



Hadoop

- Evaluation
 - TagConverter
 - Findings
 - Recommendations
- Possible Architecture



TagConverter

- Playground package
 - Converting between Hadoop file formats (csv,seq,binary,...)
 - Creating vertical partitions (groups of attributes per file)
 - Very configurable
 - Testing/benchmarking
 - Remote access API (based on ssh)

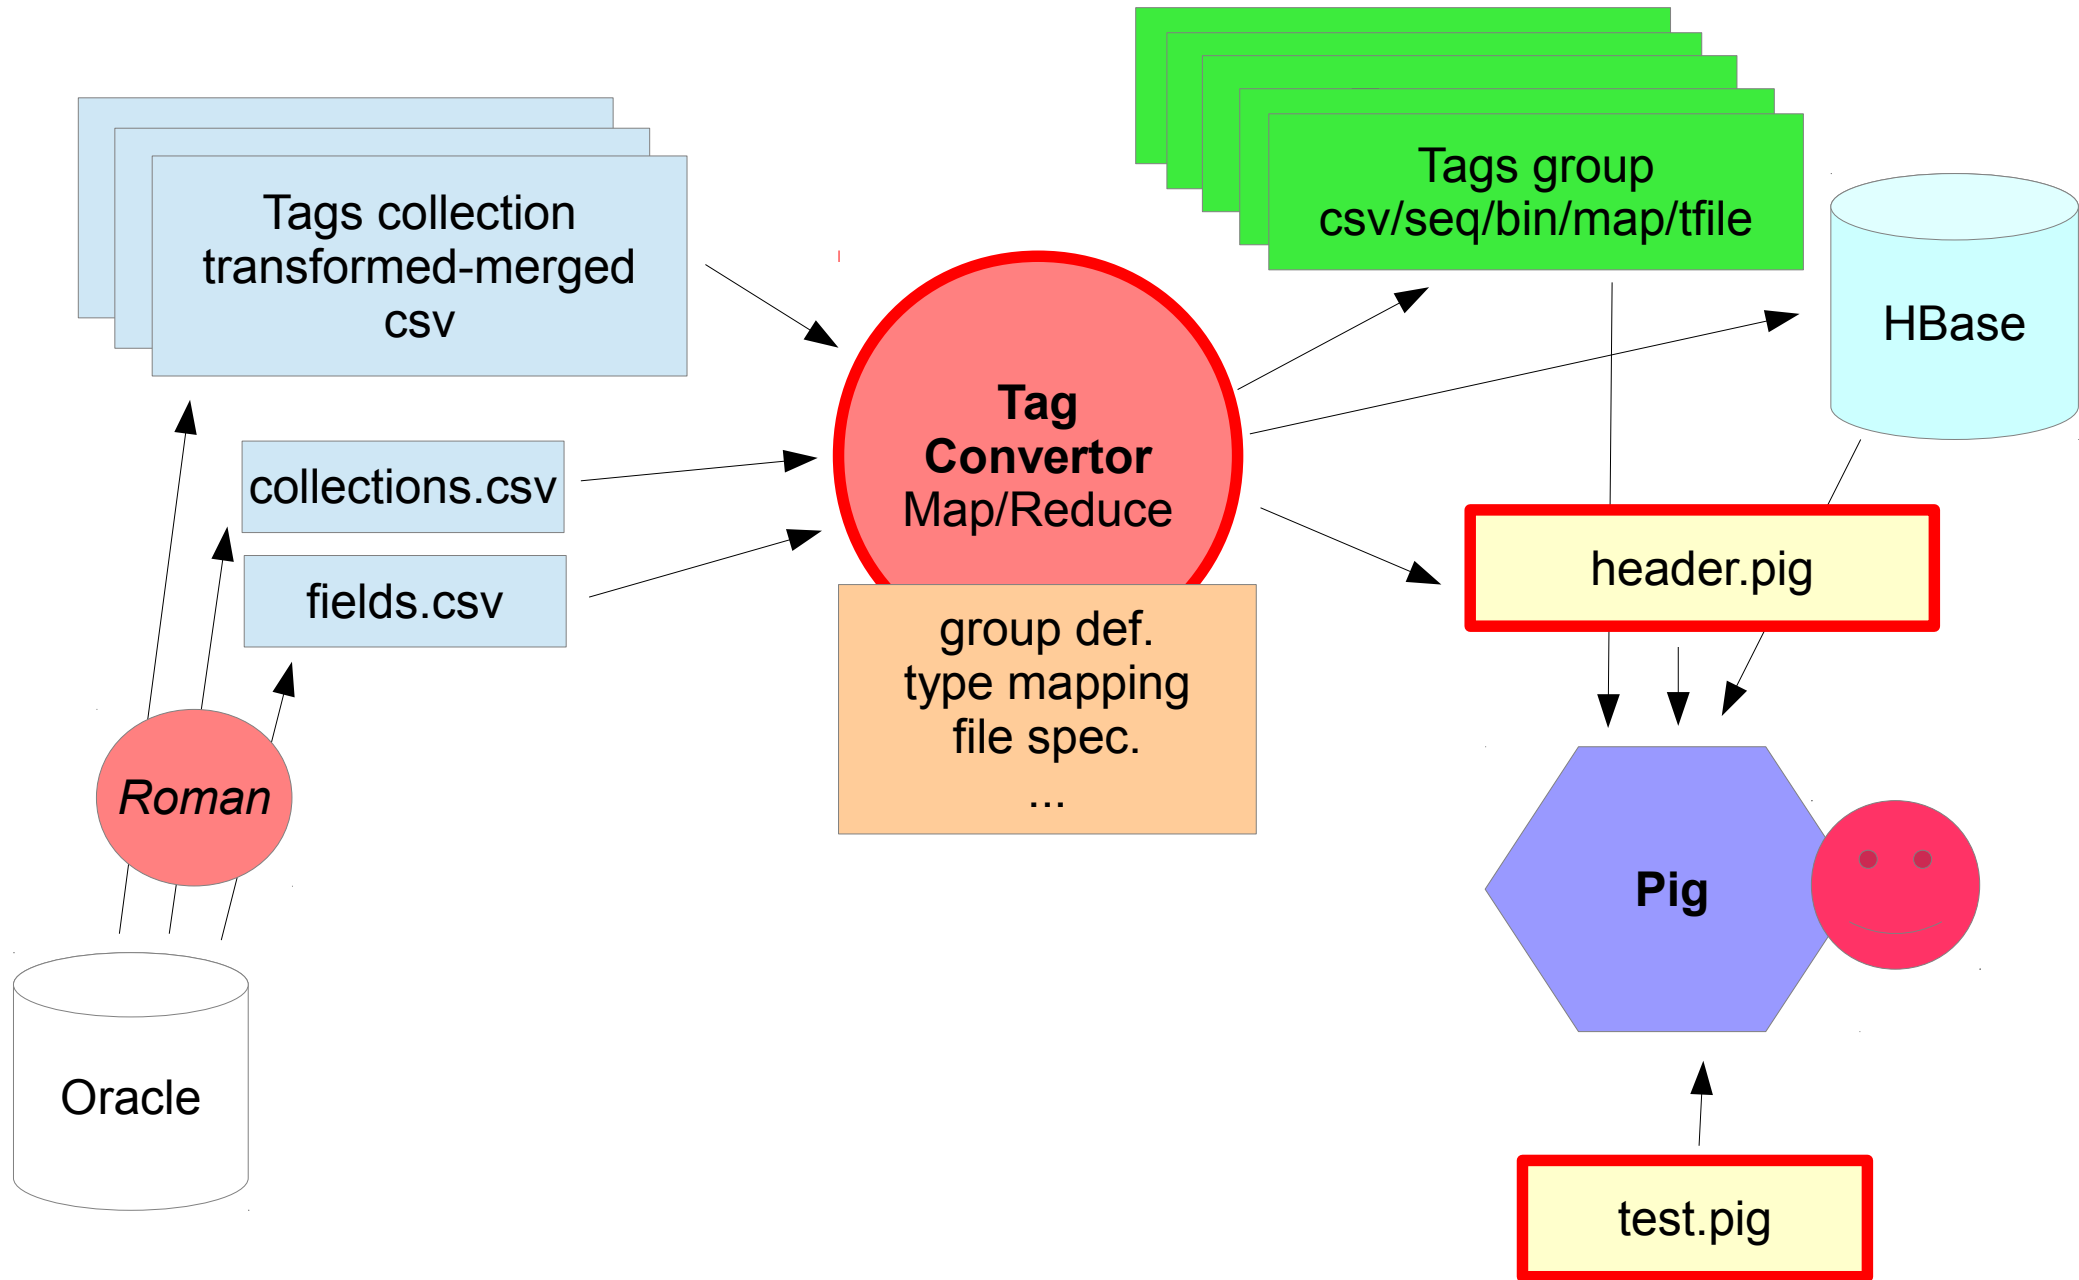
SVN

<svn+ssh://svn.cern.ch/repos/atlasoff/Database/TAGHadoop/TagConvertor>
(HowTo.txt in root)

Doc

<http://cern.ch/hrivnac/Activities/Packages/TagConvertor>

Evaluation





Findings

- Partition sensitivity:
 - The fastest read is when the partition (set of columns) exactly satisfies requested fields
 - More than 2 JOINS kills the performance
- Binary files are slower (and not automatic)
 - R/W should be implemented for non-basic file formats (csv, seq)
 - Additional encoding/decoding seems to cause a performance degradation, while doesn't bring space advantage (in comparison with standard compression)
 - We may still consider non-standard formats, but not for performance or space reasons
 - New Hadoop file types under development, supporting different data organisation (column-wise, partially column-wise, replicas with different data organisation,...) - keep eye on
- HBase useful and fast in its domain
 - Faster search if searching on keys or keys derivatives (substrings, regexp)
 - Complex (maths) search not supported; can be added, but would be slow as it requires full scan

Read/search speed is more important than write/upload speed or size.



Recommendations

- **Use HDFS for the tag files**
 - Use csv-like format
 - Use simple vertical partitions
- **Use HBase for indexes**
 - Cache-like indexes can be created later

Possible Architecture

