




Cloud Computing  
.@

**YAHOO!**®

the premier  
digital media company

.

# Yahoo's most famous cloud product?



<http://developer.yahoo.com/blogs/ymn/posts/2011/02/full-house-apache-hadoop-india-summit-2011/>

But this talk is  
not about  
Hadoop

# Overview

1. Yahoo Private Cloud Overview
  - Y! Cloud Goals
  - Y! Software stack and cloud fit
  - Y! Cloud platforms
2. Yahoo Cloud Serving Engine
  - Beyond IaaS



# Technology & Science at Global Internet Scale

## The Ultimate Sandbox for Science

» **640+ Million Users**

» **368 Million**

PEOPLE VISIT THE  
YAHOO!  
HOMEPAGE MONTHLY

» **Billions** OF ADS  
SERVED EACH DAY

» **4.5 Billion**

PAGE VIEWS  
DAILY

» **#1** IN MAIL

» **77 Billion**

MESSAGES SENT FROM  
MAIL USERS MONTHLY

» **#1** IN SPORTS, NEWS,  
FINANCE,  
AND ENTERTAINMENT

» **81 Billion**

MESSAGES SENT FROM  
112 MILLION Y! MESSENGER  
USERS

» **3 Million** FLICKR PHOTOS  
UPLOADED EVERY  
DAY

» **70 Billion**

MINUTES SPENT MONTHLY  
ON COMMUNICATIONS  
PROPERTIES



# Infrastructure: Agility & Stability

competing needs



accelerating  
innovation




increasing  
stability

# Infrastructure: Agility & Stability

competing needs



accelerating  
innovation



increasing  
stability




# Infrastructure: Agility & Stability

infrastructure to the rescue



accelerating  
innovation



increasing  
stability

# Case study : Yahoo! Mail

Enabling quick response in the spam arms race

450M mail boxes  
5B+ deliveries/day

Antispam models retrained  
every few hours on Hadoop

“  
40% less spam  
than Hotmail and  
55% less spam  
than Gmail  
”



# Yahoo!'s Cloud is behind every click

600 million users visit 11 billion times  
a month generating 98 billionpage views

Cloud allows the ability to collect,  
transform, store analyze and leverage  
big data

# Yahoo!'s cloud is behind every click.

Caching, Load Balancing, Edge

Search  
Index

Machine  
Learning  
(e.g. Spam filters)

Content  
Optimization

Advertising  
Optimization  
& Delivery

*40% less spam  
than Hotmail and  
55% less spam  
than Gmail*

RSS Feeds

Image/Video  
Storage &  
Delivery

The collage features several Yahoo! web pages:

- Homepage:** Shows the top navigation bar with links for Web, Images, Video, Local, Shopping, and More. The main content area includes a large image of an astronaut on the moon, a section for "Astronauts' pitch to President Obama", and a "POPULAR SEARCHES" list.
- Email Interface:** On the left, a screenshot of the Yahoo! Mail interface shows a list of recent emails and a "Hello, Bryan!" greeting.
- Search Results:** On the right, a screenshot of search results shows a Toyota Prius advertisement and a "Sell, market, and recruit online" section.
- Market Data:** At the bottom, a section displays market data for Dow and Nasdaq.

Overlaid callout boxes highlight the following services and features:

- Caching, Load Balancing, Edge** (top center)
- Search Index** (top right)
- Machine Learning (e.g. Spam filters)** (left side)
- Content Optimization** (center)
- Advertising Optimization & Delivery** (right side)
- RSS Feeds** (bottom center)
- Image/Video Storage & Delivery** (bottom right)

At the bottom right, there is a large yellow exclamation mark logo.

# Yahoo! Cloud is open source

## Benefits

- » Avoid technological dead ends
- » Leverage community contributions
- » Workforce already trained

### Ongoing contributions



### Yahoo!'s adoption of open source



### Future contributions

*Cloud serving  
Storage*



- So...What is Cloud Computing?

Oh! No, here we  
go again...



# The NIST definition

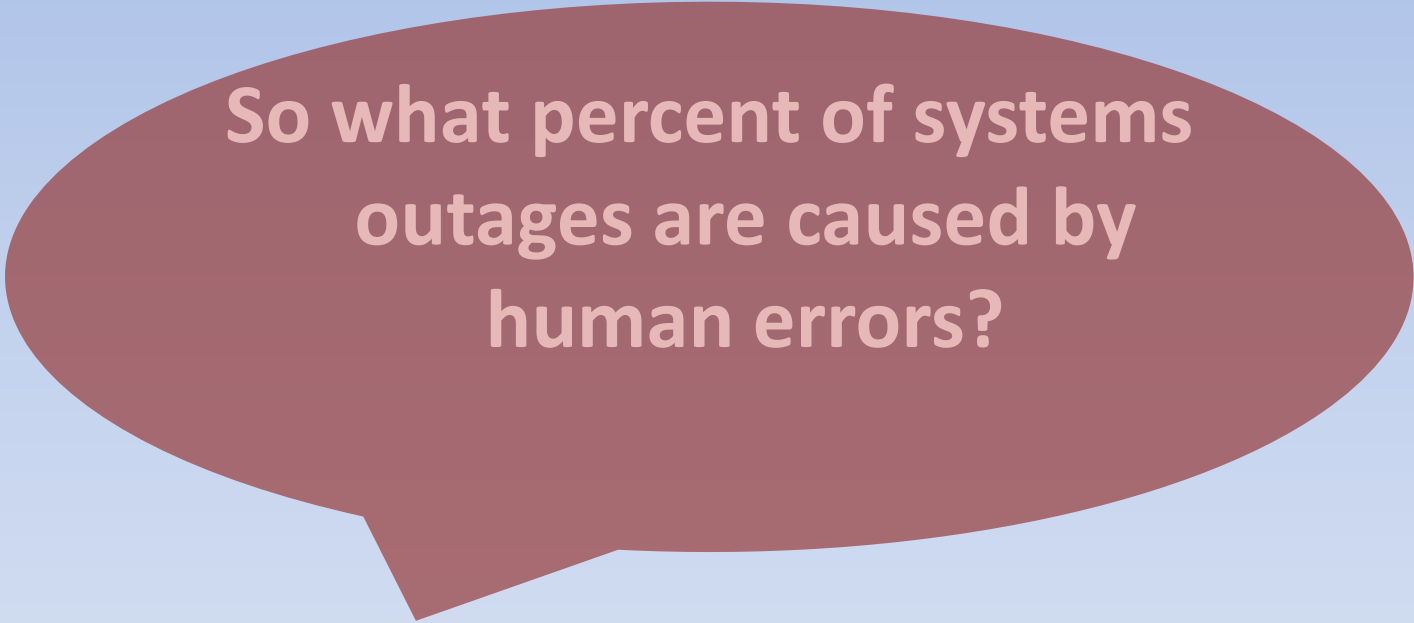
<http://csrc.nist.gov/groups/SNS/cloud-computing/index.html>

- Cloud computing is a model for enabling convenient, on-demand network access to a shared pool of configurable computing resources (e.g., networks, servers, storage, applications, and services) that can be rapidly provisioned and released with minimal management effort or service provider interaction. This cloud model promotes availability and is composed of five **essential characteristics**, three **service models**, and four **deployment models**.
- Five essential characteristics
  1. On-demand self-service
  2. Broad network access
  3. Resource pooling
  4. Rapid elasticity
  5. Measured service
- Three service models
  1. Cloud Software as a Service (SaaS)
  2. Cloud Platform as a Service (PaaS)
  3. Cloud Infrastructure as a Service (IaaS)
- Four deployment models
  1. Private cloud
  2. Community cloud
  3. Public cloud
  4. Hybrid cloud

# Yahoo Cloud Serving *“To IaaS and Beyond”*








**So what percent of systems  
outages are caused by  
human errors?**

# ACME: Lessons about human operators

- Human error is largest single failure source
  - HP HA labs: human error is #1 cause of failures (2001)
  - Oracle: half of DB failures due to human error (1999)
  - Gray/Tandem: 42% of failures from human administrator errors (1986)
  - Murphy/Gent study of VAX systems (1993):



# The problem

## **Engineers spend a lot of energy in**

- Deployment specification
  - Resource allocation, which host should run which packages
  - Interconnection between components, settings
- Deployment/upgrade process
  - How to upgrade the software from one version to the next with minimum disruption
  - How to cleanly undo the changes, if problems are discovered

## **There is no clear separation of concerns between architecture, development and deployment**

- Architect should be able to specify the architecture, which is common for all environments
- The developer should not worry about the scale, interconnection details

## **Very hard to build new environments for experiments, testing**



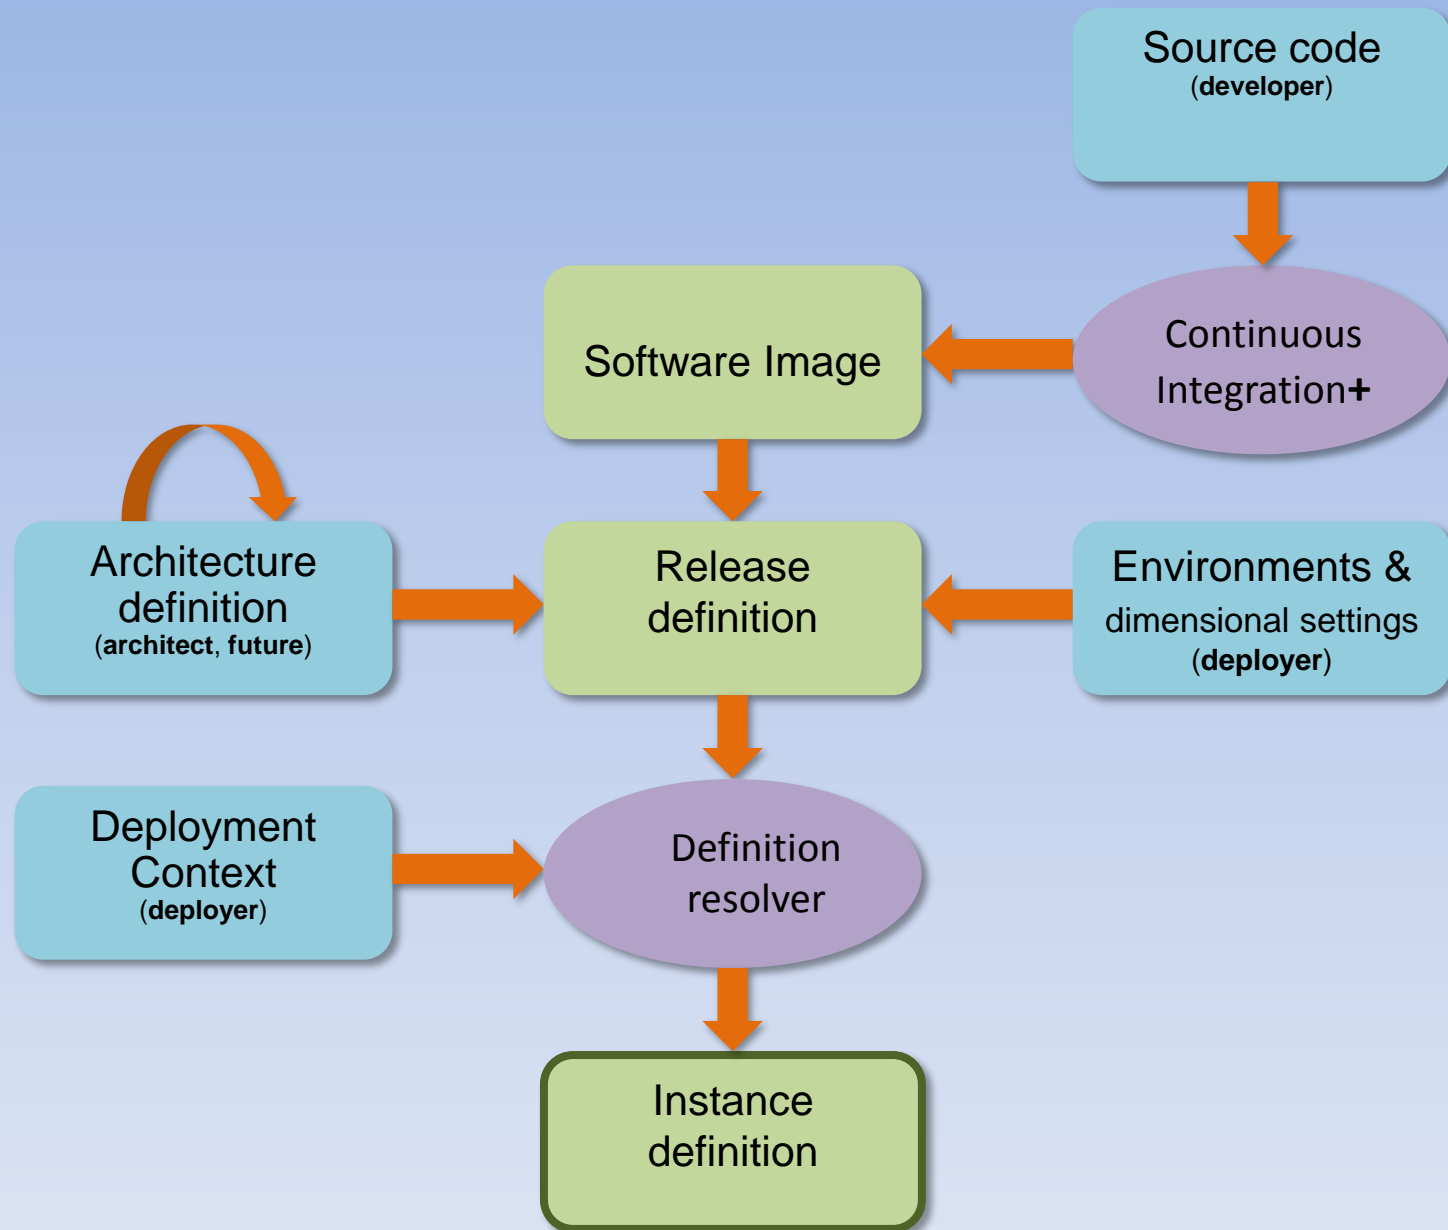
# Y! Cloud Serving

CSE does the heavy lifting in terms of

- Deployment specification
  - Release defines the definition space for instance definitions
  - Instance definition, defines a specific deployment instance
- Deployment/Upgrade process
  - Uniform deployment/upgrade process
  - Deployers can specify wait points in the process for inspection
  - Anytime in the middle, the deployer can decide to rollback, cleanly
- Tools for separation of concerns
  - Definition resolver: Release-definition  $\times$  Deployment-context  $\rightarrow$  Instance-definition
  - Tools for building architecture definition and release definition

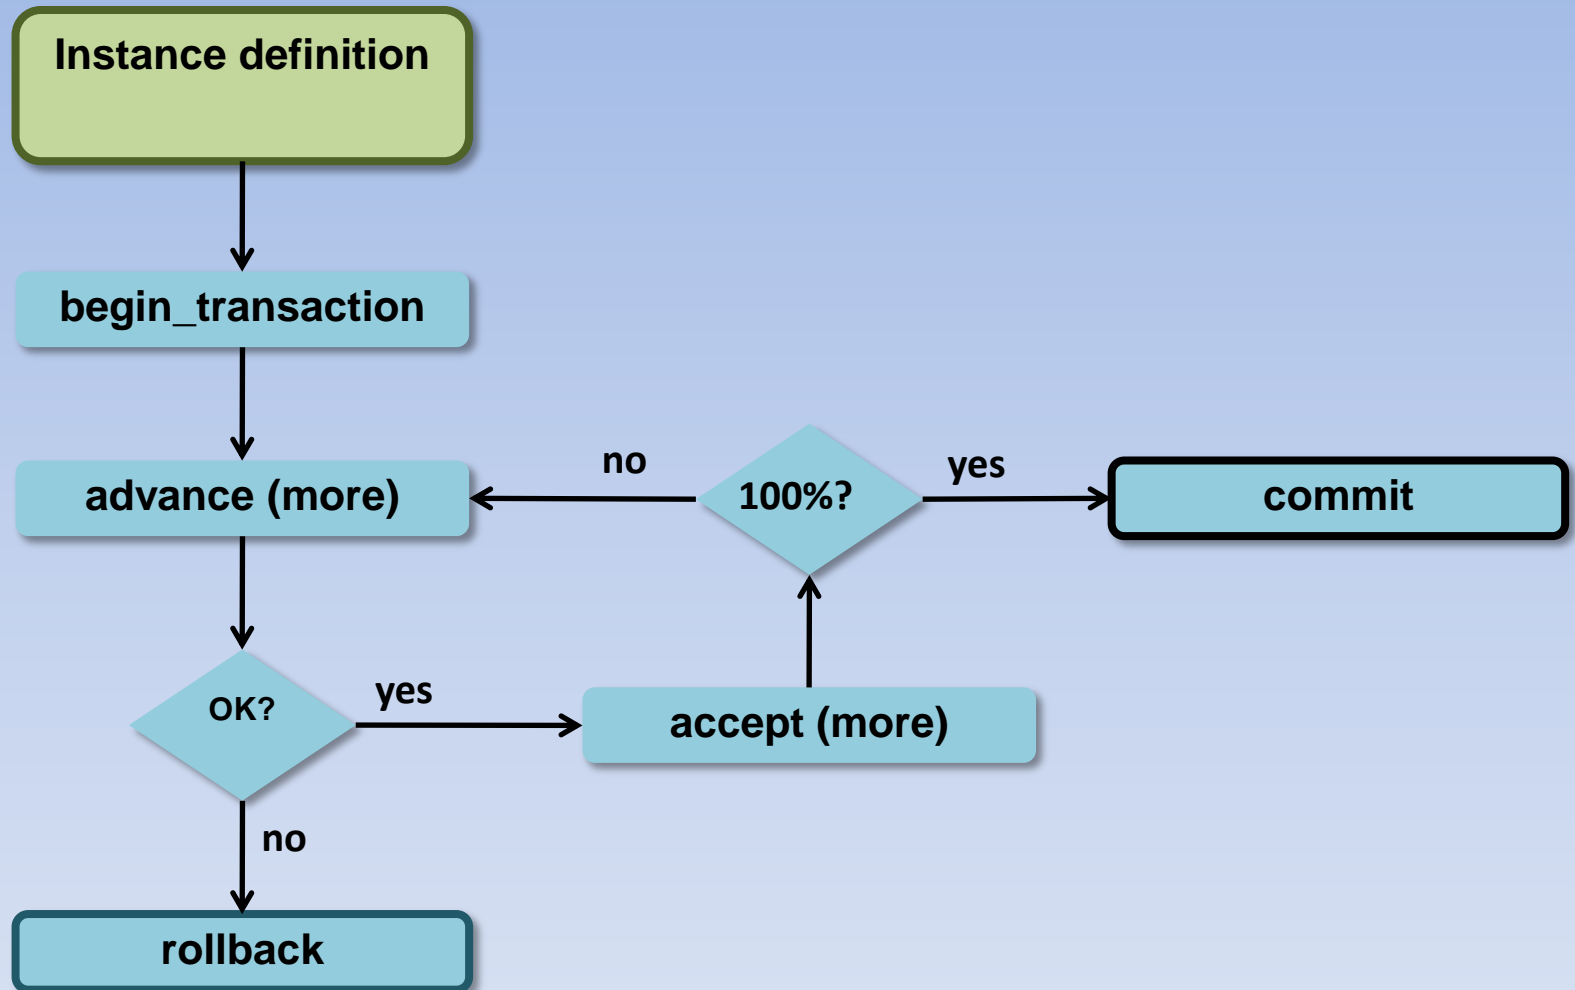
# Declarative Nirvana

- “Every non-trivial configuration system is a programming language; you might just think it that way.” – **James Gosling**
- This language (configuration system) has issues
  - Procedural – Communicate using a procedure
  - (a) Configuration files (b) Notes on Twikis/Runbooks (c) In people’s heads
- Translation issues from Architect to Dev to Testing to Production
- CSE
  - Declarative Language to describe deployments
    - Describe a 1000 line procedure in 10 lines
      - Just the desired end state
    - Automated, Uniform, reliable deployment procedure.





# Deployment process



Caveat: Approximate description

# Key Takeaways

- Think beyond IaaS
- Great savings to be had in deployment automation
- Software deployment at the touch of a button
- Multi Tier - need for bindings



# YAHOO!®

Thanks for Listening...  
Questions?



Q&A