



CMS Applications – Status and Near Future Plans

Dr Henry Nebrensky for
Bristol University
Brunel University
Imperial College London



Contents

- Project objectives
- Implementation status
- Future plans (GRIDPP1/GRIDPP2)




Overall Objective for GRIDPP

Build and test under realistic conditions* an analysis framework to enable CMS physicists to perform end user analysis of data in batch mode on the Grid.

i.e. run a private analysis algorithm over a set of reconstructed data


i.e. not interactive – submit once and leave to run.

* i.e. Data Challenges



Approach

- Build tactical solution using (as far as possible) tools that already exist.
- Iterative development – deliver first prototype quickly covering major use cases, iterating design and implementation later to refine and extend the solution.
- Involve middleware developers in ensuring their general toolkit works for our applications

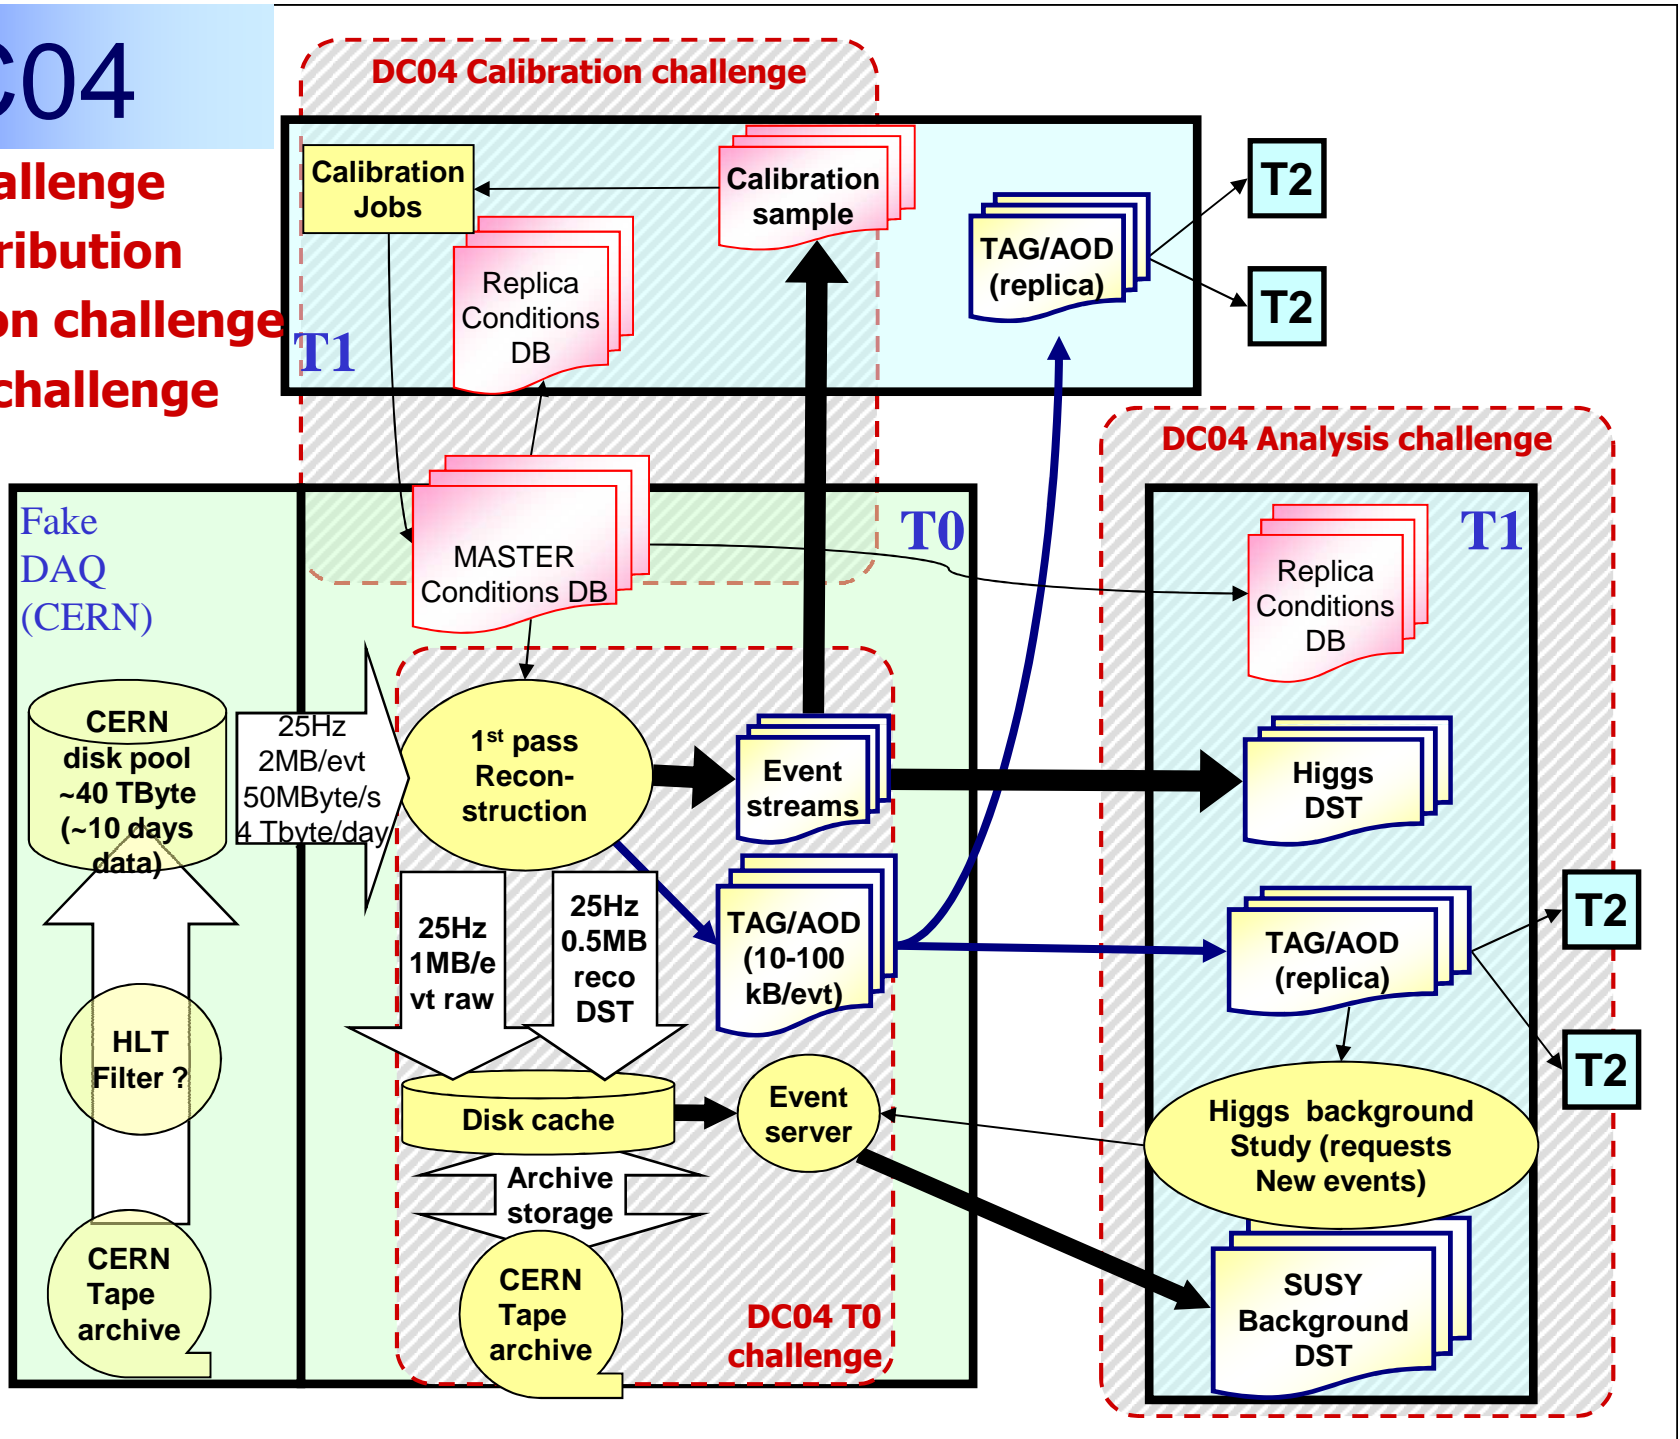



Data Challenges

- Data Challenges (DC) have been used as the prime motivator for the UK contribution to CMS Grid computing.
- DC04 is the most recent and ambitious – 75 million events simulated and 150 TB of data stored worldwide.
- We tried to bring together all the aspects of T0-T1 data transfer, replication, Grid Job submission and job status monitoring.
- We also aimed to contribute to the generation of CMS Monte Carlo data with UK resources.

DC04

Tier-0 challenge
Data distribution
Calibration challenge
Analysis challenge

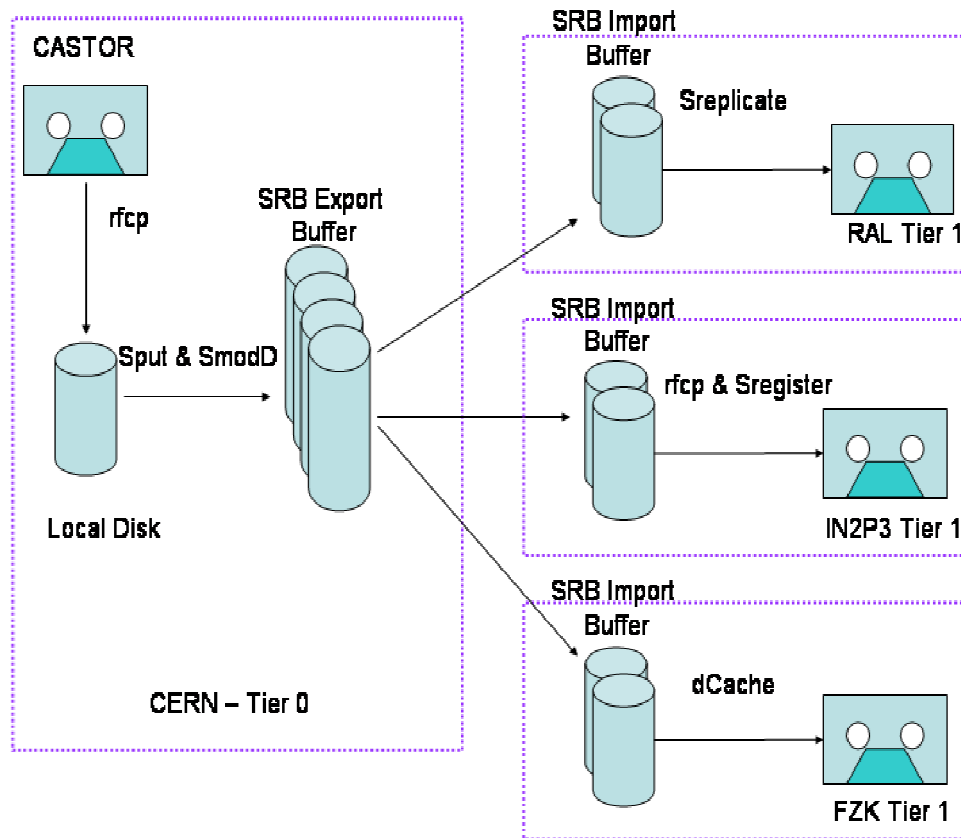





Data Challenge 04

- Processed about 30M events
- Got above 25Hz on many short occasions... but only one full day
- T1s generally kept up
- See details in slides from GridPP 10

SRB in CMS DC04



- SRB was used as the main transfer and data management tool for 3 of the 6 CMS T1 sites
- Mcat (metadata catalog) database hosted between RAL and Daresbury
- Transfers controlled by agents communicating through TMDB
- Difficulties communicating with the Mcat caused severe problems
- SRB is being phased out in favour of LCG solutions
 - may continue to be used “behind the scenes” at some sites



SRB


Tim Barrass, Simon Metson, Dave Newbold (Bristol)

SRB provided a valuable service, and useful experience.

The problems seen in the Data Challenge were a result of the short time-scale involved in setting up for the DC and a lack of a "burn in" of the system at the scale of the DC.

SRB continues to increase in impact in many areas of UK e-science: other projects beginning to use SRB, both in HEP and outside of it.


However, CMS now wants to converge on a unified solution to bulk data management. At the physical layer in each Tier-1, this will be provided products such as dcache and castor. At the top-level data management layer, Phedex is our current solution.



PhEDEx

Tim Barrass, Simon Metson, Dave Newbold (Bristol)


- CMS' system for managing the large-scale movement of data (*ex-DC04*).
 - EGEE **very** interested- no similar component yet.
- Three basic use cases:
 - Push (for vital - e.g. detector - data).
 - Subscription pull (large sites subscribe to datasets e.g. as they are produced).
 - These two are missing from existing grid use cases, but are what experiments traditionally do.
 - Random pull (“shift this file over there”).
 - Relates to grid-style optimisation.
- “User” not physicist, but replica managers, &c



PhEDEx (2)

- V2 close to deployment
 - Some teething troubles moving from ad-hoc project to a more managed one (Tim Barrass/Lassi Tuura project and technical leads).
 - Gaining experience quickly.
 - PhEDEx extended to accommodate multiple data sources- to enable transfers of MC produced around the world.
 - Will help production meet PRS goal of delivering 10M events per month from the Autumn.
- Finding/helping solve differences in approach of grid and traditional distributed computing...
- Plenty of feedback to other developer groups- LCG, EGEE, SRM, SRB ...
- Hands-on Workshop at CERN before CMS week.
- Several talks at CHEP04.

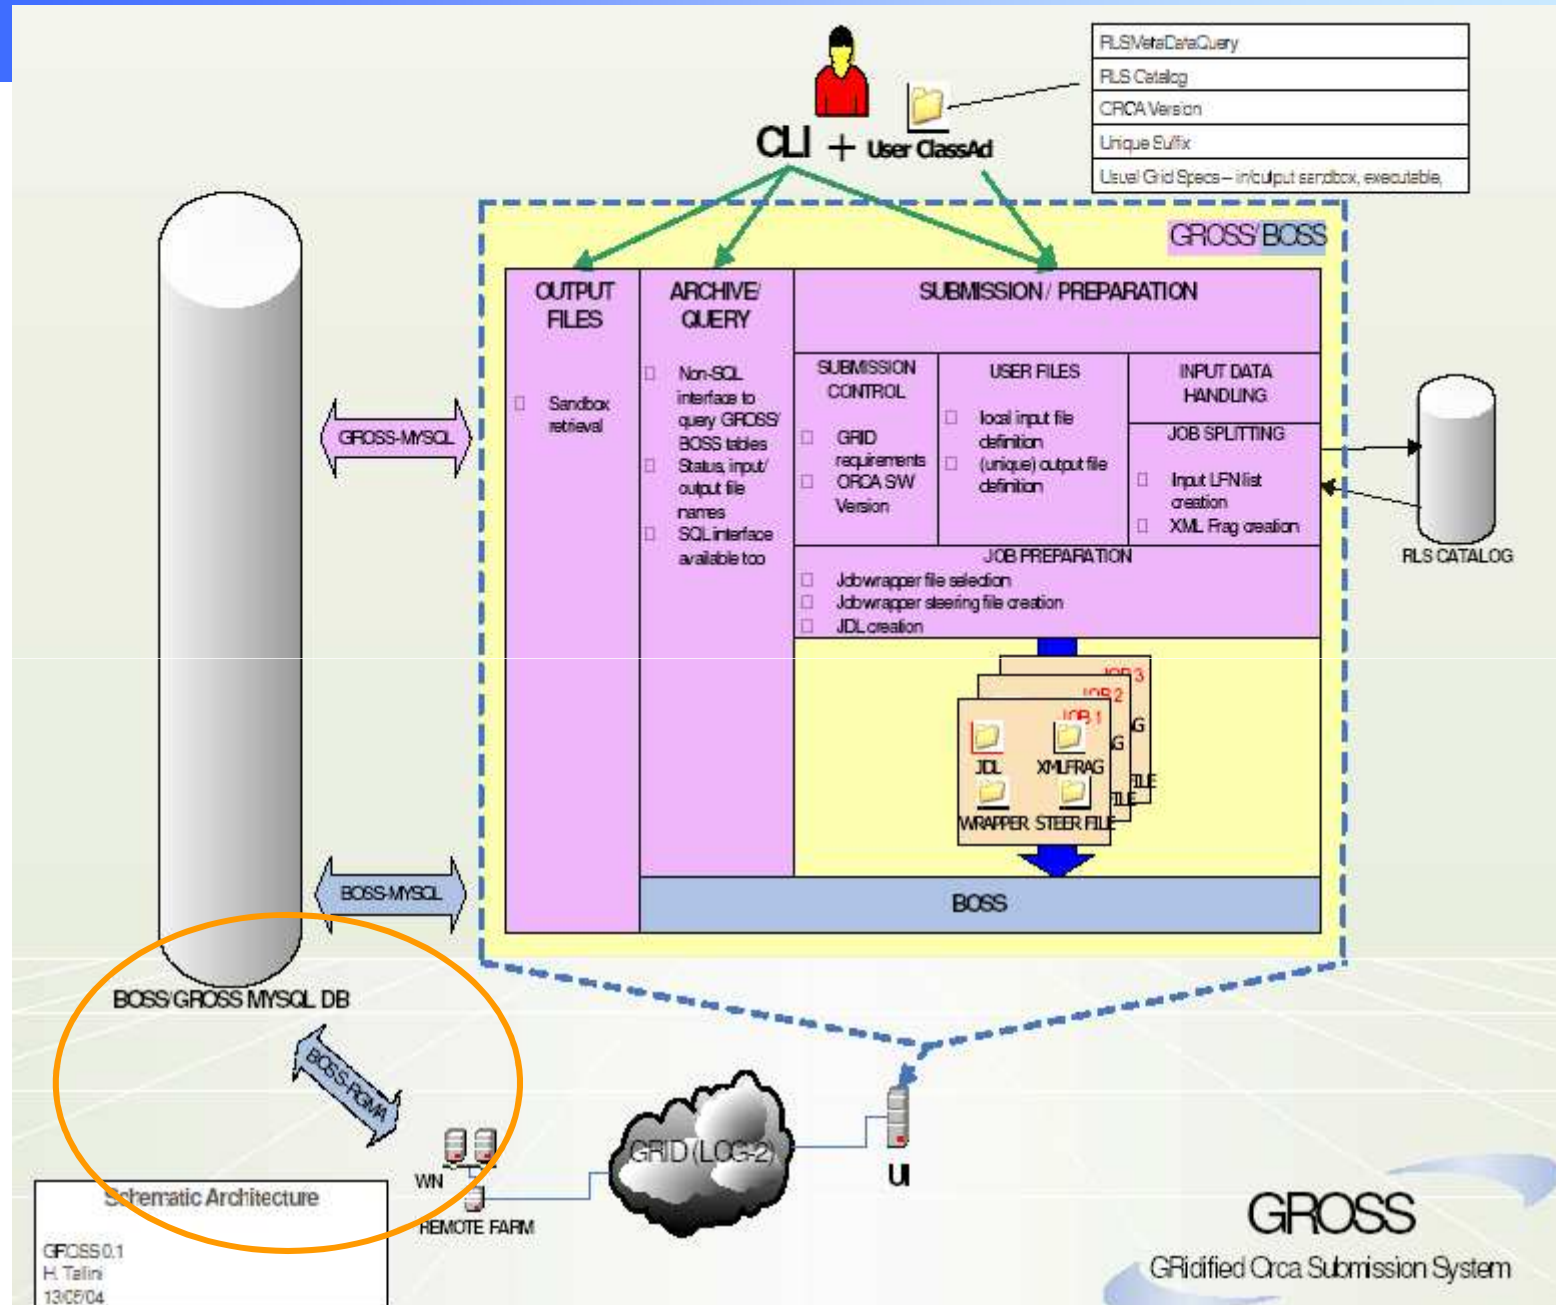
(See Tim's Data Management talk from UK CMS meeting in Bristol: <http://agenda.cern.ch/fullAgenda.php?ida=a042616>)




Job submission - GROSS

Imperial College


- GROSS is designed as an interface to the Grid for end users wishing to run distributed analysis.
- User authenticates themselves with the Grid, and then interacts solely with GROSS.
- User defines the details of an analysis task
 - GROSS will carry out the job splitting, preparation, submission, monitoring and output file collection.
 - The same functionality is implemented for local batch submission systems (were appropriate).





GROSS (2)

- Uses an Abstract Factory to build the appropriate family of classes for the specific job type which will be submitted. Currently two types have been implemented:
 - ORCA job on LCG2 (ORCA is the full reconstruction code for CMS data)
 - ORCA job using *local* batch submission system
- GROSS currently uses a CLI
- Users also have to define some configuration files.

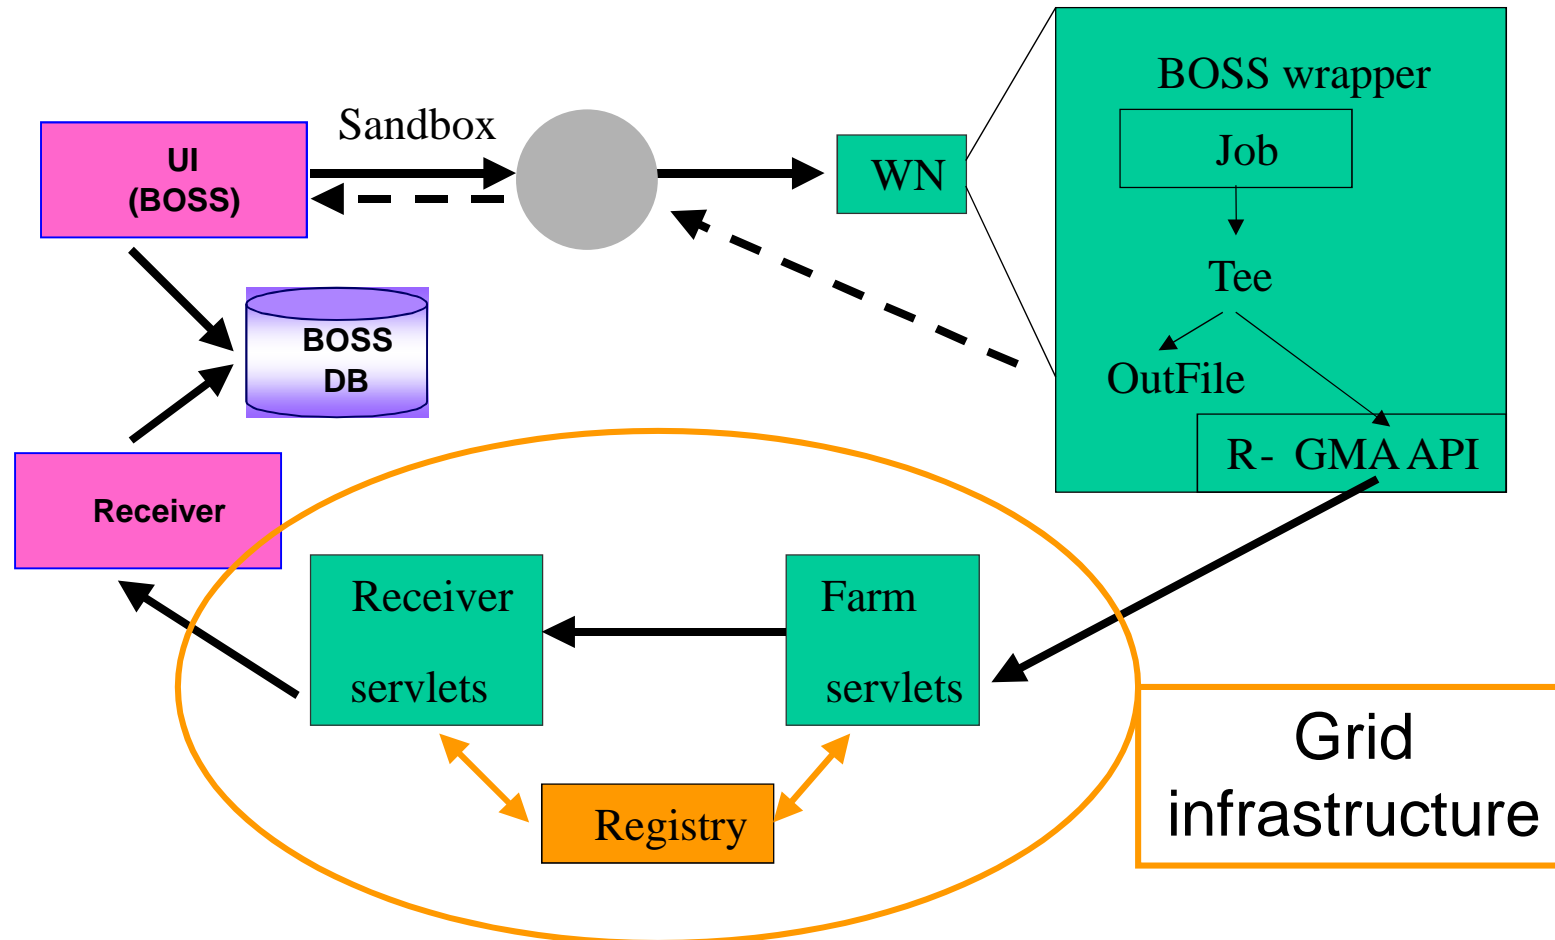


Job monitoring - RGMA


Brunel University

- R-GMA finally released in LCG 2.2.0 and deployed on the LCG
- We need to confirm that it can handle load of applications monitoring under “real-world” conditions
- Will be testing BOSS with R-GMA over next couple of months

Use of R - GMA in BOSS




R-GMA is an EDG WP3 Middleware product



Assessing R-GMA


- Use the CMSIM “emulator” from last year’s tests
- Submit jobs to the Grid; see
 - if messages get back
 - how many come back
- Don’t need to do the number crunching in between – small number of Grid jobs can put large load on R-GMA
- Results don’t apply just to BOSS – *any* applications monitoring framework would have to shift same amount of data from those jobs.



Tier 1

Bristol/RAL

- Slide Here?



“Future” Plans

CMS

- CMS' submission to the ARDA Middleware mtg is being put together Right Now.