

Chukwa

<http://wiki.apache.org/hadoop/Chukwa>

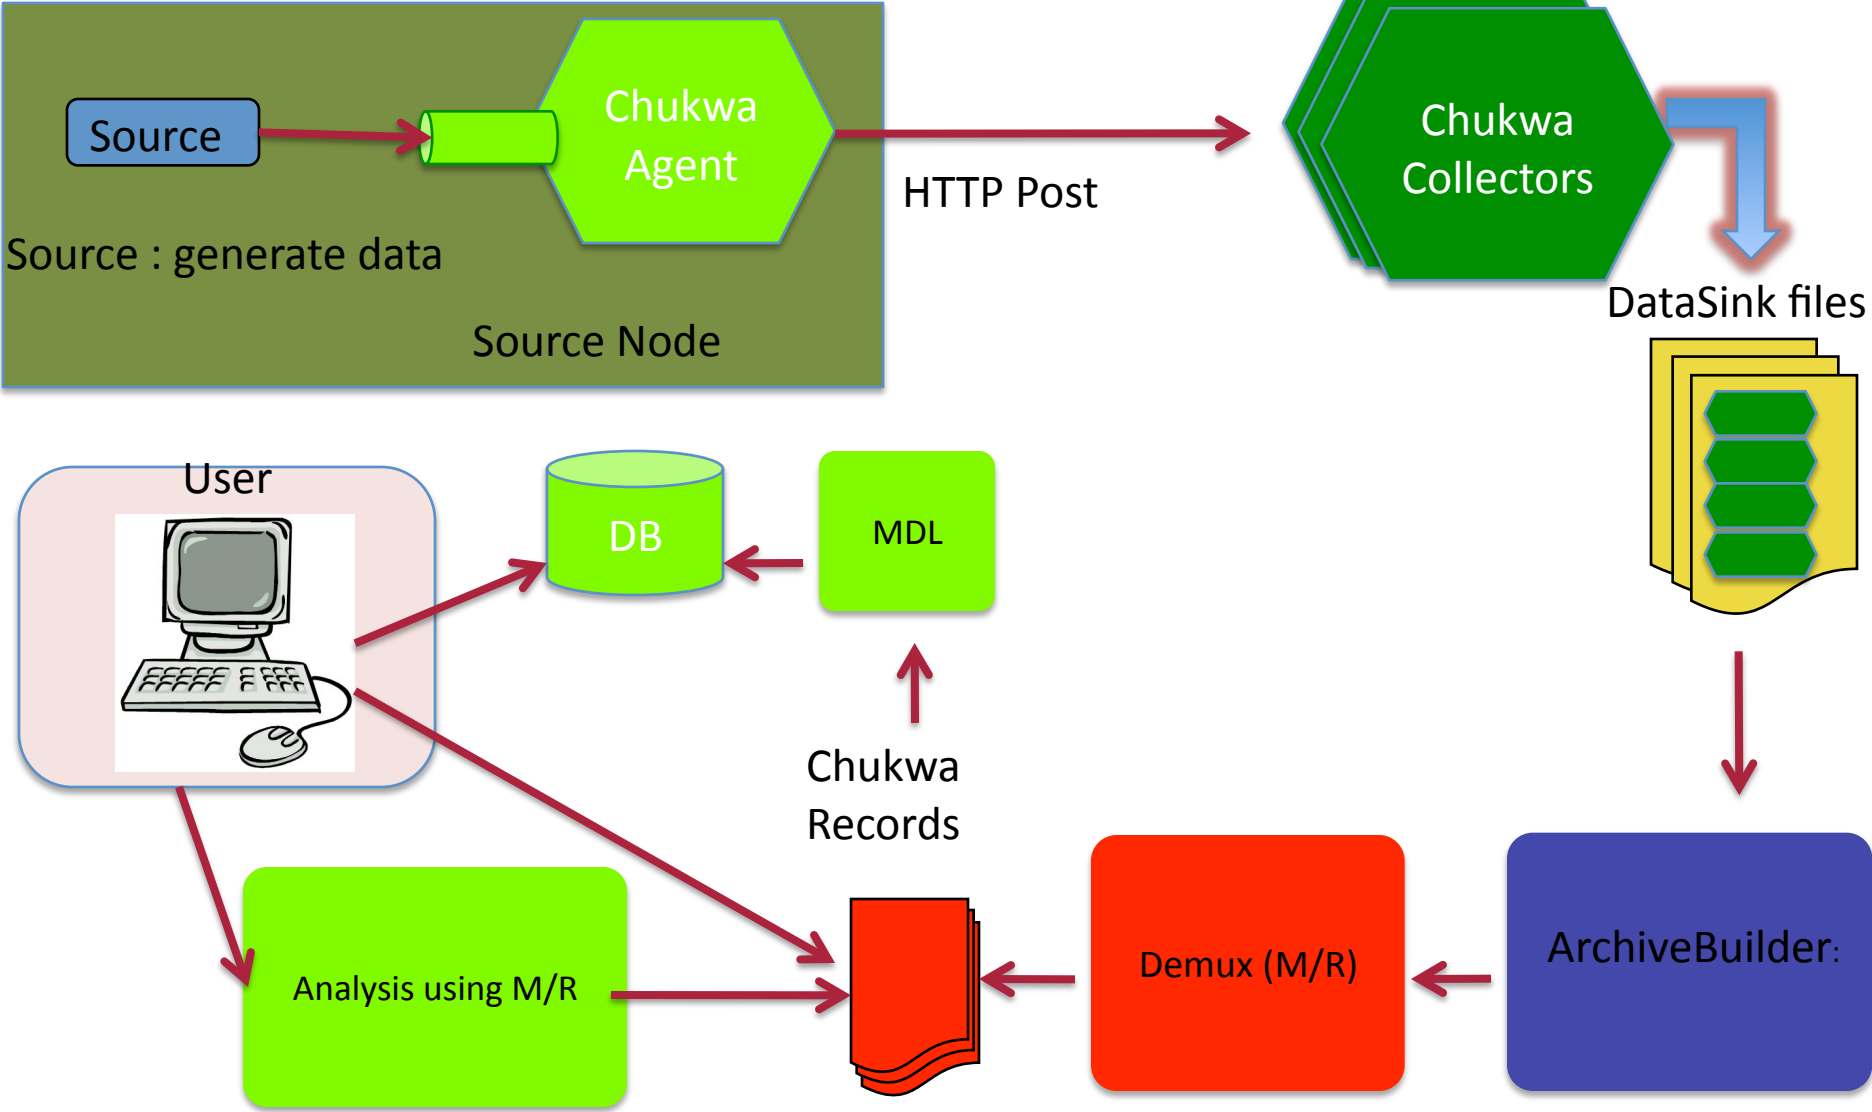
Jerome Boulon – jboulon@yahoo-inc.com



Chukwa Goals

- Collect
 - System metrics
 - Arbitrary log files (unknown format)
 - Known log files (well define format)
 - Handle log rotation
- Latency should be in minutes but not in hours
 - Freshly generated logs can be used within minutes
- Scale to large cluster
- Deal with failure of any part of the collection network
- Should use and take advantage of HDFS and M/R

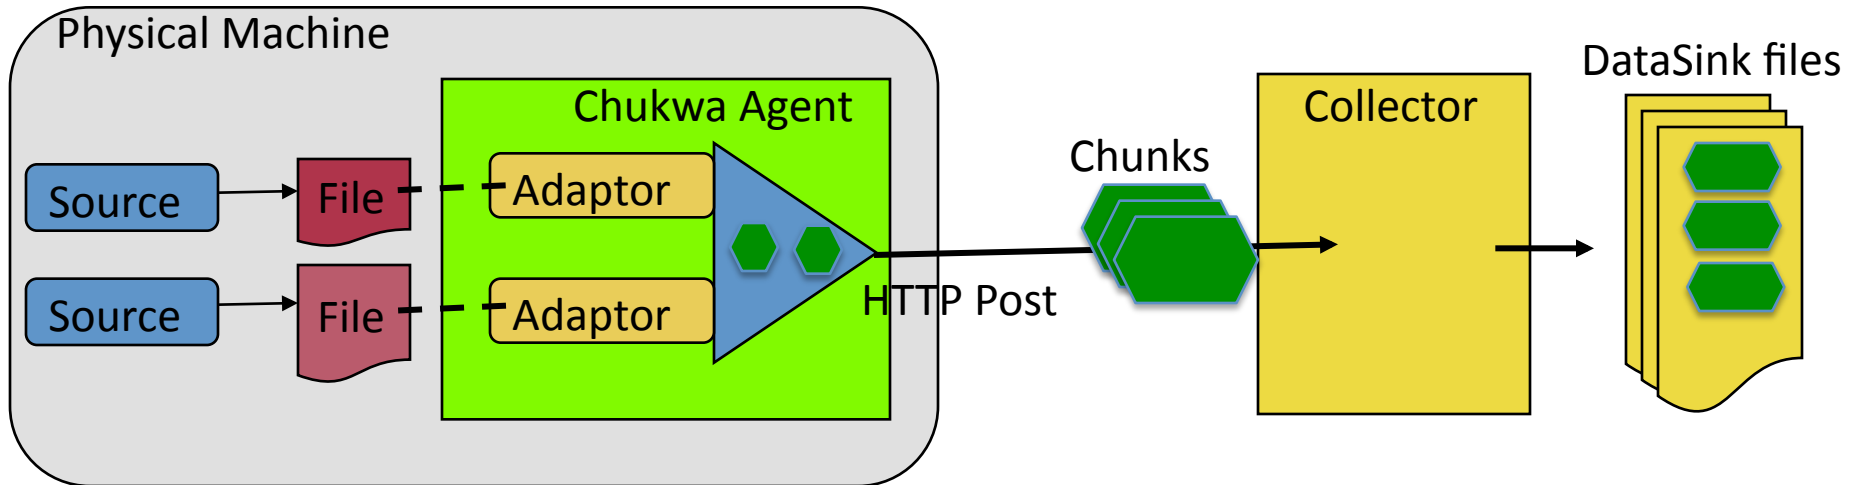
Overview




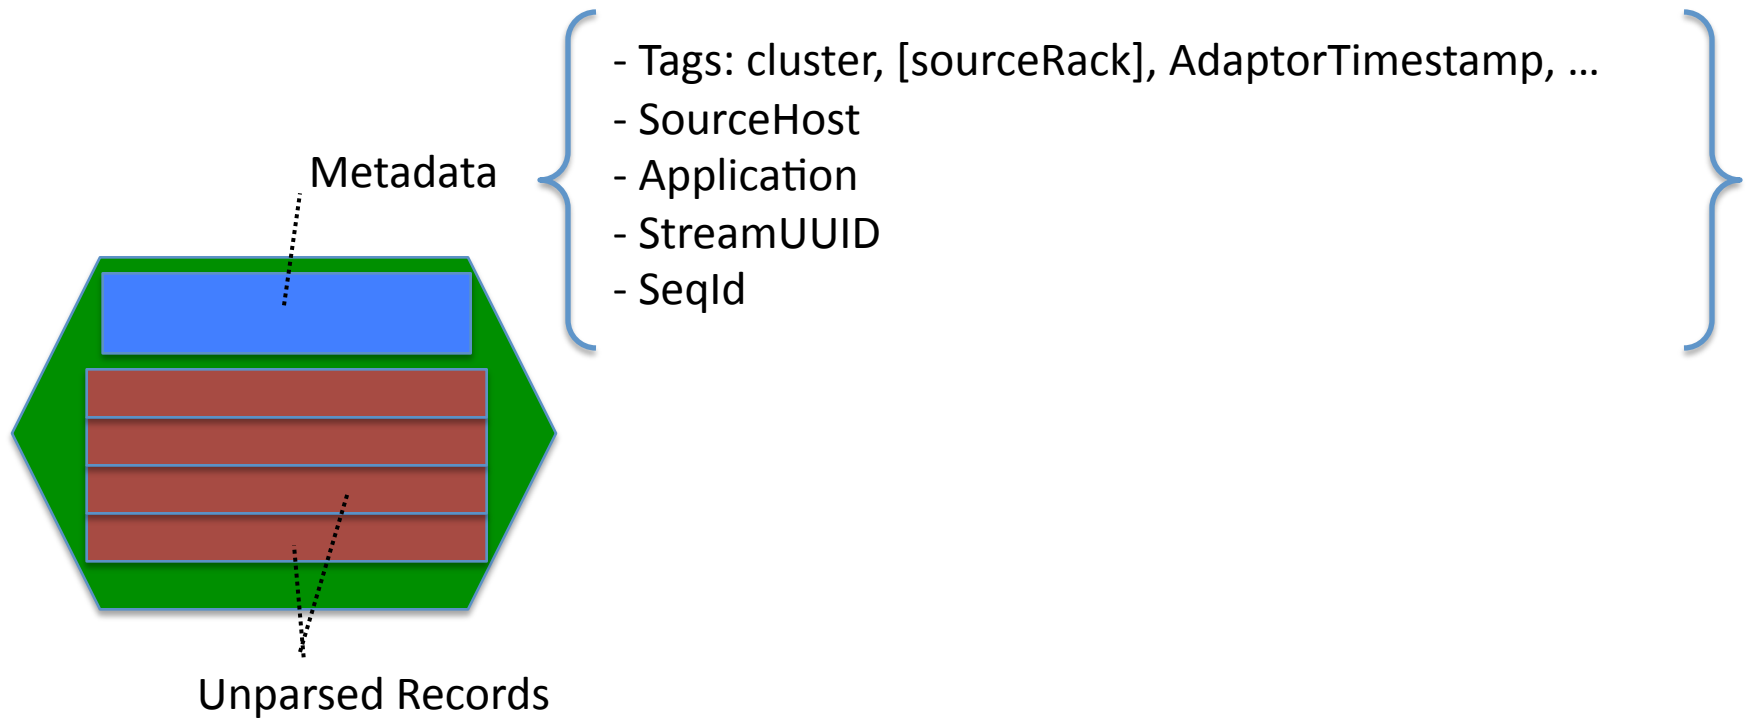
Roles

- ChukaAgent: transfer source data to collector
- Collector: continuously save data to HDFS
- DataSink: contains raw data unsorted/possible duplicates
- ArchiveBuilder: remove duplicates, merge dataSink files together & sort
- Archive File: indexed chunks ready to be parsed
- Demux: parse & convert unstructured chunks to fully or semi structured ChukwaRecord
- ChukwaRecords:
 - ready to be viewed & M/R processed
 - Prepare data for the Database
- Database loader:
- MDL: define ChukwaRecord to database mapping
 - Database: metrics data & down sampled old metrics data
 - Allows real time queries


Data collection: ChukwaAgent -> Collector



Chunk



Automatic Failover & Load Balancing




In case of error/exception, the agent pick a new collector from a predefined list and resend the same list of chunks

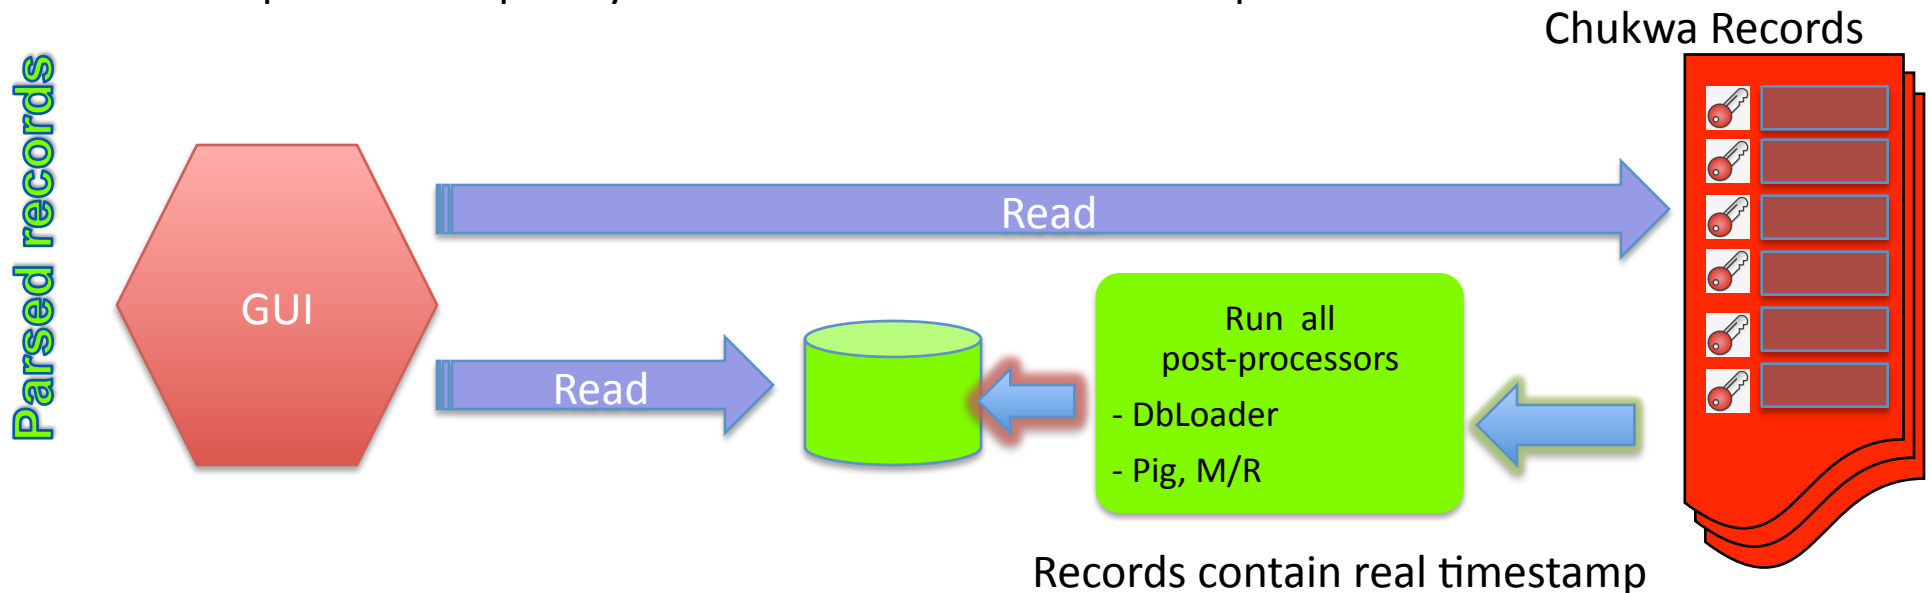
Failover 

Data flow & workflow

Backend Workflow



Steps run in seq every 5 minutes within the same script shell



Demux role

- Parse, extract & transform useful information from the record
- Convert chunks to parsed ChukwaRecords
 - *Fully structured format* (real timestamp + Key,Value pairs)
 - Df, Iostat, Sar, Top, ...
 - JobHistory, Hadoop metrics
 - PbsNodes, NodeActivity
 - *Semi structured format* (real timestamp + raw log entry)
 - Hadoop NameNode/dataNodes, ...
- Fully configurable
 - Define the RecordType-Parser relationship
 - Generally we don't drop anything

Chukwa records

- Role
 - Records with parsed fields
 - Easy post-processing and/or analysis using M/R
 - Prepare data for regular direct query
 - Used as input to database

Chukwa Record – Sequence File

Key:

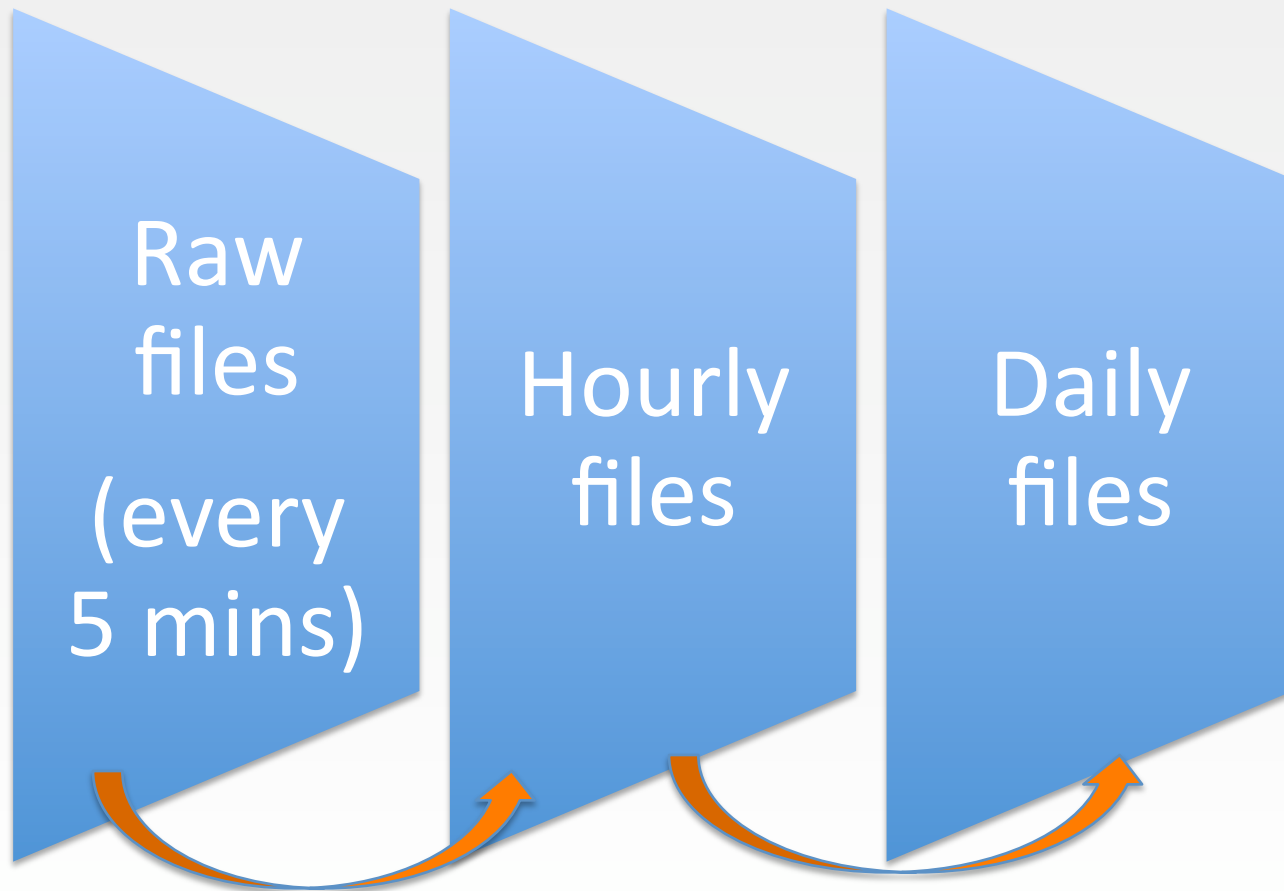
```
class ChukwaRecordKey
{
  ustring recordType;
  ustring key;
}
```

Value:

```
class ChukwaRecord
{
  long time;
  map<ustring, buffer> mapFields;
}
```

Directories structure & File rotation

File rotation for Archive and ChukwaRecords



At the end of the hour

Merge all raw files for a
datasource to an hourly
file

At the end of the day

Merge all hourly files for
a datasource to a daily
file