Integrating Kerberos into Apache Hadoop

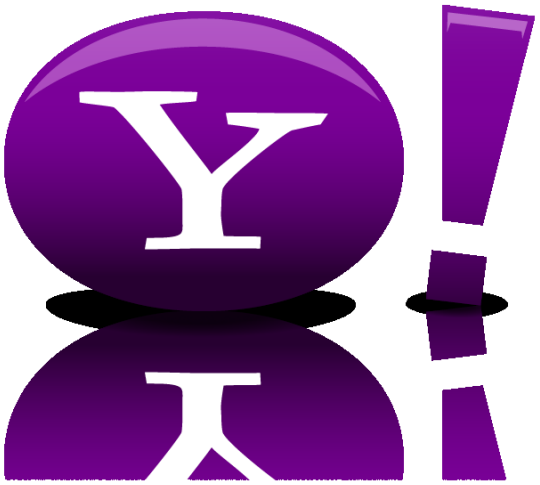

Owen O'Malley
owen@yahoo-inc.com
Yahoo's Hadoop Team

Kerberos Conference 2010

Who am I

- An architect working on Hadoop full time
  - Mainly focused on MapReduce
- Tech-lead on adding security to Hadoop
- Before Hadoop – Yahoo Search WebMap
- Before Yahoo – NASA, Sun
- PhD from UC Irvine

What is Hadoop?

- A framework for storing and processing big data on lots of commodity machines.
  - Up to 4,000 machines
  - Up to 20 PB
- Open Source Apache project
- High reliability done in software
  - Automated failover for data and computation
- Implemented in Java

What is Hadoop?

- HDFS – Distributed File System
  - Combines cluster's local storage into a single namespace.
  - All data is replicated to multiple machines.
  - Provides locality information to clients
- MapReduce
  - Batch computation framework
  - Tasks re-executed on failure
  - User code wrapped around a distributed sort
  - Optimizes for data locality of input

What is Hadoop NOT?

- Hadoop is aimed at moving large amounts of data efficiently.
- It is not aimed at doing real-time reads or updates.
- Hadoop moves data like a freight train, slow to start but very high bandwidth.
- Databases answer queries quickly, but can't match the bandwidth.

Case Study: Yahoo Front Page

Personalized

for each visitor

Result:

twice the engagement

+79% clicks
vs. randomly selected

+160% clicks
vs. one size fits all

+43% clicks
vs. editor selected

Scaling and Stability

- Yahoo is the largest contributor and user of Hadoop.
- It has become the platform of choice for big data analytics.
- Moved from Research to Science to Production

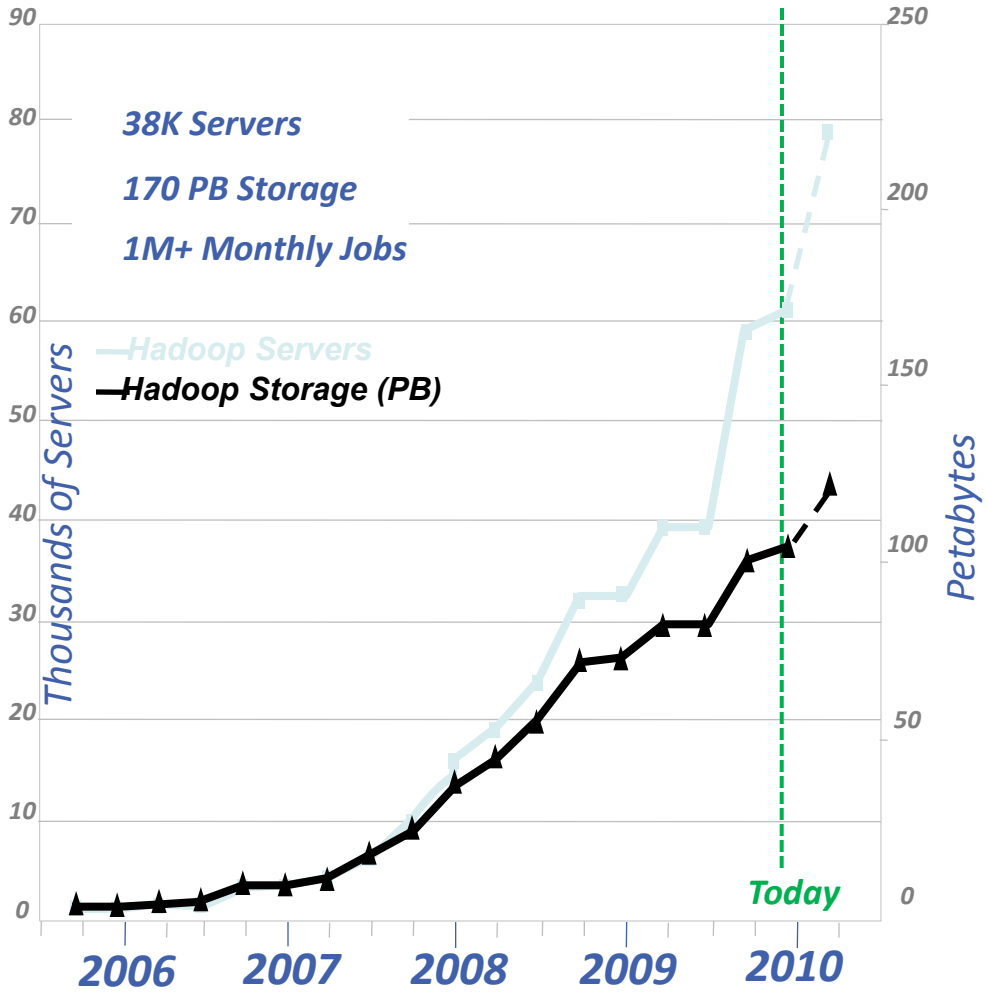

Hadoop Community

2007

YAHOO!

Powerset

last·fm

2008

Google

ableGrape

ImageShack

Cascading

IBM

facebook

ENORMO Every property. Everywhere.

A9

krugle

rackspace HOSTING

Lookery Control freaks welcome

The New York Times

Joost

Zvents Discover Things To Do

FORMATION CIENCES NSTITUTE

News Corporation

Cornell University Computing and Information Science

Visible MEASURES

LOTAME Locate, Target, & Message with Social Media

NetSeer

parc Palo Alto Research Center

SECURITY ENHANCED DOMAIN NAME SYSTEM

veoh

2009

AOL

cloudera

deepdyve BETA

cooliris

eyealike

TEXTMAP THE ENTITY SEARCH ENGINE

PSG College of Technology

iterend

tailsweep

hulu

RapLeaf

USCMS

Ning

quantcast

amazon web services

pressflip

detikSearch

WorldLingo

Systems@ETH zürich

VK SOLUTIONS Global Solutions Provider

TARAGANA Innovation + Quality + Simplicity

HOSTING HABITAT

HOLA SERVERS

Terrier

adknowledge

stampede beta

2010

SAMSUNG

rubicon PROJECT

BERKELEY LAB LAWRENCE BERKELEY NATIONAL LABORATORY

VISIBLE TECHNOLOGIES

APOLLO GROUP

ADSDAQ

rackspace HOSTING

RapLeaf

wordnik All the words.

mozigen Network Revolution

comScore

trulia real estate search

Accela COMMUNICATIONS

Forward3D

LinkedIn

Microsoft

Infochimps Find the world's data

Pharm2Phork

ADMELD

gumgum

BrainPad

Pronux

The Datagraph Blog

NETFLIX

mobileanalytics.tv

markt24.de

twitter

media6degrees

BEEBLER

SLC Security When Experience Matters...

ebay

Problem

- Yahoo! has more yahoos than clusters.
  - Hundreds of yahoos using Hadoop each month
  - 38,000 computers in ~20 Hadoop clusters.
  - Requires isolation or trust.
- Different users need different data.
  - Not all yahoos should have access to sensitive data
    - financial data and PII
- In Hadoop 0.20, easy to impersonate.
    - Segregate different data on separate clusters

Solution

- Prevent unauthorized HDFS access
  - All HDFS clients **must** be authenticated.
  - Including tasks running as part of MapReduce jobs
  - And jobs submitted through Oozie.
- Users must also authenticate servers
  - Otherwise fraudulent servers could steal credentials
- Integrate Hadoop with Kerberos
  - Provides well tested open source distributed authentication system.

Requirements

- Security must be optional.
  - Not all clusters are shared between users.
- Hadoop must not prompt for passwords
  - Makes it easy to make trojan horse versions.
  - Must have single sign on.
- Must handle the launch of a MapReduce job on 4,000 Nodes

Definitions

- **Authentication** – Determining the user
  - Hadoop 0.20 completely trusted the user
    - User passes their username and groups over wire
  - We need it on both RPC and Web UI.
- **Authorization** – What can that user do?
  - HDFS had owners, groups and permissions since 0.16.
  - Map/Reduce had nothing in 0.20.

Authentication

- Changes low-level transport
- RPC authentication using SASL
  - Kerberos (GSSAPI)
  - Token
  - Simple
- Browser HTTP secured via plugin
- Use auth_to_local name translation to map principals to user names.

Authorization

- HDFS
  - Command line and semantics unchanged
  - Web UI enforces authentication
- MapReduce added Access Control Lists
  - Lists of users and groups that have access.
  - mapreduce.job.acl-view-job – view job
  - mapreduce.job.acl-modify-job – kill or modify job
- Code for determining group membership is pluggable.

Delegation Tokens

- To prevent a flood of authentication requests at the start of a job, NameNode can create delegation tokens.
- Allows user to authenticate once and pass credentials to all tasks of a job.
- JobTracker automatically renews tokens while job is running.
- Cancels tokens when job finishes.

Primary Communication Paths

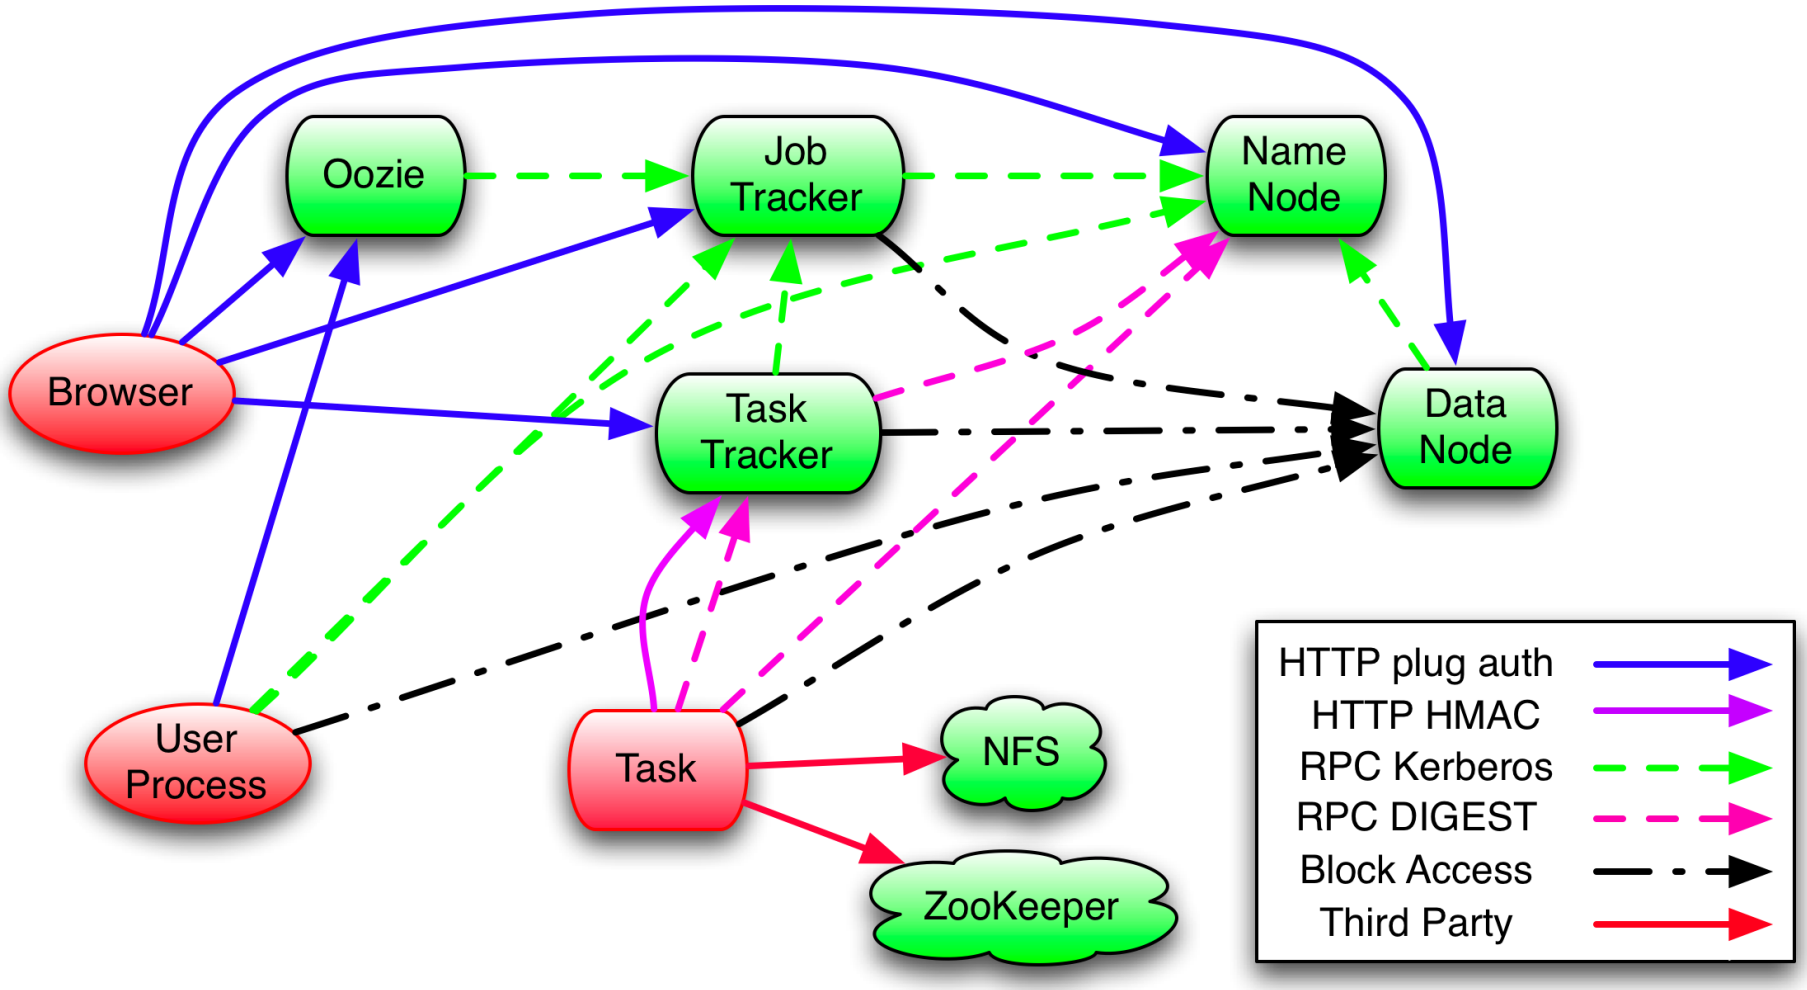

API Changes

- Very Minimal API Changes
- MapReduce added secret credentials
  - Available from JobConf and JobContext
  - Never displayed via Web UI
- Automatically get tokens for HDFS
  - Primary HDFS, File{In,Out}putFormat, and DistCp
  - Can set mapreduce.job.hdfs-servers

Web UIs

- Hadoop relies on the Web UIs.
  - These need to be authenticated also…
- Web UI authentication is pluggable.
  - Yahoo uses an internal package
  - We have written a very simple static auth plug-in
  - SPNEGO plugin being developed
- All servlets enforce permissions.

Proxy-Users

- Some services access HDFS and MapReduce as other users.
- Can’t store credentials, since they expire.
- Configure services with the proxy user:
  - Group of users that the proxy can impersonate
  - Which hosts they can impersonate from
- Provides control without over burdening operations.

Out of Scope

- Encryption
  - RPC transport – easy
  - Block transport protocol – difficult
  - On disk – difficult
- File Access Control Lists
  - Still use Unix-style owner, group, other permissions
- Non-Kerberos Authentication
  - Much easier now that framework is available

Schedule

- The security team worked hard to get security added to Hadoop on schedule.
  - Roughly 6 months of calendar time.
- Security Development team:
  - Devaraj Das, Ravi Gummadi, Jakob Homan, Owen O'Malley, Jitendra Pandey, Boris Shkolnik, Vinod Vavilapalli, Kan Zhang
- Currently on production clusters

Questions?

- Questions should be sent to:
  - common/hdfs/mapreduce-user@hadoop.apache.org
- Security holes should be sent to:
  - security@hadoop.apache.org
- Available from
  - http://developer.yahoo.com/hadoop/distribution/
  - Also a VM with Hadoop cluster with security
- Thanks!