

# Interoperability in the Apache Hive Ecosystem

Sushanth Sowmyan (Hortonworks)  
Mithun Radhakrishnan (Yahoo)  
201404081645



# Ecosystem

*“An ecosystem is a community of living organisms in conjunction with the nonliving components of their environment, interacting as a system.”*

# Hive Ecosystem

- Part of the larger Hadoop Ecosystem
- Very much a living part of it.
- Sharing data and interoperation between large scale data processing frameworks

# Hadoop Ecosystem

- HDFS
- M/R framework

# Problem : Anagrams

- Consider the following words:
  - Dictionary : ACT, BAG, CAT, DOG, GAB, GOOD, MART, TRAM
- Desired Result :
  - [ [ACT, CAT], [BAG, GAB], [DOG], [GOOD],  
[MART, TRAM]]

# HashMap!

- Where the key is something uniquely common for a group of anagrams, and the value is a list of all the words that hash to that key.
- The trick is in coming up with what's appropriate for that key.

# HashMap! (contd...)

{

“ACT” => [ “ACT” , “CAT” ],

“ABG” => [ “BAG” , “GAB” ],

“DGO” => [ “DOG” ],

“DG00” => [ “GOOD” ],

“AMRT” => [ “MART” , “TRAM” ],

}

# Memory constraints

- What if above hashmap is too large to fit into memory?
- We can flatten it so it's a list of key-values, and constantly emit them:
  - {“ACT” => “ACT”}, {“ABG” => ”BAG”},  
  {“ACT” => ”CAT”}, [“DGO” => ”DOG”},  
  {“ABG” => ”GAB”}, ...
- Then we can sort by the first key, and group.

# Key-generate-Group?

- Map , Classify, Transform
- Group, Reduce, Aggregate, FoldLeft

# Hadoop

- Moving computation to data – data locality
- Generally for jobs where the notion of fitting datasets in memory without some manner of massaging is infeasible.

# What is Hive?

- Database Engine that works on Hadoop
- Data processing/Data warehouse infrastructure on Hadoop, but with the following notions from the get-go:
  - Tabular data
  - Storage engine abstractions
  - Metadata-driven
  - based on SQL
  - Query optimizations, indexes
  - Query evaluation and M/R job generation engine

# Why Hive? (Why not Pig?)

- Pig comes from an ETL world and appeals to the scripters/programmers. Hive appeals to analysts more familiar with SQL.
- BI tools written against SQL
- Pig was written to be able to “eat anything.” Hive considers schema and structure to be central, and dines with a knife and fork.
- Ultimately, stylistic differences in communities.

# Why Hive (What about HBase?)

- Excellent for random r/w
- Excellent for Billions of rows, millions of columns sparse tables.
- These are not common use-cases for relational operations, and speak to a very different mindset and optimization requirement.
- NoSQL ?

# What is HCatalog?

- Formerly Howl, which was formerly Owl.
- Apache Incubator project that's now part of hive
- Metadata-backed storage abstraction
- Allows for interoperability between various tools – Pig or a custom M/R program can now read/write Hive tables.
- Allows for data migrations, and format-independence.

# What is HiveServer2?

- Next step in evolution of Hive usage:
  - Lone fat-client with all the logic, mostly a way to use SQL on top of HDFS data
  - Need for sharing metadata across a team - common database server for metadata
  - Need for protecting kinds of access to the database server, and not deploying database login credentials across the board – metastore server
  - Need for access control on metastore server, protected HDFS access – Hive-as-a-service
  - JDBC/ODBC connection modes, SQLLine(Beeline) support - HS2

# Project Starling

Hadoop Introspection, for make benefit Audit and Performance  
Tuning

# Going “Meta”

- ⊗ “How do we store, catalogue and analyze the logs generated from the thousands of Hadoop jobs run across the dozens of Hadoop clusters across the company?”
- ⊗ What insights do we gather from analyzing Hadoop logs?

# Example Insights

- ⊗ Metering:
  - ⊗ Which user/group uses the most compute-time?
  - ⊗ Who's hogging HDFS space?
- ⊗ Job analysis:
  - ⊗ Weapon of Choice: M/R, Pig, Hive
  - ⊗ Pig 0.10 vs. 0.11
  - ⊗ Job-failures: OOM, Auto-tune mapreduce.map.memory.mb.
- ⊗ Data analysis:
  - ⊗ Which dataset is most popular?
  - ⊗ Last-access-time:
    - ⊗ Tune data-retention periods, Replication factor, Cold Storage
- ⊗ Analyze using SQL, Graph with Excel/Tableau/MicroStrategy

# Hive

- Query large datasets on the Hadoop DFS (and elsewhere) using SQL-ish syntax
  - E.g. `SELECT page_url FROM page_views WHERE user_id = 'mithunr' AND dt='20140101';`
- Represent large datasets as databases, tables and partitions
  - HCatalog: Metadata storage
  - Wrappers for Map Reduce and Pig
- JDBC/ODBC connectivity:
  - HiveServer2
  - Business Intelligence tools: Tableau, MicroStrategy

# Example Code

❁ CREATE TABLE job\_info( job\_id STRING, job\_name STRING, user STRING, queue STRING, priority STRING, status STRING, run\_time BIGINT ) partitioned by (cluster\_name STRING, dt STRING) STORED AS RCFILE;

❁ Hive query:

```
SELECT user_id, sum(run_time) AS total_runtime
FROM job_info
WHERE cluster_name='staging3' and dt BETWEEN '20140101' AND '20140131'
GROUP BY user_id ORDER BY total_runtime DESC LIMIT 10;
```

❁ Pig:

```
job_info = LOAD 'starling.job_info' using org.apache.hive.hcatalog.pig.HCatLoader();
january_staging_job_info = FILTER job_info BY dt >= '20140101' AND dt <= '20140131' AND
cluster_name='staging3';
grouped_job_info = GROUP january_staging_job_info by user_id;
agg_runtimes = FOREACH grouped_job_info GENERATE group, SUM(run_time) total_runtime;
sorted_agg_runtimes = ORDER agg_runtimes BY total_runtime DESC;
top_users = LIMIT sorted_agg_runtimes LIMIT 10;
DUMP top_users;
```

❁ HiveServer2:

❁ jdbc:hive2://hive-server2.myth.net:10000/starling

# DistCp

- ⊗ Uses MapReduce to copy files between clusters.
- ⊗ E.g.

```
hadoop distcp -m 20 \
```

```
hftp://source-cluster:50070/jobhistory/20140331/* \
```

```
hdfs://starling-cluster:8020/jobhistory/20140331
```

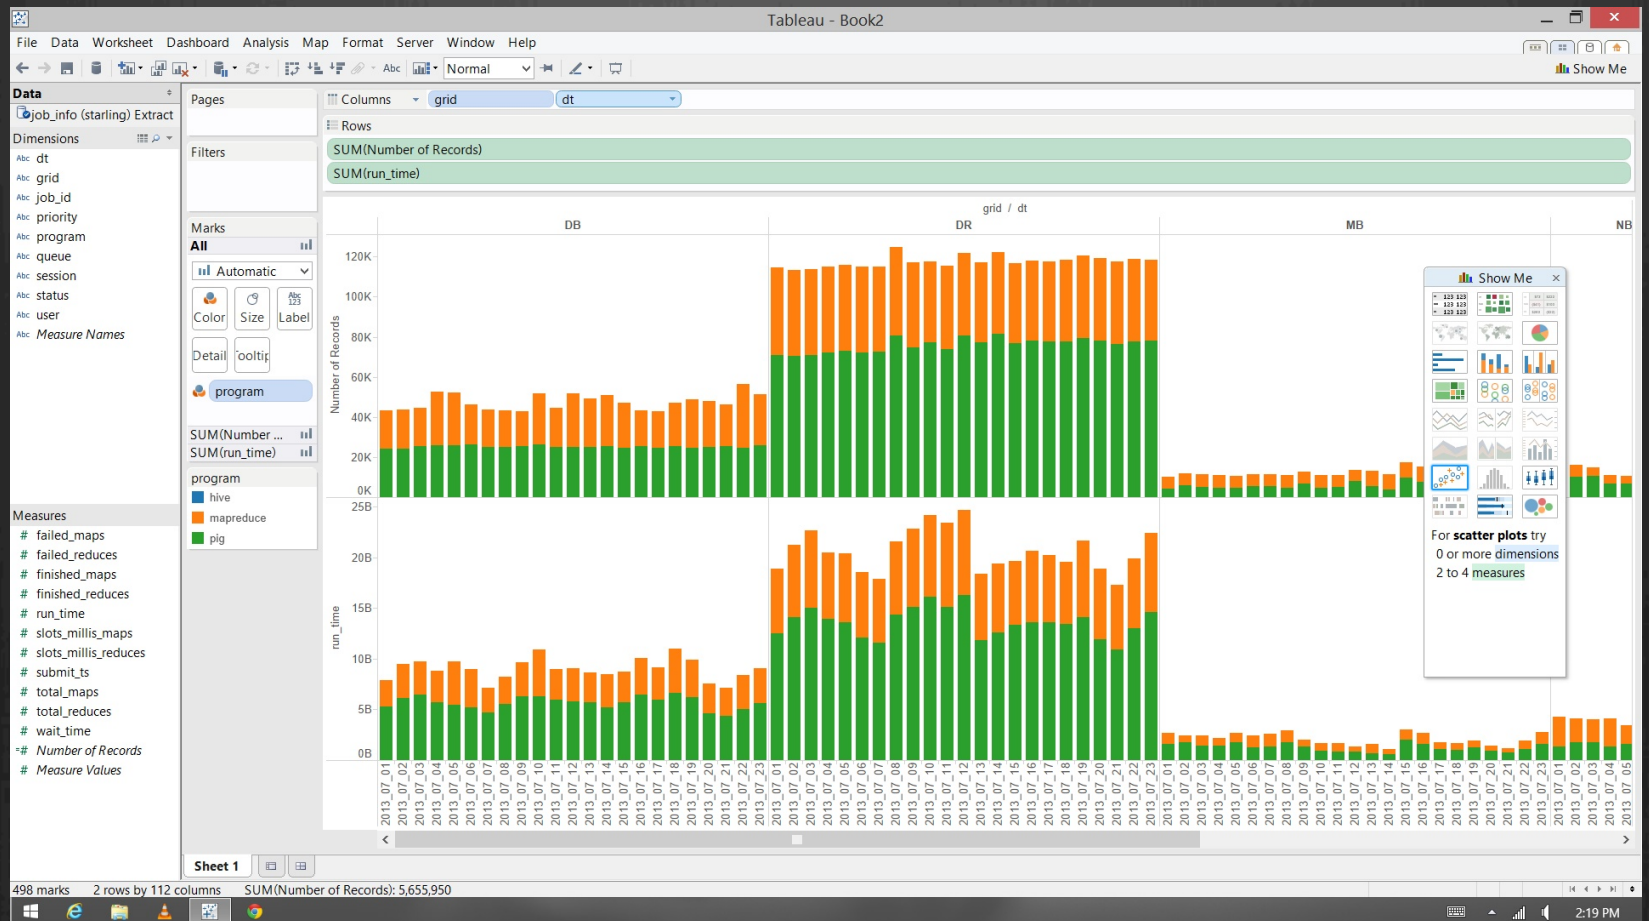
# Parser Reference

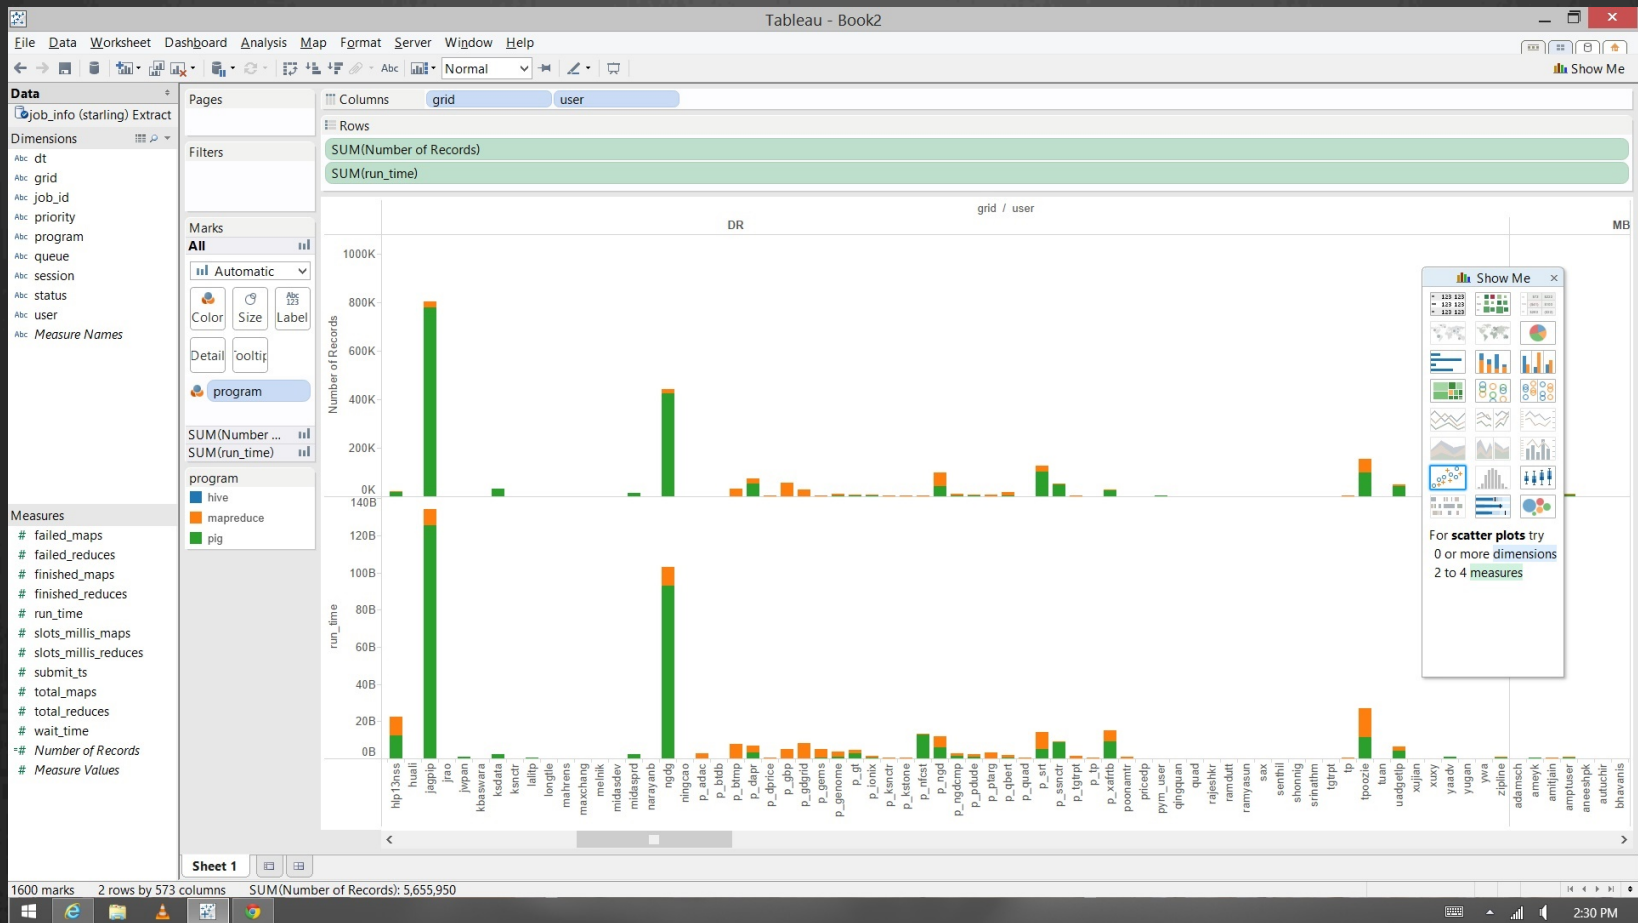
- ❶ `org.apache.hadoop.tools.rumen.JobHistoryParser`
- ❷ Hadoop File System Offline Image View (HDFS OIV)
- ❸ Convert from Raw log-lines to HiveRecords.

# Oozie

## ⌘ cron for Hadoop

```
<workflow-app xmlns='uri:oozie:workflow:0.1' name='DoWhatIMeanTo'>
  <action name='CopyRawLogs'>
    <java>
      <main-class>org.apache.hadoop.tools.DistCp</main-class>
      <arg>hdfs://source-cluster/jobhistory/20140331</arg>
      <arg>hdfs://starling-cluster/jobhistory/20140331</arg>
    </java>
  </action>
  <action name='ProcessLogs'>
    <java>
      <main-class>net.myth.starling.ProcessJHistLogs</main-class>
      <arg>hdfs://starling-cluster/jobhistory/20140331</arg>
    </action>
  </workflow-app>
```





Questions?

EOF