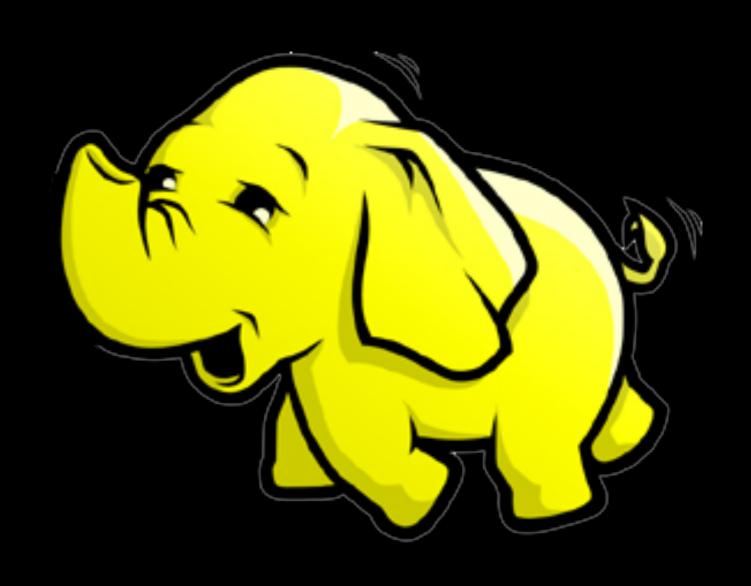
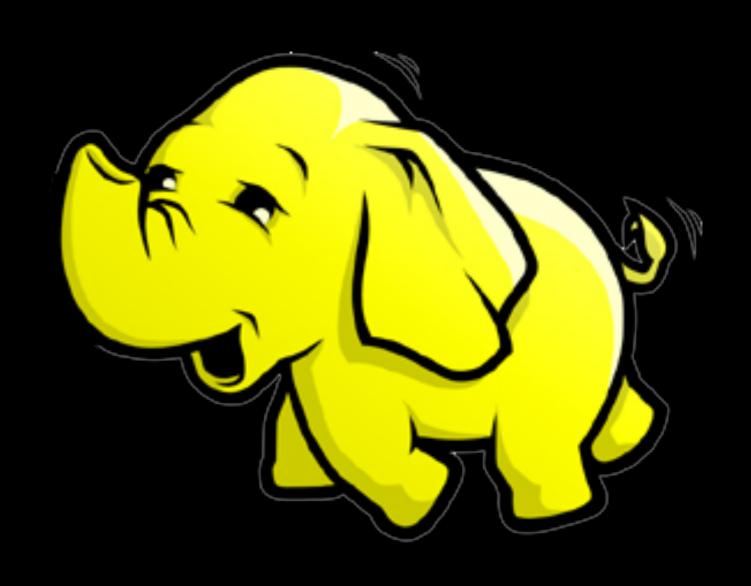
HADOOP

divide and conquer petabyte-scale data











METADATA

Matthew McCullough

Ambient Ideas, LLC

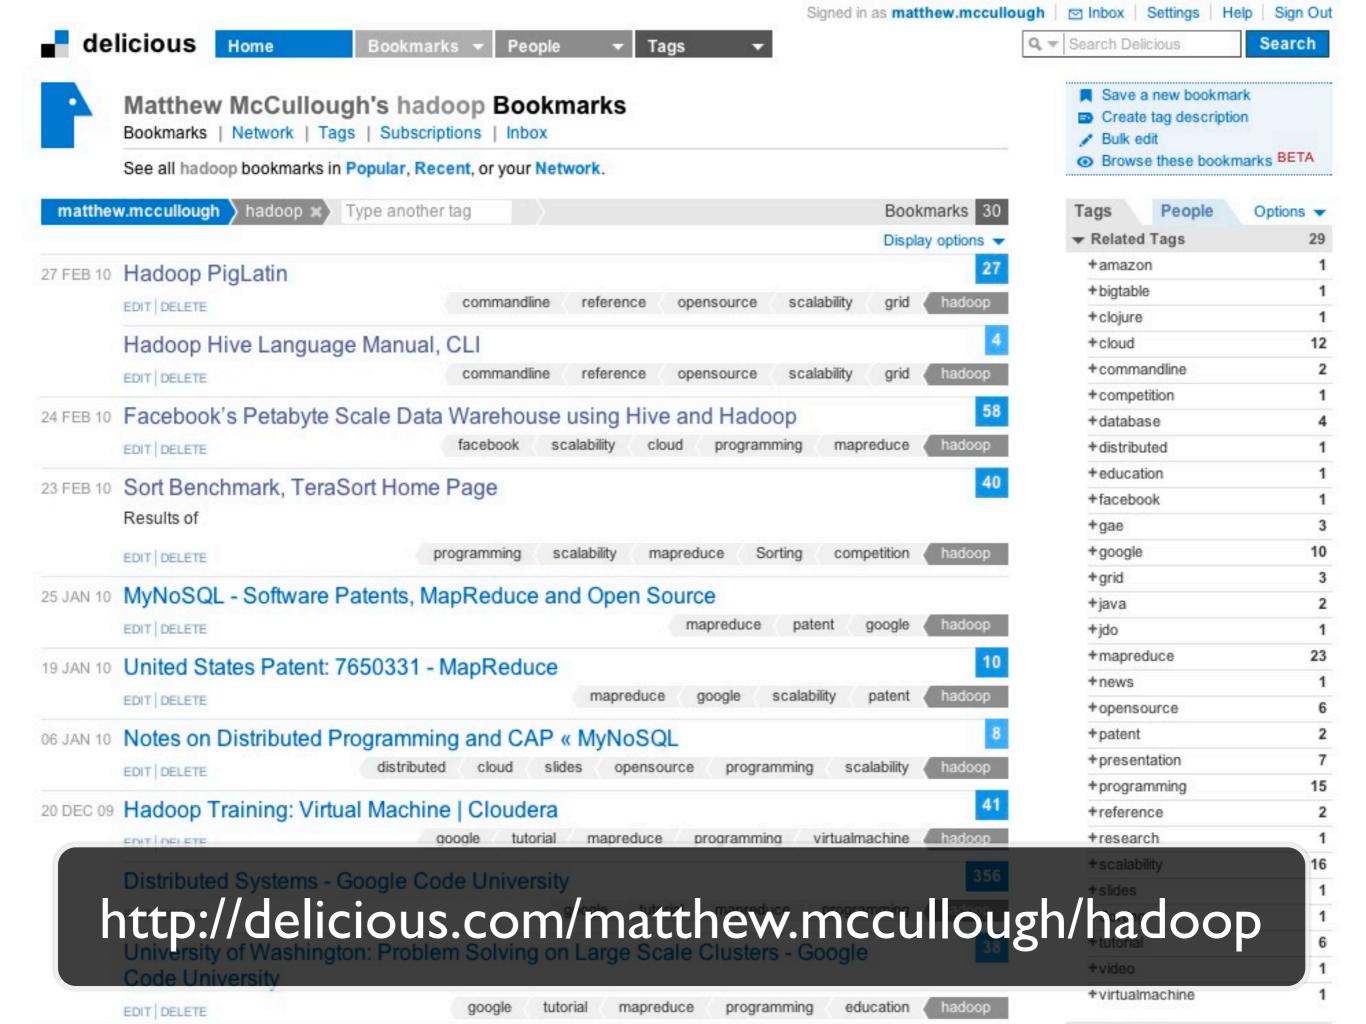
matthewm@ambientideas.com

http://ambientideas.com/blog

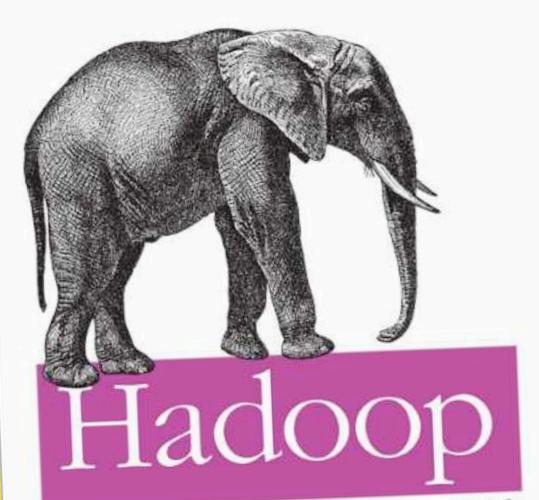
@matthewmccull

Code

http://github.com/matthewmccullough/hadoop-intro



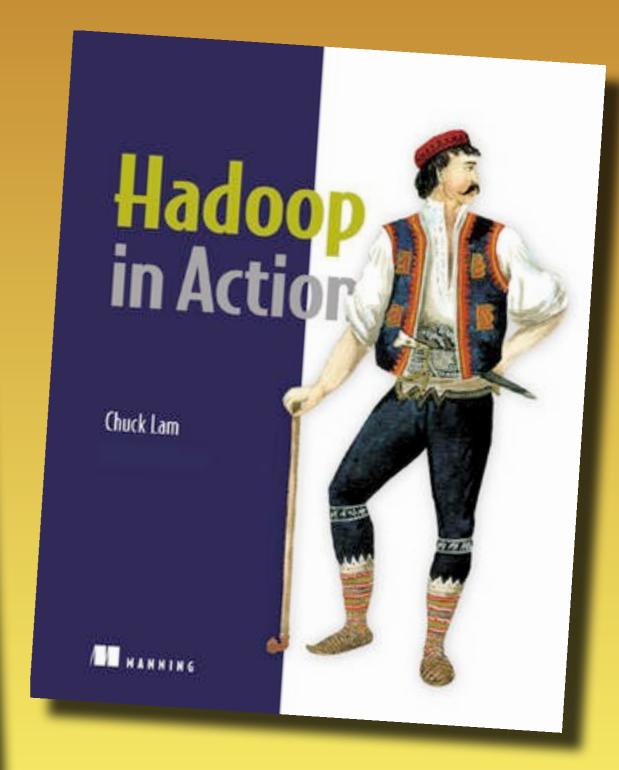
MapReduce for the Cloud



The Definitive Guide

O'REILLY" YAHOO! PRESS

Tom White



Why Hadoop? Manageduce

YOUR BIG DATA

- ➡ Data type?
- ** Business need?
- Processing?
- How large?
- Growth projections?
- What if you saved everything?

Mhy Hacoop? HDFS MapReduce

Mhy Hadoop?

I use Hadoop often.

I use Hadoop often.

That's the sound my Tivo That's the sound my Tivo makes every time I skip a commercial.

-Brian Goetz, author of Java Concurrency in Practice

BIG DATA

NO SQL

NO SQL

NO SQL

Not SQL

Not SQL

Not Only SQL

NOSQL

- **Applies to data
 - No strong schemas
 - No foreign keys
- **Applies to processing
 - No SQL-99 standard
 - No execution plan

The computer you are using right now may very well have the fastest GHZ processor you'll ever own





Scale up?

Scale out



WEB CRAWLING

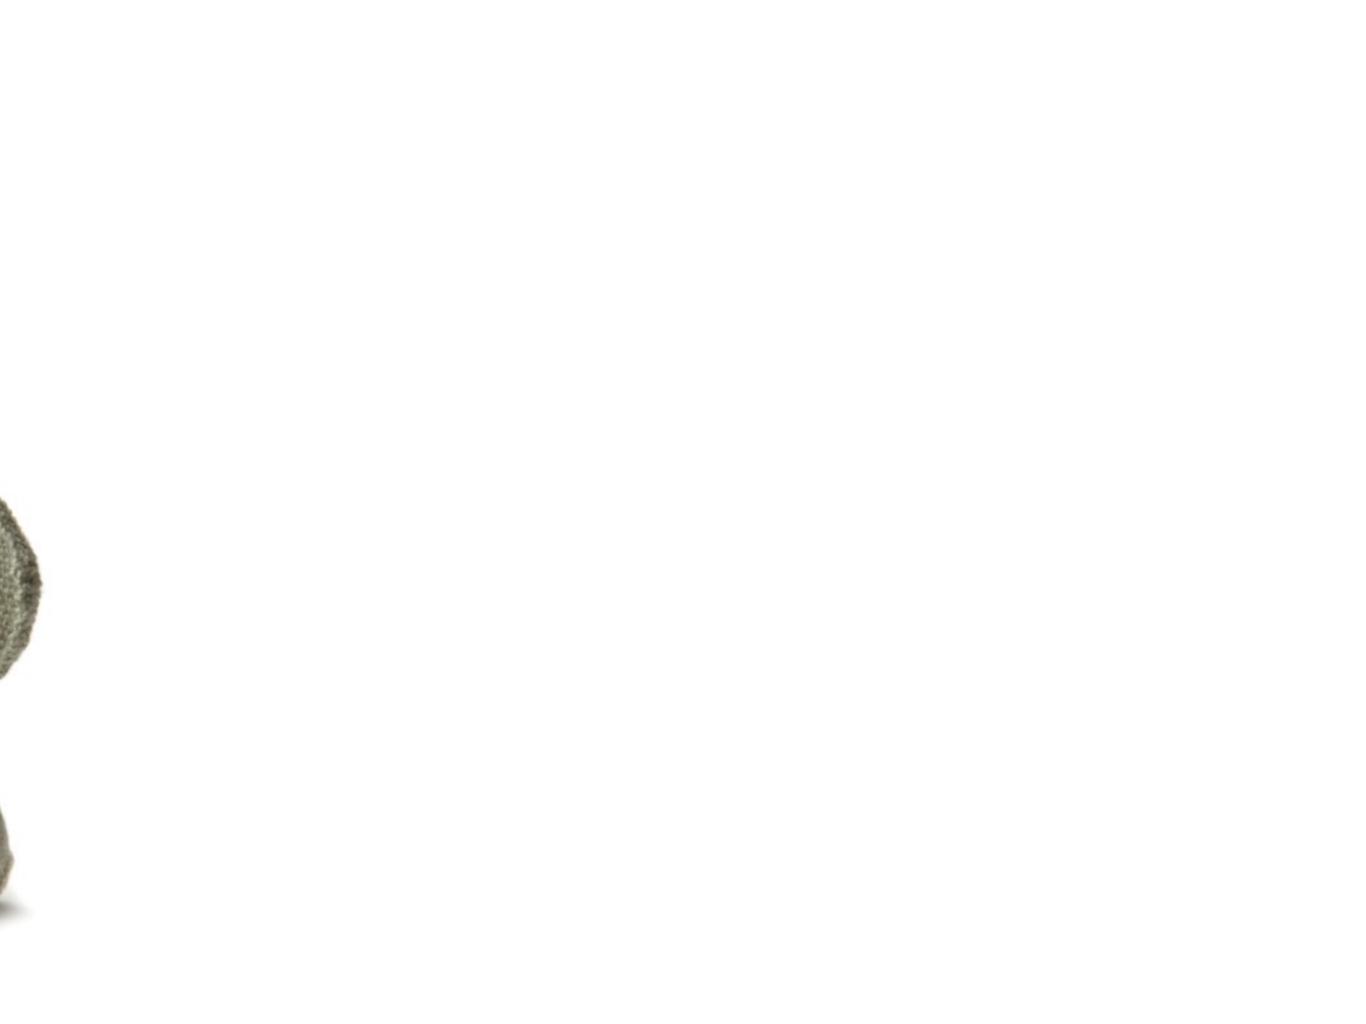


open source



MARKETINGS

NAME RESEARCHS



DAUGHTER'S TUFFED TOY



Lucene

Lucene V Nutch

Lucene Nutch Hadoop

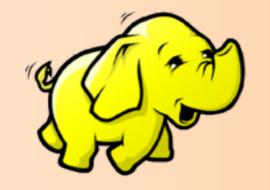
Lucene Nutch x Hadoop-

Lucene Mahout Nutch Hadoop



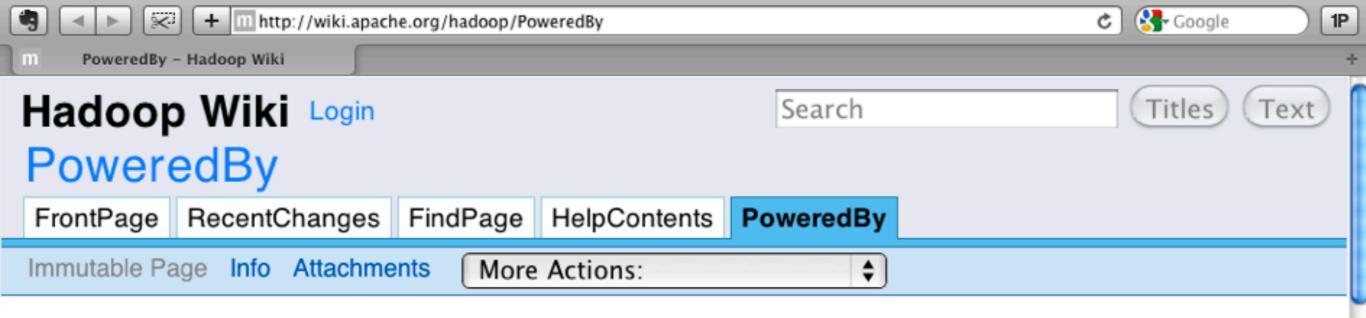


0.21.0 current version



Dozens of companies contributing

Hundreds of companies using



Applications and organizations using Hadoop include (alphabetically):

- A9.com Amazon
 - We build Amazon's product search indices using the streaming API and preexisting C++, Perl, and Python tools.
 - We process millions of sessions daily for analytics, using both the Java and streaming APIs.
 - Our clusters vary from 1 to 100 nodes.

Adobe

- We use Hadoop and HBase in several areas from social services to structured data storage and processing for internal use.
- We currently have about 30 nodes running HDFS, Hadoop and HBase in clusters ranging from 5 to 14 nodes on both production and development. We plan a deployment on an 80 nodes cluster.
- We constantly write data to HBase and run MapReduce jobs to process then store it back to HBase or external systems.
- Our production cluster has been running since Oct 2008.
- Able Grane Vertical search engine for trustworthy wine information

- We have one of the world's smaller hadoop clusters (2 hodes & 6 Cros/hode)
 - Hadoop and Nutch used to analyze and index textual information

Adknowledge - Ad network

- Hadoop used to build the recommender system for behavioral targeting, plus other clickstream analytics
- We handle 500MM clickstream events per day
- Our clusters vary from 50 to 200 nodes, mostly on EC2.
- Investigating use of R clusters atop Hadoop for statistical analysis and modeling at scale.

Salaba

- A 15-node cluster dedicated to processing sorts of business data dumped out of database and joining them together. These data will then be fed into iSearch, our vertical search engine.
- Each node has 8 cores, 16G RAM and 1.4T storage.

Amazon Web Services

- We provide Amazon Elastic MapReduce. It's a web service that provides a hosted Hadoop framework running on the web-scale infrastructure of Amazon Elastic Compute Cloud (Amazon EC2) and Amazon Simple Storage Service (Amazon S3).
- Our customers can instantly provision as much or as little capacity as they like to perform data-intensive tasks for applications such as web indexing, data mining, log file analysis, machine learning, financial analysis, scientific simulation, and bioinformatics research.

• AOL

 We use hadoop for variety of things ranging from ETL style processing and statistics generation to running advanced algorithms for doing behavioral We are one of six universities participating in IBM/Google's academic cloud computing initiative. Ongoing research and teaching efforts include projects in machine translation, language modeling, bioinformatics, email analysis, and image processing.

University of Nebraska Lincoln, Research Computing Facility

We currently run one medium-sized Hadoop cluster (200TB) to store and serve up physics data for the computing portion of the Compact Muon Solenoid (CMS) experiment. This requires a filesystem which can download data at multiple Gbps and process data at an even higher rate locally. Additionally, several of our students are involved in research projects on Hadoop.

- Veoh
 - We use a small Hadoop cluster to reduce usage data for internal metrics, for search indexing and for recommendation data.
- Visible Measures Corporation uses Hadoop as a component in our Scalable Data Pipeline, which ultimately powers VisibleSuite and other products. We use Hadoop to aggregate, store, and analyze data related to in-stream viewing behavior of Internet video audiences. Our current grid contains more than 128 CPU cores and in excess of 100 terabytes of storage, and we plan to grow that substantially during 2008.
- VK Solutions
 - We use a small Hadoop cluster in the scope of our general research activities at VK Labs to get a faster data access from web applications.
 - We also use Hadoop for filtering and indexing listing, processing log analysis, and for recommendation data.
- Wuelos baratos
 - We use a small Hadoop

- W Vuelos baratos
 - We use a small Hadoop

WorldLingo

- Hardware: 44 servers (each server has: 2 dual core CPUs, 2TB storage, 8GB RAM)
- Each server runs Xen with one Hadoop/HBase instance and another instance with web or application servers, giving us 88 usable virtual machines.
- We run two separate Hadoop/HBase clusters with 22 nodes each.
- Hadoop is primarily used to run HBase and Map/Reduce jobs scanning over the HBase tables to perform specific tasks.
- HBase is used as a scalable and fast storage back end for millions of documents.
- Currently we store 12million documents with a target of 450million in the near future.

Symbol

- More than 100,000 CPUs in >25,000 computers running Hadoop
- Our biggest cluster: 4000 nodes (2*4cpu boxes w 4*1TB disk & 16GB RAM)
 - Used to support research for Ad Systems and Web Search
 - Also used to do scaling tests to support development of Hadoop on larger clusters
- Our Blog Learn more about how we use Hadoop.
- >40% of Hadoop Jobs within Yahoo are Pig jobs.

Szvents

- 10 node cluster (Dual-Core AMD Opteron 2210, 4GB RAM, 1TB/node storage)
- Run Naive Bayes classifiers in parallel over crawl data to discover event information



Bing

Why Hadoop? MapReduce

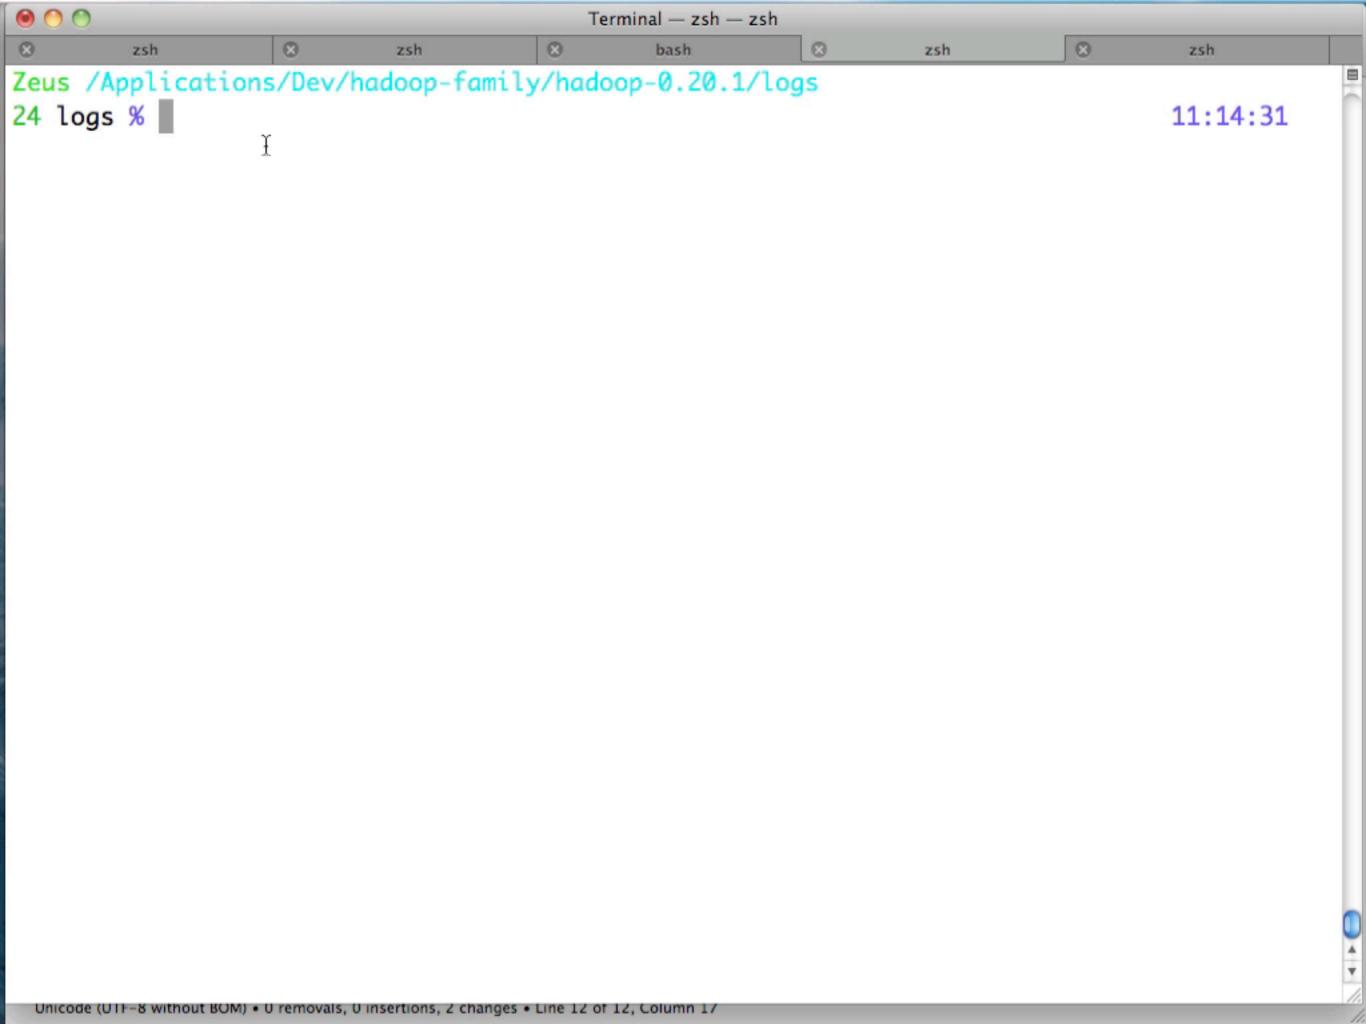
VIRTUAL MACHINE

VM SOURCES

- **Yahoo**
 - True to the OSS distribution
- **Cloudera**
 - Desktop tools
- Both VMWare based

STARTING UP

Tailing the logs



FILESYSTEM

The Google File System

Sanjay Ghemawat, Howard Gobioff, and Shun-Tak Leung Google

ABSTRACT

We have designed and implemented the Google File System, a scalable distributed file system for large distributed data-intensive applications. It provides fault tolerance while running on inexpensive commodity hardware, and it delivers high aggregate performance to a large number of clients.

While sharing many of the same goals as previous distributed file systems, our design has been driven by observations of our and selection of the same goals as previous disronment, beyond that reflect departure as earlier assumption of the same goals as previous disdeparture as earlier assumption is and cadically dissign points.

The em has successfull reds. loyed within Googl for m on and processing sersearch and develop The largest clust ds of of storage acr sands of a the hines, and acurrently a undreds

a this paper,
designed to support distributed applications, discuss many
aspects of our design, and report measurements from both
micro-benchmarks and real world use.

Categories and Subject Descriptors

D [4]: 3—Distributed file systems

General Terms

Design, reliability, performance, measurement

Keywords

Fault tolerance, scalability, data storage, clustered storage

*The authors can be reached at the following addresses: {sanjay,hgobioff,shuntak}@google.com.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires price specific permission and/or a fee.

SOSP'03, October 19-22, 2003, Bolton Landing, New York, USA. Copyright 2003 ACM 1-58113-757-5/03/0010 _-\$5.00.

1. INTRODUCTION

We have designed and implemented the Google File tem (GFS) to meet the rapidly growing demands of Good data processing needs. GFS shares many of the same g as previous distributed file systems such as performascalability, reliability, and availability. However, its dehas been driven by key observations of our application w and technological enviont and that reflect a p zie est design assy nined to tions. and ex adically di pints in design

First. are the norr than exception consists of s or e thousand a built from modity p naive co number client ma The q. of the come nents vi uarantee 1... ot functional any gir and some will no om their We have seen proble applicatio sting system b e failures .exs, memory, co plies. Therefore, cy ver sage--aitoring, error a, fault tolerance, and au recovery must be in to the system.

Second, files are traditional stands application objects such as a regularly working with fast growing billions of objects, it is unwietay to manage billions of approximately KB-sized files even when the file system could support it. As a result, design assumptions and parameters such as I/O operation and block sizes have to be revisited.

Third, most files are mutated by appending new data rather than overwriting existing data. Random writes within a file are practically non-existent. Once written, the files are only read, and often only sequentially. A variety of data share these characteristics. Some may constitute large repositories that data analysis programs scan through. Some may be data streams continuously generated by running applications. Some may be archival data. Some may be intermediate results produced on one machine and processed on another, whether simultaneously or later in time. Given this access pattern on huge files, appending becomes the focus of performance optimization and atomicity guarantees, while caching data blocks in the client loses its appeal.

Fourth, co-designing the applications and the file system API benefits the overall system by increasing our flexibility. "scalable distributed file System for large distributed data-intensive applications. It provides fault tolerance while running on inexpensive commodity hardware"

HDFS BASICS

- Open Source implementation of Google BigTable
- Replicated data store
- Stored in 64MB blocks

HDFS

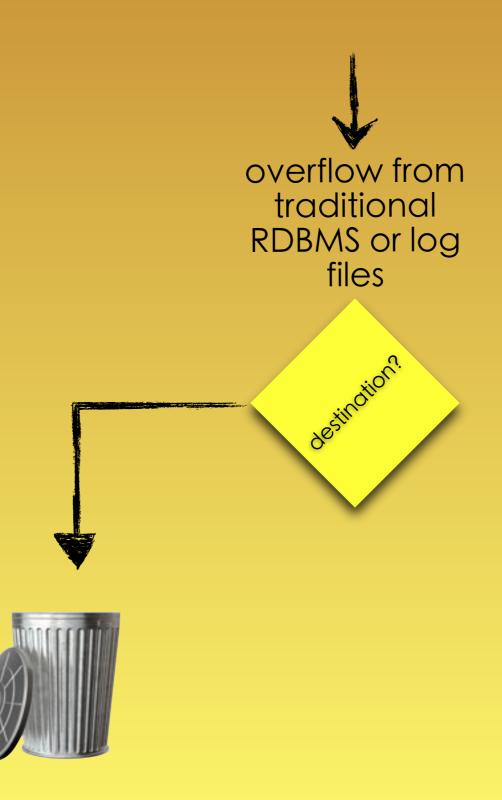
- Rack location aware
- Configurable redundancy factor
- Self-healing
- Looks almost like *NIX filesystem

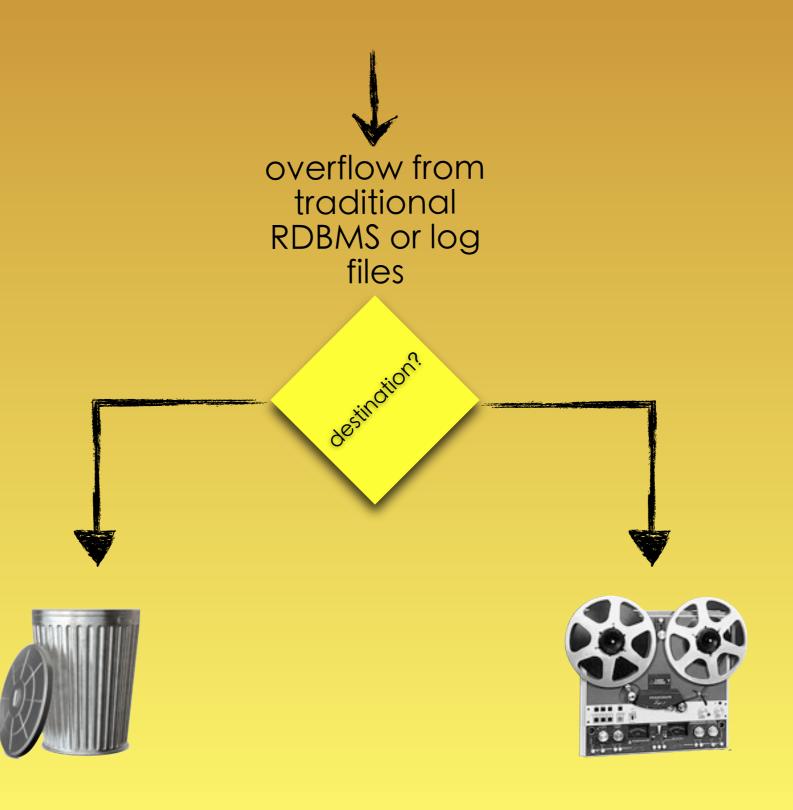
WHY HDFS!

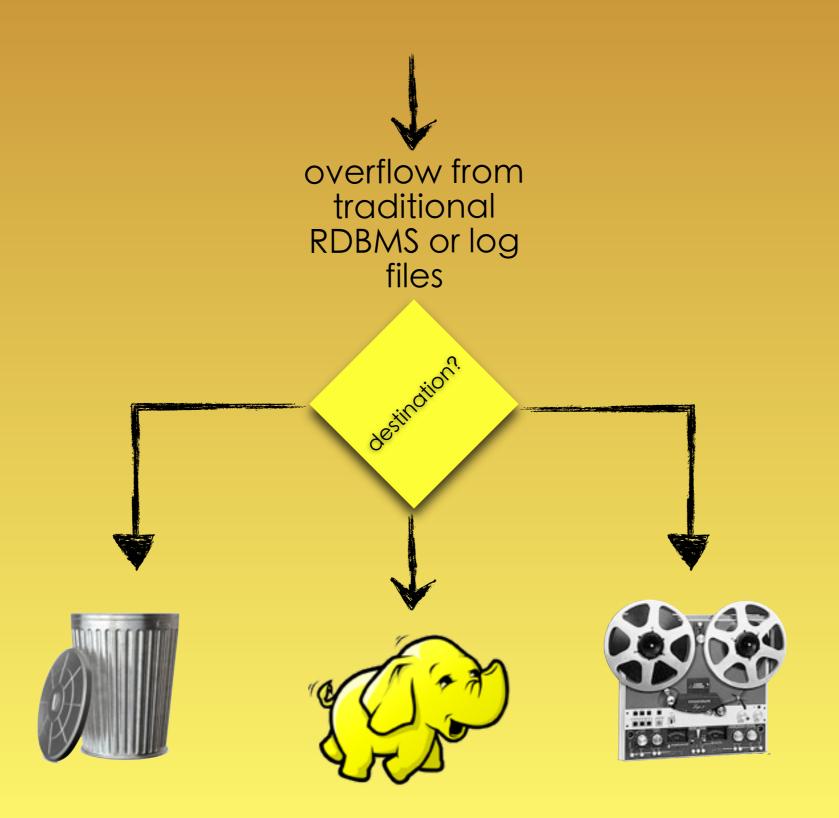
- Random reads
- Parallel reads
- Redundancy



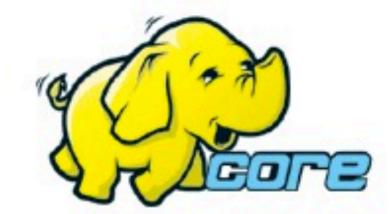












Project

Wiki

Hadoop 0.20 Documentation

- Getting Started
- Programming Guides
- **▼ HDFS**
 - User Guide
 - Architecture
 - File System Shell Guide
 - Permissions Guide
 - Quotas Guide
 - Synthetic Load Generator Guide
 - C API libhdfs
- **▶** HOD
- Miscellaneous

HDFS File System Shell Guide

Overview

- cat
- chgrp
- chmod
- chown
- copyFromLocal
- copyToLocal
- count
- cp
- du
- dus dus
- expunge
- get
- getmerge
- Is
- Isr
- mkdir

TEST HDFS

- Upload a file
- List directories



HDFS UPLOAD

- Show contents of file in HDFS
- Show vocabulary of HDFS



HDFS CHALLENGES

- Writes are re-writes today
 - Append is planned
- Block size, alignment
- Small file inefficiencies
- NameNode SPOF

Why Hadoop? HDFS Mankeduce

Mapreduce

MapReduceis functional programming distributed processing platform



HADOOP'S PREMISE

Geography-aware distribution of brute-force processing

MAPREDUCE the algorithm

MapReduce: Simplified Data Processing on Large Clusters

Jeffrey Dean and Sanjay Ghemawat

jeff@google.com, sanjay@google.com

Google, Inc.

Abstract

MapRe ramming model and an ated ssing and gene function the es a e pair to general nterme value nd a reduce function nediate rges associated with the Many me orld tasks are express shown aper.

ams written in this f style matilelized and exect arge clu mm whines. The r stem takes a
deta. Jata, scheduling a
gram's c. and managing the required inter-machine
communication. This allows programmers without any
experience with parallel and distributed systems to easily utilize the resources of a large distributed system.

Our implementation of MapReduce runs on a large cluster of commodity machines and is highly scalable: a typical MapReduce computation processes many terabytes of data on thousands of machines. Programmers find the system easy to use: hundreds of MapReduce programs have been implemented and upwards of one thousand MapReduce jobs are executed on Google's clusters every day.

1 Introduction

Over the past five years, the authors and many others at Google have implemented hundreds of special-purpose computations that process large amounts of raw data, such as crawled documents, web request logs, etc., to compute various kinds of derived data, such as inverted indices, various representations of the graph structure of web documents, summaries of the number of pages crawled per host, the set of most frequent queries in a

given day, etc. Most such computations are conceptutraightforward. However he computation ousands les in or sh in a reasc ount a The issues of parallelize utati bute the data, ille failures c e original sim tation wit mplex code t ith these issue

As a re this co a new abstraction ows us to ea, computations w ing to perform but I essy deization, fault-tolerance tion alancing in a 131 pace by the map and and many other fun guages. We realize most of our compu volved applying a ma eration to each log ord" in our input in or compute a set of int key/value pairs, applying a reduce ope the same key, in order to propriately. Our use of a tune. specified map and reduce operations allows us to parallelize large computations easily and to use re-execution as the primary mechanism for fault tolerance.

The major contributions of this work are a simple and powerful interface that enables automatic parallelization and distribution of large-scale computations, combined with an implementation of this interface that achieves high performance on large clusters of commodity PCs.

Section 2 describes the basic programming model and gives several examples. Section 3 describes an implementation of the MapReduce interface tailored towards our cluster-based computing environment. Section 4 describes several refinements of the programming model that we have found useful. Section 5 has performance measurements of our implementation for a variety of tasks. Section 6 explores the use of MapReduce within Google including our experiences in using it as the basis

To appear in OSDI 2004

"A programming model and implementation for processing and generating large data sets"

VERIFY SETUP

- Launch "cloudera-training" VMWare instance
- Open a terminal window
- Verify hadoop





THE PROCESS

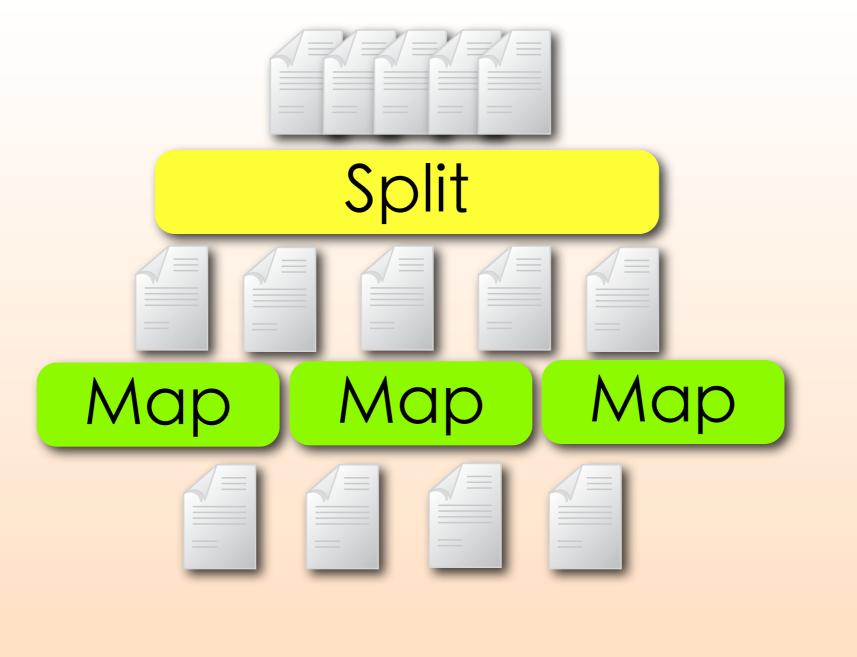
- Map(k1,v1) -> list(k2,v2)
 - Every item is parallel candidate for Map
- Shuffle (group) pairs from all lists by key
- Reduce(k2, list (v2)) -> list(v3)
 - Reduce in parallel on each group of keys



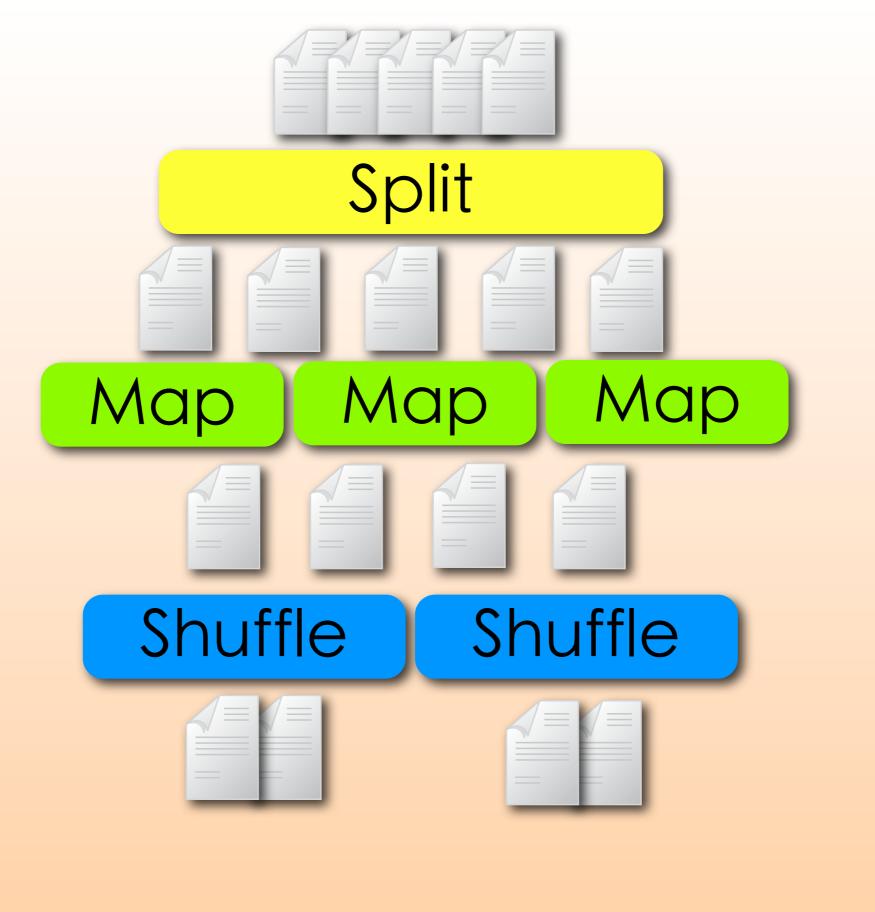




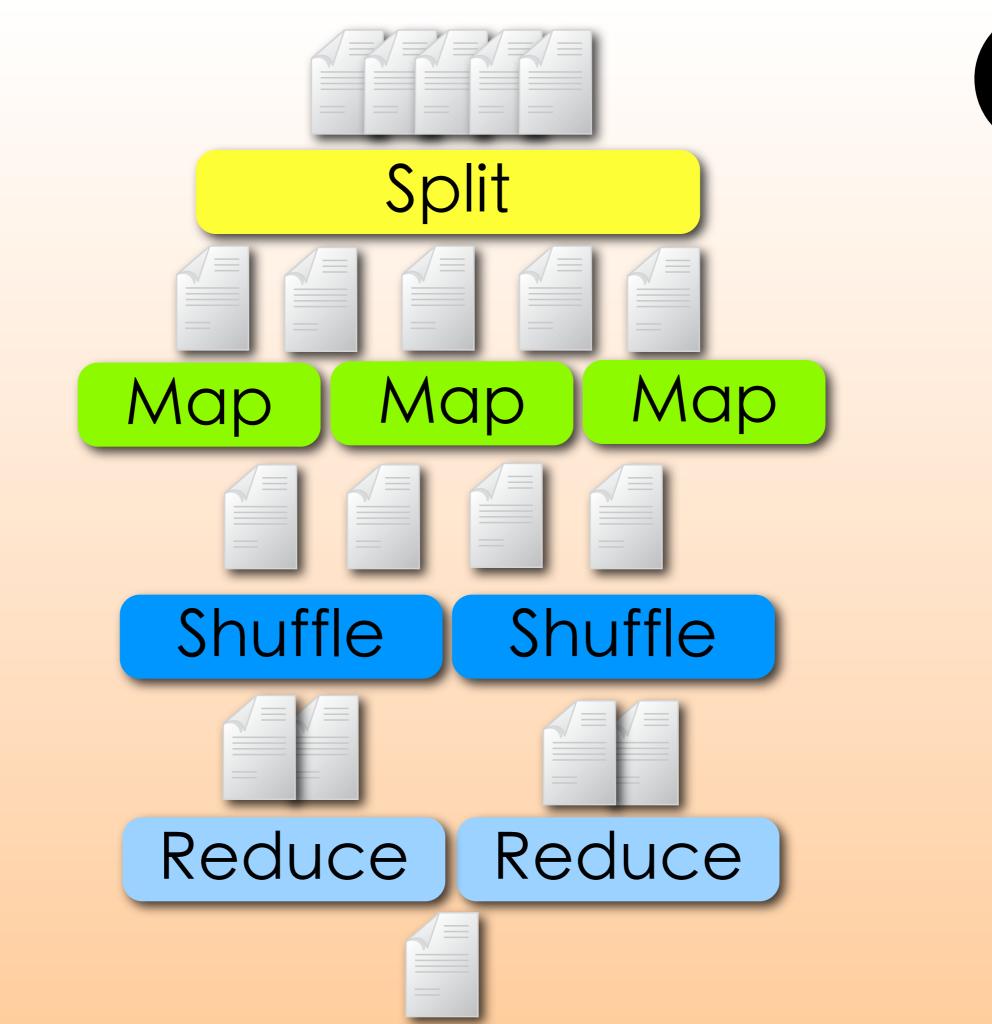












MAPREDUCE

- Run the grep Shakespeare job
- View the status consoles



MAPREDUCE

a word counting conceptual example

THE GOAL

Provide the occurrence count of each distinct word across all documents

RAW DATA



a folder of documents

mydoc1.txt

At four years old I acted out

mydoc2.txt

Glad I am not four years old

mydoc3.txt

Yet I still act like someone four



break documents into words

four

years

old

acted

out

glad

am

not

four

years

old

yet

still

act

like

someone

four

SHUFFLE

physically group (relocate) by key

glad at yet not four four still act four old like am old years someone acted out years

REDUCE

count word occurrences

$$at = 1$$

$$four = 3$$

$$old = 2$$

$$glad = 1$$

$$1 = 3$$

$$am = 1$$

$$years = 2$$

$$yet = 1$$

$$not = 1$$

$$still = 1$$

$$act = 1$$

$$like = 1$$

$$out = 1$$

REDUCE AGAIN

sort occurrences alphabetically

$$act = 1$$

$$acted = 1$$

$$at = 1$$

$$am = 1$$

$$four = 3$$

$$glad = 1$$

$$1 = 3$$

$$like = 1$$

$$not = 1$$

$$out = 1$$

$$old = 2$$

$$still = 1$$

$$yet = 1$$

GREP-JAVA

```
package org.apache.hadoop.examples;
import java.util.Random;
import org.apache.hadoop.conf.Configuration;
import org.apache.hadoop.conf.Configured;
import org.apache.hadoop.fs.FileSystem;
import org.apache.hadoop.fs.Path;
import org.apache.hadoop.io.LongWritable;
import org.apache.hadoop.io.Text;
import org.apache.hadoop.mapred.*;
import org.apache.hadoop.mapred.lib.*;
import org.apache.hadoop.util.Tool;
import org.apache.hadoop.util.ToolRunner;
/* Extracts matching regexs from input files and counts them. */
public class Grep extends Configured implements Tool {
 private Grep() {}
                                                // singleton
 public int run(String[] args) throws Exception {
   if (args.length < 3) {</pre>
     System.out.println("Grep <inDir> <outDir> <regex>
[<group>]");
```

```
/* Extracts matching regexs from input files and counts them. */
public class Grep extends Configured implements Tool {
  private Grep() {}
                                                   // singleton
  public int run(String[] args) throws Exception {
    if (args.length < 3) {</pre>
      System.out.println("Grep <inDir> <outDir> <regex>
[<group>]");
      ToolRunner.printGenericCommandUsage(System.out);
      return -1;
    Path tempDir =
      new Path("grep-temp-"+
          Integer.toString(new Random().nextInt
(Integer.MAX_VALUE)));
    JobConf grepJob = new JobConf(getConf(), Grep.class);
    try {
      grepJob.setJobName("grep-search");
      FileInputFormat.setInputPaths(grepJob, args[0]);
```

```
grepJob.setJobName("grep-search");
FileInputFormat.setInputPaths(grepJob, args[0]);
grepJob.setMapperClass(RegexMapper.class);
grepJob.set("mapred.mapper.regex", args[2]);
if (args.length == 4)
  grepJob.set("mapred.mapper.regex.group", args[3]);
grepJob.setCombinerClass(LongSumReducer.class);
grepJob.setReducerClass(LongSumReducer.class);
FileOutputFormat.setOutputPath(grepJob, tempDir);
grepJob.setOutputFormat(SequenceFileOutputFormat.class);
grepJob.setOutputKeyClass(Text.class);
grepJob.setOutputValueClass(LongWritable.class);
JobClient.runJob(grepJob);
JobConf sortJob = new JobConf(Grep.class);
sortJob.setJobName("grep-sort");
FileInputFormat.setInputPaths(sortJob, tempDir);
sortJob.setInputFormat(SequenceFileInputFormat.class);
```

```
FileOutputFormat.setOutputPath(grepJob, tempDir);
      grepJob.setOutputFormat(SequenceFileOutputFormat.class);
      grepJob.setOutputKeyClass(Text.class);
      grepJob.setOutputValueClass(LongWritable.class);
      JobClient.runJob(grepJob);
      JobConf sortJob = new JobConf(Grep.class);
      sortJob.setJobName("grep-sort");
      FileInputFormat.setInputPaths(sortJob, tempDir);
      sortJob.setInputFormat(SequenceFileInputFormat.class);
      sortJob.setMapperClass(InverseMapper.class);
      sortJob.setNumReduceTasks(1);
                                                     // write a
single file
      FileOutputFormat.setOutputPath(sortJob, new Path(args)
[1]));
      sortJob.setOutputKeyComparatorClass
                                                     // sort by
decreasing freq
      (LongWritable.DecreasingComparator.class);
      JobClient.runJob(sortJob);
    finally 5
```

```
JobClient.runJob(grepJob);
      JobConf sortJob = new JobConf(Grep.class);
      sortJob.setJobName("grep-sort");
      FileInputFormat.setInputPaths(sortJob, tempDir);
      sortJob.setInputFormat(SequenceFileInputFormat.class);
      sortJob.setMapperClass(InverseMapper.class);
                                                     // write a
      sortJob.setNumReduceTasks(1);
single file
      FileOutputFormat.setOutputPath(sortJob, new Path(args
[1]));
      sortJob.setOutputKeyComparatorClass
                                                     // sort by
decreasing freq
      (LongWritable.DecreasingComparator.class);
      JobClient.runJob(sortJob);
    finally {
      FileSystem.get(grepJob).delete(tempDir, true);
    return 0;
```

REGEXMAPPER.JAVA

```
/**
 * Licensed to the Apache Software Foundation (ASF) under one
 * or more contributor license agreements. See the NOTICE file
 * distributed with this work for additional information
 * regarding copyright ownership. The ASF licenses this file
 * to you under the Apache License, Version 2.0 (the
 * "License"); you may not use this file except in compliance
 * with the License. You may obtain a copy of the License at
 *
       http://www.apache.org/licenses/LICENSE-2.0
 *
 *
 * Unless required by applicable law or agreed to in writing, software
 * distributed under the License is distributed on an "AS IS" BASIS,
 * WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or
implied.
 * See the License for the specific language governing permissions and
 * limitations under the License.
 */
package org.apache.hadoop.mapred.lib;
import java.io.IOException;
import java.util.regex.Matcher;
import java.util.regex.Pattern;
```

```
/** A {@link Mapper} that extracts text matching a regular expression.
*/
public class RegexMapper<K> extends MapReduceBase
    implements Mapper<K, Text, Text, LongWritable> {
 private Pattern pattern;
 private int group;
  public void configure(JobConf job) {
    pattern = Pattern.compile(job.get("mapred.mapper.regex"));
    group = job.getInt("mapred.mapper.regex.group", 0);
  public void map(K key, Text value,
                  OutputCollector<Text, LongWritable> output,
                  Reporter reporter)
    throws IOException {
    String text = value.toString();
    Matcher matcher = pattern.matcher(text);
    while (matcher.find()) {
      output.collect(new Text(matcher.group(group)), new LongWritable
(1));
```

MAPREDUCE

- View the Shakespeare output "part-XXXX" file
- Why "part-XXXXX" naming?



HAVE CODE, WILL TRAVEL

- Code travels to the data
- Opposite of traditional systems

STREAMING

UNIX VIA MAPREDUCE

- Any UNIX command
- Any shell-invokable script
 - Perl
 - Python
 - **Ruby**

UNIX VIA MAPREDUCE

- Line at a time
- Tab separator

```
hadoop jar contrib/streaming/
hadoop-0.20.1-streaming.jar
-input people.csv
-output outputstuff
-mapper 'cut -f 2 -d ,'
-reducer 'uniq'
```

STREAMING

Run a streaming job



The Grad DSLS Tools

The Grid

MOTIVATIONS

1 Gigabyte

1 Terabyte

1 Petabyte

16 Petabytes

Near-Linear Hardware Scalability

APPLICATIONS

- Protein folding pharmaceutical research
- Search Engine Indexing walking billions of web pages
- Product Recommendations
 based on other customer purchases
- Sorting terabytes to petabyes in size
- Classification government intelligence

CONTEXTUAL ADS

CONTEXTUAL ADS

Shopper

looking at

Product \

is told that bblo of other customers who bought X

Customer B

0/50 bought x

Customer

did not buy

Product

Customer

SELECT reccProd.name, reccProd.id FROM products reccProd WHERE purchases.customerld =

SELECT reccProd.name, reccProd.id FROM products reccProd WHERE purchases.customerId =

(SELECT customerId FROM customers WHERE purchases.productId = thisProd)

LIMIT 5

GRID BENEFITS

SCALABLE

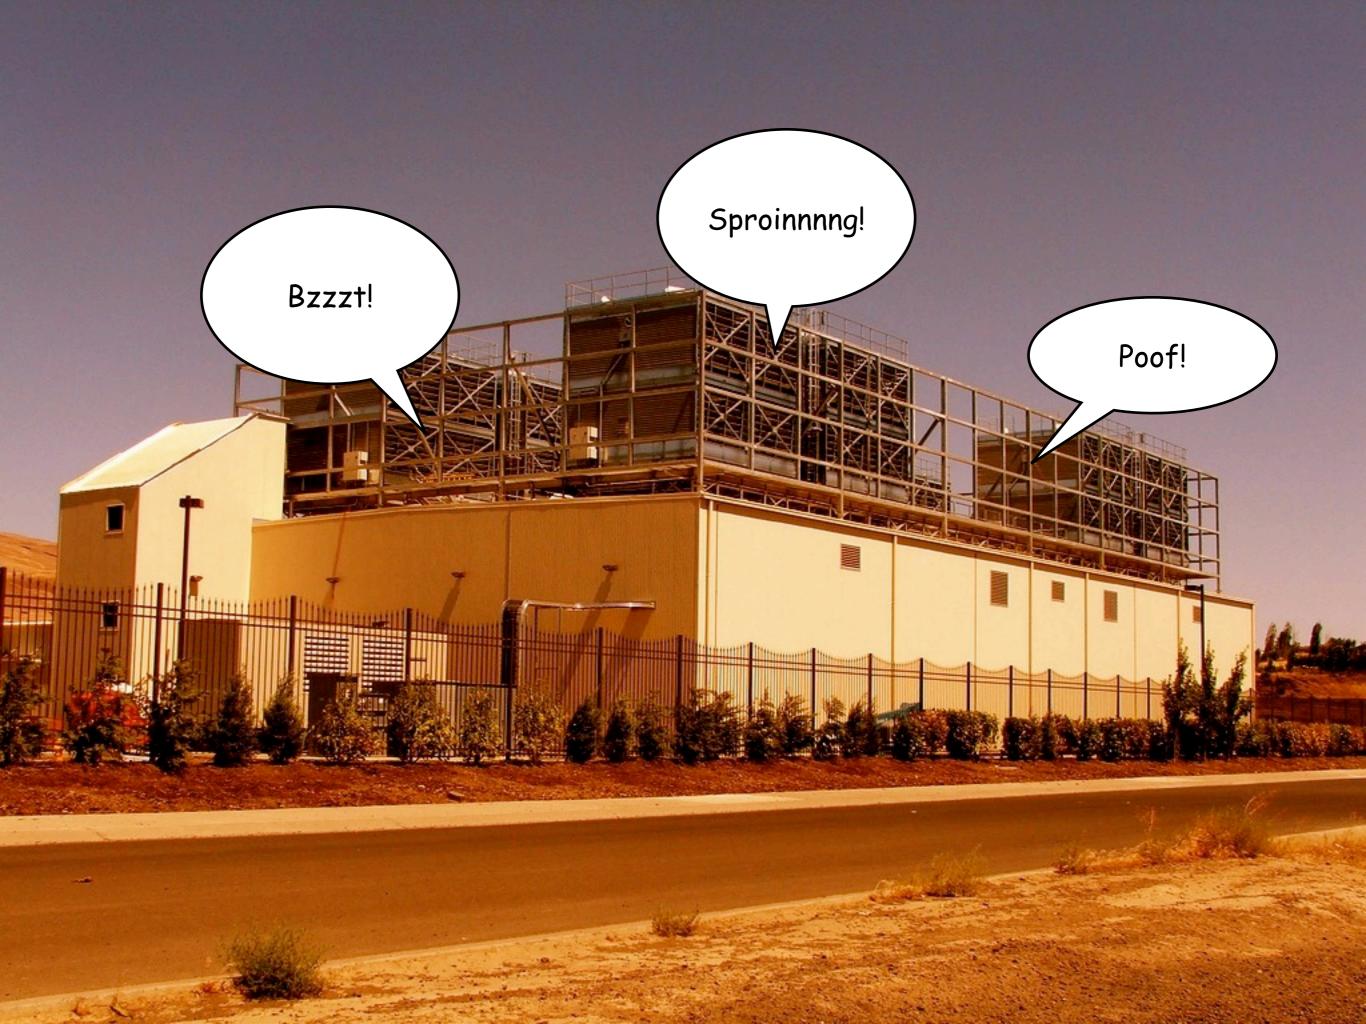
- Data storage is pipelined
- Code travels to data
- Near linear hardware scalability

OPTIMIZED

- Preemptive execution
- Maximizes use of faster hardware
- New jobs trump P.E.

FAULT TOLERANT

- Configurable data redundancy
- Minimizes hardware failure impact
- Automatic job retries
- Self healing filesystem



SERVER FUNERALS

- No pagers go off when machines die
- Report of dead machines once a week
 - Clean out the carcasses

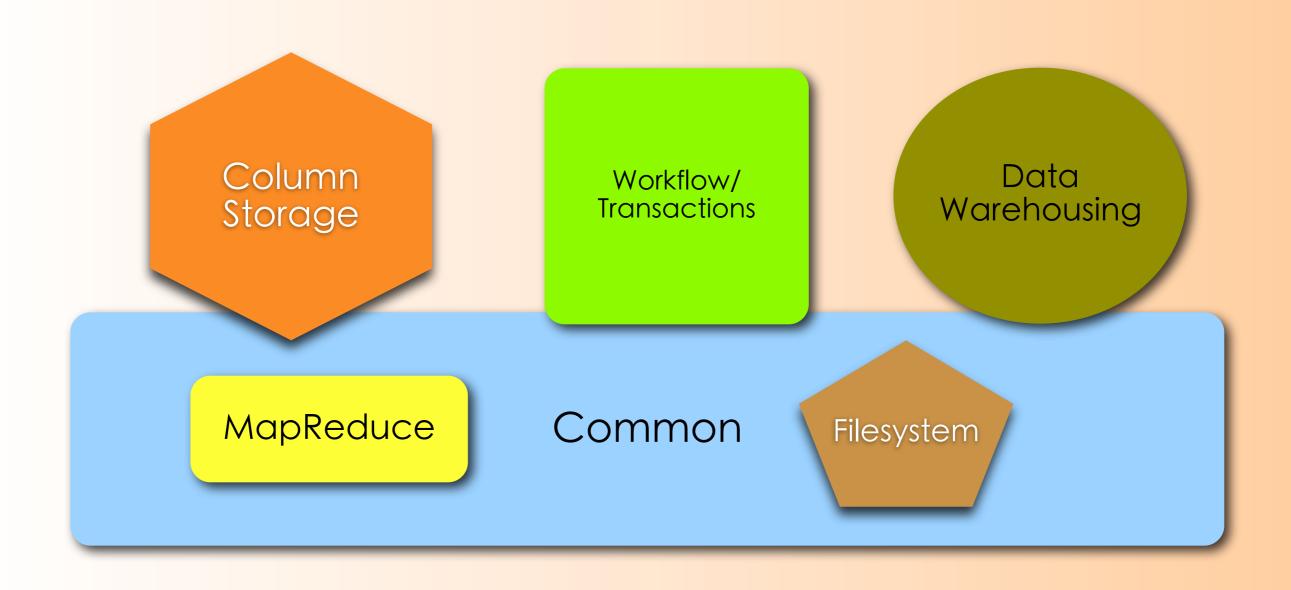


- Tata redundancy
- Node death
- **Retries**
- Data geography
- Parallelism
- Scalability

COMPONENTS

HADOOP COMPONENTS

HADOOP COMPONENTS



HADOOP COMPONENTS

Tool	Purpose
Common	MapReduce
HDFS	Filesystem
Pig	Analyst MapReduce language
HBase	Column-oriented data storage
Hive	SQL-like language for HBase
ZooKeeper	Workflow & distributed transactions
Chukwa	Log file processing

The Grid Tools

SYNC, ASYNC

- Pig is asynchronous
- Hive is asynchronous
- HBase is near realtime
- RDBMS is realtime



PIG BASICS

- Yahoo-authored add-on
- High-level language for authoring data analysis programs
- Console

PIG SAMPLE

```
Person = LOAD 'people.csv' using PigStorage(',');
Names = FOREACH Person GENERATE $2 AS name;
OrderedNames = ORDER Names BY name ASC;
GroupedNames = GROUP OrderedNames BY name;
NameCount = FOREACH GroupedNames
GENERATE group, COUNT(OrderedNames);
store NameCount into 'names.out';
```

[pig-example (master)/]>



```
[pig-example (master)/] > pig -x local get-names-from-people.pi
2010-04-20 05:58:16,851 [main] INFO org.apache.pig.Main - Log
ging error messages to: /Users/mccm06/Documents/Teach/Courses/
Hadoop-Intro/hadoop-examples_git/pig-example/pig_1271764696851
.log
2010-04-20 05:58:52,460 [main] INFO org.apache.pig.backend.lo
cal.executionengine.LocalPigLauncher - Successfully stored res
ult in: "file:/Users/mccm06/Documents/Teach/Courses/Hadoop-Int
ro/hadoop-examples_git/pig-example/names.out"
2010-04-20 05:58:52,460 [main] INFO org.apache.pig.backend.lo
cal.executionengine.LocalPigLauncher - Records written : 5158
2010-04-20 05:58:52,462 [main] INFO org.apache.pig.backend.lo
cal.executionengine.LocalPigLauncher - Bytes written : 0
2010-04-20 05:58:52,462 [main] INFO
                                     org.apache.pig.backend.lo
cal.executionengine.LocalPigLauncher
                                     - 100% complete!
2010-04-20 05:58:52,462 [main] INFO
                                     org.apache.pig.backend.lo
cal.executionengine.LocalPigLauncher - Success!!
[pig-example (master) / ] >
```



HIVE BASICS

- Authored by facebook.
- SQL interface to HBase
- Hive is low-level
- Hive-specific metadata
- Data warehousing

```
SELECT * FROM shakespeare
WHERE freq > 100
SORT BY freq ASC
LIMIT 10;
```

SYNC, ASYNC

- RDBMS SQL is realtime
- Hive is primarily asynchronous



Hbase/PoweredBy FrontPage RecentChanges FindPage HelpContents Hbase/PoweredBy

More Actions:

Immutable Page Info Attachments

Adobe - We currently have about 30 nodes running HDFS, Hadoop and HBase in clusters ranging from 5 to 14 nodes on both production and development. We plan a deployment on an 80 nodes cluster. We are using HBase in several areas from social services to structured data and processing for internal use. We constantly write data to HBase and run mapreduce jobs to process then store it back to HBase or external systems. Our production cluster has been running since Oct 2008.

Drawn to Scale Consulting consults on HBase, Hadoop, Distributed Search, and Scalable architectures.

*

- Filmweb is a film web portal with a large dataset of films, persons and movie-related entities. We have just started a small cluster of 3 HBase nodes to handle our web cache persistency layer. We plan to increase the cluster size, and also to start migrating some of the data from our databases which have some demanding scalability requirements.
- Flurry provides mobile application analytics. We use HBase and Hadoop for all of our analytics processing, and serve all of our live requests directly out of HBase on our 16-node production cluster with billions of rows over several tables.
- GumGum is an in-image ad network. We use HBase 0.20 on a 4-node Amazon EC2 Large Instance (m1.large) cluster for both real-time data and analytics. Our production cluster has been running since June 2010.
- Skalooga is a discovery service for image galleries. We use Hadoop, Hbase, Chukwa and Pig on a 20-node cluster for our crawling, analysis and events processing.
- Lily is an open source content repository backed by HBase and SOLR from Outerthought scalable content applications.
- Mahalo, "...the world's first human-powered search engine". All the markup that powers the wiki is stored in HBase. It's been in use for a few months now. MediaWiki the same software that power Wikipedia has version/revision control. Mahalo's in-house editors produce a lot of revisions per day, which was not working well in a RDBMS. An hbase-based solution for this was built and tested, and the data migrated out of MySQL and into HBase. Right now it's at something like 6 million items in HBase. The upload tool runs every hour from a shell script to back up that data, and on 6 nodes takes about 5-10 minutes to run and does not slow down production at all.
- Meetup is on a mission to help the world's people self-organize into local groups. We use Hadoop and HBase to power a site-wide, real-time activity feed system for all of our members and groups. Group activity is written directly to HBase, and indexed per member, with the member's custom feed served directly from HBase for incoming requests. We're running HBase 0.20.0 on a 11 node cluster.
- Ning uses HBase to store and serve the results of processing user events and log files, which allows us to provide near-real time analytics and reporting. We use a small cluster of commodity machines with 4 cores and 16GB of RAM per machine to handle all our analytics and reporting needs.

- Stumbleupon and Su.pr use HBase as a real time data storage and analytics platform. Serving directly out of HBase, various site features and statistics are kept up to date in a real time fashion. We also use HBase a map-reduce data source to overcome traditional query speed limits in MySQL.
- SubRecord Project is an Open Source project that is using HBase as a repository of records (persisted map-like data) for the aspects it provides like logging, tracing or metrics. HBase and Lucene index both constitute a repo/storage for this platform.
- Shopping Engine at Tokenizer is a web crawler; it uses HBase to store URLs and Outlinks (AnchorText + LinkedURL): more than a billion. It was initially designed as Nutch-Hadoop extension, then (due to very specific 'shopping' scenario) moved to SOLR + MySQL(InnoDB) (ten thousands queries per second), and now to HBase. HBase is significantly faster due to: no need for huge transaction logs, column-oriented design exactly matches 'lazy' business logic, data compression, MapReduce support. Number of mutable 'indexes' (term from RDBMS) significantly reduced due to the fact that each 'row::column' structure is physically sorted by 'row'. MySQL InnoDB engine is best DB choice for highly-concurrent updates. However, necessity to flash a block of data to harddrive even if we changed only few bytes is obvious bottleneck. HBase greatly helps: not-so-popular in modern DBMS 'delete-insert', 'mutable primary key', and 'natural primary key' patterns become a big advantage with HBase.
- Trend Micro uses HBase as a foundation for cloud scale storage for a variety of applications. We have been developing with HBase since version 0.1 and production since version 0.20.0.
- Twitter runs HBase across its entire Hadoop cluster. HBase provides a distributed, read/write backup of all mysql tables in Twitter's production backend, allowing engineers to run MapReduce jobs over the data while maintaining the ability to apply periodic row updates (something that is more difficult to do with vanilla HDFS). A number of applications including people search rely on HBase internally for data generation. Additionally, the operations team uses HBase as a timeseries database for cluster-wide monitoring/performance data.
- Veoh Networks uses HBase to store and process visitor(human) and entity(non-human) profiles which are used for behavioral targeting, demographic detection, and personalization services. Our site reads this data in real-time (heavily cached) and submits updates via various batch map/reduce jobs. With 25 million unique visitors a month storing this data in a traditional RDBMS is not an option. We currently have a 24 node Hadoop/HBase cluster and our profiling system is sharing this cluster with our other Hadoop data pipeline processes.
- SideoSurf "The video search engine that has taught computers to see". We're using Hbase to persist various large graphs of data and other statistics. Hbase was a real win for us because it let us store substantially larger datasets without the need for manually partitioning the data and it's column-oriented nature allowed us to create schemas that were substantially more efficient for storing and retrieving data.
- Visible Technologies We use Hadoop, HBase, Katta, and more to collect, parse, store, and search hundreds of millions of Social Media content. We get incredibly fast throughput and very low latency on commodity hardware. HBase enables our business to exist.
- WorldLingo The WorldLingo Multilingual Archive. We use HBase to store millions of documents that we scan using Map/Reduce jobs to machine translate them into all or selected target languages from our set of available machine translation languages. We currently store 12 million documents but plan to eventually reach the 450 million mark. HBase allows us to scale out as we need to grow our storage capacities. Combined with Hadoop to keep the data replicated and therefore fail-safe we have the backbone our service can rely on now and in the future. WorldLingo is using HBase since December 2007 and is along with a few others one of the longest running HBase installation. Currently we are running the latest HBase 0.20 and serving directly from it: MultilingualArchive.

- Stumbleupon and Su.pr use HBase as a real time data storage and analytics platform. Serving directly out of HBase, various site features and statistics are kept up to date in a real time fashion. We also use HBase a map-reduce data source to overcome traditional query speed limits in MySQL.
- SubRecord Project is an Open Source project that is using HBase as a repository of records (persisted map-like data) for the aspects it provides like logging, tracing or metrics. HBase and Lucene index both constitute a repo/storage for this platform.
- Shopping Engine at Tokenizer is a web crawler; it uses HBase to store URLs and Outlinks (AnchorText + LinkedURL): more than a billion. It was initially designed as Nutch-Hadoop extension, then (due to very specific 'shopping' scenario) moved to SOLR + MySQL(InnoDB) (ten thousands queries per second), and now to HBase. HBase is significantly faster due to: no need for huge transaction logs, column-oriented design exactly matches 'lazy' business logic, data compression, MapReduce support. Number of mutable 'indexes' (term from RDBMS) significantly reduced due to the fact that each 'row::column' structure is physically sorted by 'row'. MySQL InnoDB engine is best DB choice for highly-concurrent updates. However, necessity to flash a block of data to harddrive even if we changed only few bytes is obvious bottleneck. HBase greatly helps: not-so-popular in modern DBMS 'delete-insert', 'mutable primary key', and 'natural primary key' patterns become a big advantage with HBase.
- Trend Micro uses HBase as a foundation for cloud scale storage for a variety of applications. We have been developing with HBase since version 0.1 and production since version 0.20.0.
- Twitter runs HBase across its entire Hadoop cluster. HBase provides a distributed, read/write backup of all mysql tables in Twitter's production backend, allowing engineers to run MapReduce jobs over the data while maintaining the ability to apply periodic row updates (something that is more difficult to do with vanilla HDFS). A number of applications including people search rely on HBase internally for data generation. Additionally, the operations team uses HBase as a timeseries database for cluster-wide monitoring/performance data.
- Veoh Networks uses HBase to store and process visitor(human) and entity(non-human) profiles which are used for behavioral targeting, demographic detection, and personalization services. Our site reads this data in real-time (heavily cached) and submits updates via various batch map/reduce jobs. With 25 million unique visitors a month storing this data in a traditional RDBMS is not an option. We currently have a 24 node Hadoop/HBase cluster and our profiling system is sharing this cluster with our other Hadoop data pipeline processes.
- VideoSurf "The video search engine that has taught computers to see". We're using Hbase to persist various large graphs of data and other statistics. Hbase was a real win for us because it let us store substantially larger datasets without the need for manually partitioning the data and it's column-oriented nature allowed us to create schemas that were substantially more efficient for storing and retrieving data.
- Visible Technologies We use Hadoop, HBase, Katta, and more to collect, parse, store, and search hundreds of millions of Social Media content. We get incredibly fast throughput and very low latency on commodity hardware. HBase enables our business to exist.
- WorldLingo The WorldLingo Multilingual Archive. We use HBase to store millions of documents that we scan using Map/Reduce jobs to machine translate them into all or selected target languages from our set of available machine translation languages. We currently store 12 million documents but plan to eventually reach the 450 million mark. HBase allows us to scale out as we need to grow our storage capacities. Combined with Hadoop to keep the data replicated and therefore fail-safe we have the backbone our service can rely on now and in the future. WorldLingo is using HBase since December 2007 and is along with a few others one of the longest running HBase installation. Currently we are running the latest HBase 0.20 and serving directly from it: MultilingualArchive.

HBASE BASICS

- Map-oriented storage
- Key value pairs
- Column families
- Stores to HDFS
- **Fast**
- Usable for synchronous responses

```
hbase>
help

create 'mylittletable', 'mylittlecolumnfamily'
describe 'mylittletable'

put 'mylittletable', 'r2', 'mylittlecolumnfamily', 'x'

get 'mylittletable', 'r2'
scan 'mylittletable'
```





What is Hadoop?

Why Hadoop? Downloads Learn Hadoop Get Support

Events Blog Twitter



Products & Services Customers

Resources

Downloads

Company Contact

Search

Developer Center

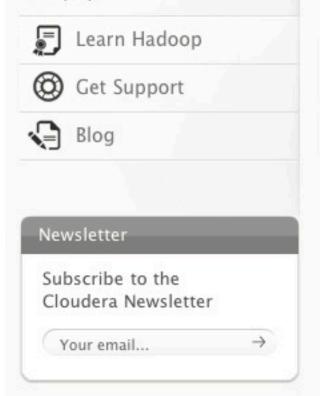


Overview



Oownloads Downloads

- Cloudera's Hadoop Distribution
- Get Virtual Machine
- CDH for Amazon
- Cloudera Desktop
- Sqoop





Sqoop

SQL-to-Hadoop database import tool

Sqoop is a tool designed to import data from relational databases into Hadoop. Sqoop uses JDBC to connect to a database. It examines each table's schema and automatically generates the necessary classes to import data into the Hadoop Distributed File System (HDFS). Sgoop then creates and launches a MapReduce job to read tables from the database via DBInputFormat, the JDBC-based InputFormat. Tables are read into a set of files in HDFS. Sgoop supports both SequenceFile and text-based target and includes performance enhancements for loading data from MySQL.

Getting Sqoop

Sgoop is an open-source program contributed to the Apache Hadoop project. The most recent release of Cloudera's Distribution for Hadoop contains Sgoop.

If you're already using our RPM packages, running the command yum update hadoop ale and all leaders are some the alake Delates are seen

Instructions

The instructions below describe how to get started using Sgoop to import your data into Hadoop:

- Example Usage
- Connecting to a Database Server
- Listing Available Tables

SQOOP

- A Cloudera utility
- Imports from RDBMS
- Outputs plaintext, SequenceFile, or Hive

sqoop --connect jdbc:mysql://database.example.com/

The Grid DSLS

MONITORING

WEB STATUS PANELS

**NameNode

http://localhost:50070/

JobTracker

http://localhost:50030/

EXTENDED FAMILY MEMBERS

Cascading

ANOTHER DIL

- Java-based
- Different vocabulary
- Abstraction from MapReduce
- Uses Hadoop
- Cascalog for Clojure

Cascading

Search

Search

Site

Home

About

Features

News

Powered By

Contact

Recent News

September 22, 2010

Memcached, Membase, and ElasticSearch Integration

September 16, 2010

O'Reilly Strata Conference -Last Call

August 26, 2010

Bixo Hackathon

August 25, 2010

O'Reilly Strata Conference

August 2, 2010

Cascading 1.1.2

Recent Articles

BigDataCamp 2010

Documentation

Diving In

Welcome

Cascading is a Query API and Query Planner used for defining and executing complex, scale-free, and fault tolerant deprocessing workflows on a Hadoop cluster.

Cascading is a thin Java library that sits on top of Hadoop's MapReduce layer. It is not a new text based query syntax (like Pig) or another complex system that must be installed on a cluster and maintained (like Hive). Though Cascading both complimentary to and is a valid alternative to either application.

Cascading is simply a query processing API that lets the developer quickly assemble complex distributed processes without having to "think" in MapReduce. And to efficiently schedule them based on their dependencies. Obviously sim data processing applications are supported as well, as complex applications tend to start simple.

Cascading is Open Source and dual licensed under the GPL and OEM/Commercial Licenses. OEM/Commercial Licenses and Developer Support can be obtained through Concurrent, Inc.

Cascading has a strong community of users and contributors, see our Cascading modules page for related projects a extensions.

Cascading, extensions, and related libraries are also hosted in the Conjars maven repository maintained by Concurrer Inc. The repository is open to the public.

Read more about <u>Cascadings features</u> or thumb through the <u>Cascading User Guide</u>.

Recent Events

Memcached, Membase, and ElasticSearch Integration

September 22, 2010 12:40 PM | Permalink

We have added a link to the Cascading. Memecached project on GitHub to the Modules and Extensions page.

This sub-project provides Memcached API integration allowing Cascading Flows to push data into various memcached API compliant applications like ElasticSearch and Membase.

thoughts from the red planet

Blog

About

Archives

Contact

Follow Nathan on



git GitHub

LinkedIn

■ Blog RSS



Search



Featured Posts

You should blog even if you have no readers

« New Cascalog features: outer joins, combiners, sorting, and more | Main | Fun with equality in Clojure »

Introducing Cascalog: a Clojure-based query language for Hadoop

WEDNESDAY, APRIL 14, 2010 W CASCADING, CASCALOG, CLOJURE, HADOOP

I'm very excited to be releasing Cascalog as open-source today. Cascalog is a Clojurebased query language for Hadoop inspired by Datalog.

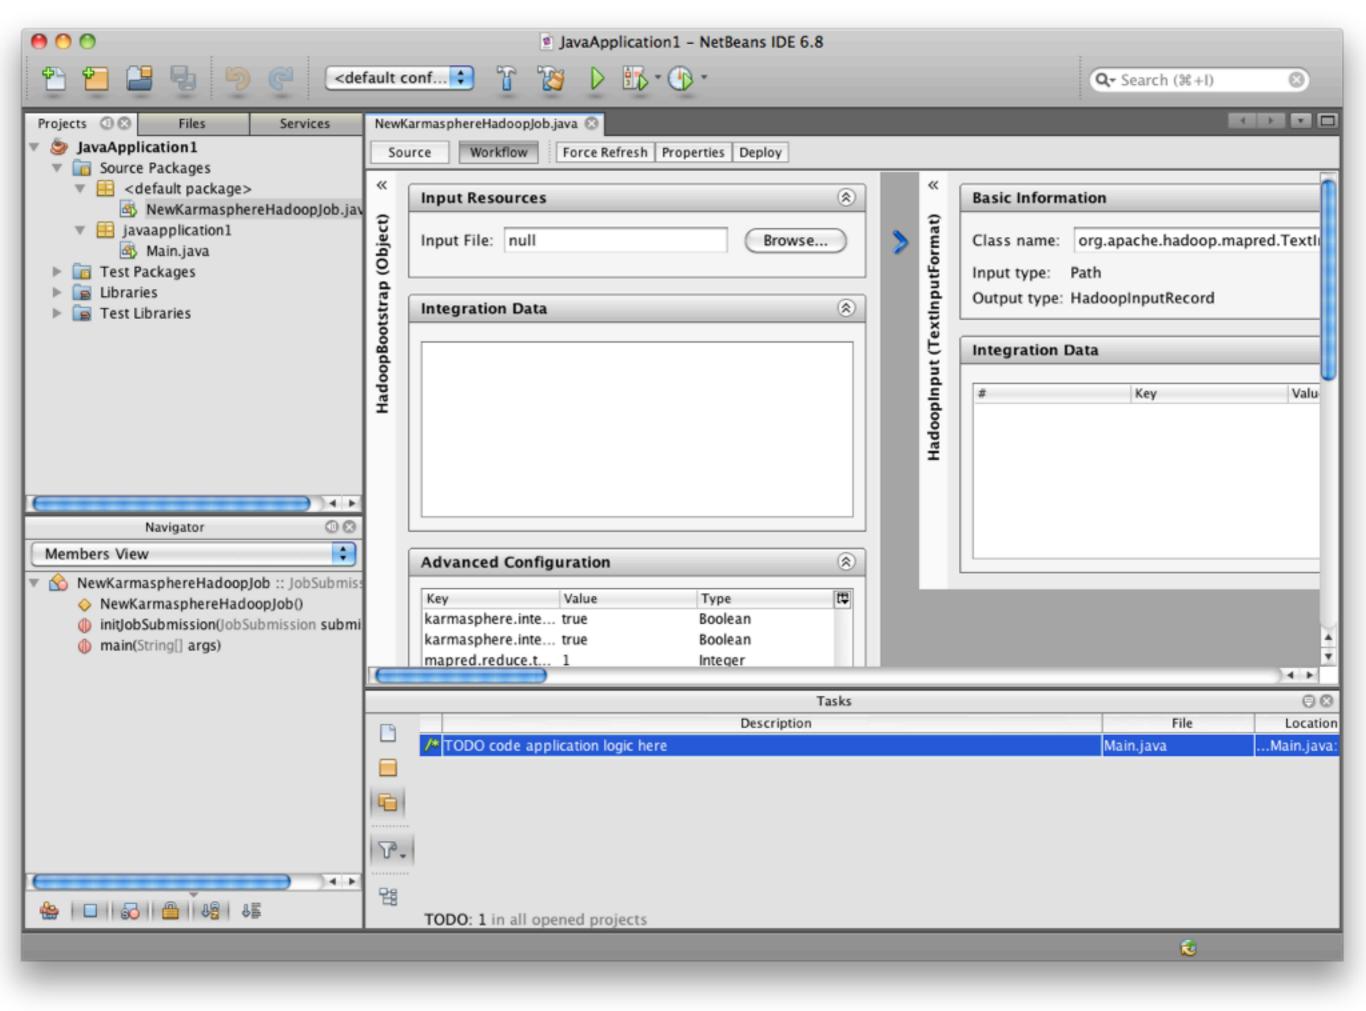
Highlights

- Simple Functions, filters, and aggregators all use the same syntax. Joins are implicit and natural.
- Expressive Logical composition is very powerful, and you can run arbitrary
 Clojure code in your query with little effort.
- Interactive Run queries from the Clojure REPL.
- Scalable Cascalog queries run as a series of MapReduce jobs.
- Query anything Query HDFS data, database data, and/or local data by making use of Cascading's "Tap" abstraction
- Careful handling of null values Null values can make life difficult. Cascalog
 has a feature called "non-nullable variables" that makes dealing with nulls

Karmasphere

HADOOP STUDIO

- Free and commercial offerings
- Workflow designer
- TIDE based testing



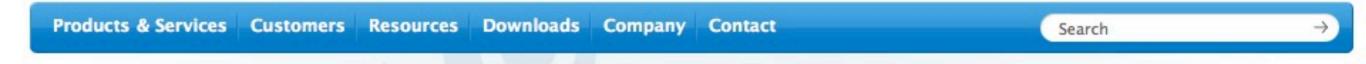
Cloudera Desktop



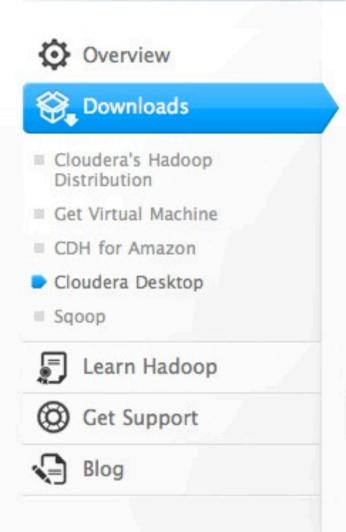
What is Hadoop?

Why Hadoop? Downloads Learn Hadoop Get Support Blog
- Archive
Twitter





Developer Center

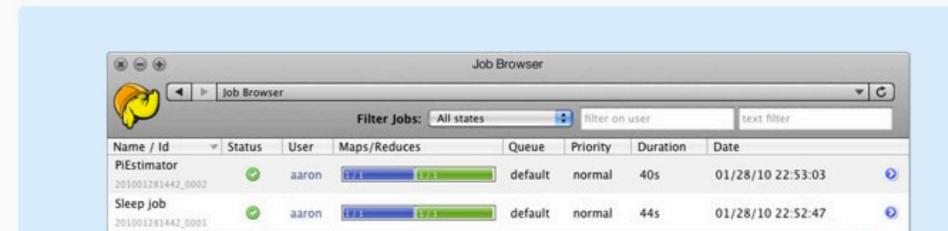




Cloudera Desktop

A unified user interface for users and operators of Hadoop clusters.

Cloudera Desktop is a graphical user interface for the many tools required to operate and develop for Hadoop. These tools are collected into a desktop environment and delivered as a Web application that simplifies cluster administration and job development. With Cloudera Desktop, the world-class performance and scalability of Apache Hadoop is now accessible to anyone in your organization.



GRID ON THE CLOUD



ELASTIC MAP REDUCE

- Cluster MapReduce by the hour
- Easiest grid to set up
- Storage on S3

EMR PRICING

US - N. Virginia US -	N. California EU – Ireland	
Standard On-Demand Instances	Amazon EC2 Price per hour (On-Demand Instances)	Amazon Elastic MapReduce Price per hour
Small (Default)	\$0.085 per hour	\$0.015 per hour
Large	\$0.34 per hour	\$0.06 per hour
Extra Large	\$0.68 per hour	\$0.12 per hour
High Memory On-Demand Instances	Amazon EC2 Price per hour (On-Demand Instances)	Amazon Elastic MapReduce Price per hour
m2.2xlarge	\$1.20 per hour	\$0.21 per hour
m2.4xlarge	\$2.40 per hour	\$0.42 per hour
High CPU On-Demand Instances	Amazon EC2 Price per hour (On-Demand Instances)	Amazon Elastic MapReduce Price per hour
Medium	\$0.17 per hour	\$0.03 per hour





SAPIR-WHORF HYPOTHESIS.... REMIXED



SAPIR-WHORF HYPOTHESIS.... REMIXED

The **ability** to store and process **massive data** influences what you decide to store.



In the next several years, the scale of data problems we are aiming to solve will grow near-exponentially.

HADOOP

divide and conquer petabyte-scale data



METADATA

Matthew McCullough

Ambient Ideas, LLC

matthewm@ambientideas.com

http://ambientideas.com/blog

@matthewmccull

Code

http://github.com/matthewmccullough/hadoop-intro

REFERENCES

REFERENCES

- http://labs.google.com/papers/mapreduce.html
- http://labs.google.com/papers/gfs.html
- http://labs.google.com/papers/bigtable.html
- http://wiki.apache.org/hadoop/Hbase/PoweredBy
- http://mahout.apache.org/
- http://mahout.apache.org/taste.html
- http://www.cascading.org/

REFERENCES

- http://hadoop.apache.org
- http://delicious.com/matthew.mccullough/hadoop
- http://code.google.com/edu/parallel/
- http://www.infoworld.com/d/open-source/whats-new-york-times-doing-hadoop-392?r=788
- http://atbrox.com/2009/10/01/mapreduce-and-hadoop-academic-papers/
- http://www.cloudera.com/

IMAGE CREDITS

- http://www.fontspace.com/david-rakowski/tribeca
- http://www.cern.ch/
- http://www.robinmajumdar.com/2006/08/05/google-dalles-data-centre-has-serious-cooling-needs/
- http://www.greenm3.com/2009/10/googles-secret-to-efficient-data-center-design-ability-to-predict-performance.html
- http://upload.wikimedia.org/wikipedia/commons/f/fc/CERN_LHC_Tunnel1.jpg
- http://www.flickr.com/photos/mandj98/3804322095/
- http://www.flickr.com/photos/8583446@N05/3304141843/
- http://www.flickr.com/photos/joits/219824254/
- http://www.flickr.com/photos/streetfly_jz/2312194534/
- http://www.flickr.com/photos/sybrenstuvel/2811467787/
- http://www.flickr.com/photos/lacklusters/2080288154/
- http://www.flickr.com/photos/sybrenstuvel/2811467787/
- http://www.flickr.com/photos/robryb/14826417/sizes/l/
- http://www.flickr.com/photos/mckaysavage/1037160492/sizes/l/
- http://www.flickr.com/photos/robryb/14826486/sizes/l/
- All others, iStockPhoto.com