

Finding the Needle in a Big Data Haystack

Wolfgang Hoschek (@whoschek)
JAX 2014

About Wolfgang

- Software Engineer @ Cloudera Search Platform Team
- Previously CERN, Lawrence Berkeley National Laboratory, Skytide
- Very Large Scale Analytics Platforms, Distributed OLAP Servers, Search Servers, BI
- Committer on Apache Solr/Lucene, Flume, Lily HBase Indexer, Kite, Morphlines

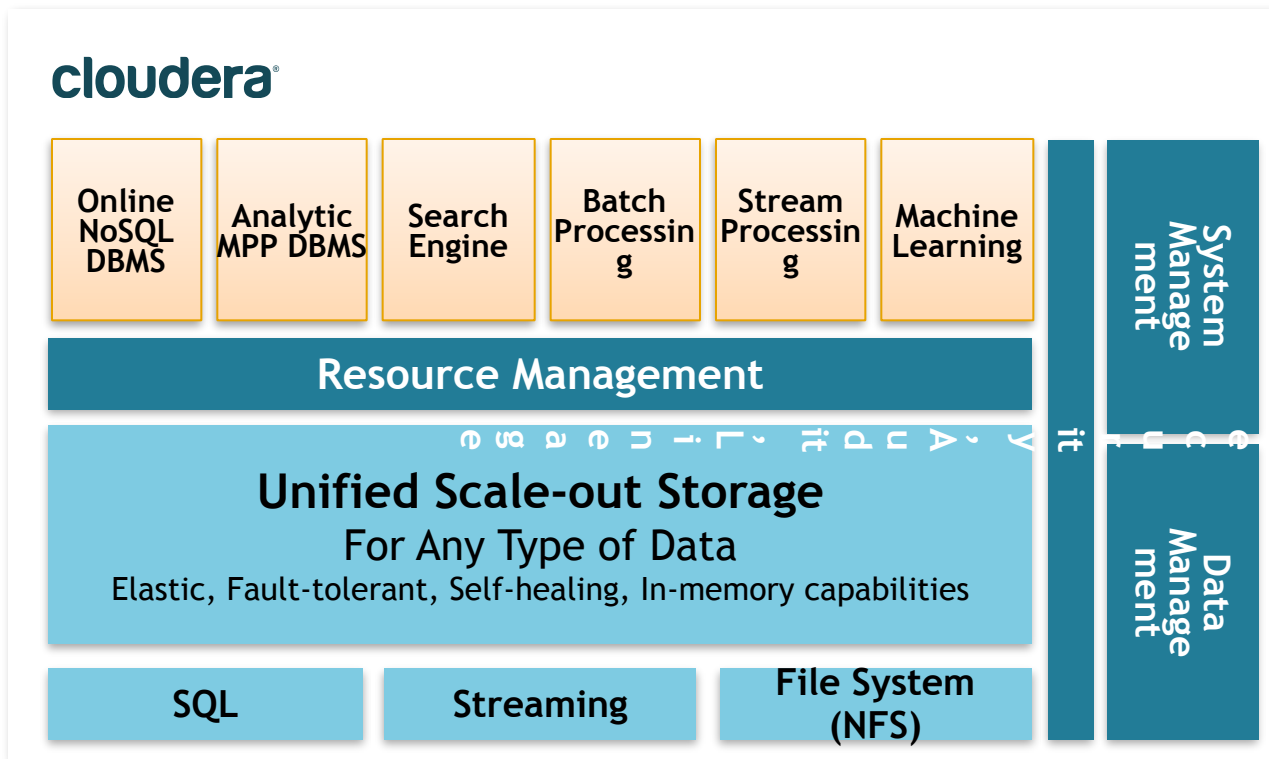
About Mark Miller

- Lucene / Solr committer.
- Previously, Core Engineering Manager at LucidWorks.
- Currently, Software Engineer at Cloudera.
- Co-creator of SolrCloud - Solr's distributed / clustering features.

Agenda

- Big Data and Search - setting the stage
- Cloudera Search's Architecture
- Near Real Time and Batch Use Cases
- Conclusion and Q&A

The Enterprise Data Hub



- Multiple processing frameworks
- One pool of data
- One set of system resources
- One management interface
- One security framework

Search Simplifies Interaction - to Everyone!



Explore



Navigate



Correlate

Experts know MapReduce. Savvy people know SQL.
Everyone knows Search!

What is Cloudera Search?

Interactive search for Hadoop

- Full-text and faceted navigation
- Batch, near real-time, and on-demand indexing

Apache Solr integrated with CDH

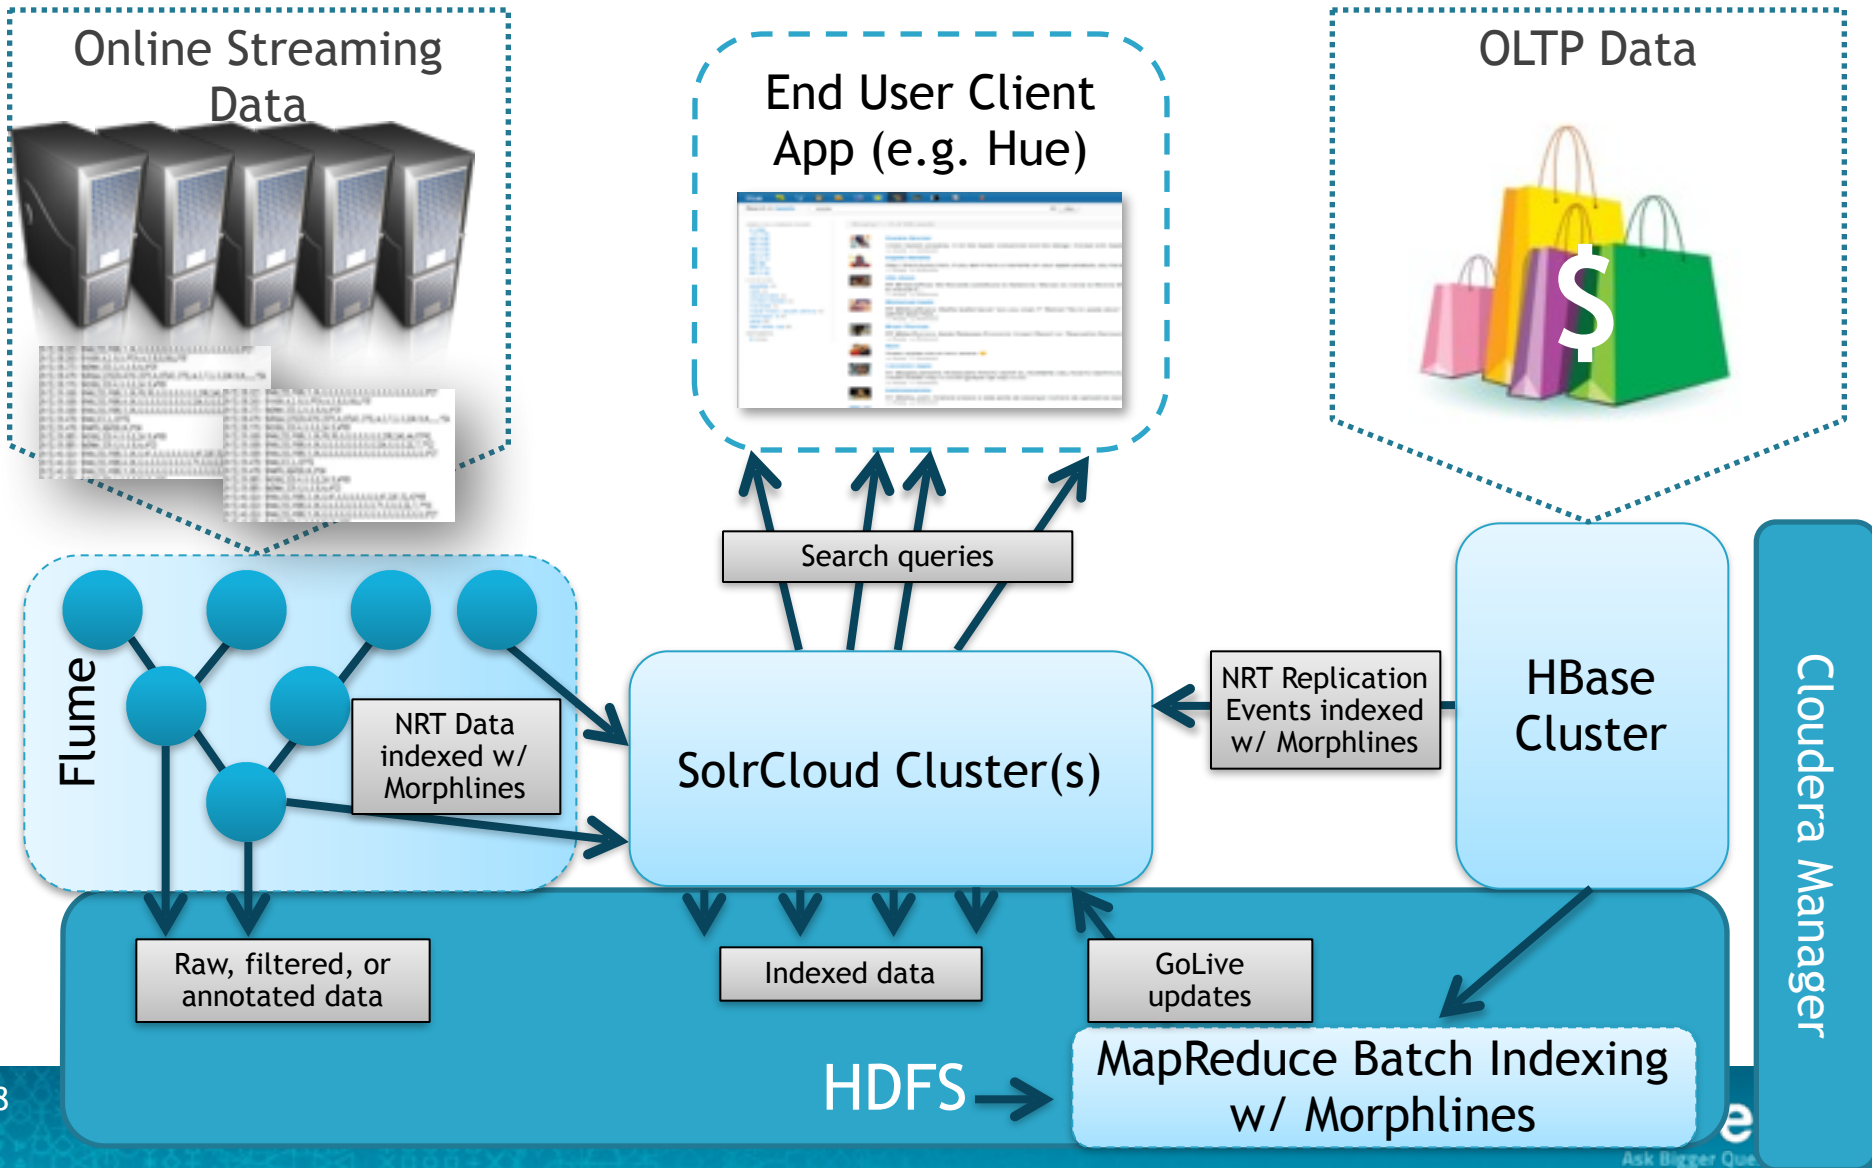
- Established, mature search with vibrant community
- Incorporated as part of the Hadoop ecosystem
 - Apache Flume, Apache HBase
 - Apache MapReduce, Kite Morphlines

Open Source

- 100% Apache, 100% Solr
- Standard Solr APIs



Cloudera Search Architecture Overview



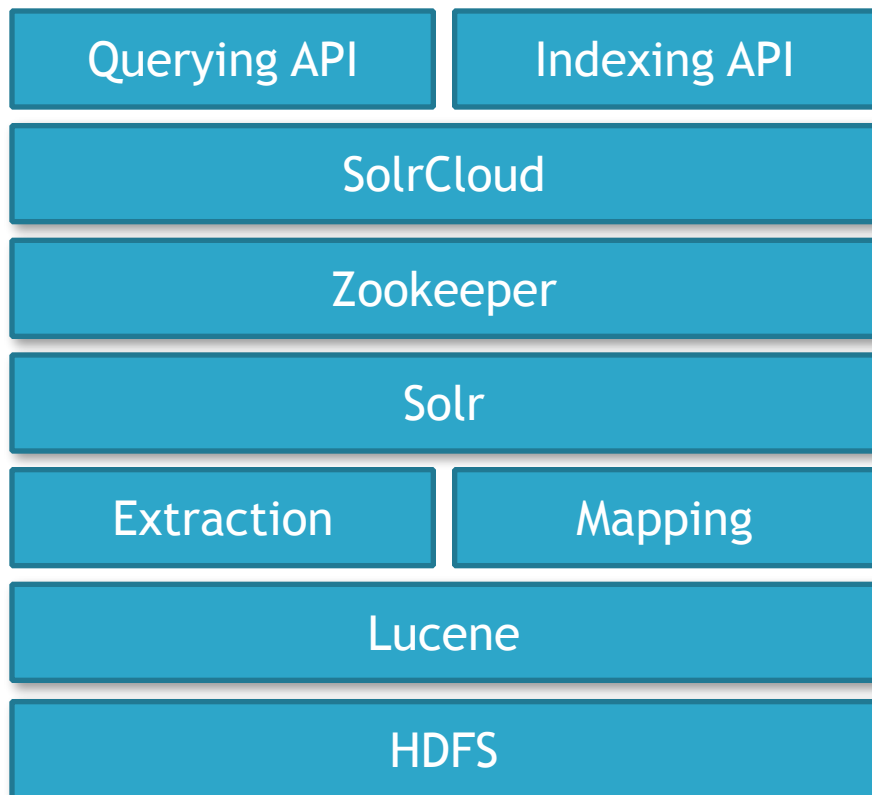
Challenges

- Scalable/Reliable Index Storage
- Near Real Time (NRT) indexing
- Scalable Batch Indexing
- Usability

Apache Lucene/Solr

- Lucene - full text search library
- Solr - search service on Lucene
- SolrCloud - distributed search
 - Partitioned and replicated inverted index
 - Low latency, scalable, reliable, HA, secure

Scalable and Robust Index Storage



Solr and HDFS

- Scalable, cost-efficient index storage
- Higher availability
- Search *and* process data in *one* platform

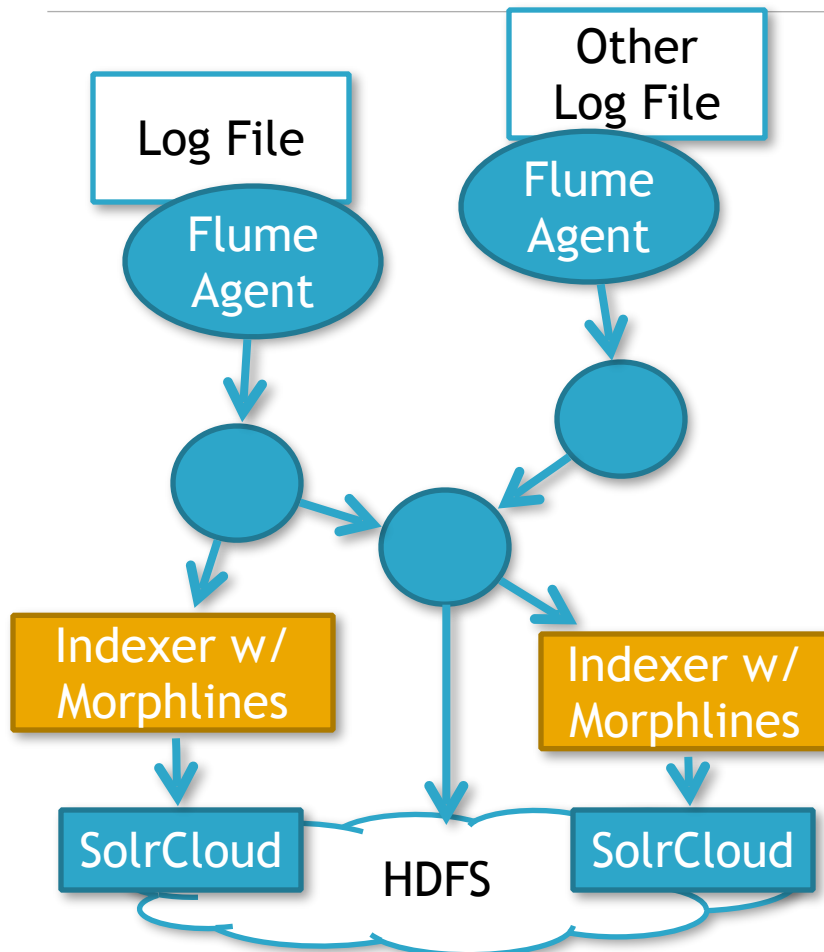
Integrate Solr/Lucene with HDFS

- Lucene Directory Abstraction
 - Implemented HDFSDirectory using HDFS client library
 - Read/Write index files directly to HDFS
- Solr DirectoryFactory Abstraction
 - HDFSDirectoryFactory plugs HDFSDirectory into Solr
 - Configuration - Solr and HDFS
- Solr HdfsTransactionLog
 - Puts transaction log into hdfs so that local file storage is not a concern and for replica failover on hdfs.

Auto Replica Failover

- When indexes are in HDFS, if a node goes down, you can start serving the index for the downed machine from a machine that is still up.
- Currently under development.
- Adds a level of failover and availability that is not available when storing indexes on the local filesystem.

Near Real Time Indexing with Apache Flume



Solr and Flume

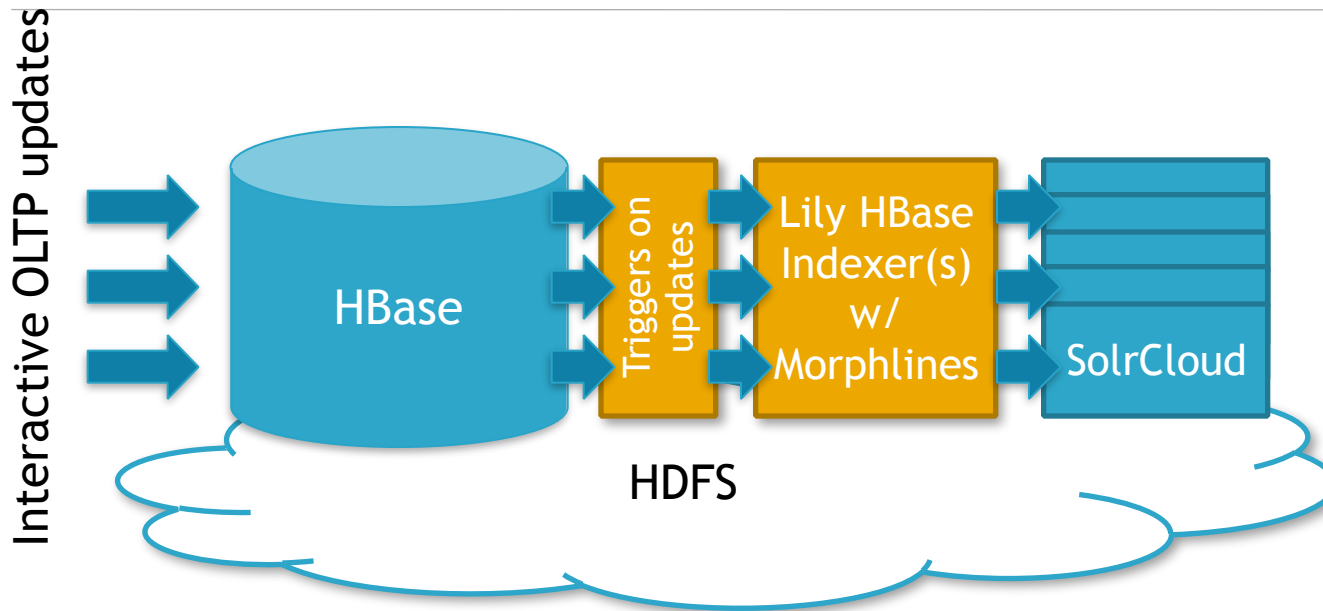
- Reliable streaming data ingestion at scale
- Indexing at data ingest
- Flexible ETL via Morphlines
- Packaged as “Flume Morphline Solr Sink”

Flume.conf

```
agent.sinks.solrSink.type = org.apache.flume.sink.solr.morphline.MorphlineSolrSink
agent.sinks.solrSink.morphlineFile = /etc/flume-ng/conf/morphline.conf
```

Searchable Real-Time Data

Indexing HBase



- Updates Solr index immediately on HBase cell update
- Non-intrusive, Flexible, Scalable, Reliable
- Partitioned and Replicated Indexing w/ Failover, HA


hbase-indexer.conf

```
<indexer table="myHBaseTable"  
mapper="com.ngdata.hbaseindexer.morphline.MorphlineResultToSolrMapper">  
  <param name="morphlineFile" value="/path/to/extractHBaseCell.conf"/>
```

Lily HBase Indexer

- Collaboration between NGData & Cloudera
 - NGData are creators of the Lily data management platform
- Lily HBase Indexer
 - Service which acts as a HBase replication listener
 - Replication updates trigger indexing of updates (rows/cells)
 - Integrates Kite Morphlines library for flexible ETL of rows/cells
 - ASL licensed on github <http://github.com/ngdata/hbase-indexer>

Scalable Batch ETL & Indexing



Solr and MapReduce

- Flexible, scalable, reliable batch indexing
- On-demand indexing, cost-efficient re-indexing
- Start serving new indices without downtime
- “MapReduceIndexerTool”
- “HBaseMapReduceIndexerTool”

```
hadoop ... MapReduceIndexerTool --morphline-file morphline.conf ...
```

Scalable Batch Indexing

Morphline Mapper:
Parse HDFS input into
indexable document

Morphline Mapper:
Parse HDFS input into
indexable document

Morphline Mapper:
Parse HDFS input into
indexable document

Arbitrary reduce steps of indexing & merging using Solr/Lucene
(can use all reducer slots for scalability even if #reducers >> #solrShards)

SolrCloud
on HDFS

Solr Shard 1

Solr Shard N

GOLIVE

GOLIVE

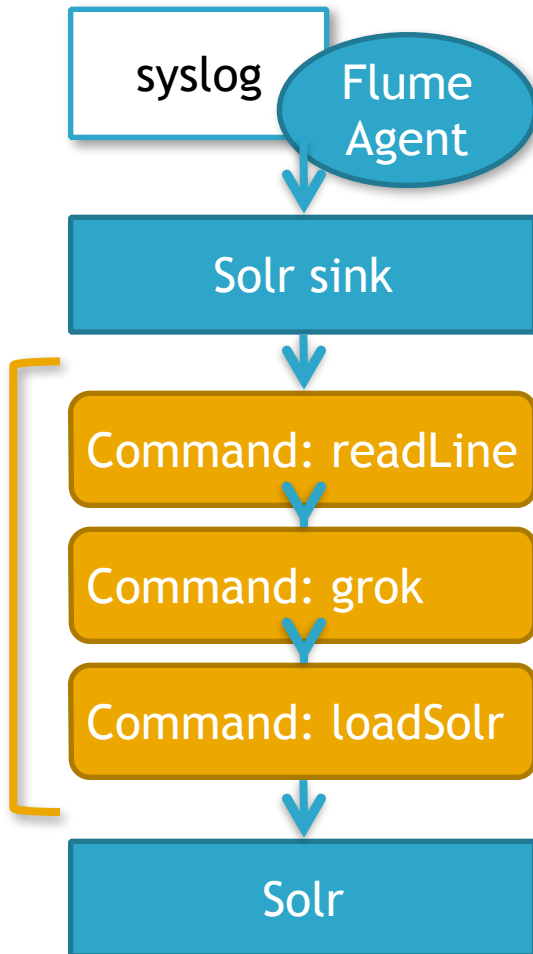
Solr Server 1

Solr Server N

MapReduce Indexer

- MapReduce Job with two parts
 - 1) Scan HDFS (or HBase) for files to be indexed
 - 2) Mapper/Reducer indexing step
 - Mapper extracts content via Kite Morphlines
 - Reducer uses Lucene to index documents directly to HDFS
- “golive”
 - Cloudera created this to bridge the gap between NRT (low latency, expensive) and Batch (high latency, cheap at scale) indexing
 - Results of MR indexing operation are immediately merged into a live SolrCloud serving cluster
 - No downtime for users
 - No NRT expense
 - Linear scale out to the size of your MR cluster

Streaming ETL (Extract, Transform, Load)

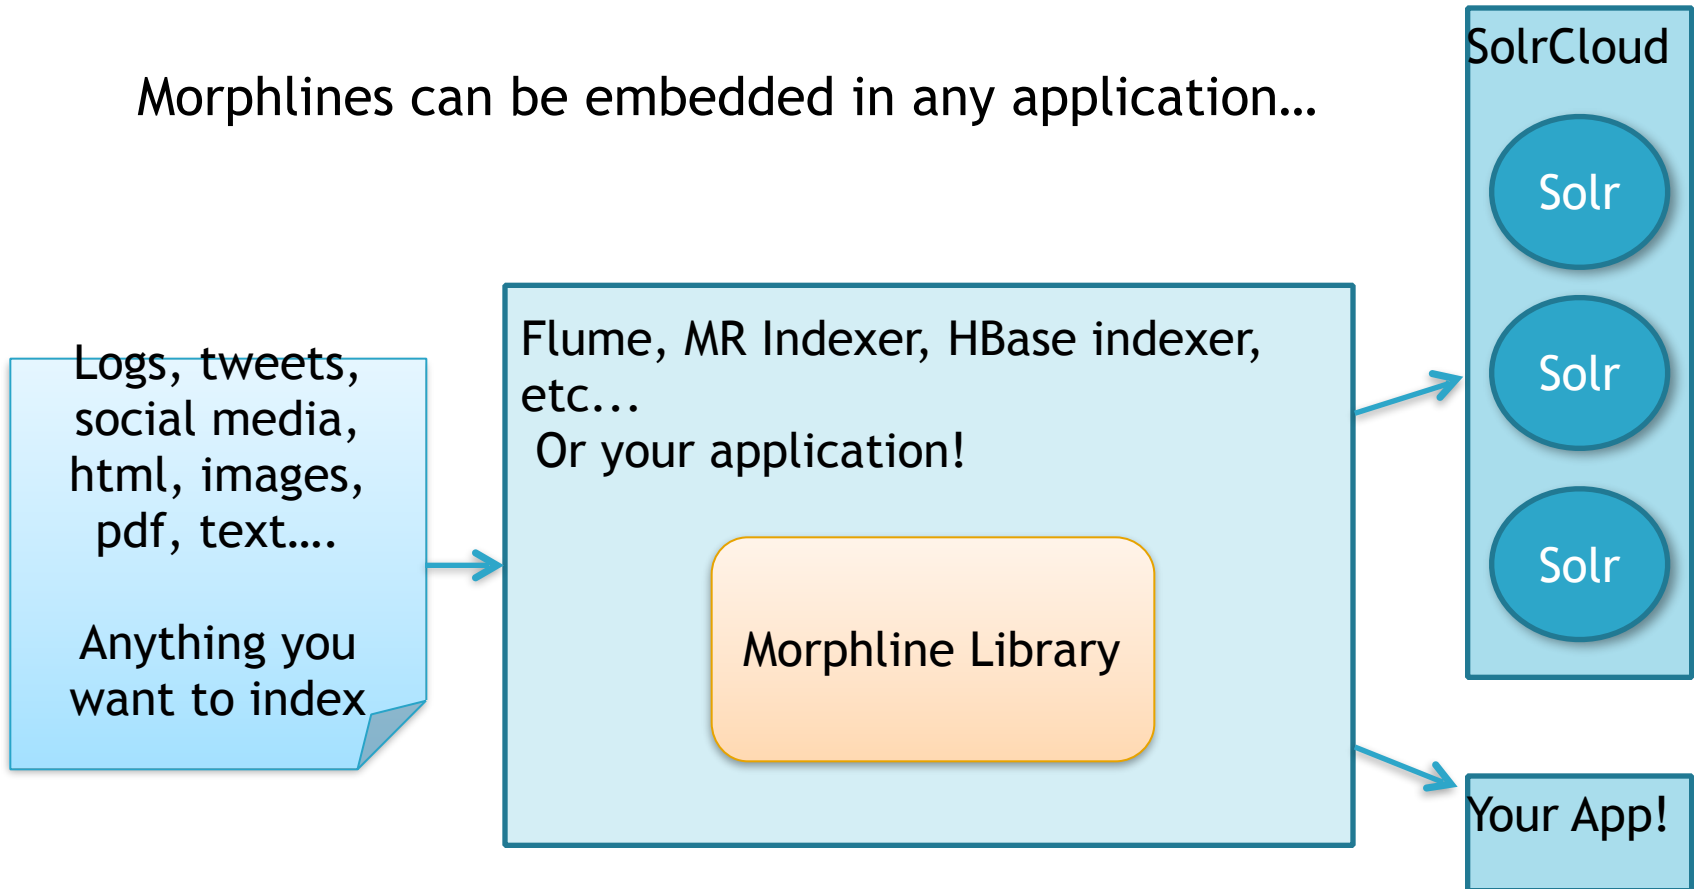


Kite Morphlines

- Event
 - Consume any kind of data from any kind of data source, process and load into Solr, HDFS, HBase or anything else
- Record
 - Simple and flexible data transformation
- Record
 - Extensible set of transf. commands
 - Reusable across multiple workloads
- Record
 - For Batch & Near Real Time
 - Configuration over coding
- Document
 - reduces time & skills
- ASL licensed on github <https://github.com/kite-sdk/kite>

Kite Morphlines Architecture

Morphlines can be embedded in any application...



Morphline Example - syslog with grok

```
morphlines : [  
  {  
    id : morphline1  
    importCommands : ["org.kitesdk.**", "org.apache.solr.**"]  
    commands : [  
      { readLine {} }  
      {  
        grok {  
          dictionaryFiles : [/tmp/grok-dictionaries]  
          expressions : {  
            message : """"<{%POSINT:syslog_pri}>{%SYSLOGTIMESTAMP:syslog_timestamp} %  
{SYSLOGHOST:syslog_hostname} {%DATA:syslog_program}(?:\[{%POSINT:syslog_pid}\])?: %  
{GREEDYDATA:syslog_message}""""  
          }  
        }  
      }  
      { loadSolr {} }  
    ]  
  }  
]
```

Example Input

<164>Feb 4 10:46:14 syslog sshd[607]: listening on 0.0.0.0
port 22

Output Record

syslog_pri:164
syslog_timestamp:Feb 4 10:46:14
syslog_hostname:syslog
syslog_program:sshd
syslog_pid:607
syslog_message:listening on 0.0.0.0 port 22.

Example: Indexing log4j w/ stacktraces

<code>juil. 25, 2012 10:49:40 AM hudson.triggers.SafeTimerTask run ok</code>	Record 1
<code>juil. 25, 2012 10:49:46 AM hudson.triggers.SafeTimerTask run failed</code>	Record 2
<code>com.amazonaws.AmazonClientException: Unable to calculate a request signature at com.amazonaws.auth.AbstractAWSSigner.signAndBase64Encode(AbstractAWSSigner.java:71) at java.util.TimerThread.run(Timer.java:505) Caused by: com.amazonaws.AmazonClientException: Unable to calculate a request signature at com.amazonaws.auth.AbstractAWSSigner.sign(AbstractAWSSigner.java:90) at com.amazonaws.auth.AbstractAWSSigner.signAndBase64Encode(AbstractAWSSigner.java:68) ... 14 more Caused by: java.lang.IllegalArgumentException: Empty key at javax.crypto.spec.SecretKeySpec.<init>(SecretKeySpec.java:96) at com.amazonaws.auth.AbstractAWSSigner.sign(AbstractAWSSigner.java:87) ... 15 more</code>	
<code>juil. 25, 2012 10:49:54 AM hudson.slaves.SlaveComputer tryReconnect</code>	Record 3

Example: Indexing log4j w/ stacktraces

```
morphlines : [  
  {  
    id : morphline1  
    importCommands : ["org.kitesdk.**", "org.apache.solr.**"]  
  
    commands : [  
      {  
        readMultiLine {  
          regex : "(^.+Exception: .+)|(^\\s+at .+)|(^\\s+... \\d+ more)|(^\\s*Caused by:.+)"  
          what : previous  
          charset : UTF-8  
        }  
      }  
      { logDebug { format : "output record: {}", args : ["@{}"] } }  
      { loadSolr {}  
    }  
  ]  
}
```


Example: Escape to Java Code

```
morphlines : [  
  {  
    id : morphline1  
    importCommands : ["org.kitesdk.**"]  
    commands : [  
      { java  
        {  
          code: """  
            List tags = record.get("tags");  
            if (!tags.contains("hello")) {  
              return false;  
            }  
            tags.add("world");  
            return child.process(record);  
          """  
        }  
      }  
    ]  
  }  
]
```

Current Morphline Command Library

- Supported Data Formats
 - Text: Single-line record, multi-line records, CSV, CLOB
 - Apache Avro, Parquet files
 - Apache Hadoop Sequence Files
 - Apache Hadoop RCFiles
 - JSON
 - XML, XPath, XQuery
 - Via Apache Tika: HTML, PDF, MS-Office, Images, Audio, Video, Email
 - HBase rows/cells
 - Via pluggable commands: Your custom data formats
- Regex based pattern matching and extraction
- Flexible log file analysis
- Integrate with and load data into Apache Solr
- Auto-detection of MIME types from binary data using Apache Tika

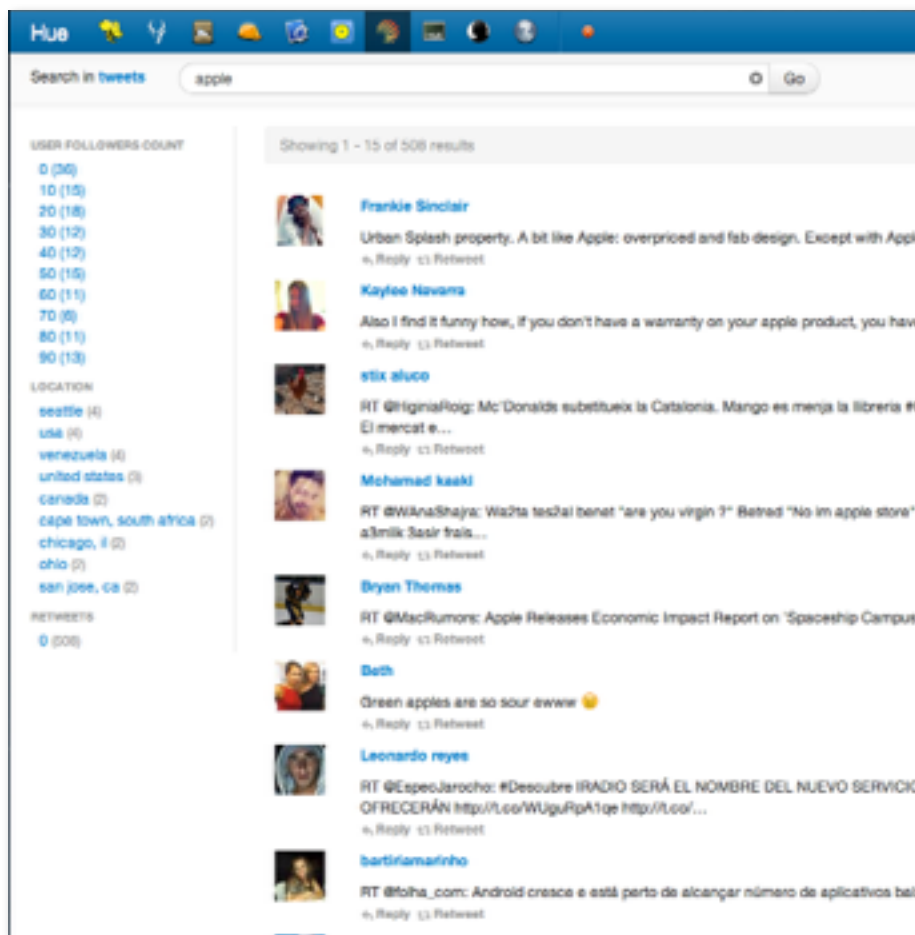
Current Morphline Command Library (cont'd)

- Scripting support for dynamic Java code
- Geolocation info for an IP address (Maxmind)
- Operations on fields for assignment and comparison
- Operations on fields with list and set semantics
- if-then-else conditionals
- A small rules engine (tryRules)
- String and timestamp conversions
- slf4j logging
- Yammer metrics and counters
- Decompression and unpacking of arbitrarily nested container file formats
- etc, etc

Morphline Performance and Scaling

- The runtime compiles morphline on the fly
- The runtime processes all commands of a given morphline in the same thread
- Piping a record from one command to another is fast
 - just a cheap Java method call
 - no queues, no handoffs among threads, no context switches, and no serialization between commands
- For scalability, deploy many morphline instances on a cluster in many Flume agents and MapReduce tasks

Simple, Customizable Search Interface



Hue

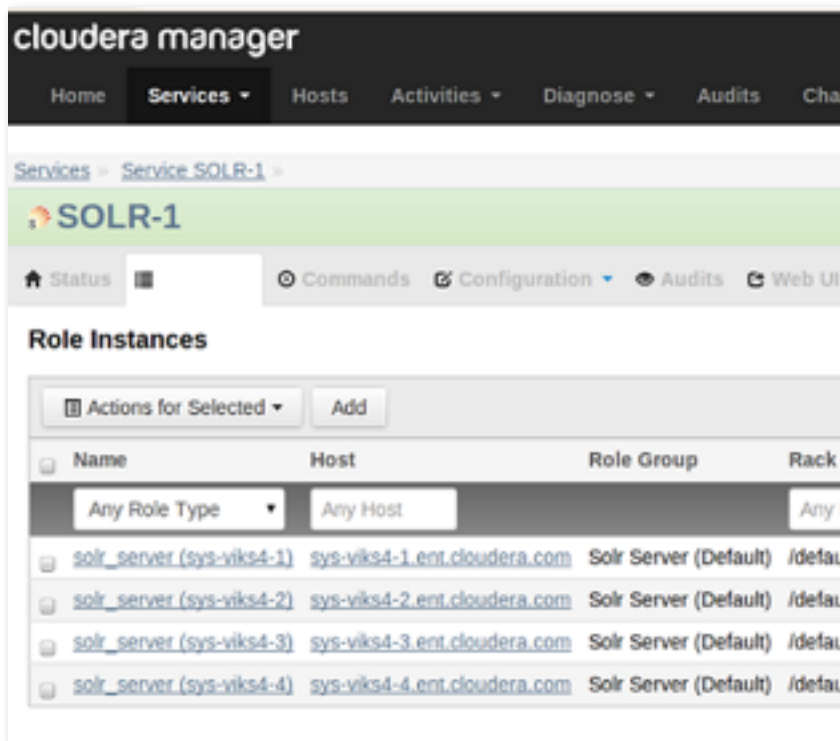
- Simple UI
- Navigated, faceted drill down
- Customizable display
- Full text search, standard Solr API and query language

Security

Cloudera Search + Sentry

- Cluster level access control
- Index level access control
- Document level access control
- Field level access control

Simplified Management



Cloudera Manager

- Install, configure, deploy Solr services on the cluster
- Unified management and monitoring
- Resource management

Cloudera Manager Flume Morphline GUI

The screenshot displays the Cloudera Manager interface for configuring a Flume Morphline. The top navigation bar includes 'Services', 'Hosts', 'Activities', 'Diagnose', 'Audits', 'Charts', 'Reports', and 'Administration'. The main header shows 'flume1' with a 'Concerning Health' status and an 'Actions' menu. The left sidebar contains navigation options: 'Status', 'Instances', 'Commands', 'Configuration', 'Audits', and 'Charts Library'. The 'Configuration' tab is active, showing a search bar and a 'Notes' field with a 'Save Changes' button.

Category	Property	Value	Description
Default	Morphlines File	<pre>SOLR_LOCATOR : { # Name of solr collection collection : collection3 # ZooKeeper ensemble zkHost : "\$ZK_HOST" # The maximum number of documents to send to Solr per network batch (throughput knob) # batchSize : 100 } morphlines : [{ id : morphline1 importCommands : ["com.cloudera.***", "org.apache.solr.***"] commands : [</pre>	Text that goes into morphlines.conf file used by the Flume-NG Solr sink. The text goes verbatim into the config file except that \$ZK_HOST is replaced by the ZooKeeper quorum of the Solr service.

Conclusion and Getting Started

- Cloudera Search for CDH 4 & CDH5
 - Free [Download](#)
 - Extensive [documentation](#)
 - Send your questions and feedback to search-user@cloudera.org
 - Take the Search [online training](#)
- Get started with a Live Instance @ [Cloudera Live](#)
 - No download, no installation, no waiting!
- QuickStart VM also [available](#)



cloudera[®]

Ask Bigger Questions