Status and Plans for **Applications** Monitoring

Henry Nebrensky (Brunel University)
CMS UK

Bristol, 1-2 July 2004

Some slides stolen from Claudio Grandi's talk at CHEP'03

BOSS

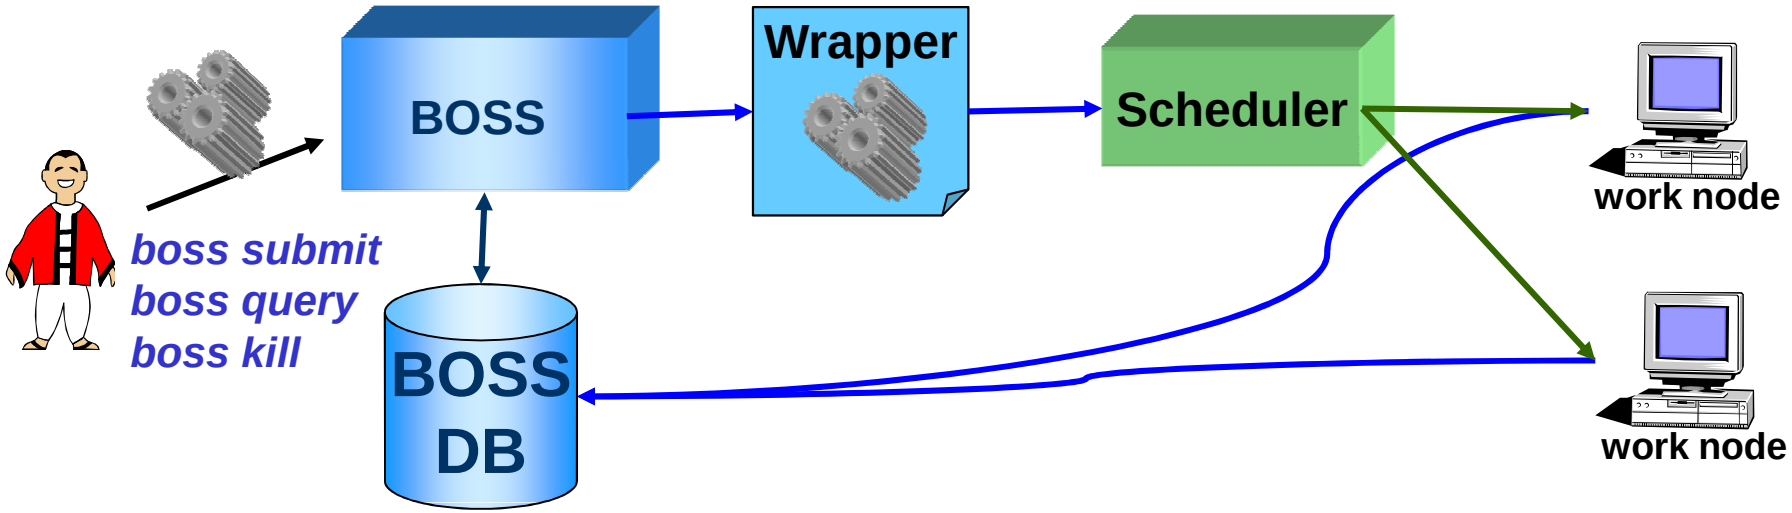

Accepts job submission from users

Stores info about job in a DB

Builds a wrapper around the job (*jobExecutor*)

Sends the wrapper to the local scheduler

The wrapper sends to the DB info about the job

BOSS on the Grid

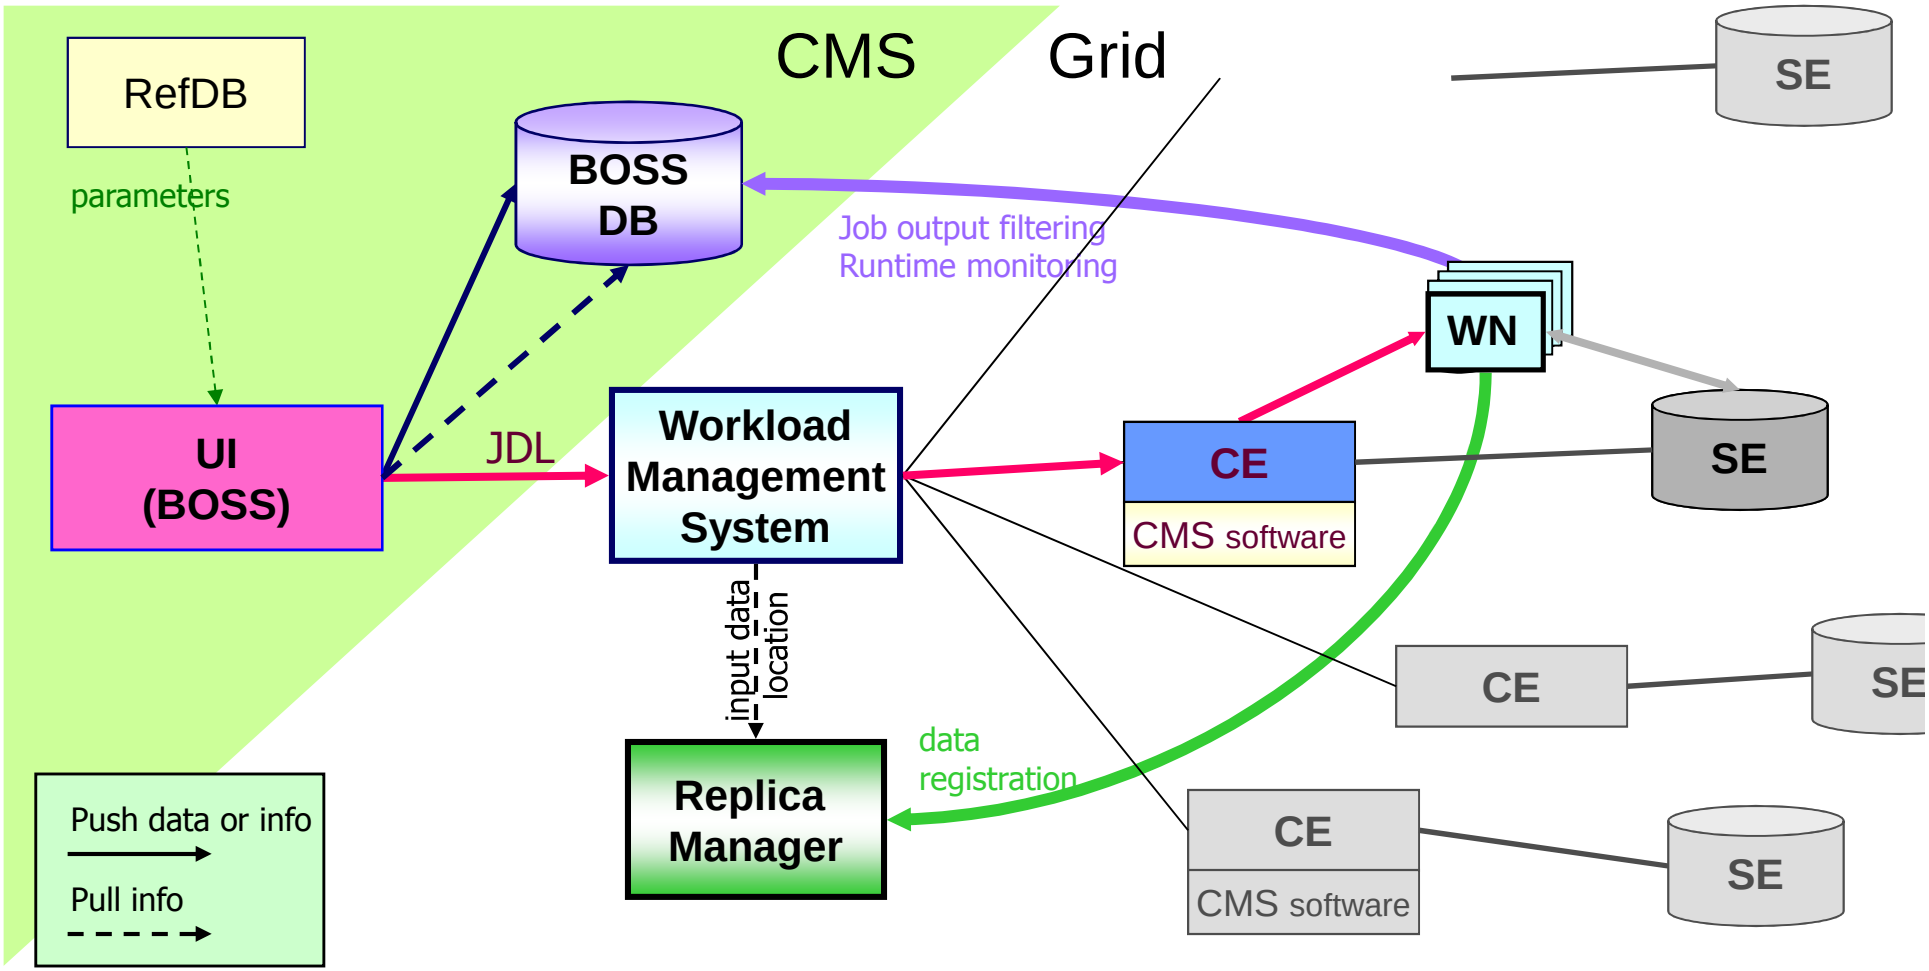

Use of R-GMA in BOSS (the idea)

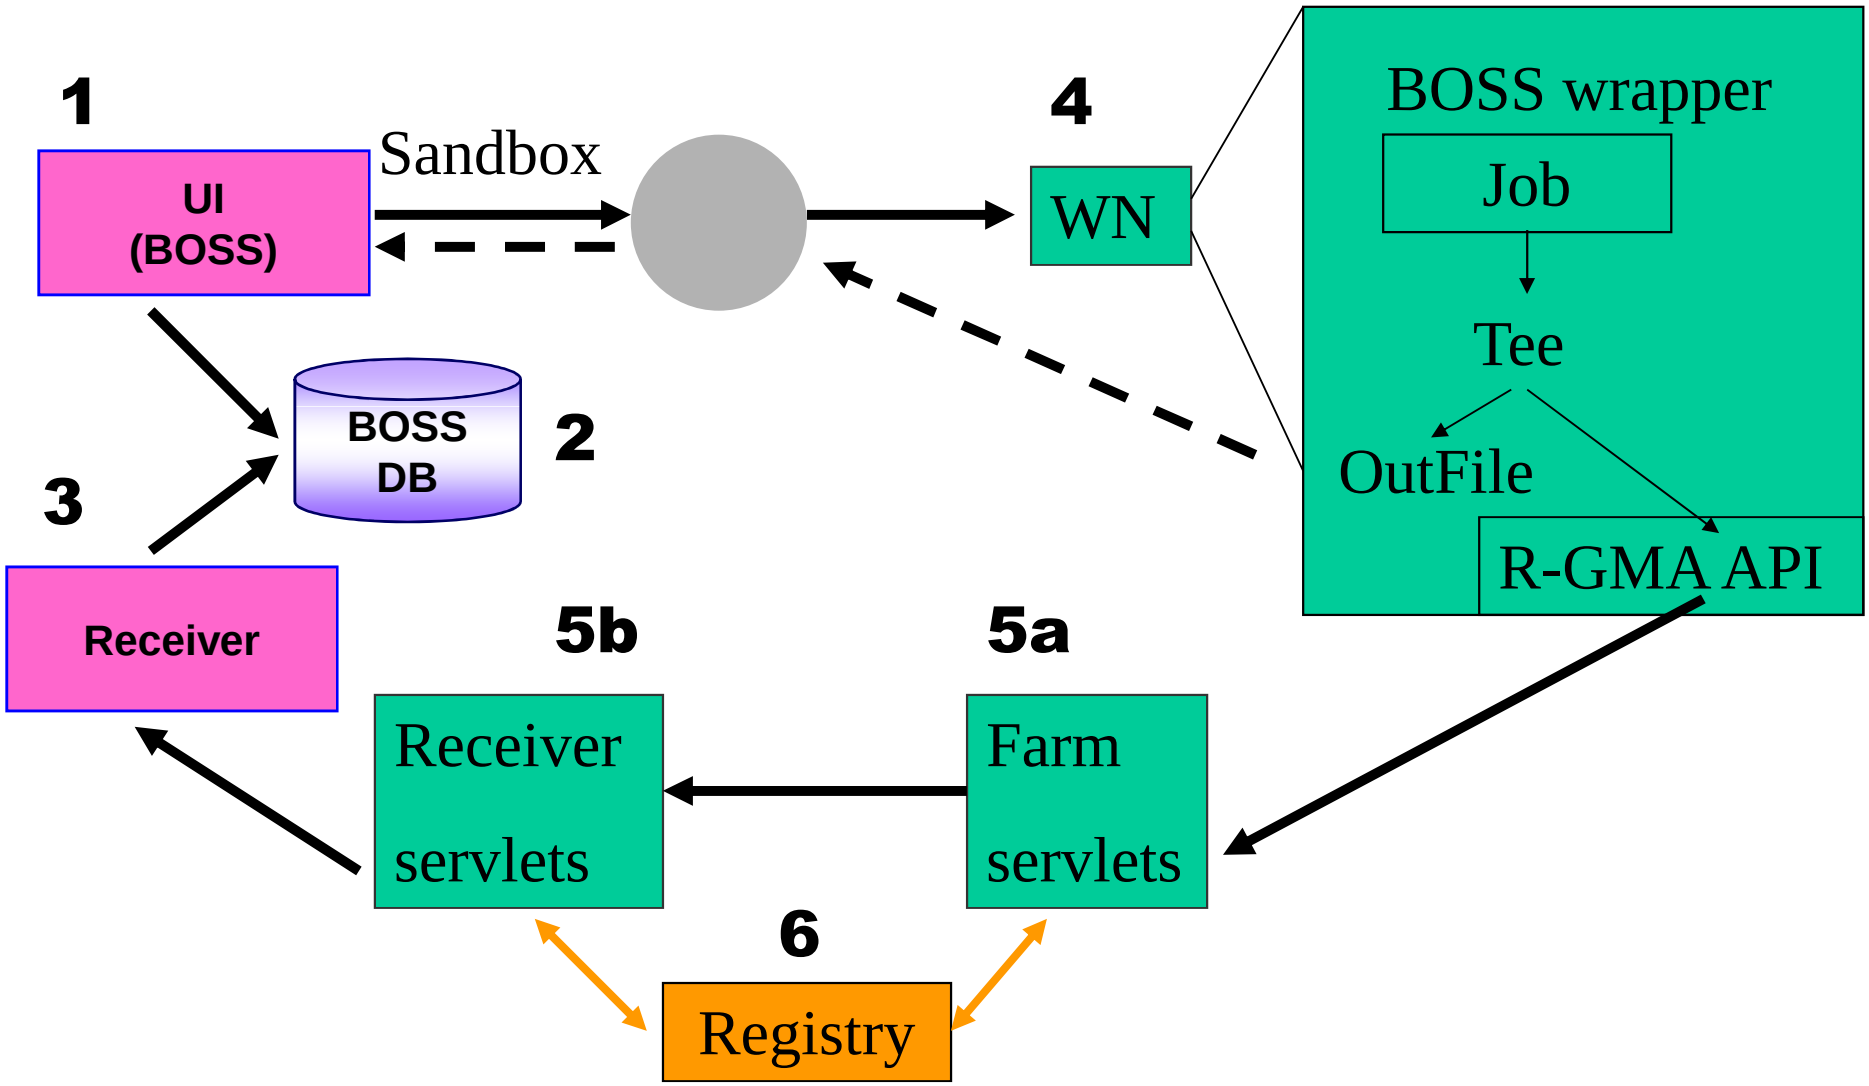

Use of R-GMA in BOSS

- Publish each update into R-GMA as a separate message – separate row
- Each producer gives host and name of “home” BOSS DB, and jobId; this identifies it uniquely
- Receiver looks for all rows relating to its DB; uses jobId and jobType to do SQL **UPDATE**

Use of R-GMA in BOSS

- R-GMA smoothes “firewall” issues
- Consumer can watch many producers; producers can feed multiple consumers.
- Provides uniform access to range of monitoring data (WP7 network, etc.)
- Can define minimum retention period but no guarantees

Scalability Tests With CMS, Boss and R-GMA

Stolen from Rob Byrom's slides at

http://agenda.cern.ch/fullAgenda.php?ida=a036755

(Presented at 2003 IEEE/NSS mtg, sub. to Trans. Nuc. Sci.)

Test Motivation

- Want to ensure R-GMA can cope with volume of expected traffic and is scalable.
- CMS production load estimated at around 2000 jobs.
- Initial tests with v3-3-28 only managed about 400 - could do better L .

Test Design

- A simulation of the CMS production system was created.
  - A Java MC simulation was designed to represent a typical job.
  - Each job creates a stream producer.
  - Each job publishes a number of tuples depending on the job phase.
  - Each job contains 3 phases with varying time delays.
  - Emulates “CMSIM” message publishing pattern, but so far with 10 hour run time compressed into a minute …
  - … so actually have fewer simultaneous jobs than real case, but overall a much higher rate of message production.

Test Design

- An Archiver scoops up published tuples.
  - The Archiver db used is a representation of the BOSS db, but stores history of received messages, rather than just a cumulative update.
  - Archived tuples are compared with published tuples to verify the test outcome.

Topology

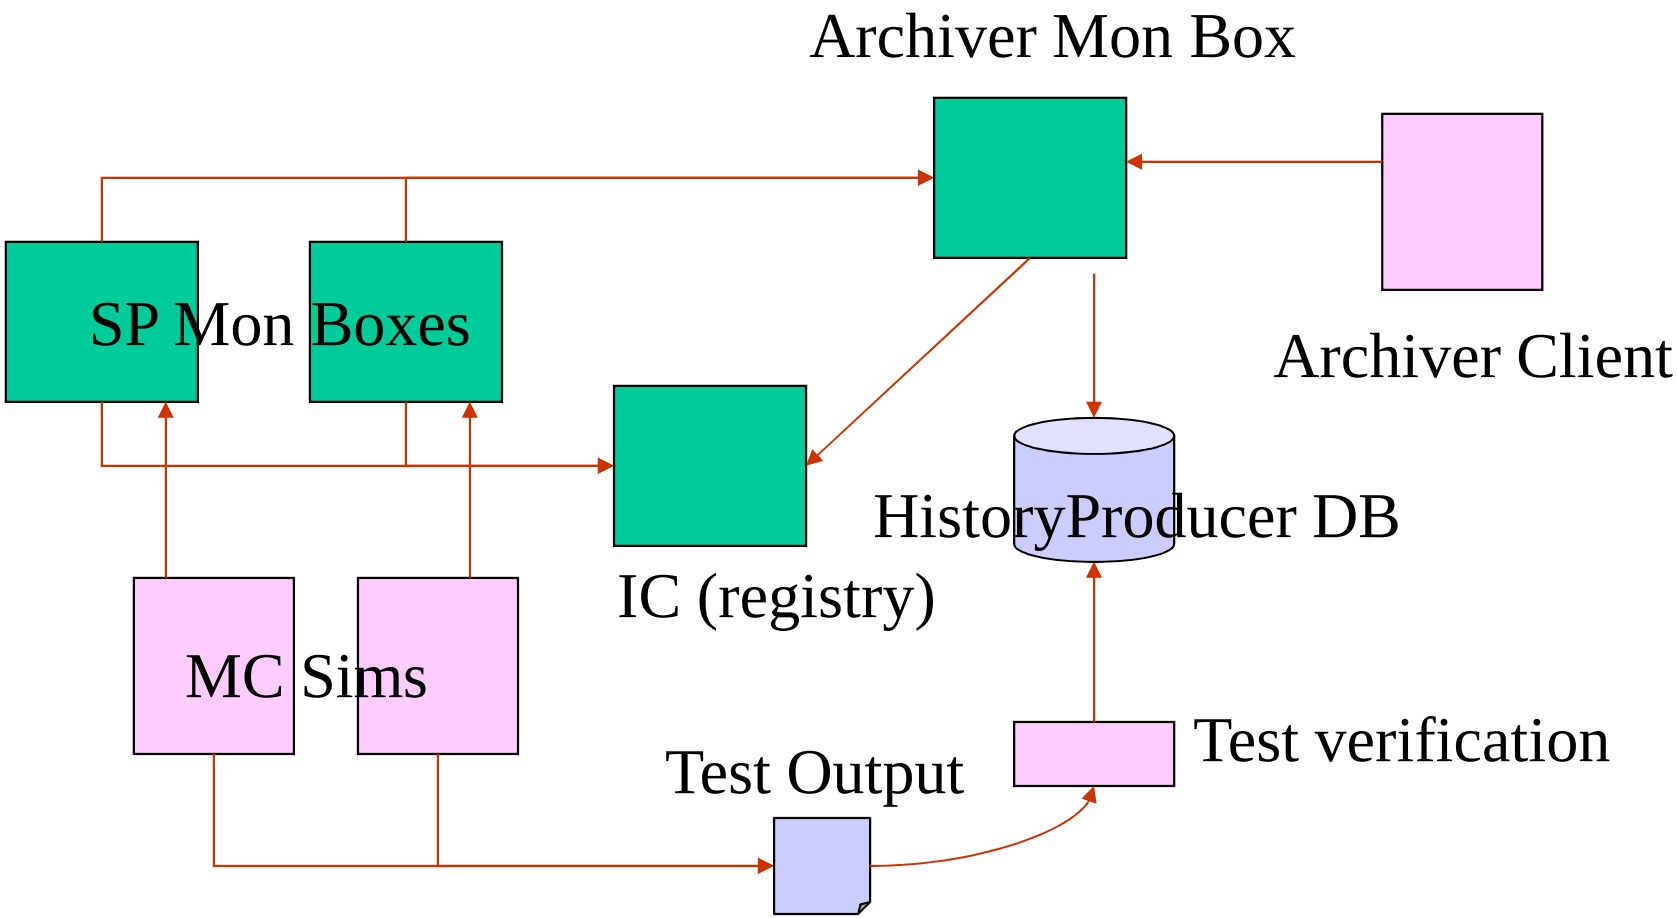
Archiver Mon Box
SP Mon Boxes
Archiver Client
IC (registry)
HistoryProducer DB
MC Sims
Test Output
Test verification

Test Setup

- Archiver & SP mon box setup at Imperial.
- SP mon box & IC setup at Brunel.
- Archiver and MC sim clients positioned at various nodes within both sites.
- Tried 1 MC sim and Archiver with variable Job submissions.
- Also setup similar test on WP3 testbed using 2 MC sims and 1 Archiver.

Results (i.e. Status)

- 1 MC sim creating 2000 jobs and publishing 7600 tuples proven to work without glitch (R-GMA v3.4.13)
- Demonstrated 2 MC sims each running 4000 jobs (with 15200 published tuples) on the WP3 testbed.
- We believe that R-GMA will scale to number of jobs required…

Plans

- But these have not been real jobs – need to confirm performance with full-length and real jobs, at full scale and under real conditions: i.e. with R-GMA infrastructure shared with rest of Grid monitoring
- Check R-GMA functionality in LCG2
- Properly integrate into BOSS source tree