.conf2013

YOUR DATA
NO LIMITS

Technical Deep Dive: Hunk: Splunk Analytics for Hadoop Beta

Ledion Bitincka

Splunk Inc

splunk>

Legal Notices

During the course of this presentation, we may make forward-looking statements regarding future events or the expected performance of the company. We caution you that such statements reflect our current expectations and estimates based on factors currently known to us and that actual events or results could differ materially. For important factors that may cause actual results to differ from those contained in our forward-looking statements, please review our filings with the SEC. The forward-looking statements made in this presentation are being made as of the time and date of its live presentation. If reviewed after its live presentation, this presentation may not contain current or accurate information. We do not assume any obligation to update any forward-looking statements we may make. In addition, any information about our roadmap outlines our general product direction and is subject to change at any time without notice. It is for informational purposes only and shall not, be incorporated into any contract or other commitment. Splunk undertakes no obligation either to develop the features or functionality described or to include any such feature or functionality in a future release.

About Me

- Sr Architect
- 6+ years at Splunk
- Mainly involved in search time stuff, like:
  - Key-value pair extraction
  - Scheduler & Alerting
  - Transactions, eventtypes , tags etc
  - MySQLConnect, HadoopConnect
  - Hunk
- @ledbit

The Problem

- Large amounts of data in Hadoop
  - Relatively easy to get the data in
  - Hard & time-consuming to get **value** out

- Splunk has solved this problem before
  - Primarily for event timeseries

- Wouldn't if be great if Splunk could be used to analyze Hadoop data?

Hadoop + Splunk = Hunk

The Goals

- A viable solution must:
  - Process the data in place
  - Maintain support for Splunk Processing Language (SPL)
  - True schema on read
  - Report previews
  - Ease of setup & use

Support SPL

- Naturally suitable for MapReduce
- Reduces adoption time
- Challenge: Hadoop “apps” written in Java & all SPL code is in C++
- Porting SPL to Java would be a daunting task
- Reuse the C++ code somehow
  - use “splunkd” (the binary) to process the data
  - JNI is not easy nor stable

Schema on Read

- Apply Splunk's index-time schema at search time
  - Event breaking, time stamping etc.
- Anything else would be brittle & maintenance nightmare
- Extremely flexible
- Runtime overhead (manpower >>$ computation)
- Challenge: Hadoop "apps" written in Java & all index-time schema logic is implemnted in C++

Previews

- No one likes to stare at a blank screen!
- Challenge: Hadoop is designed for batch-like jobs

Ease of Setup & Use

- Users should just specify:
  - Hadoop cluster they want to use
  - Data within the cluster they want to process
- Immediately be able to explore & analyze their data

Deployment Overview

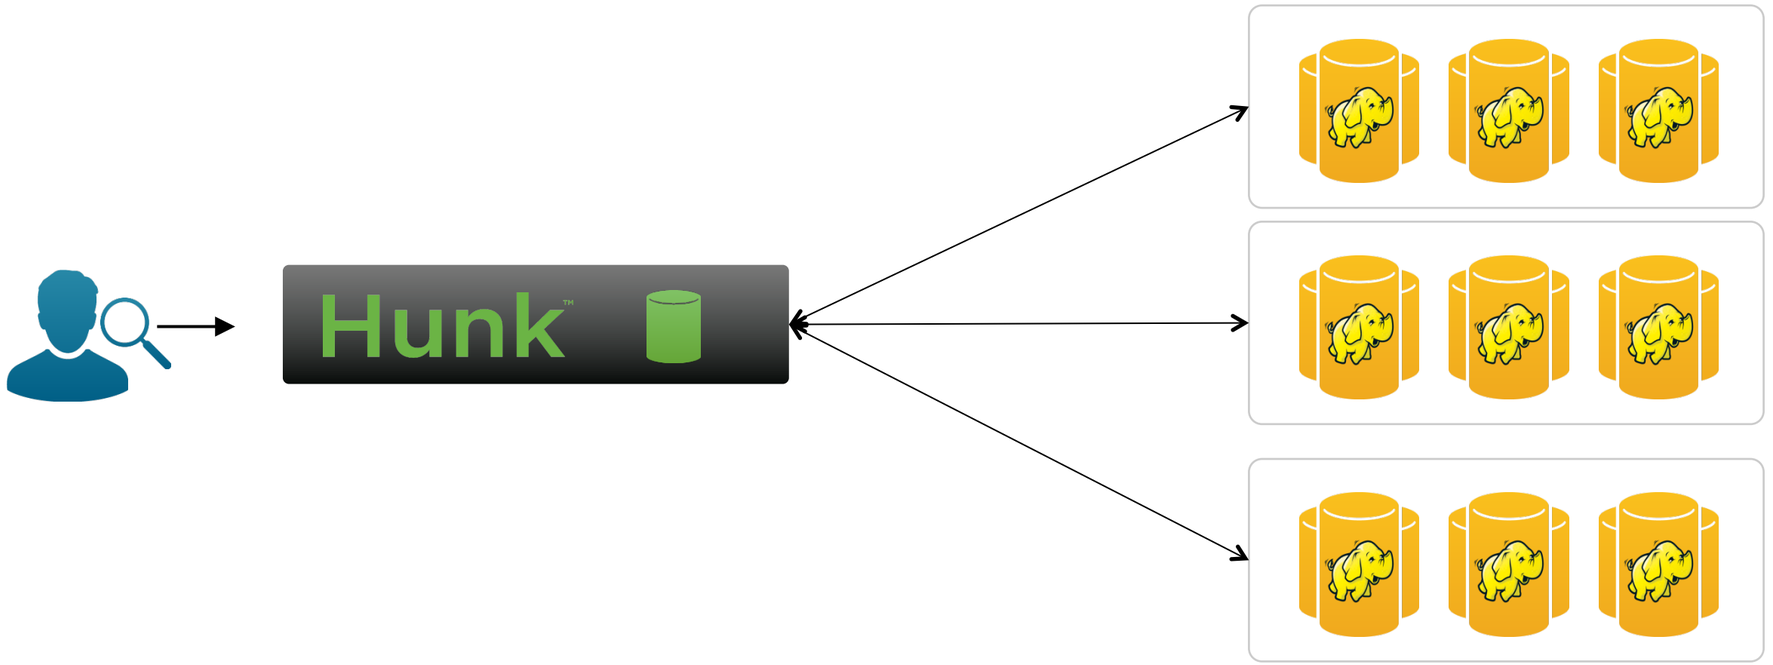

Move Data to Computation (stream)

- Move data from HDFS to SH
- Process it in a streaming fashion
- Visualize the results

- Problem?

Move Computation to Data (MR)

- Create and start a MapReduce job to do the processing
- Monitor MR job & collect its results
- Merge the results and visualize

- Problem?

Mixed Mode

- Use **both** computation models concurrently

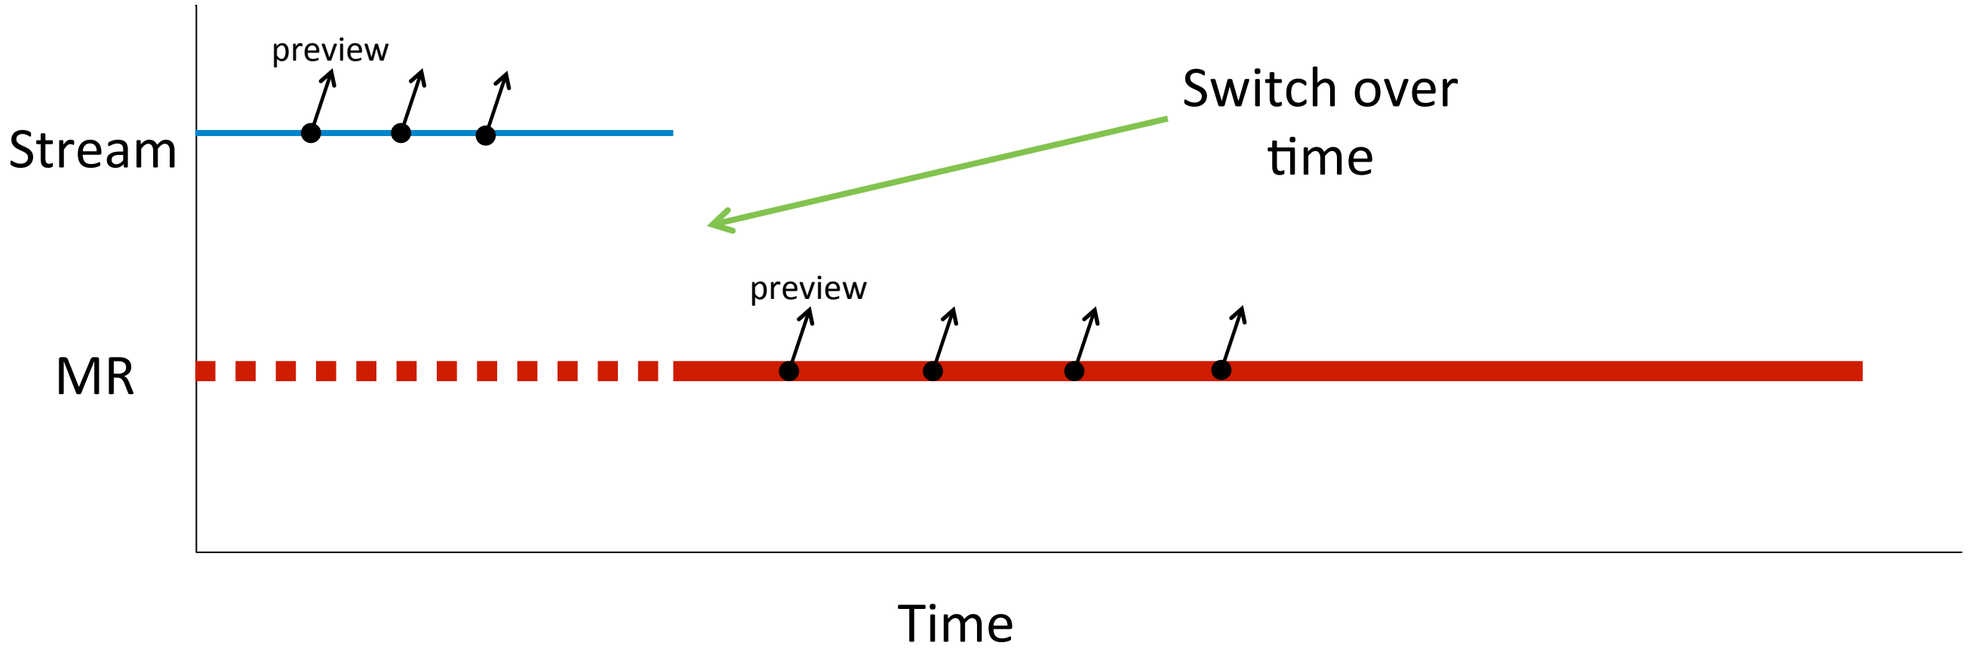

First Search Setup

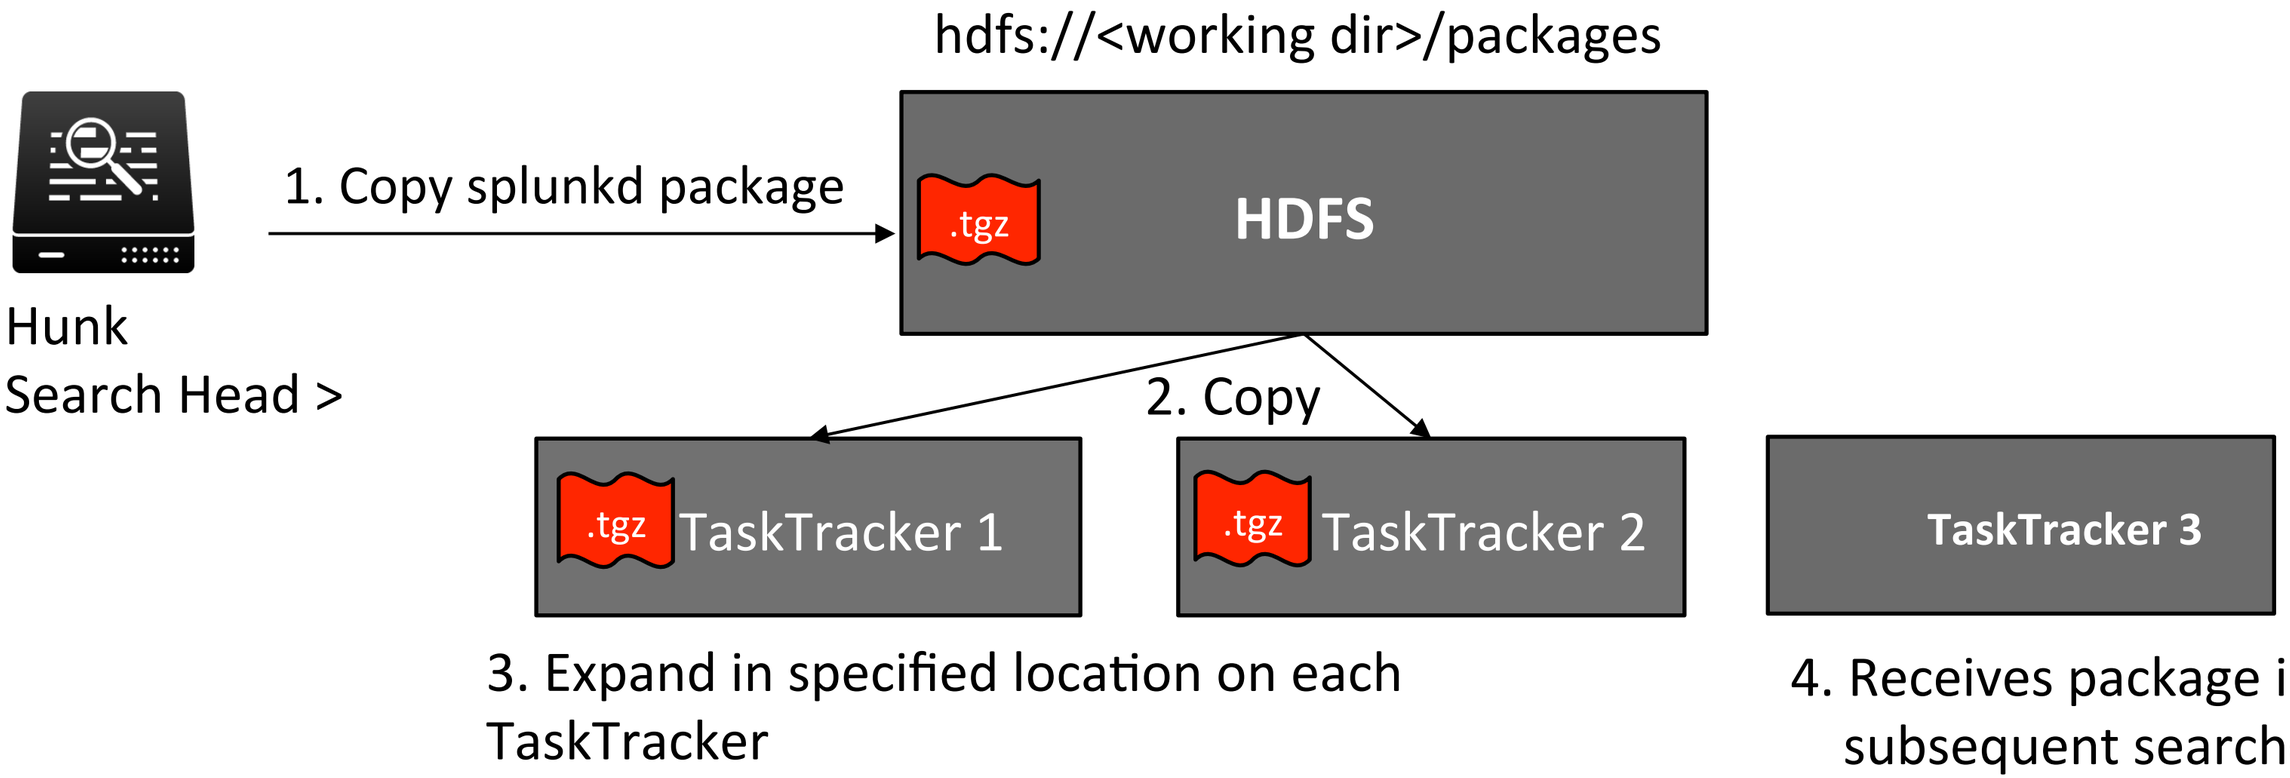

Data Flow

Streaming Data Flow

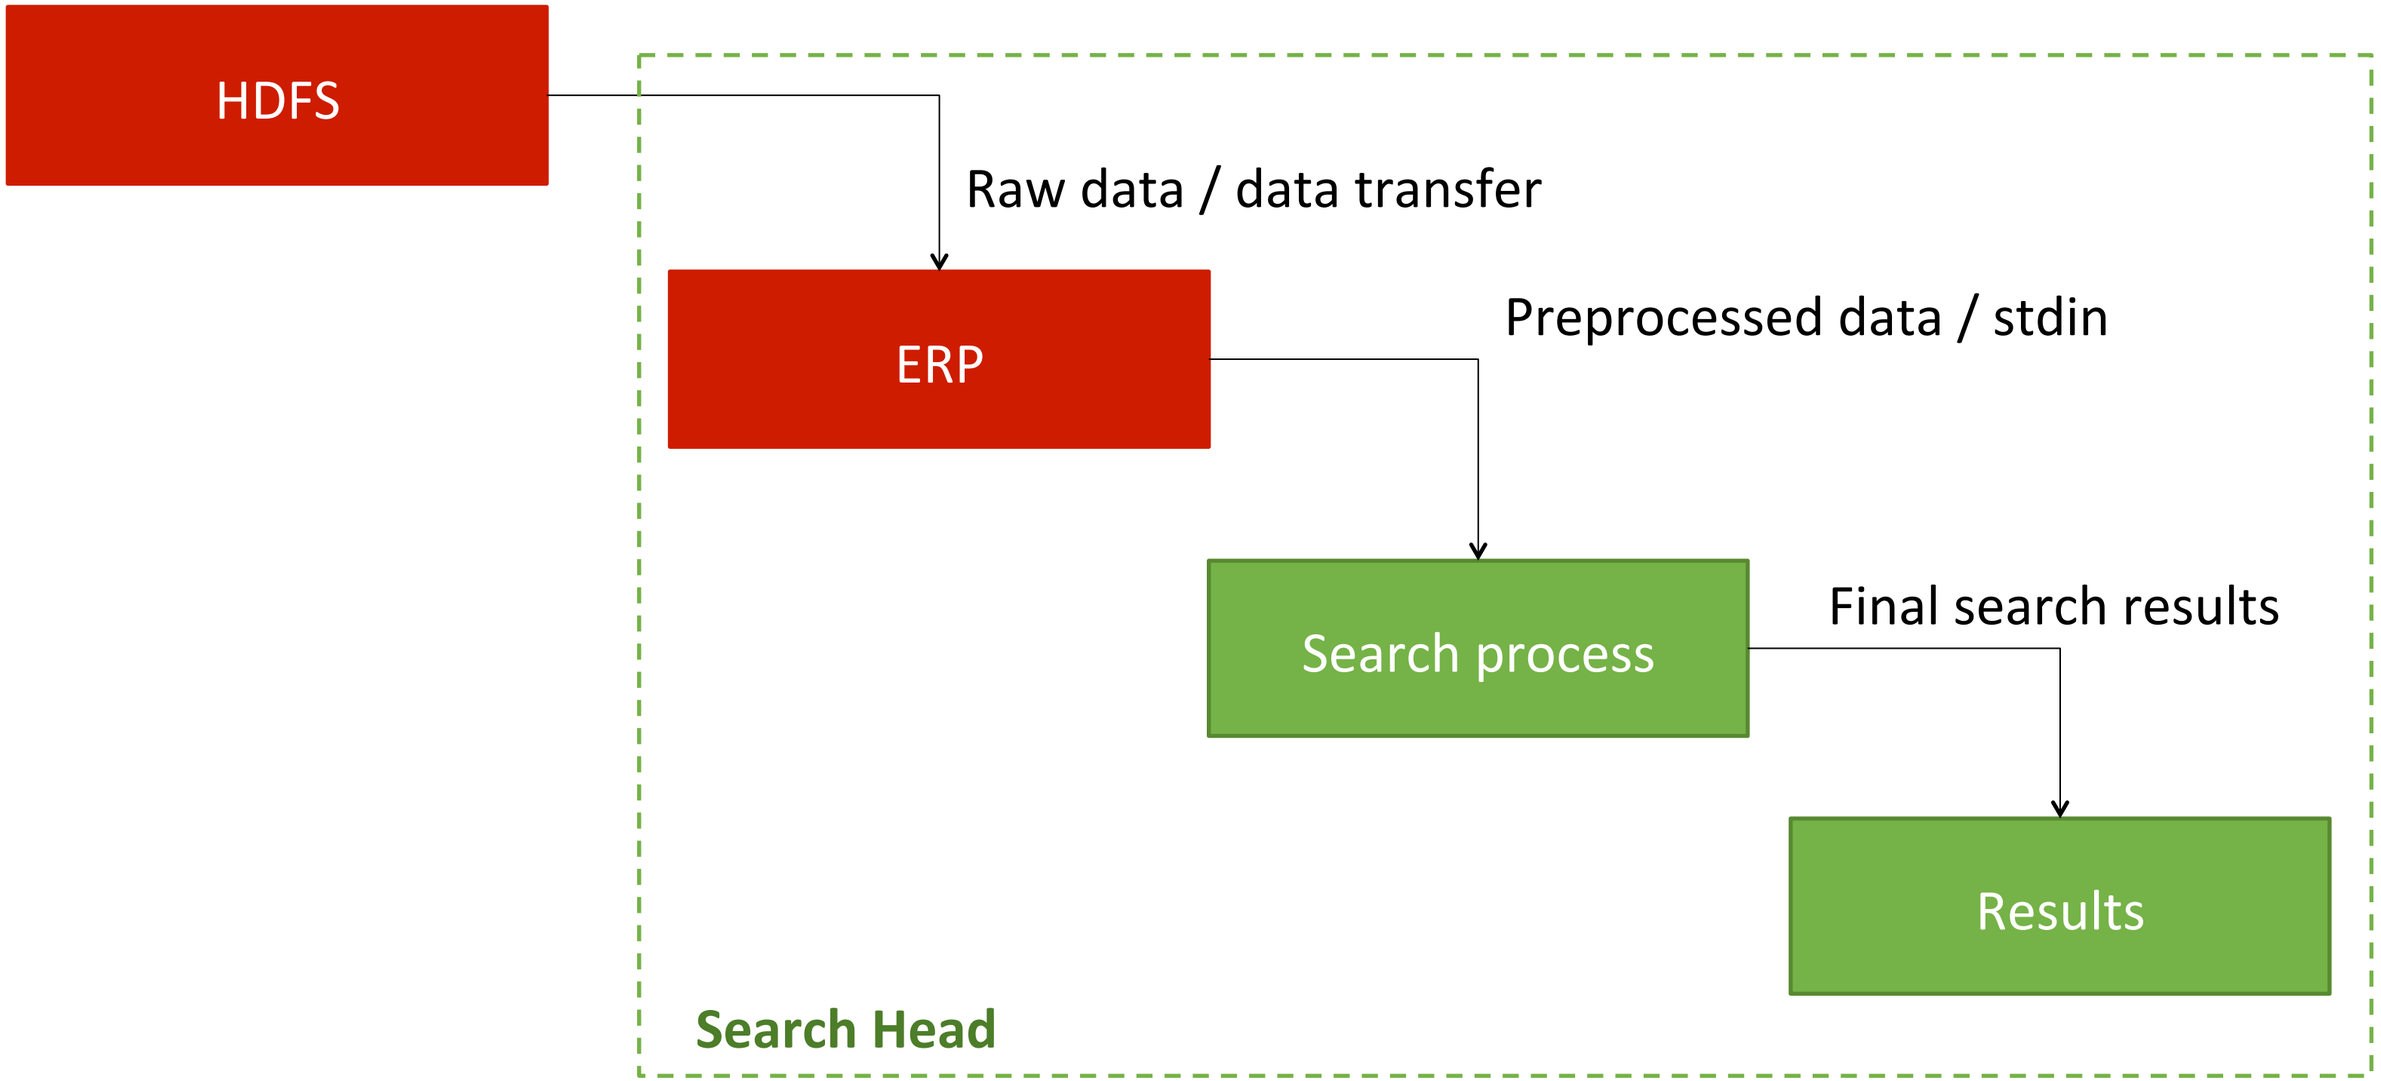

Reporting Data Flow

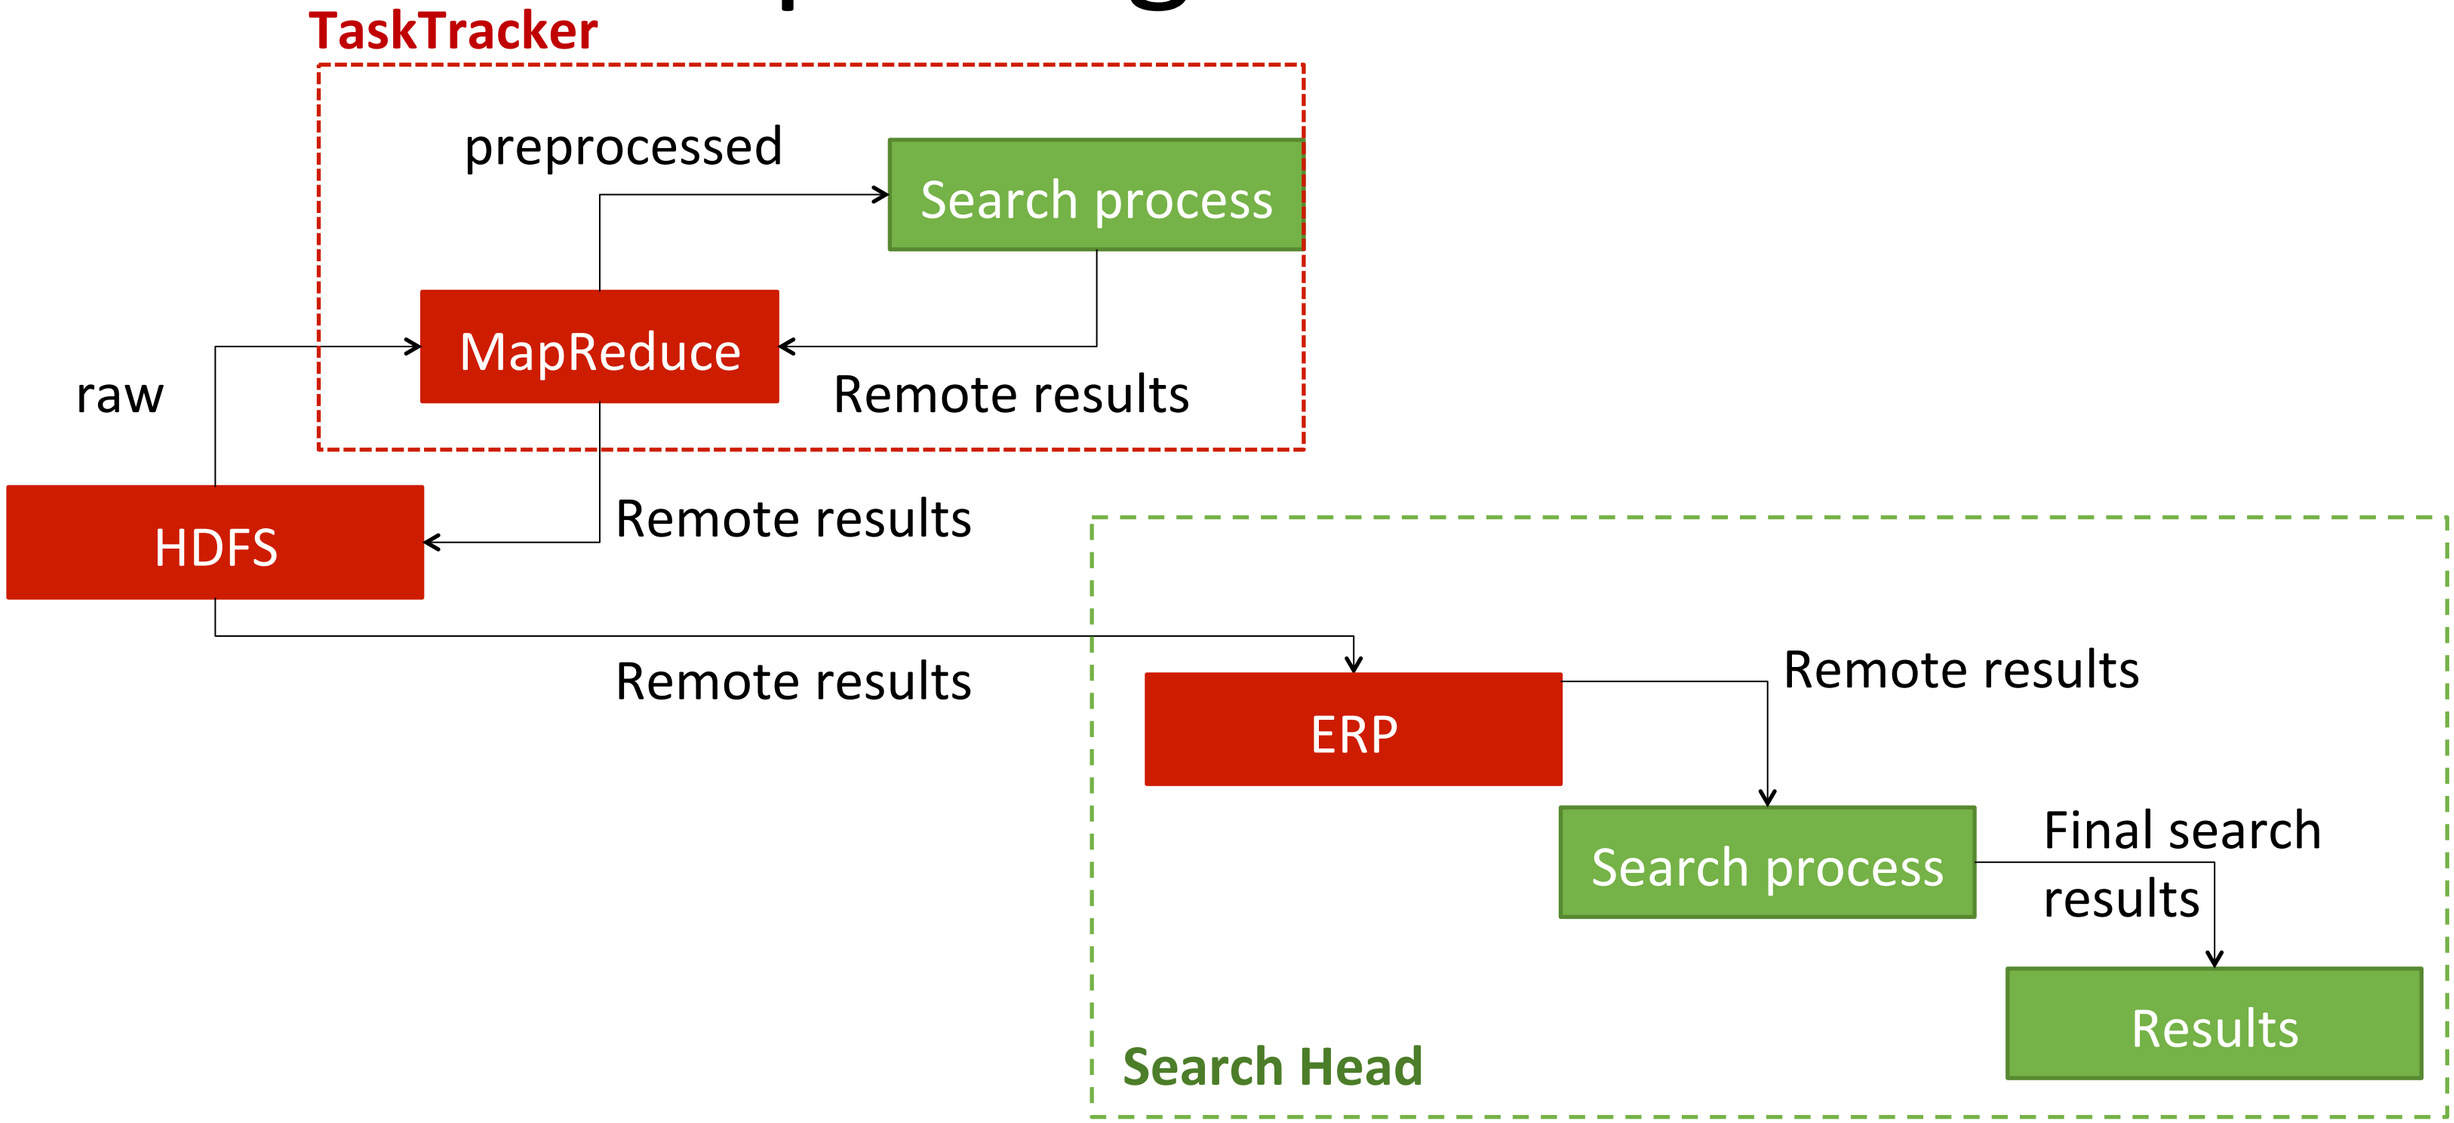

Data Processing

DataNode/TaskTracker

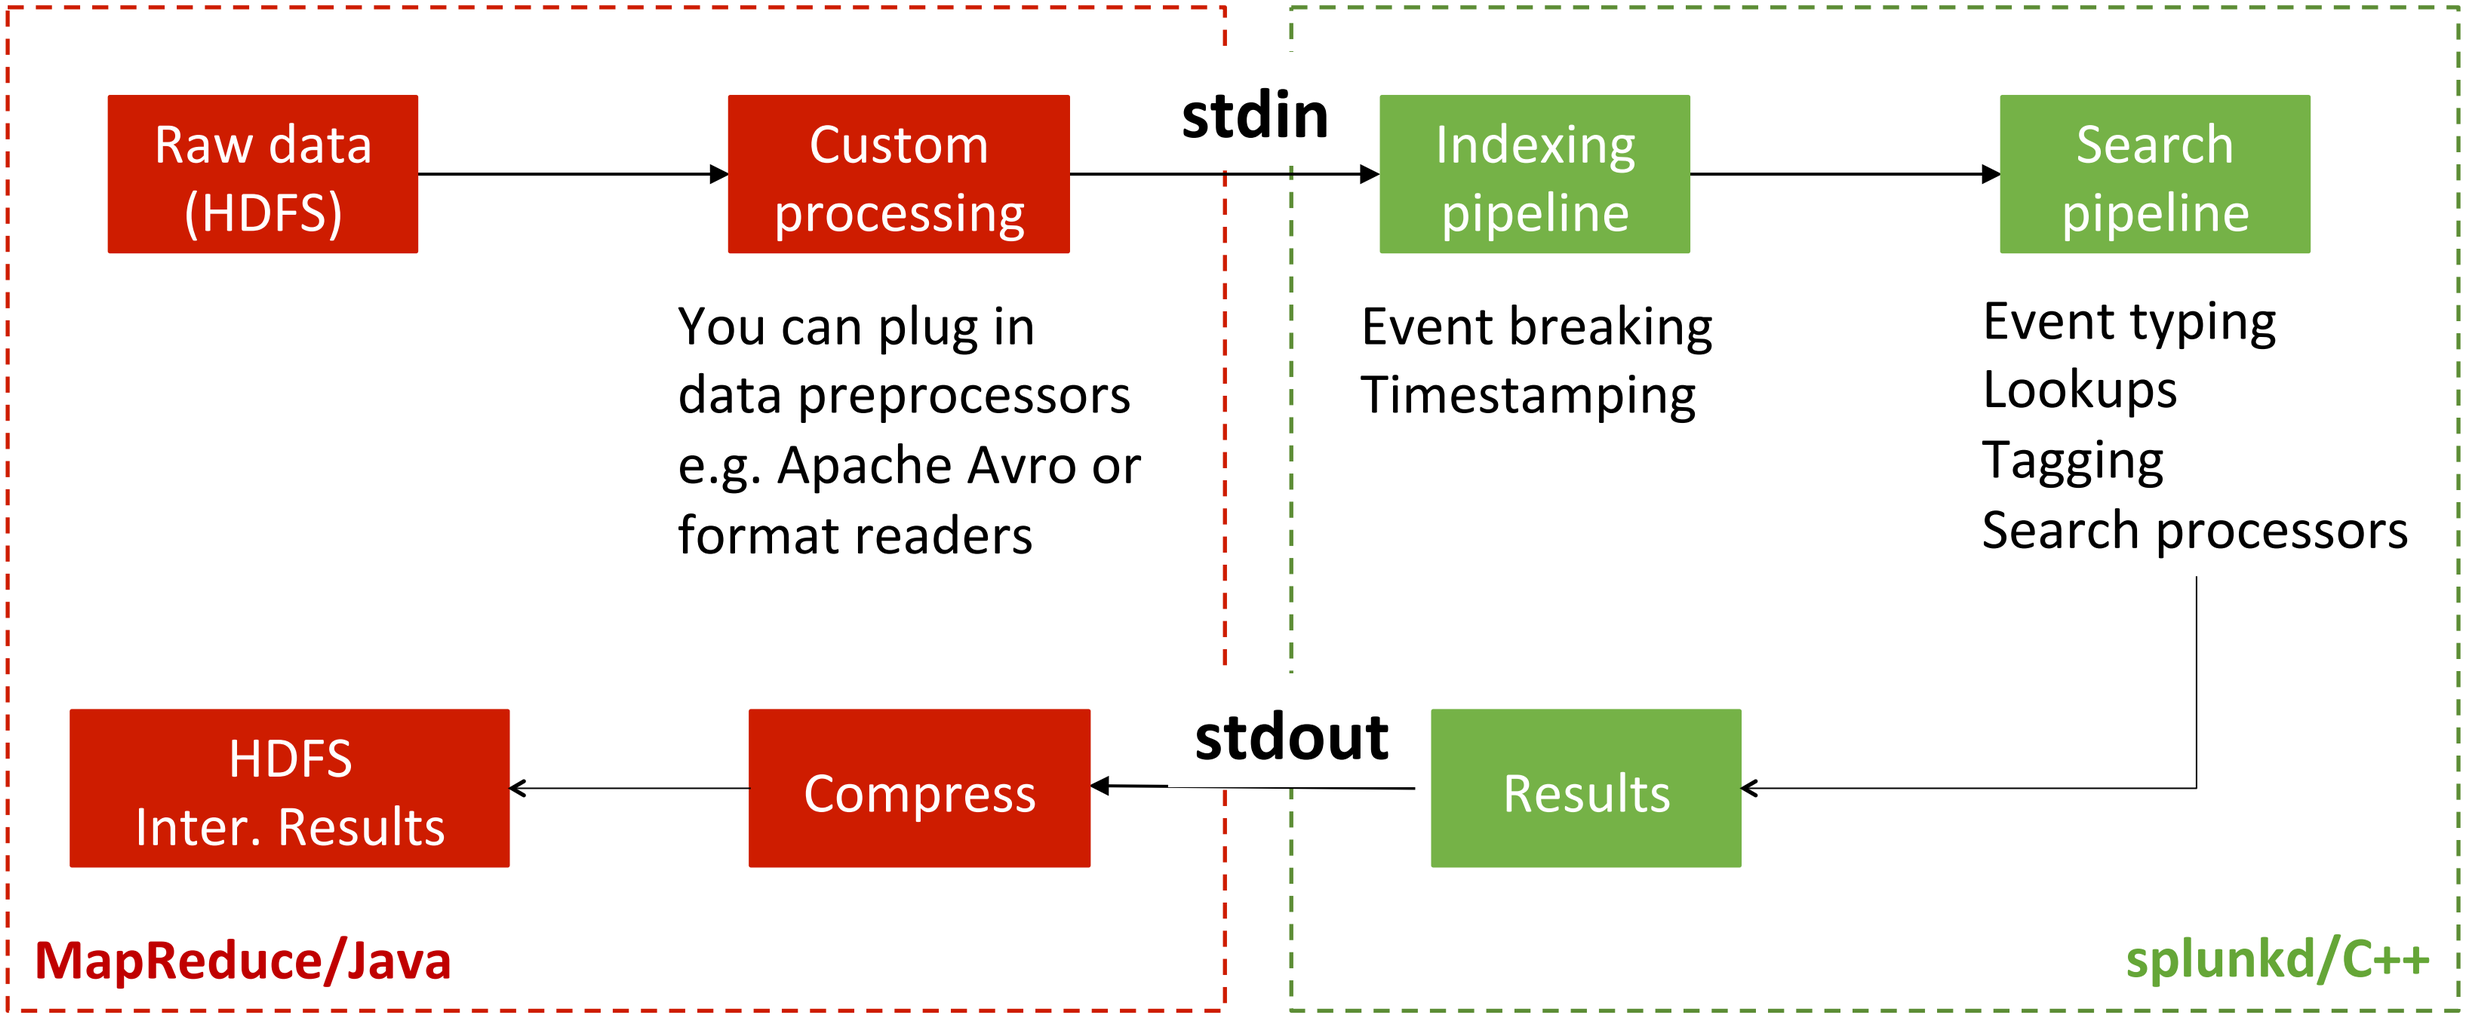

Schematization

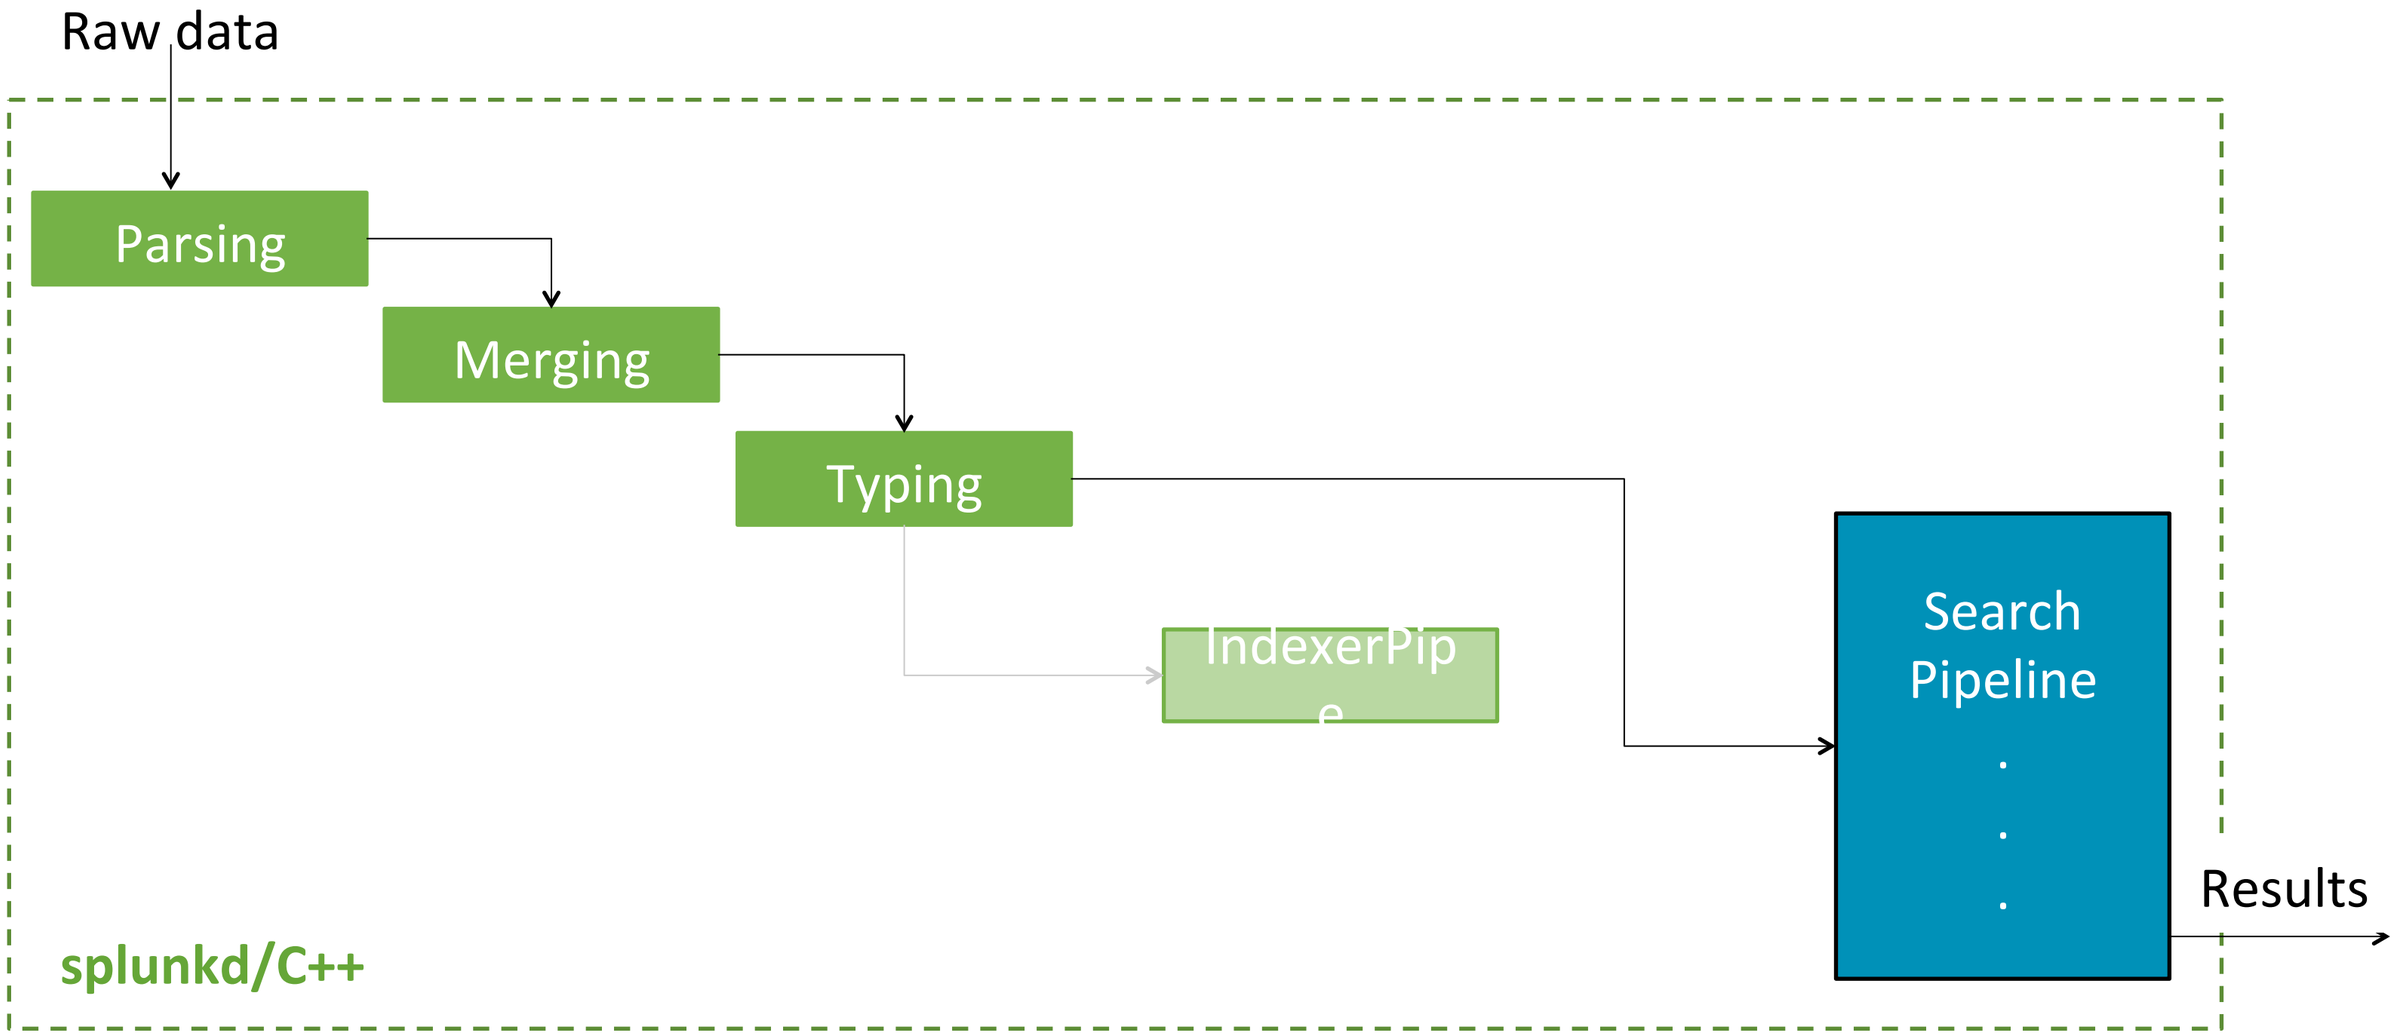

Search – Search Head

- Responsible for:
  - Orchestrating everything
  - Submitting MR jobs (optionally splitting bigger jobs into smaller ones)
  - Merging the results of MR jobs
    - Potentially with results from other VIXes or native indexes
  - Handling high level optimizations

Partition Pruning

- Data is usually organized into hierarchical dirs, eg. /<base_path>/<date>/<hour>/<hostname>/somefile.log
- Hunk can be instructed to extract fields and time ranges from a path
- Ignores directories that cannot possibly contain search results

Partition Pruning, e.g.

Paths in a VIX:
/home/hunk/20130610/01/host1/access_combined.log
/home/hunk/20130610/02/host1/access_combined.log
/home/hunk/20130610/01/host2/access_combined.log
/home/hunk/20130610/02/host2/access_combined.log

Search: `index=hunk` **`server=host1`**

Paths searched:
/home/hunk/20130610/01/**host1**/access_combined.log
/home/hunk/20130610/02/**host1**/access_combined.log

Partition Pruning, e.g.

Paths in a VIX:
/home/hunk/20130610/01/host1/access_combined.log
/home/hunk/20130610/02/host1/access_combined.log
/home/hunk/20130610/01/host2/access_combined.log
/home/hunk/20130610/02/host2/access_combined.log

Search: `index=hunk` **`earliest_time="2013-06-10T01:00:00" latest_time="2013-06-10T02:00:00"`**

Paths searched:
/home/hunk/**20130610/01**/host1/access_combined.log
/home/hunk/**20130610/01**/host2/access_combined.log

Best Practices

- Partition data in FS using fields that:
  - Are commonly used
  - Relatively low cardinality
- For new data, use formats that are well defined, e.g.
  - Avro, json etc
  - Avoid columnar formats, like csv/tsv (hard to split)
- Use compression, gzip, snappy etc
  - I/O becomes a bottleneck at scale

Troubleshooting

- Search.log is your friend !!!
- Log lines annotated with ERP.<name> …
- Links for spawned MR job(s)
- Follow these links to troubleshoot MR issues
- hdfs://<base_path>/dispatch/<sid>/<num>/<dispatch_dirs> contains the dispatch dir content of searches ran on TaskTracker

Job Inspector

31.435	erp.charlie.MR	17	60	60
31.435	erp.charlie.MR.SPLK_ronnie.sv.splunk.com_1374191982.3_0	17	60	60
6.224	erp.charlie.report.bytes	58	55,260	174,669
43.836	erp.charlie.report.delay	1	-	-
0.288	erp.charlie.setup	1	-	-
0.268	erp.charlie.setup.bundles	1	-	-
0.004	erp.charlie.setup.splunk	1	-	-
42.751	erp.charlie.stream.bytes	2	67,732,337	575,250,866
1.045	erp.charlie.stream.delay	1	-	-
0	erp.charlie.vix.hunk.files.filter.search	120	-	-
0	erp.charlie.vix.hunk.files.filter.time	1,095	-	-
1.923	erp.charlie.vix.hunk.files.listed	1,275	-	-

Common Problems

- User running Splunk does not have permission to write to HDFS or run MapReduce jobs
- HDFS SPLUNK_HOME not writable
- DN/TT SPLUNK_HOME not writable, out of disk
- Data reading permission issues

Demo

Helpful Resources

- Download
  - http://www.splunk.com/bigdata
- Help & Docs
  - http://docs.splunk.com/Documentation/Hunk/6.0/Hunk/MeetHunk
- Resource
  - http://answers.splunk.com

Next Steps

1. **Download the .conf2013 Mobile App**
   If not iPhone, iPad or Android, use the Web App

2. Take the survey & **WIN A PASS FOR .CONF2014**… Or one of these bags!

3. **View the other “Using” sessions**
   All sessions are available on the Mobile App
   Videos will be available shortly


.conf2013
YOUR DATA
NO LIMITS
THANK YOU
splunk>