



Data Processing in the Cloud at Yahoo!

Sanjay Radia

Chief Architect, Hadoop/Grid

Sradia@yahoo-inc.com

Cloud Computing & Data Infrastructure

Yahoo Inc.





Outline

- What are clouds and their benefits
- Clouds at Yahoo
- Hadoop





Types of Cloud Services

- Two kinds of cloud services:
 - Horizontal (“Platform”) Cloud Services
 - Functionality enabling tenants to build applications or new services on top of the cloud
 - Functional Cloud Services
 - Functionality that is useful in and of itself to tenants. E.g., various SaaS instances, such as Salesforce.com; Google Analytics and Yahoo!’s IndexTools; Yahoo! properties aimed at end-users and small businesses, e.g., flickr, Groups, Mail, News, Shopping
 - Could be built on top of horizontal cloud services or from scratch
 - Yahoo! has been offering these for a long while (e.g., Mail for SMB, Groups, Flickr, BOSS, Ad exchanges)





Cloud Characteristics

- Elastic, horizontal, capacity
 - Grow and shrink as needed
 - Large scale possible
- Utility – pay-per-usage, often no long term commitment needed
- Multi-tenant
- Availability, operations management, and connectivity built in
 - Customer focuses on his problem not IT management
- Private/Public
 - Private (within an organization): - e.g. Yahoo!, Google
 - Public: e.g. EC2, ElasticMapReduce App Engine
- Different abstraction levels
 - EC2 – provides “bare” machines (VMs)
 - Higher level: Hadoop (Yahoo), ElasticMapReduce, App Engine





Benefits of utility cloud computing

- Business agility, business innovation, technical innovation
 - Key computations solved in days and not months
 - Try out new ideas
 - Projects move from research to production in days
 - Easy to learn, even our rocket scientists use it!
- Scale, performance, availability
 - More robust, more global, more complete, better growth, for a given budget
- Cost
 - Lower cost – infrastructure is shared
- Major factors
 - You don't need to find new hardware to experiment
 - You can work with all your data!
 - No need for R&D to do IT (it just works)
 - Focus on the problem not IT





Yahoo!'s Unique Cloud: Unprecedented Scale

- **Massive user base and engagement**
 - 500M+ unique users per month
 - Hundreds of petabyte of storage
 - Hundreds of billions of objects
 - Hundred of thousands of requests/sec
- **Global**
 - Tens of globally distributed data centers
 - Serving each region at low latencies
- **Challenging Users**
 - Rapidly extracting value from voluminous data
 - Downtime is not an option (outages cost \$millions)
 - Variable usage patterns






Technology Strategy


- **Open Source**
 - Hadoop, PIG, Xen, ...
- **High performance**
 - Throughput
 - Latency
- **Multi-data center**
 - N-way replication
 - Consistency/availability tradeoffs
 - BCP
- **Supporting technologies used at Yahoo!**
- **Flat, flexible infrastructure**





Hadoop: A Key Cloud Technology

- A framework for storing & processing Petabyte of data using commodity hardware and storage
- Storage: HDFS, HBase
- Processing: MapReduce, Pig
- Distributed coordination: Zookeeper





Hadoop Characteristics

- Commodity HW + Horizontal scaling
 - Add inexpensive servers with JBODS
 - Storage servers and their disks are *not* assumed to be highly reliable and available
- Use replication across servers to deal with unreliable storage/servers
- Metadata-data separation - simple design
 - Storage scales horizontally
 - Metadata scales vertically (today)
- Slightly Restricted file semantics
 - Focus is mostly sequential access
 - Single writers
 - No file locking features
- Support for moving computation close to data
 - i.e. servers have 2 purposes: data storage and computation
- MapReduce Data processing framework

Simplicity of design

why a small team could build such a large system in the first place





A growing user base - "powered by"

Year: 2007



Year: 2008



Year: 2009





Hadoop is critical to Yahoo's business

My Yahoo! | May 12, 2009

Sign In | New here? Sign Up

The screenshot shows the Yahoo! homepage interface with several key areas annotated with blue callouts:

- Search Index:** Located in the top right, it lists top search terms such as "Carrie Prejean", "Leonard Nimoy", "Mission", "Farrah Fawcett", "Dean", "U.S. Stamps", and "TV Recaps".
- Machine Learning (e.g. Spam filters):** Annotated over the "Check your Yahoo! Mail" section, which includes a "Sign In" button and a "Don't have a Yahoo! ID?" link.
- Ads Optimization:** Annotated over a Toyota advertisement for a "PERFECT TIMING EVENT" featuring a silver car.
- Content Feed Processing:** Annotated over a news feed section showing headlines like "White House officials say no decision has been..." and "Light from an exploding star races across..."

Other visible elements include the "MY FAVORITES" sidebar with links to sites like eBay, Finance, and Games; the "Web Search" bar; and a "Yahoo! Mail" preview window showing an email list with recipients like Jack Roberts and Macy's.





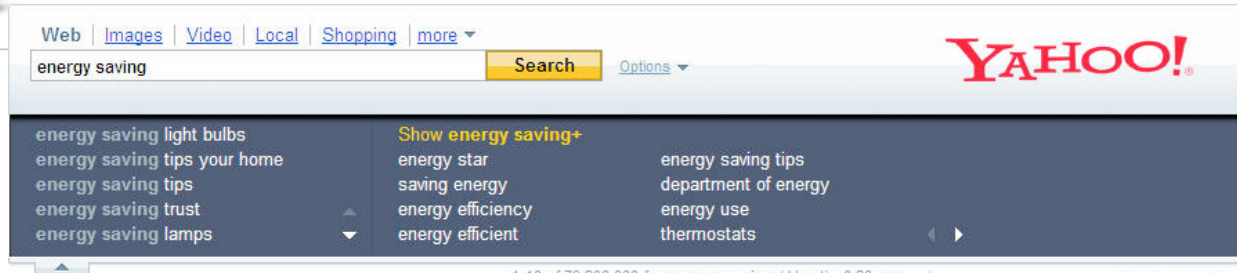
Massive Hadoop-Based Application @ Yahoo!

	2008	2009
Webmap	~70 hours runtime ~300 TB shuffling ~200 TB output	~73 hours runtime ~490 TB shuffling ~280 TB output +55% hardware
Terasort	209 seconds 1 Terabyte sorted 900 nodes	62 seconds 1 Terabyte, 1500 nodes 16.25 hours 1 Petabyte, 3700 nodes
Largest cluster	2000 nodes •6PB raw disk •16TB of RAM •16K CPUs	4000 nodes •16PB raw disk •64TB of RAM •32K CPUs •(40% faster CPUs too)






Hadoop Applications: Search Assist™



- » Database for Search Assist™ is built using Hadoop.
 - » 3 years of log-data
 - » 20-steps of map-reduce

	Before Hadoop	After Hadoop
Time	26 days	20 minutes
Language	C++	Python
Development Time	2-3 weeks	2-3 days





Collaborations around the globe

- M45 - Yahoo!'s supercomputing cluster
 - Carnegie Mellon University
 - The University of California at Berkeley
 - Cornell University
 - The University of Massachusetts at Amherst joined
- Partners in India
 - Computational Research Laboratories (CRL), India's Tata Group
 - Universities – IIT, IISc, IIIT-H, PSG
- Open Cirrus™ - cloud computing research & education
 - The University of Illinois at Urbana-Champaign
 - Infocomm Development Authority (IDA) in Singapore
 - The Karlsruhe Institute of Technology (KIT) in Germany
 - HP, Intel
 - The Russian Academy of Sciences, Electronics & Telecomm.
 - Malaysian Institute of Microelectronic Systems





More Info



- Main Web sites

- <http://hadoop.apache.org/core/>
- <http://wiki.apache.org/hadoop/>
- <http://wiki.apache.org/hadoop/GettingStartedWithHadoop>
- <http://wiki.apache.org/hadoop/HadoopMapReduce>





THANKS

Sanjay Radia

sradia@yahoo-inc.com

Cloud Computing & Data Infrastructure

Yahoo!

