


GridMix3



WBDB 2012
May 8, 2012

Srigurunath Chakravarthi
Director, Engineering
Yahoo! Inc.
sriguru@yahoo-inc.com

What is GridMix3?

1. A workload emulation benchmark

- Uses Hadoop *JobHistory* logs as source of workload info.
- Re-plays jobs from traces on small-sized test-bed cluster using dummy data.
- Mimics workload on a real Hadoop cluster using real hardware/software.

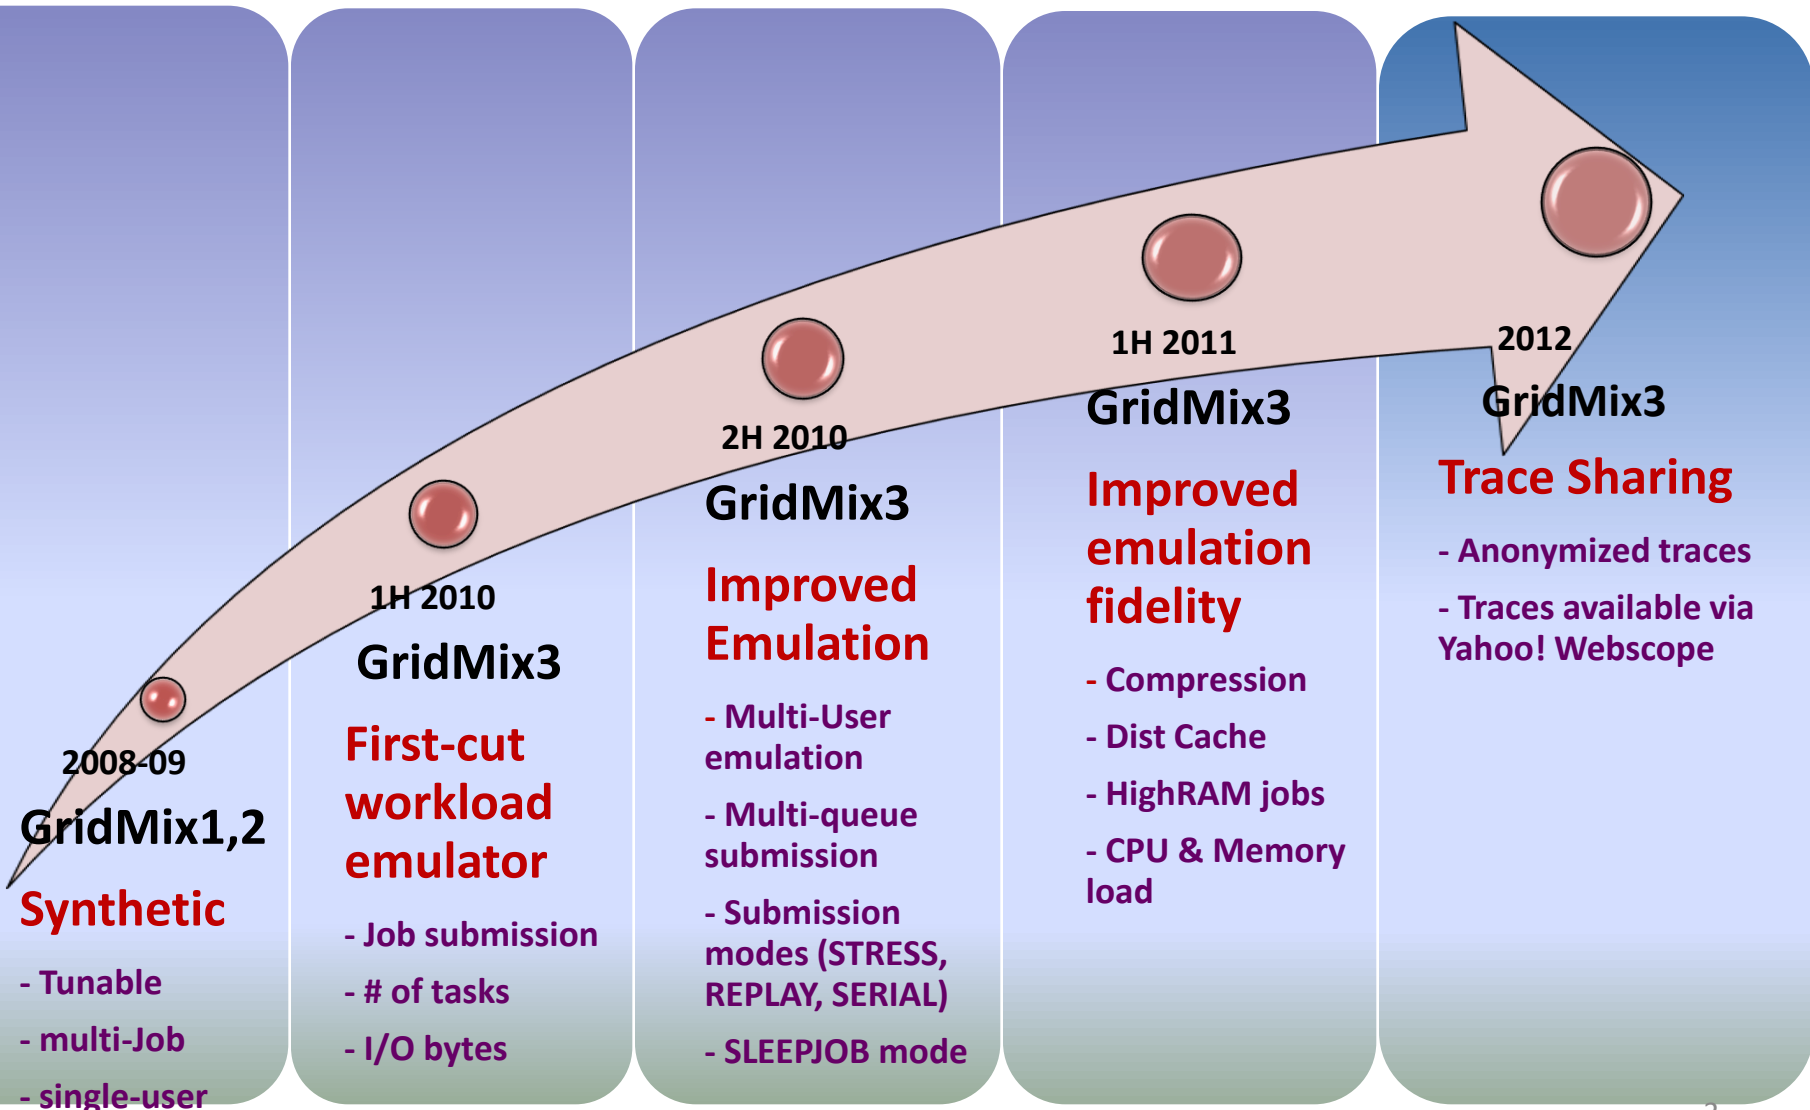
2. Mimics a bunch of things

- *Matches* number of tasks (Maps & Reduces).
- *Emulates* resource usage (Disk, Network, CPU, Memory).
- *Supports* multiple users. Submits to multiple queues.
- *Control* over job submission rate. Three modes: STRESS, REPLAY & SERIAL.
- *Supports* core Hadoop features like *HighRAM*, *Dist. Cache*, etc.

3. Value proposition

- Study workload behavior using a *small test-bed* (10x smaller than original cluster).
- Benchmark influenced by real-world apps (via traces).
- Available as open source in Apache Hadoop.

GridMix Evolution



Benchmarking Methodology Proposal

- ✓ Contribute your workload trace
 - ① Run your Hadoop workload(s) in a “production” environment.
 - ② Collect *JobHistory* logs.
 - ③ Optionally scale & filter to produce *workload trace file*.
 - ④ Anonymize via *Rumen* and share with community.

- ✓ Workgroup constructs a set of “golden traces” representing workloads across industry & academia.
- ✓ Optionally publish base lines on popular system configs.

- ✓ For each system configuration of interest, anyone can
 - ① Replay trace file via GridMix3 on a small test-bed cluster (10x smaller than original cluster *or* 10s of nodes).
 - ② Observe execution metrics.
 - ③ Compare execution metrics across system configs and/or well known baselines to benchmark your system.

Use Cases met by GridMix3



Technique:
Reproduce workload
characteristics on a smaller
controlled test-bed.

Quantify &
Analyze
Cluster
Performance

- ✓ Optimize Hadoop code
- ✓ Tune Hadoop Clusters

Reproduce
Scalability
Issues

- ✓ Analyze & Fix Hadoop
scalability issues

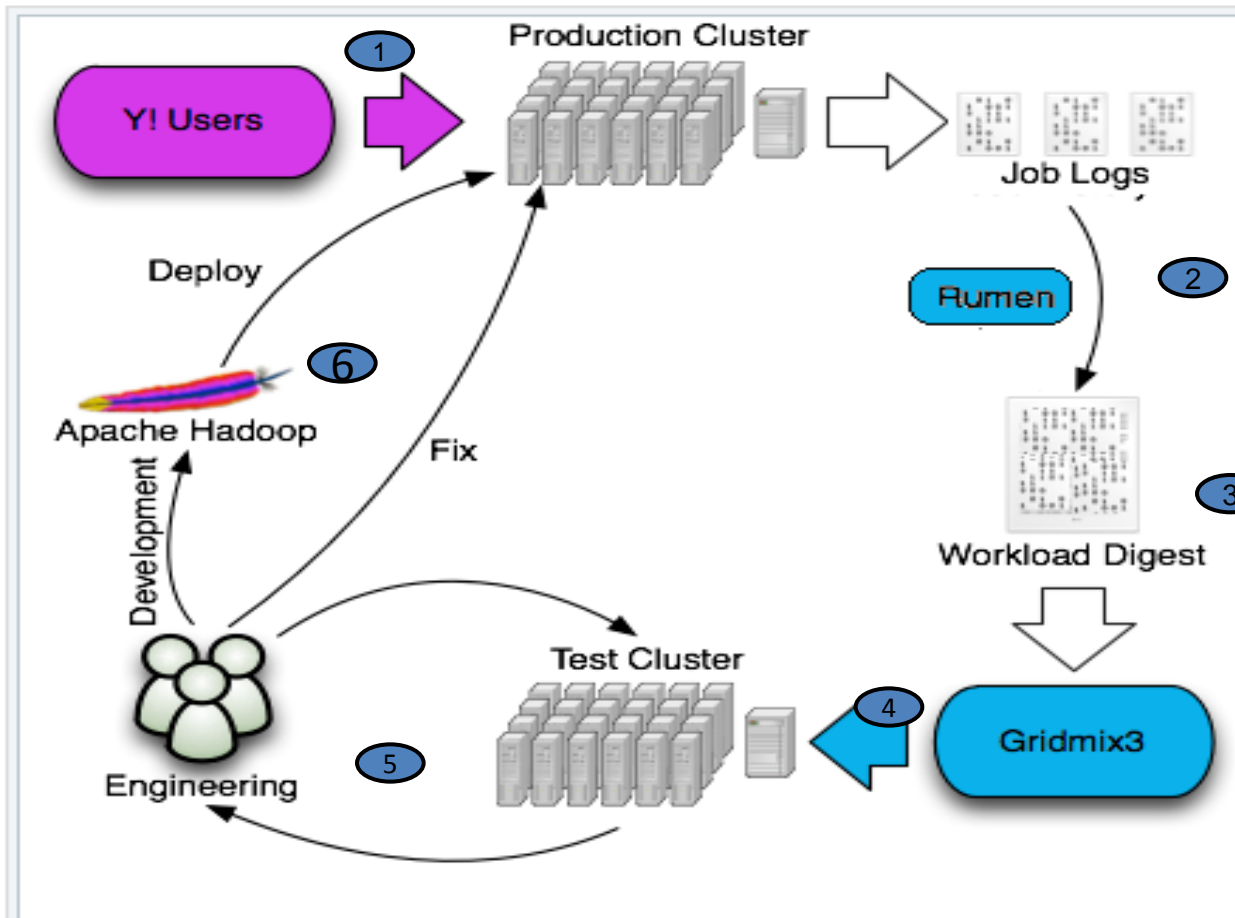
Certify
Software
Performance

- ✓ Preserve performance
across releases
- ✓ Detect performance
regressions nightly

Hardware
Selection

- ✓ Select Optimal Configuration
- ✓ Optimize Cluster Capacity

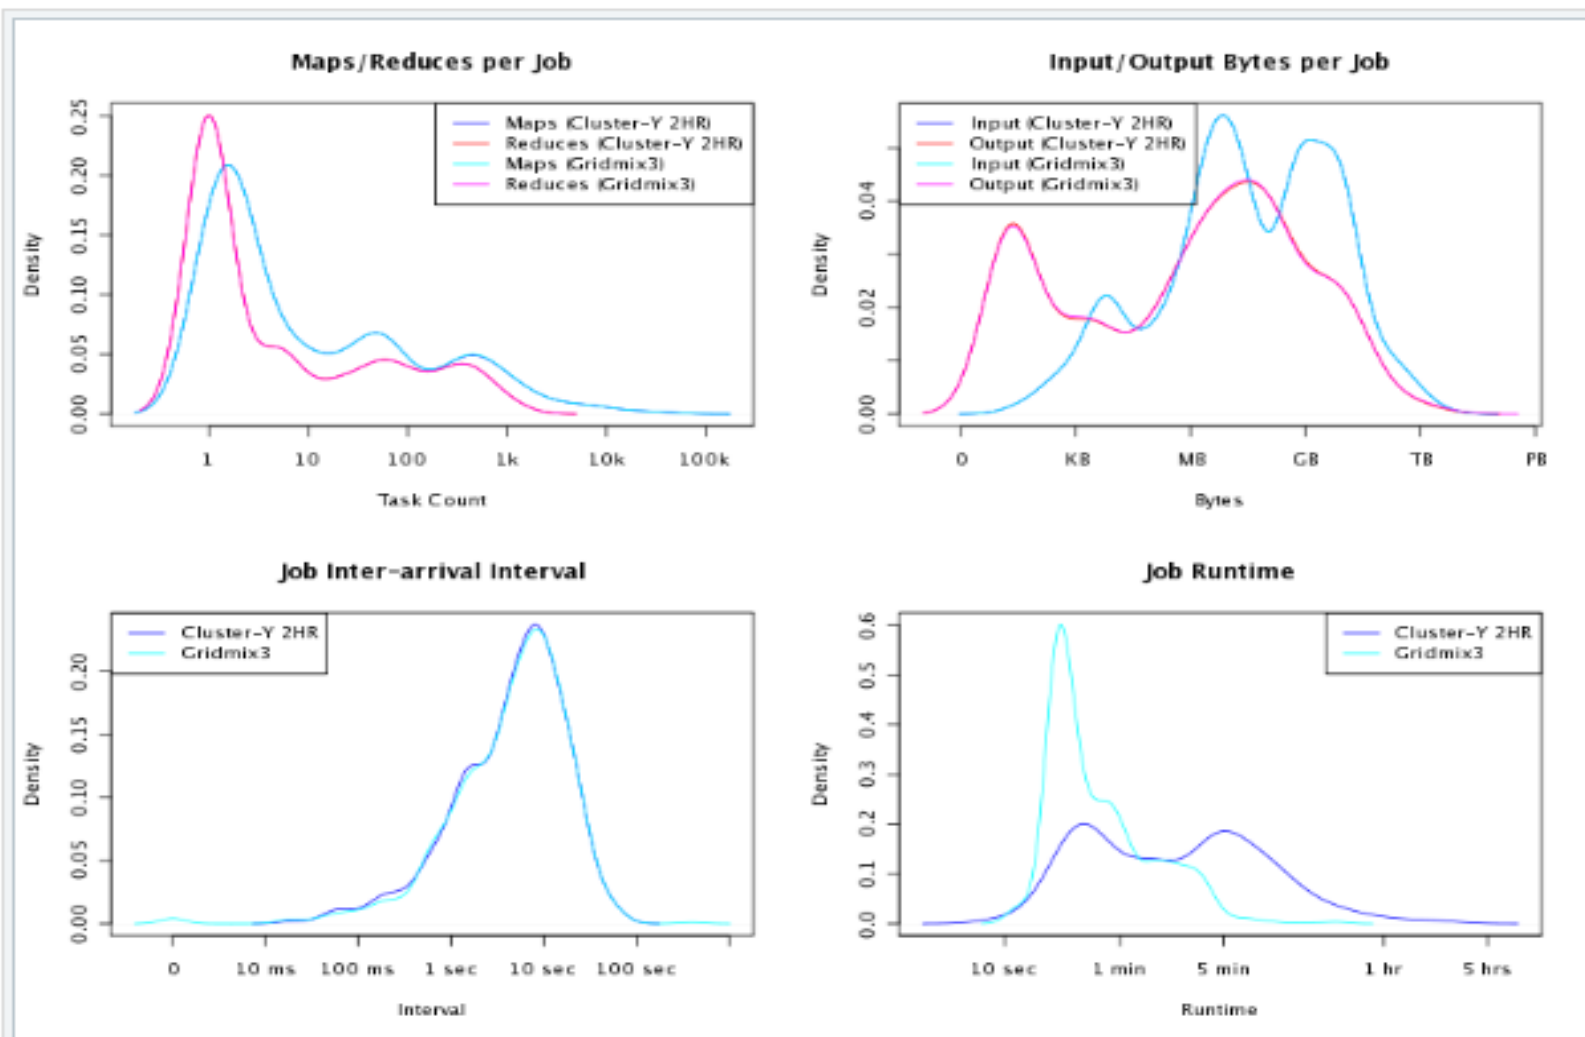
GridMix3 Usage Workflow: A closer look



- ① Run workload on production cluster
- ② The job trace is derived using *Rumen*.
- ③ *Rumen* provides basic means to adjust *density* of the trace to match test cluster size.
- ④ GridMix3 takes job trace as input and plays back on test cluster.
- ⑤ Engineers analyze results & iterate on test cluster with patches / fixes / tunables using GridMix3 .
- ⑥ Verified and validated changes are rolled out to Hadoop source and ultimately to Y! production clusters.

Benchmarking & Optimization Workflow with GridMix3

GridMix3 v/s Original – Sample workload characteristics



Comparing Workload Characteristics - Gridmix3 vs Production.

GridMix3 Source Base & Contributors

Where it lives

Gridmix3:

<http://svn.apache.org/repos/asf/hadoop/common/trunk/mapreduce/src/contrib/gridmix/>

Rumen:

<http://svn.apache.org/repos/asf/hadoop/common/trunk/mapreduce/src/tools/org/apache/hadoop/tools/rumen/>

To get workload traces from Yahoo!, contact:

yahoo-hadoop-traces@yahoo-inc.com

Yahoo! & Ex-Yahoo! Contributors to GridMix3 & Rumen

Amar Kamat

Chris Douglas

Dick King

Hong Tang

Rahul Singh

Ravi Gummadi

Vinay Thota

- END -