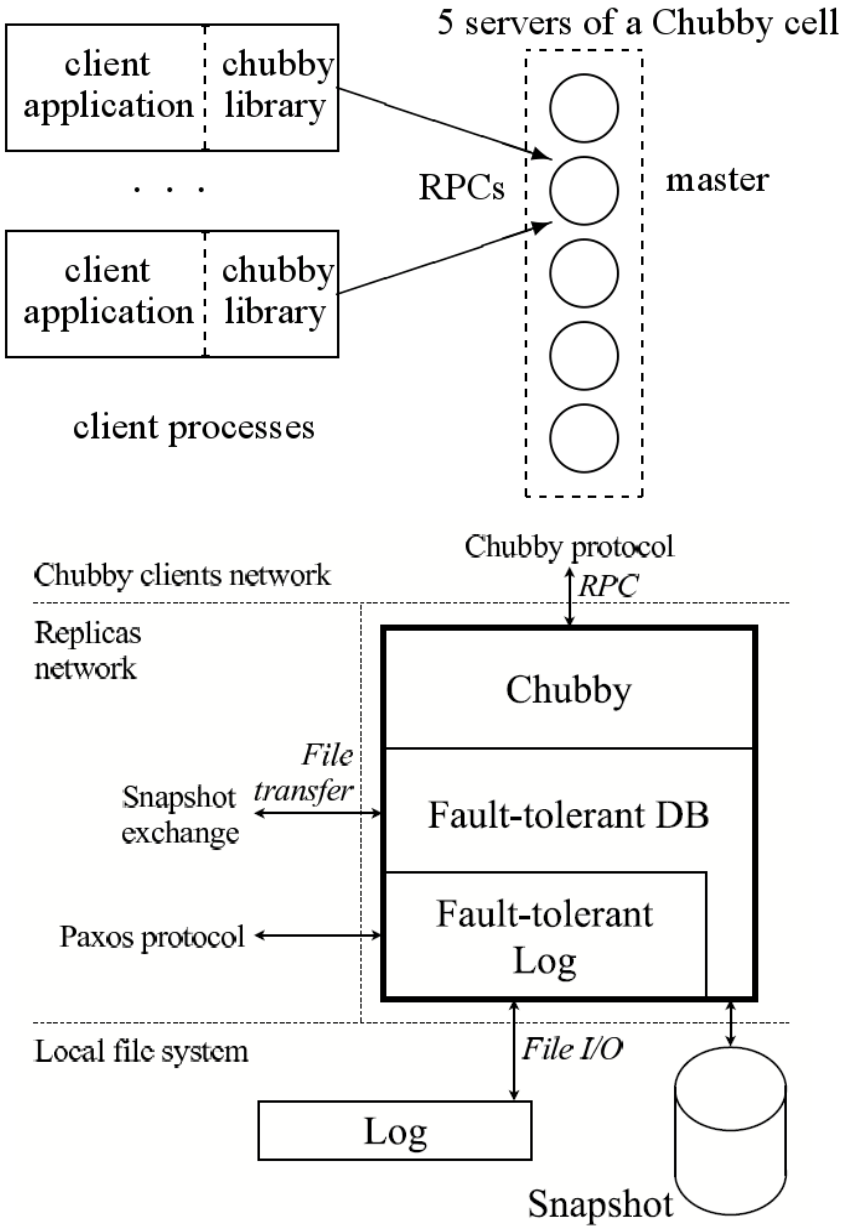

- No-SQL Database:
  Make database more efficient.
- No relational, but only key-value.
- No index, but algorithm.
- No SQL language.
- Easier to add machine.

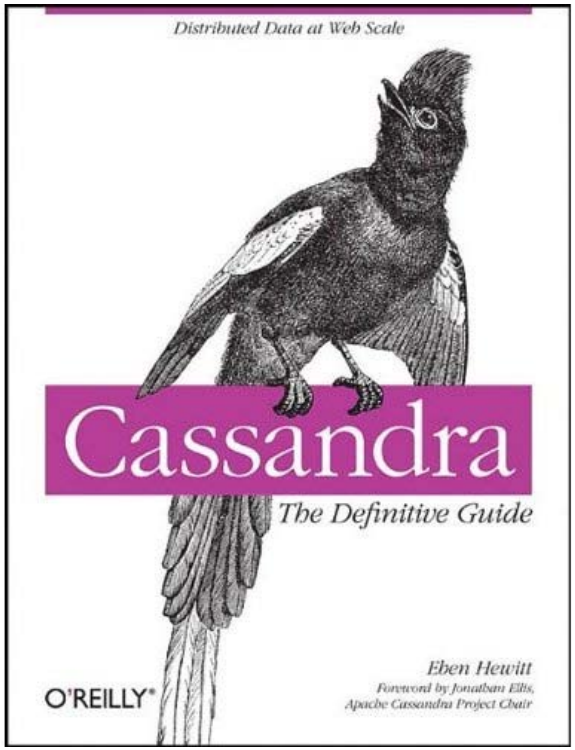

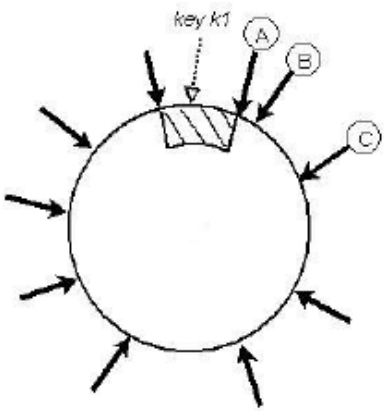

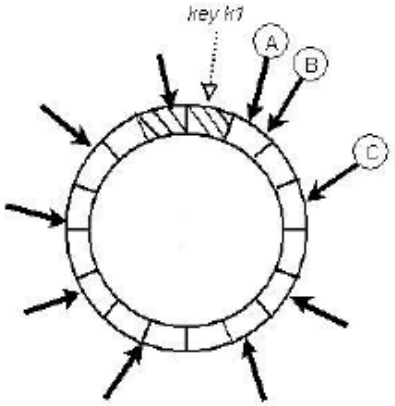

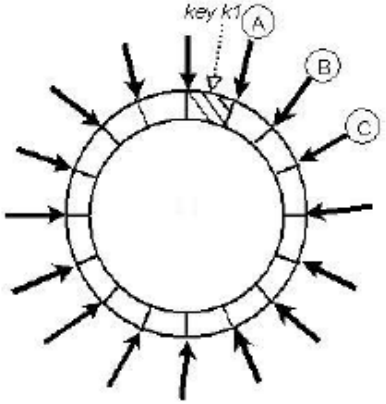

- Paralleled computing:

  Process big data by divide and conquer
- Google's MapReduce
- Not a panacea, case study.

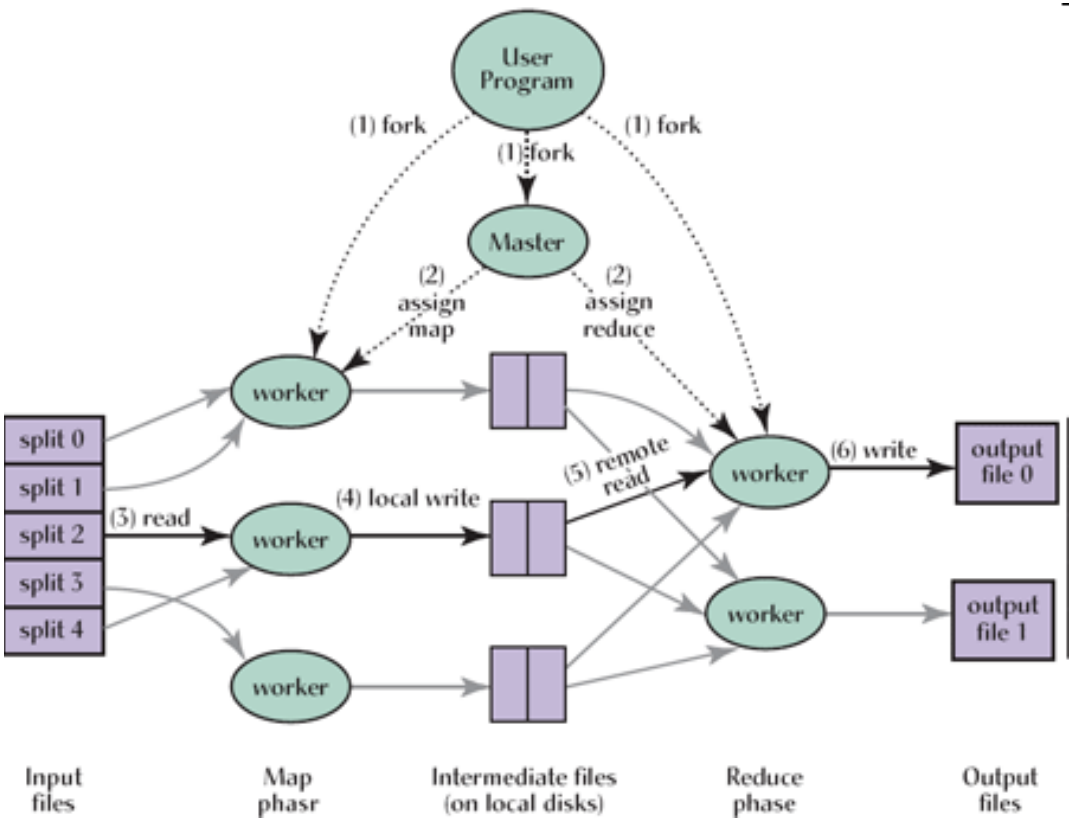

- Virtual Machine:
  Run multiple OSes on single machine.
- Separate modules,
  Present bugs and virus from infecting.
- Dynamically allocate resource.

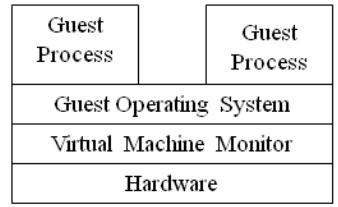

Type I VMM

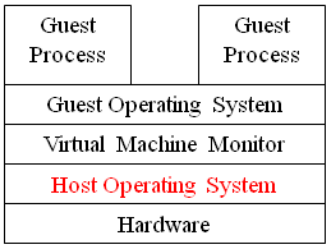

Type II VMM

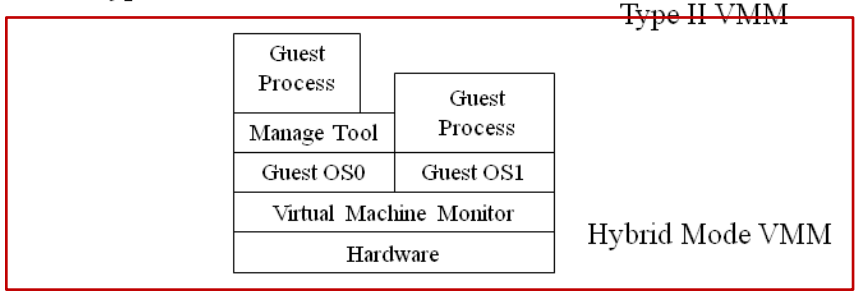

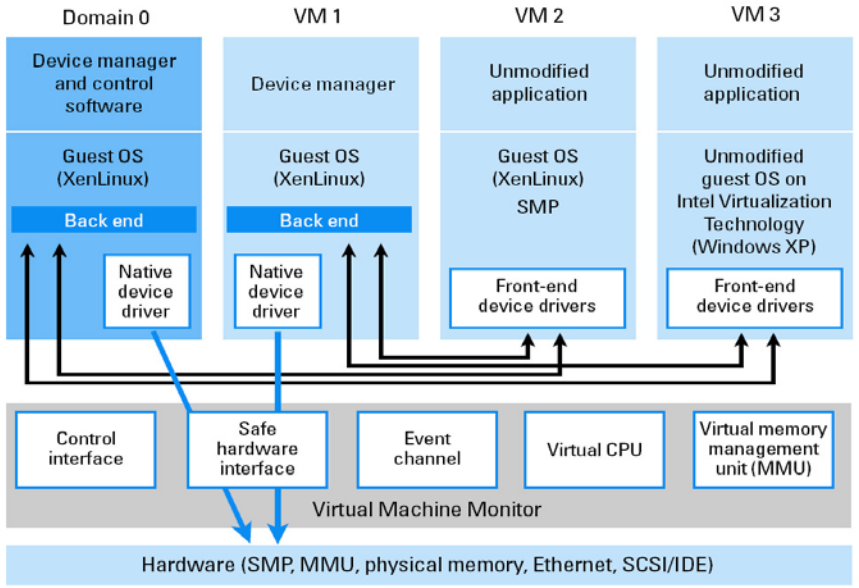

- VLAN:
  Regardless physical locations,
  multiple machine operate as if in the same network domain.
- VLAN vs. VPN
- Group machines in different regions
  as in one LAN.
- Separate machines in the same LAN,
  into different groups,
  mutually inaccessible.

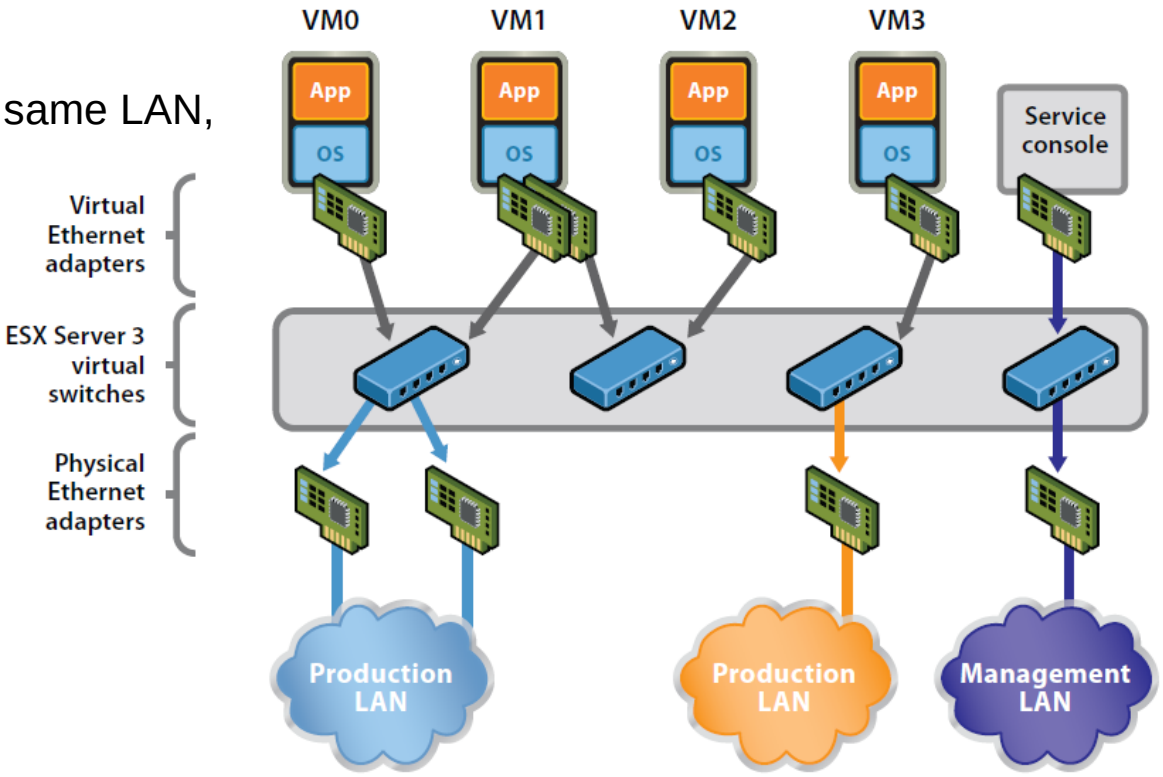

- Traffic Monitoring and Network Topology.

  Construct the entire cloud system.

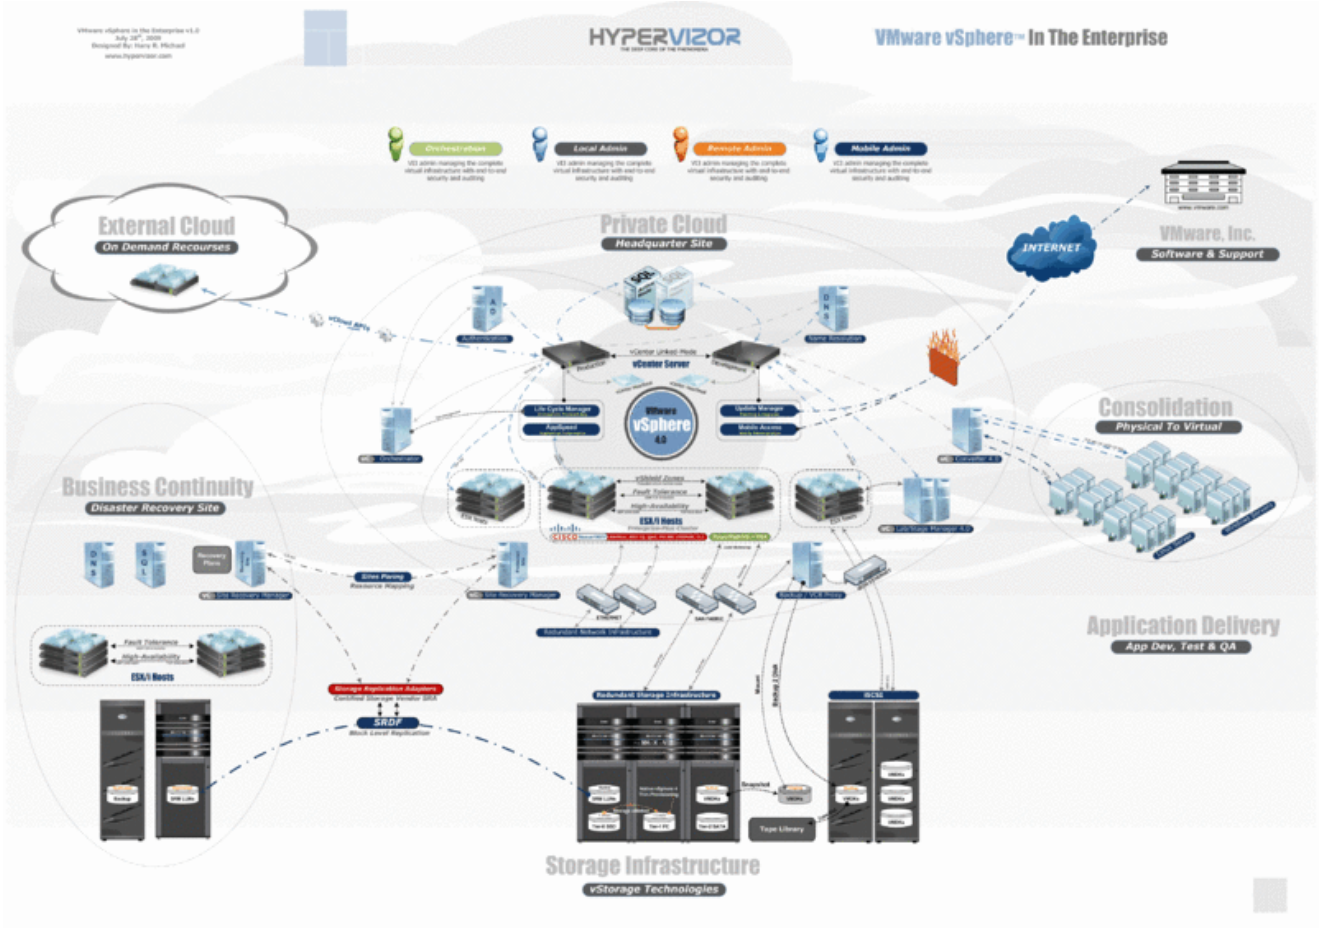

- Future trends:
  smaller, bigger, faster, easier.
- One chip with 48 CPUs.
- Data-center TCP.
- Cloud in RAM.
- Erlang, PigLatin:
  languages for cloud computing.

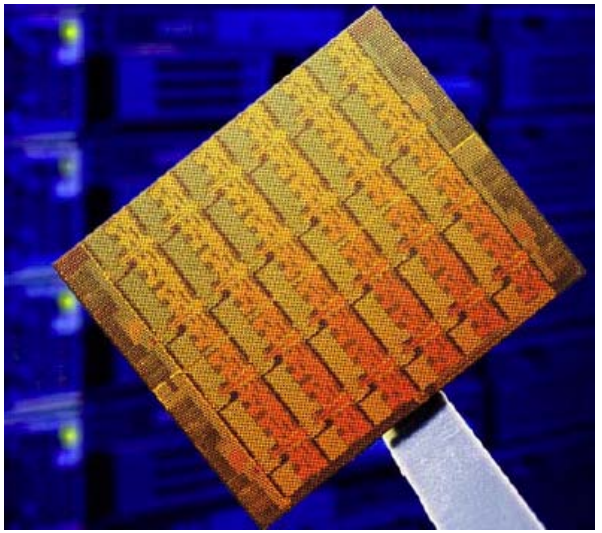

Syllabus
Invited Seminars
Homework

1. 2/28, 18:00pm - 20:00pm, Tuesday,
Introduction to clouding computer? Why cloud, what to do, and how to do?
Homework: Construct a simple 3-tier website.

2. 3/6, 18:00pm - 20:00pm, Tuesday,
Cluster-based scalable network services, SOA.
Homework: Learn to use THRIFT and MemCached to implement a messaging system.

3. 3/13, 18:00pm - 20:00pm, Tuesday,
Scalable file system, Google file system.
Homework: Learn to use SWIFT file system.

4. 3/20, 18:00pm - 20:00pm, Tuesday,
Distributed RDBMS database, Google Bigtable.
Homework: Learn to use Hadoop HBase.

5. 3/27, 18:00pm - 20:00pm, Tuesday,
Invited seminar: Baidu.

6. 4/3, 18:00pm - 20:00pm, Tuesday,
Distributed Locking system, Paxos and Google Chubby.
Homework: Learn to use Hadoop ZooKeeper

7. 4/10, 18:00pm - 20:00pm, Tuesday,
Distributed NO-SQL Database.
Homework: Learn to use Facebook Cassandra.

8. 4/17, 18:00pm - 20:00pm, Tuesday,
Paralleled computation, Google MapReduce.
Homework: Learn to use Hadoop MapReduce.

- Syllabus.
- Core cloud techniques.
  Understand principles,
- Learn how to use,
  but not re-implement.
  (that is for advanced courses)

9. 4/24, 18:00pm - 20:00pm, Tuesday,
Invited Seminar: Taobao.

10. 5/1, 18:00pm - 20:00pm, Tuesday,
Virtual Machine for dynamic resource allocation.
Homework: Learn to use KVM.

11. 5/8, 18:00pm - 20:00pm, Tuesday,
Cloud security and VLAN.
Homework: TBD

12. 5/15, 18:00pm - 20:00pm, Tuesday,
Invited seminar: EMC/VMWare.

13. 5/22, 18:00pm - 20:00pm, Tuesday,
Datacenter network topology and traffic management.
Homework: Learn to use Zookeeper.

14. 5/29, 18:00pm - 20:00pm, Tuesday,
Invited seminar: Google.

15. 6/5, 18:00pm - 20:00pm, Tuesday,
Future Trend:
Bigger: Datacenter as a warehouse-scale computer, Datacenter needs an OS.
Smaller: Multicore CPU and GPUs.
Faster: In-Memory Framework, Piccolo an Spark.
Easier: Erlang, PigLatin language.

16. 6/12, 18:00pm - 20:00pm, Tuesday,
Invited seminar: CloudValley.

- Syllabus.
- Core cloud techniques.
  Understand principles,
- Learn how to use,
  but not re-implement.
  (that is for advanced courses)

- Invited seminars.
- Top cloud players will be your teachers.
- Diverse opinions, also deviated from theory, and why?
- Scheduled for mid-term & final exam periods, and no homework!

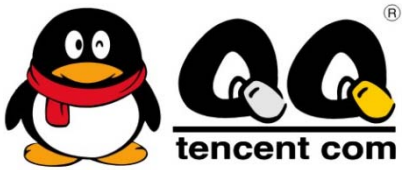

- Homework.
- Homework: 50%

  Mid-term exam: 20%

  Final exam: 30%
- You will be able to build a cloud!

  Not just Hadoop, and beyond.

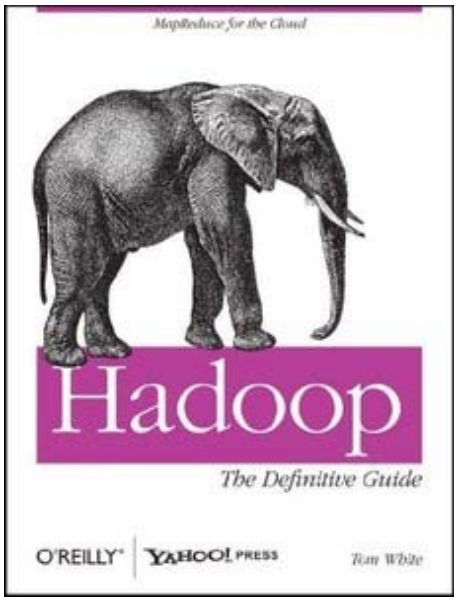

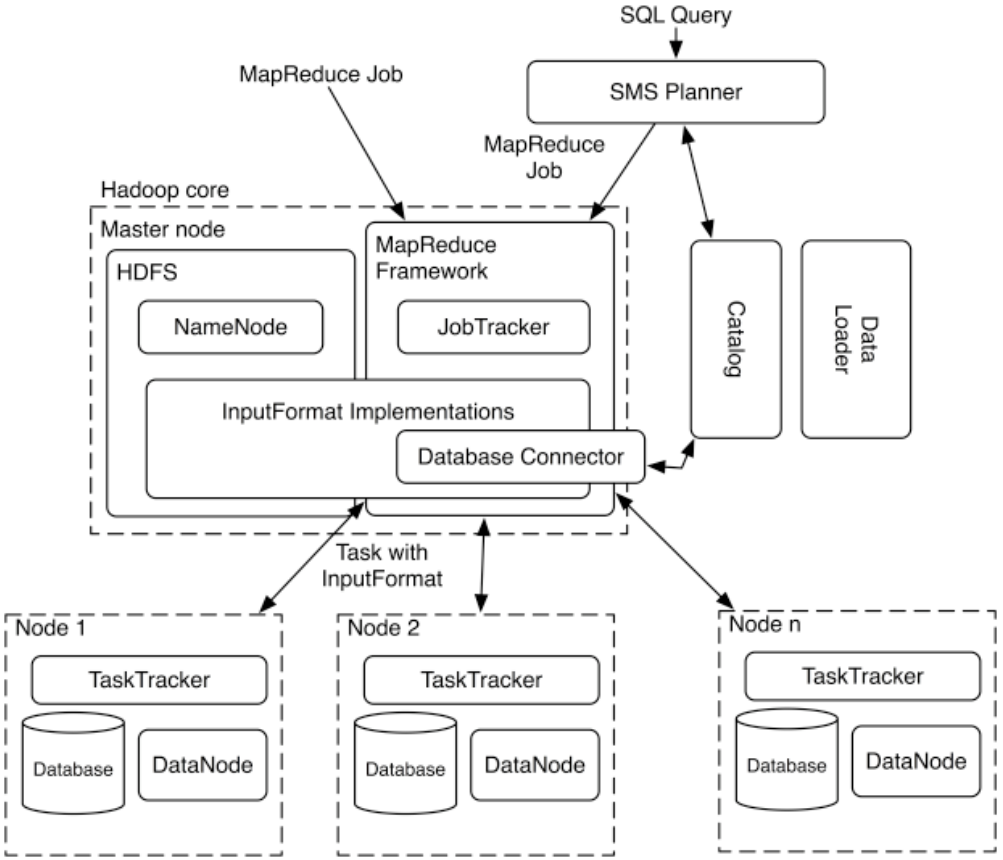

Q&A

- No stupid questions, but it is stupid if not ask!
- Ask a good question, and impress your professor and classmates!