ARDA Workshop, CERN October, 21st, 2004

"ALICE @ ARDA"
Interactive Analysis on a GRID

P. Cerello on behalf of the ALICE Offline Team

Enabling Grids for E-science in Europe

Outline

1. The ALICE Physics Data Challenge
2. Summary of phase I / II
3. Phase III
   1. What?
   2. How?
   3. When?
4. Summary & Outlook

The ALICE Physics Data Challenge

- Phase 1 - production of underlying events using heavy ion MC generators
  - Status: 100% complete (Mar-May 2004)
  - Basic statistics - ~ 1.3 million files, 26 TB data volume
- Phase 2 – mixing of signal events in the underlying events
  - Status: 100% complete (Jun-Sep 2004)
- Phase 3 – analysis of signal + underlying events
  - Goal – to test the data analysis model of ALICE
  - Status – to be started

Phase I & II layout: a "Meta-Grid"

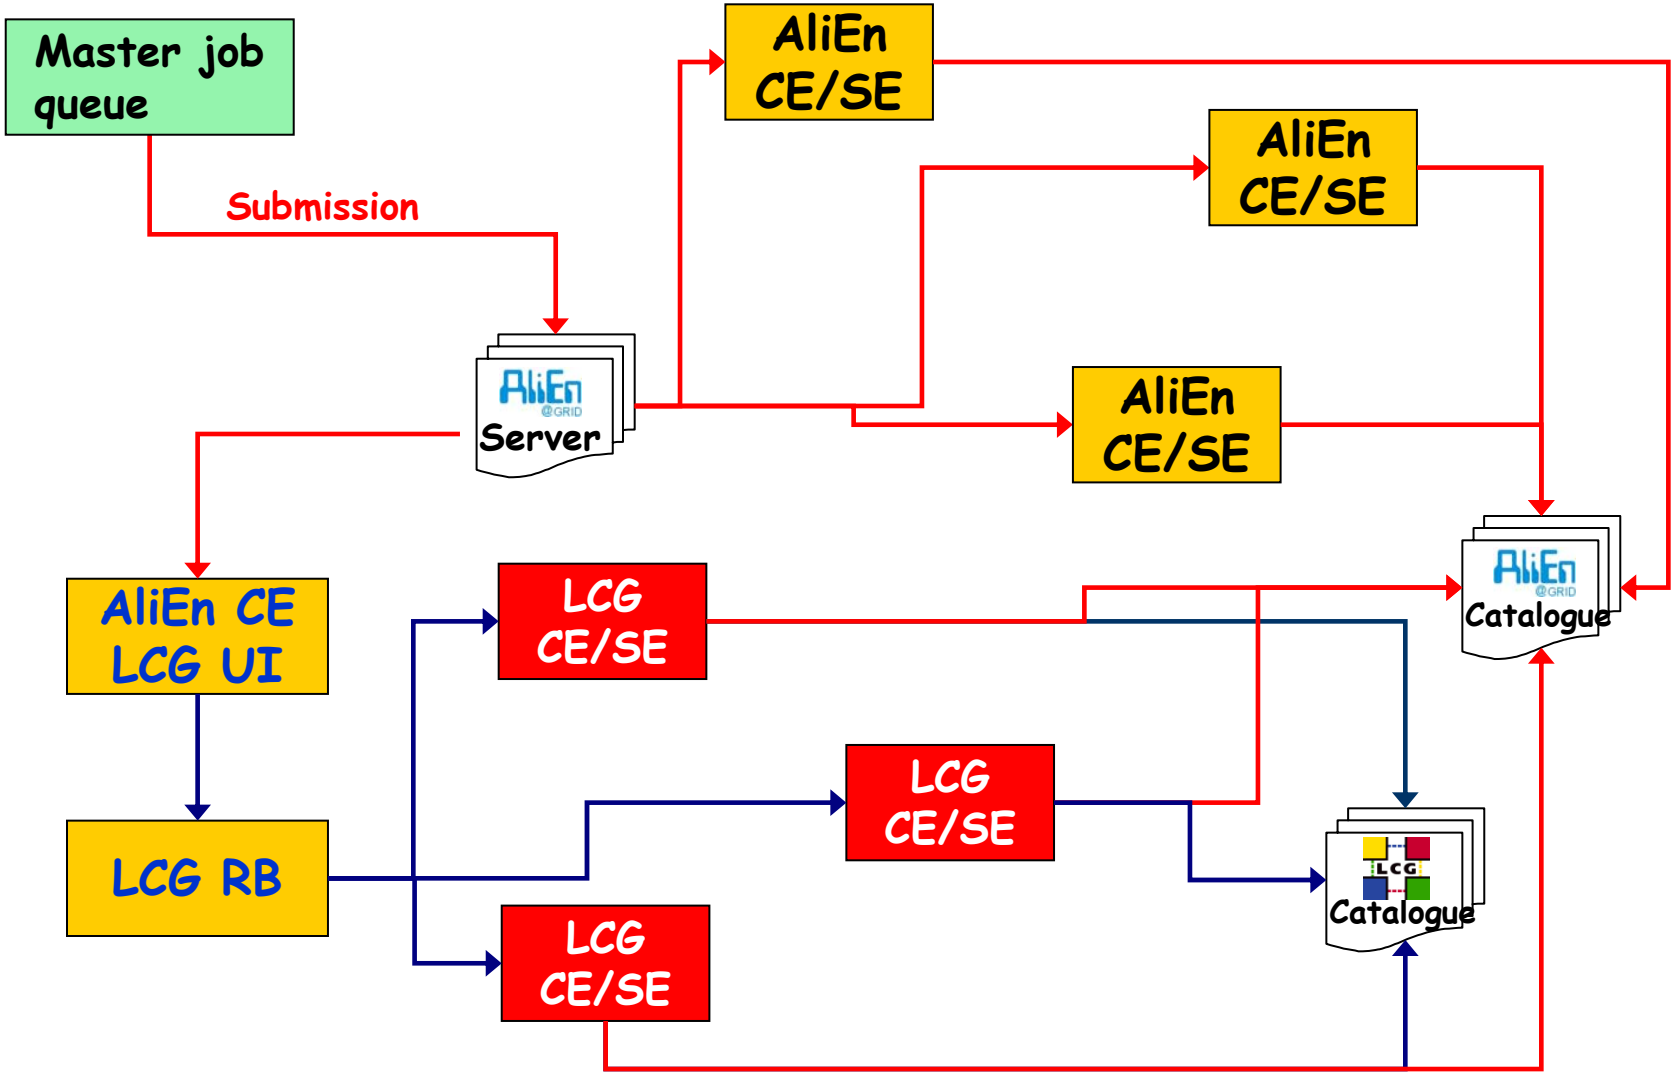

Interfacing AliEn and LCG

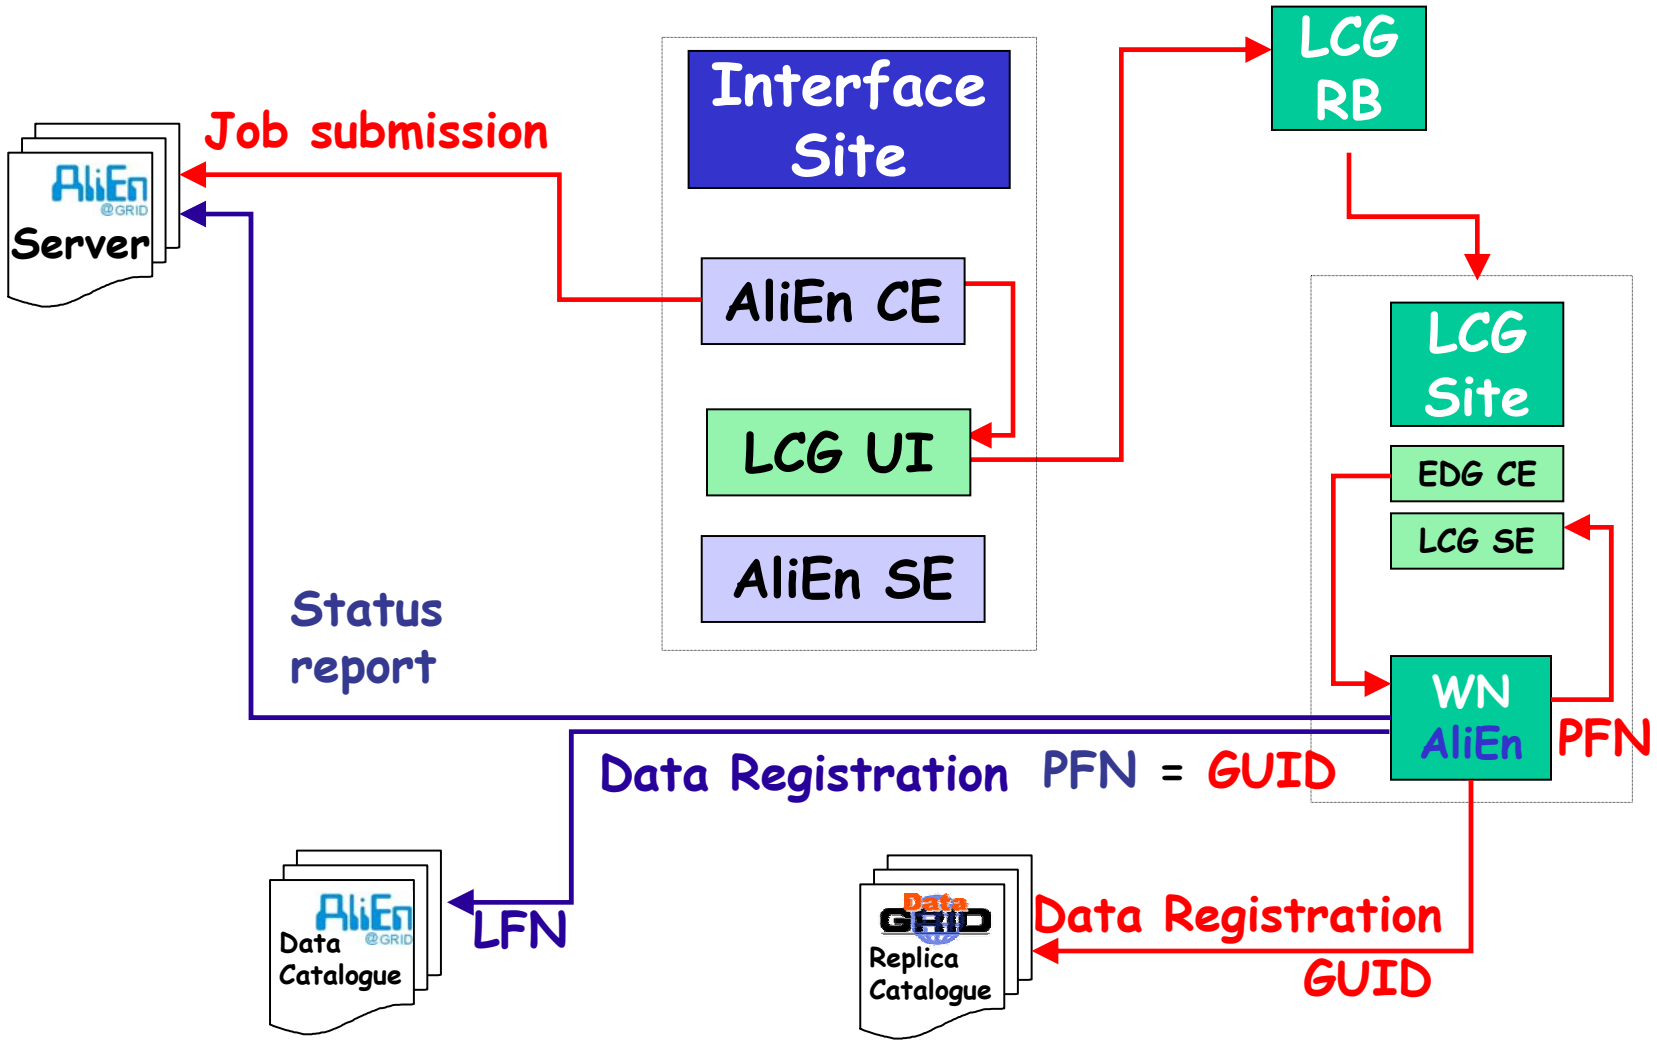

Phase I - Simulation

Small input with interaction conditions

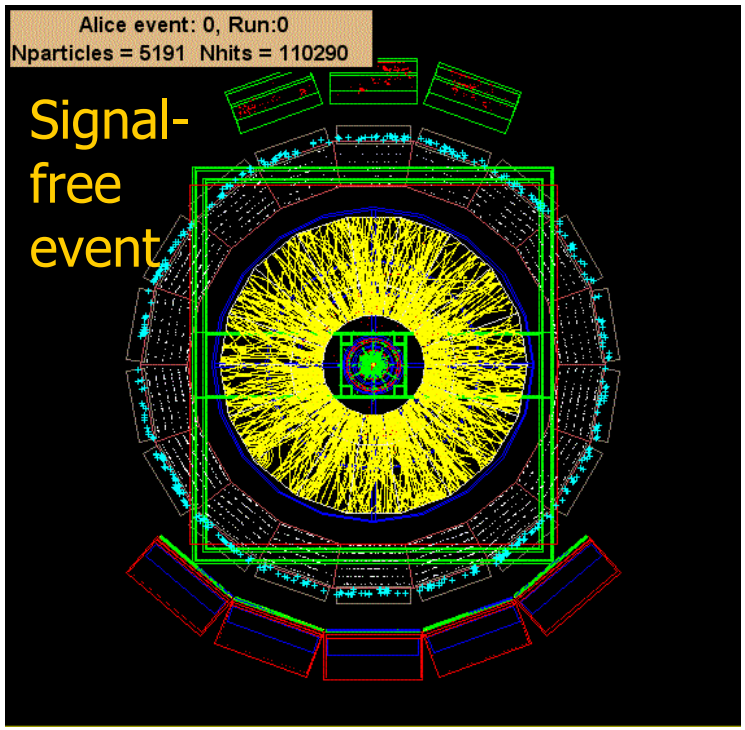

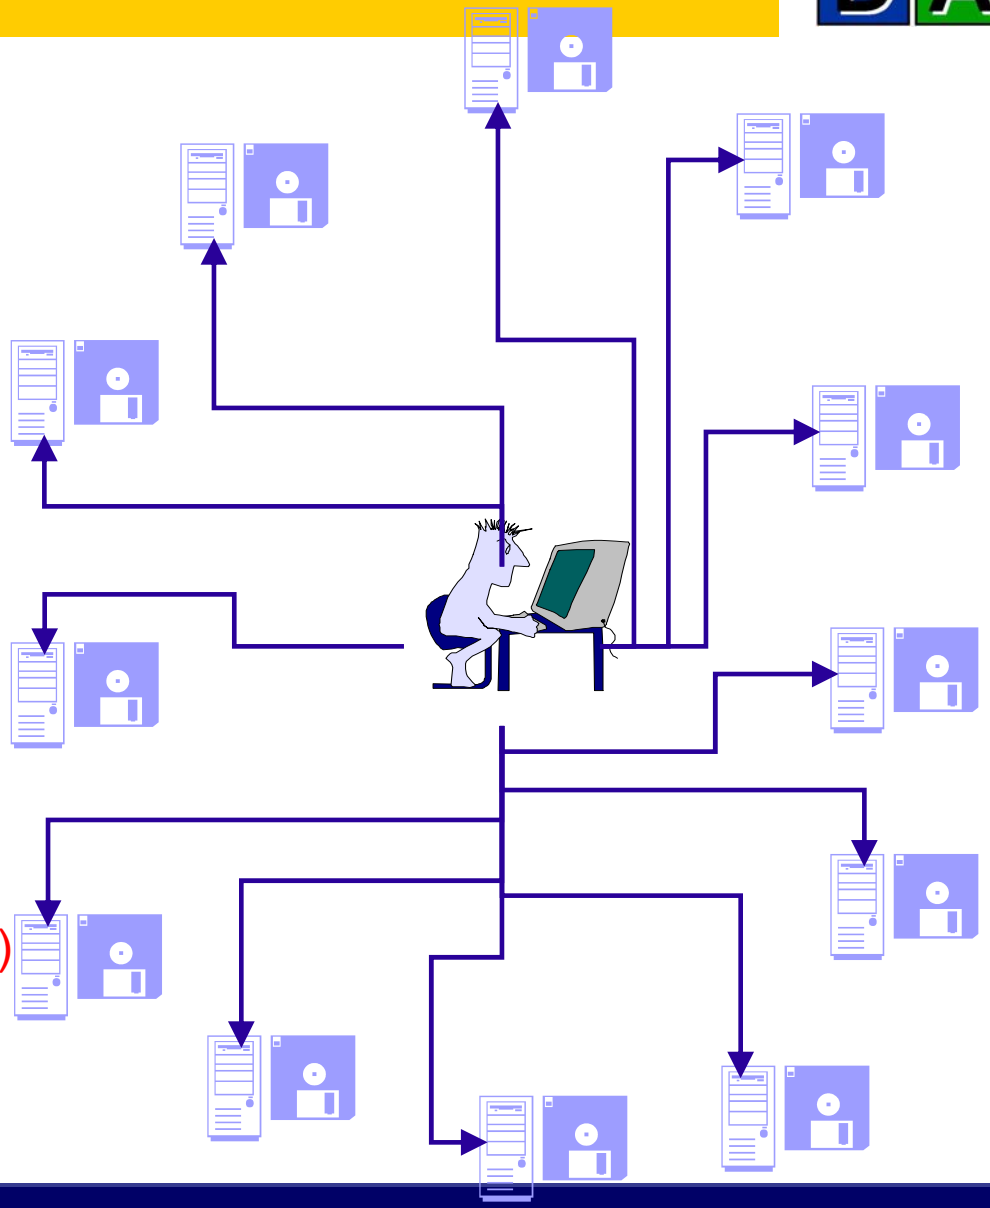

Large **distributed** output (1 GB/event)
with simulated detector response
Long execution time (10 hours/event)

Phase I results

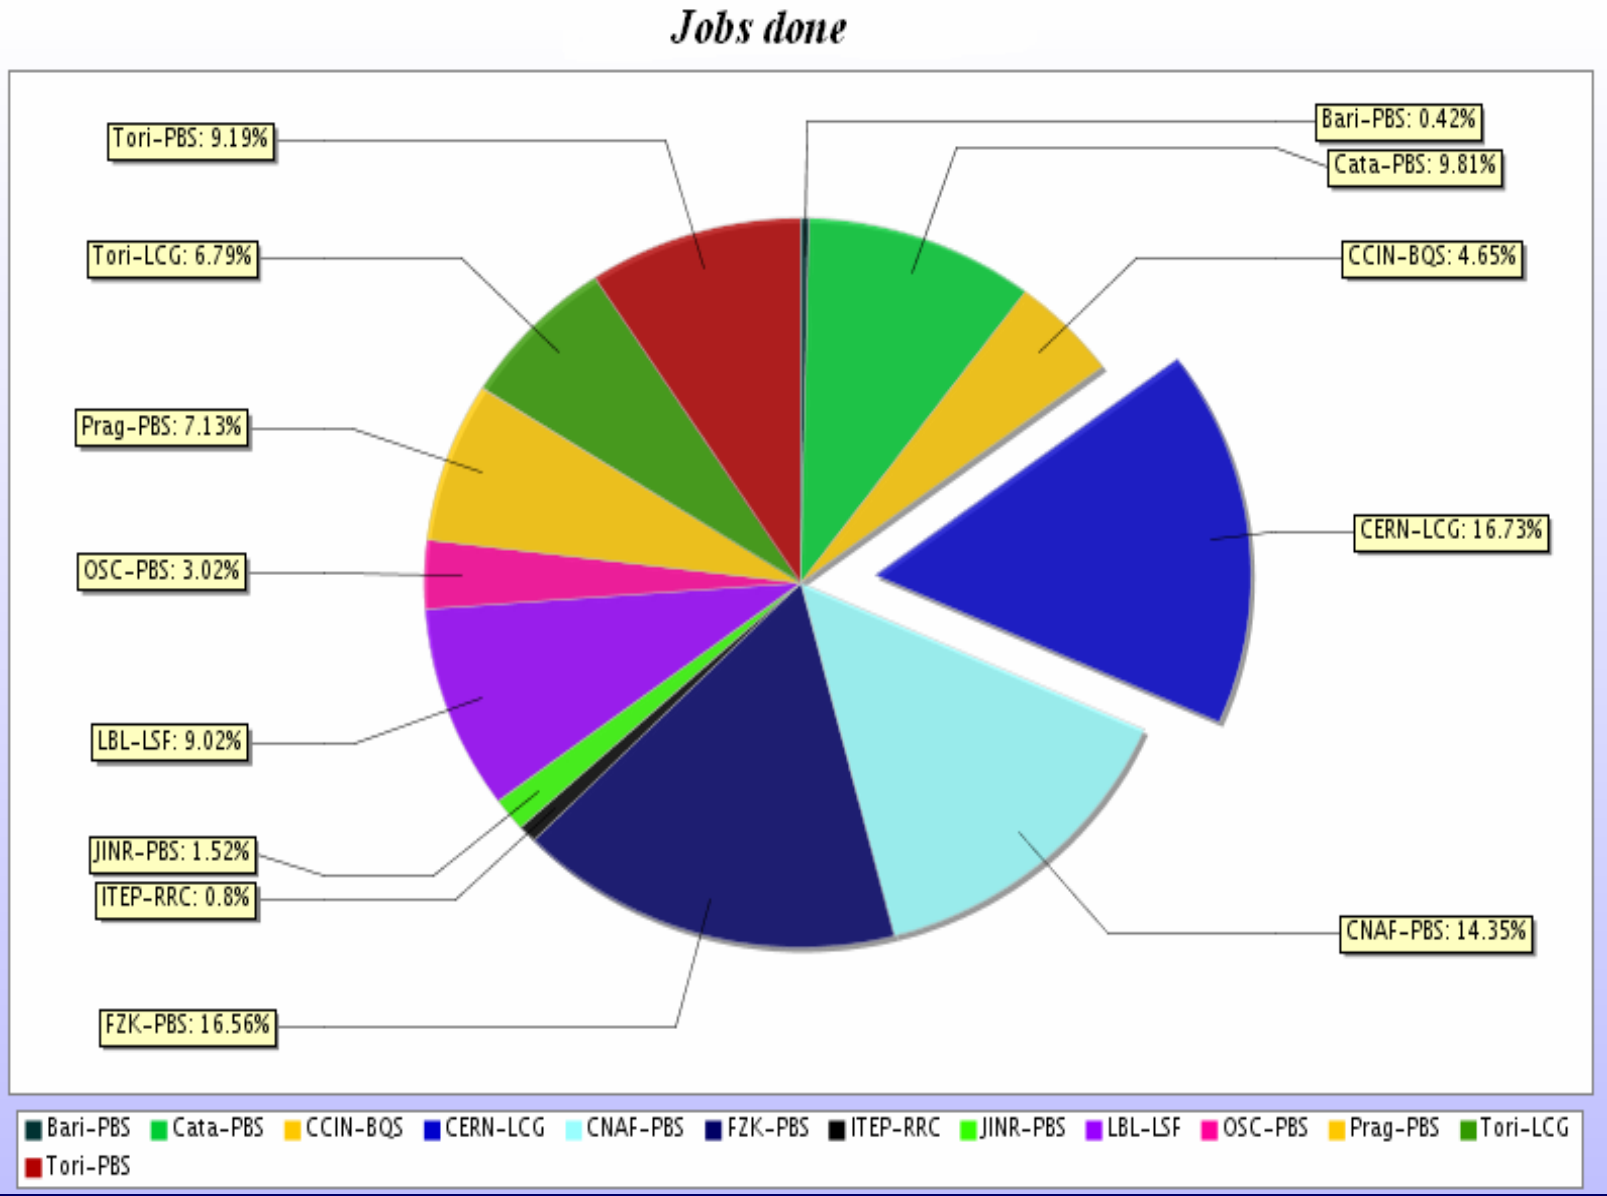

Phase II - Merging & Reconstruction

Large **distributed** input (1 GB/event)

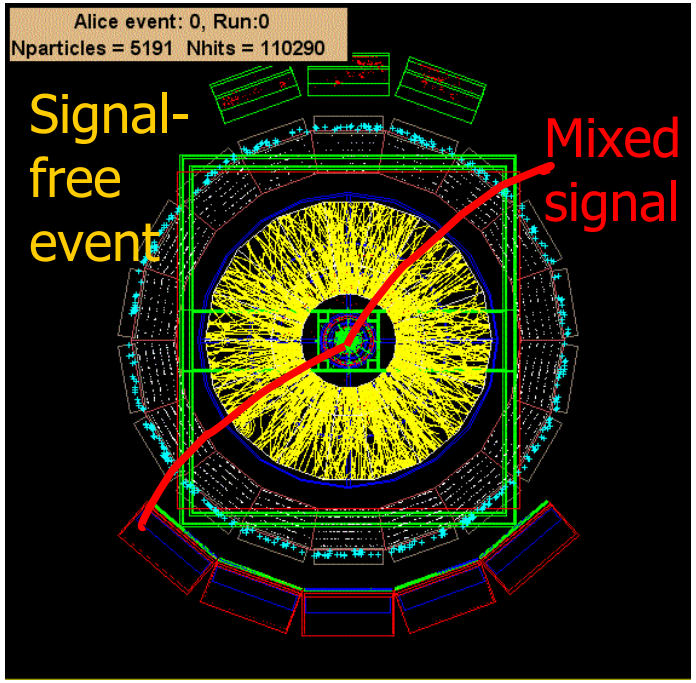

Fairly Large **distributed** output (100 MB/event, 7MB files) with reconstructed events

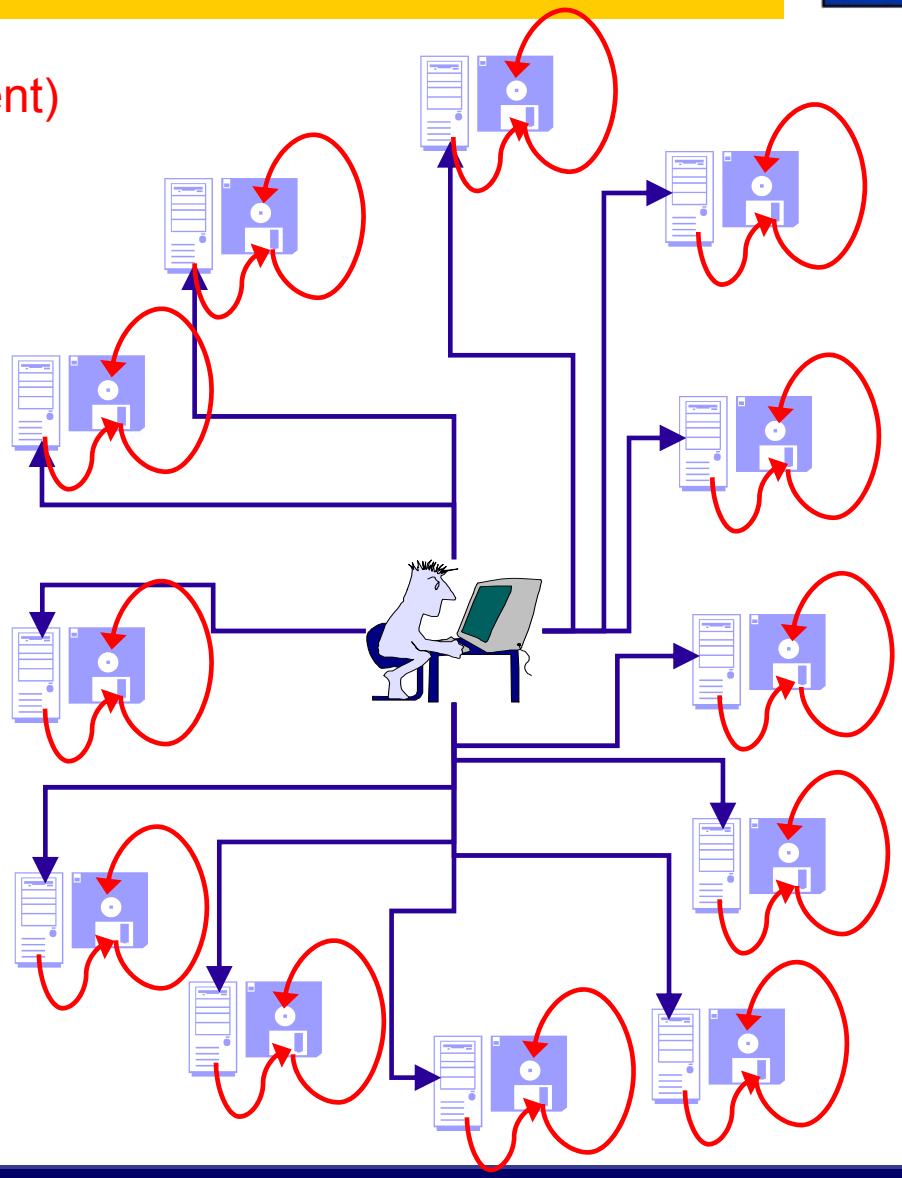

Phase II results

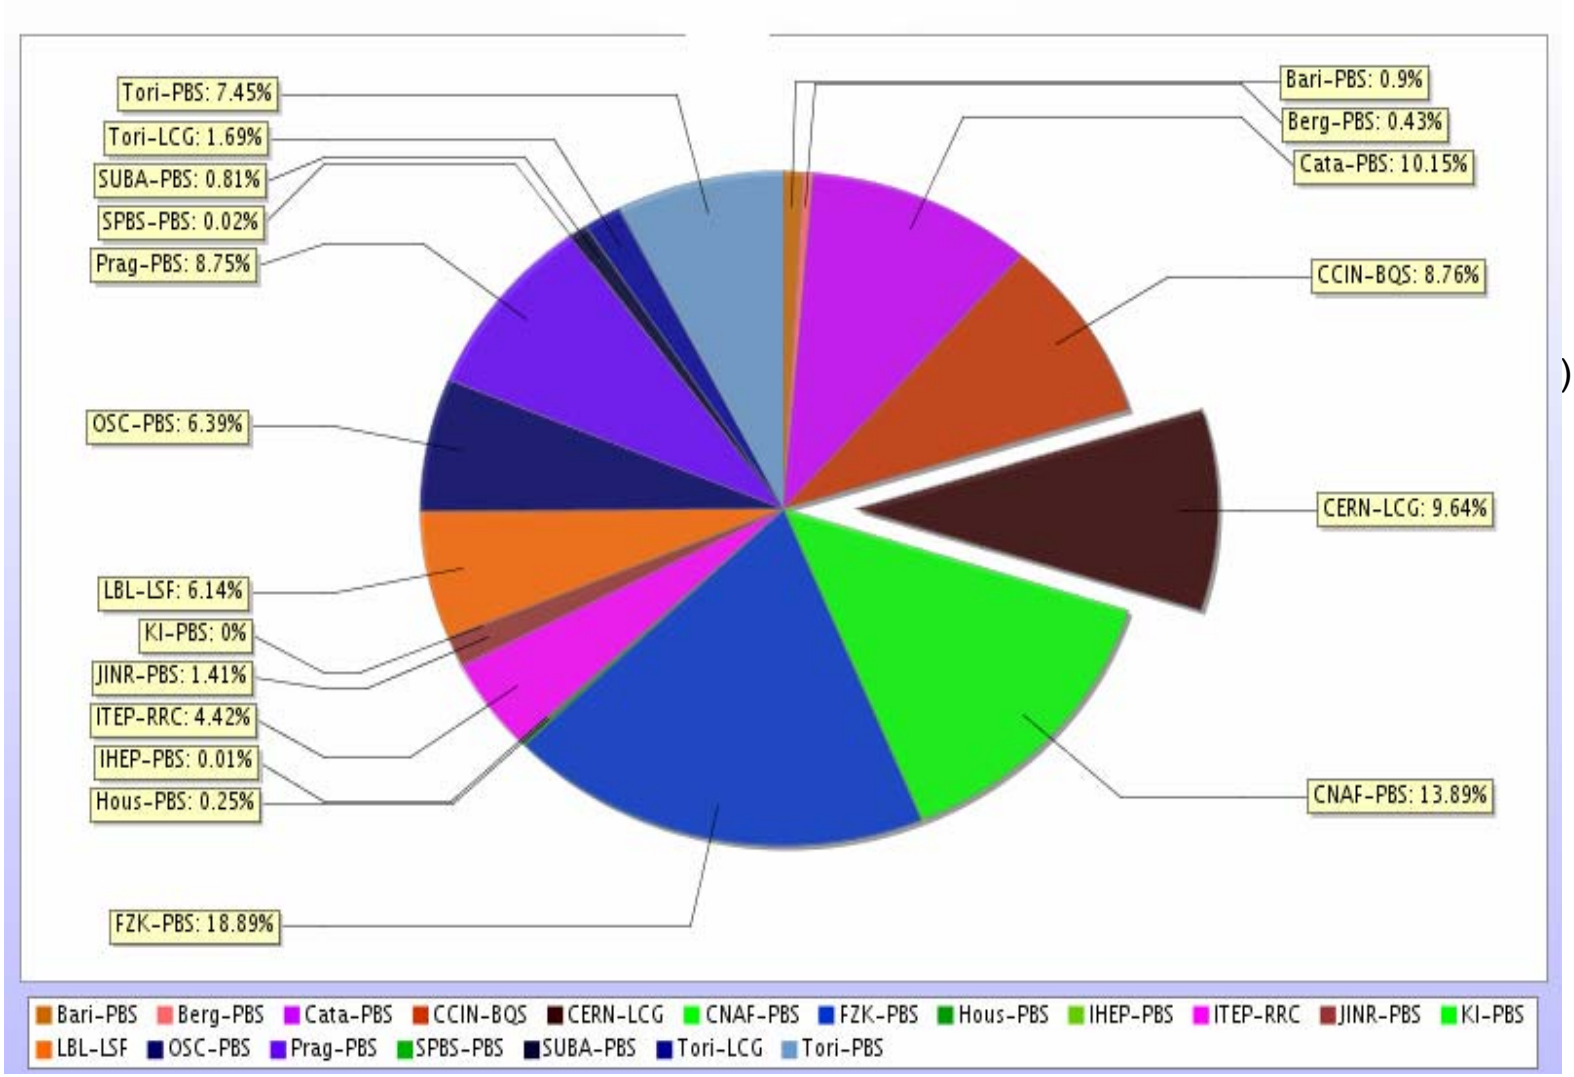

Phase III – (Interactive) Analysis

Large **distributed** input (100 MB/event)

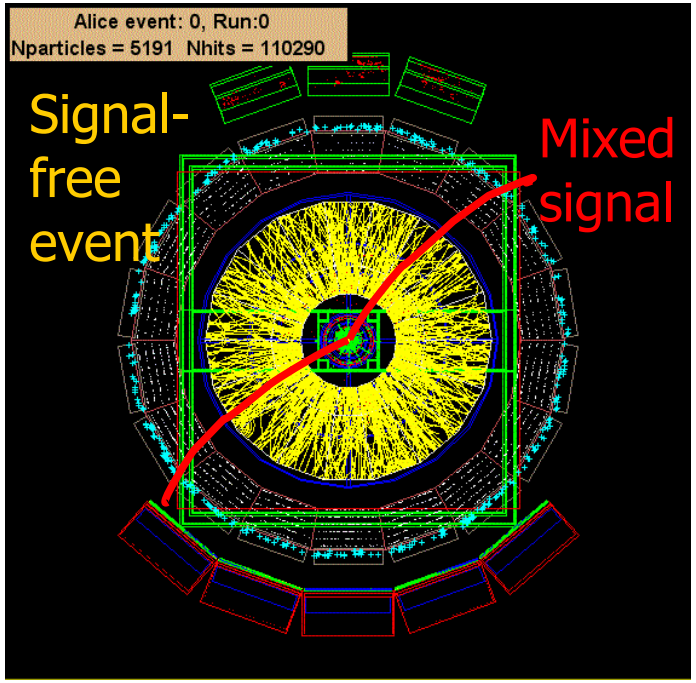

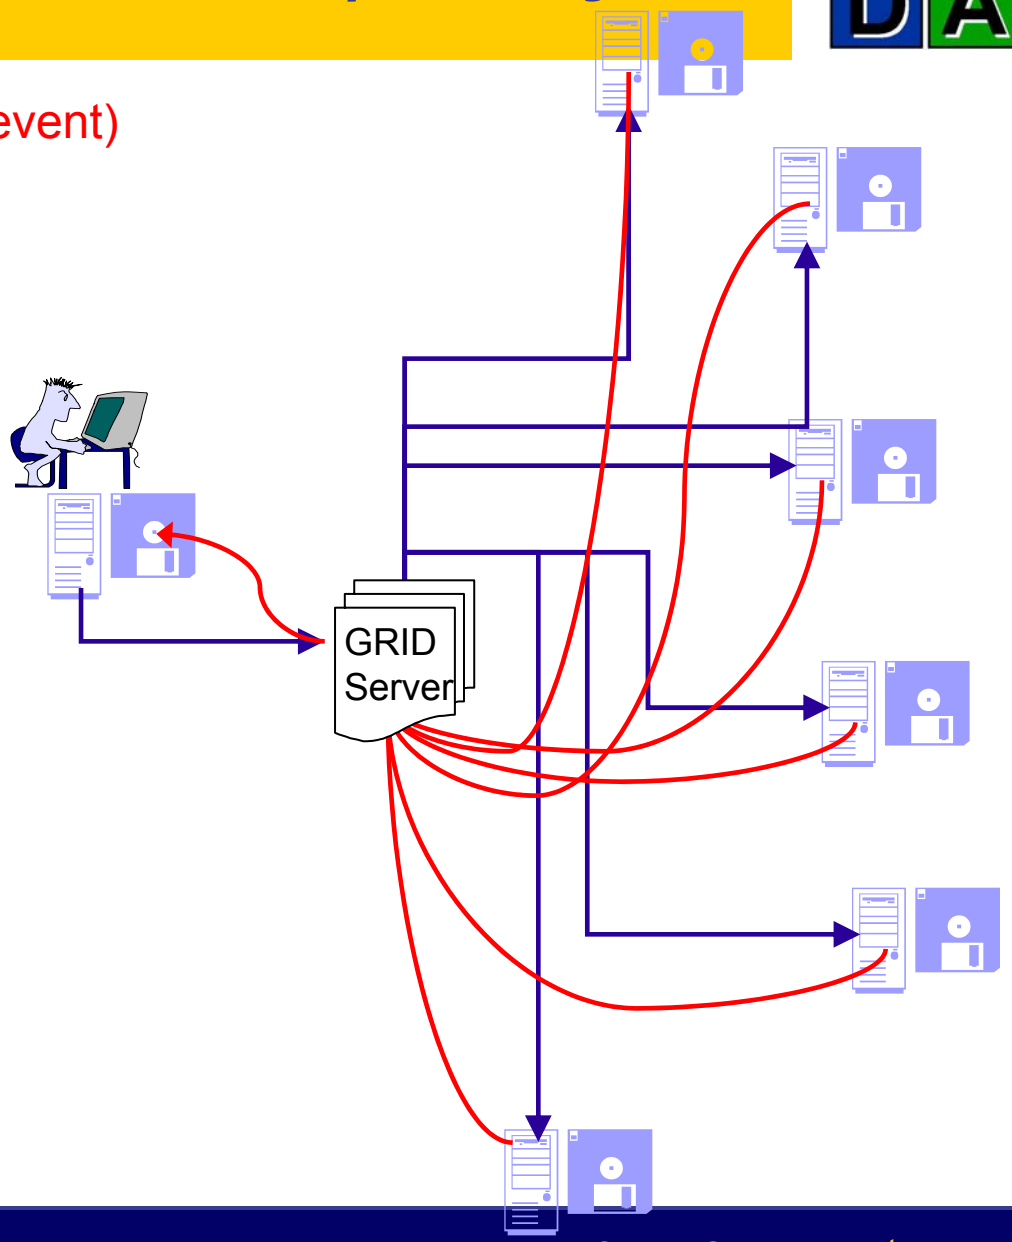

Fairly small **merged** output

Phase III - Design strategy

- Distributed Input
  - 5.4 M files, on about 30 different SEs
- Do not move the input
  - Algorithm definition by a user on site S
  - On site S, Input Query to the Data Catalogue, based on selected MetaData
  - Input files typically from many SEs, so
    - Split them in N subgroups defined by files stored on a given SE
    - Split the task into N sub-tasks, to be run in parallel on the CEs associated to SEs containing a fraction of the input files
    - Run the N sub-tasks in parallel
    - Merge the output on the user's site S
- How?
  - From a ROOT shell on the user's site
  - Hopefully interactively with PROOF
  - Refer to Derek's and Andreas' presentation yesterday

Phase III - Original Plan

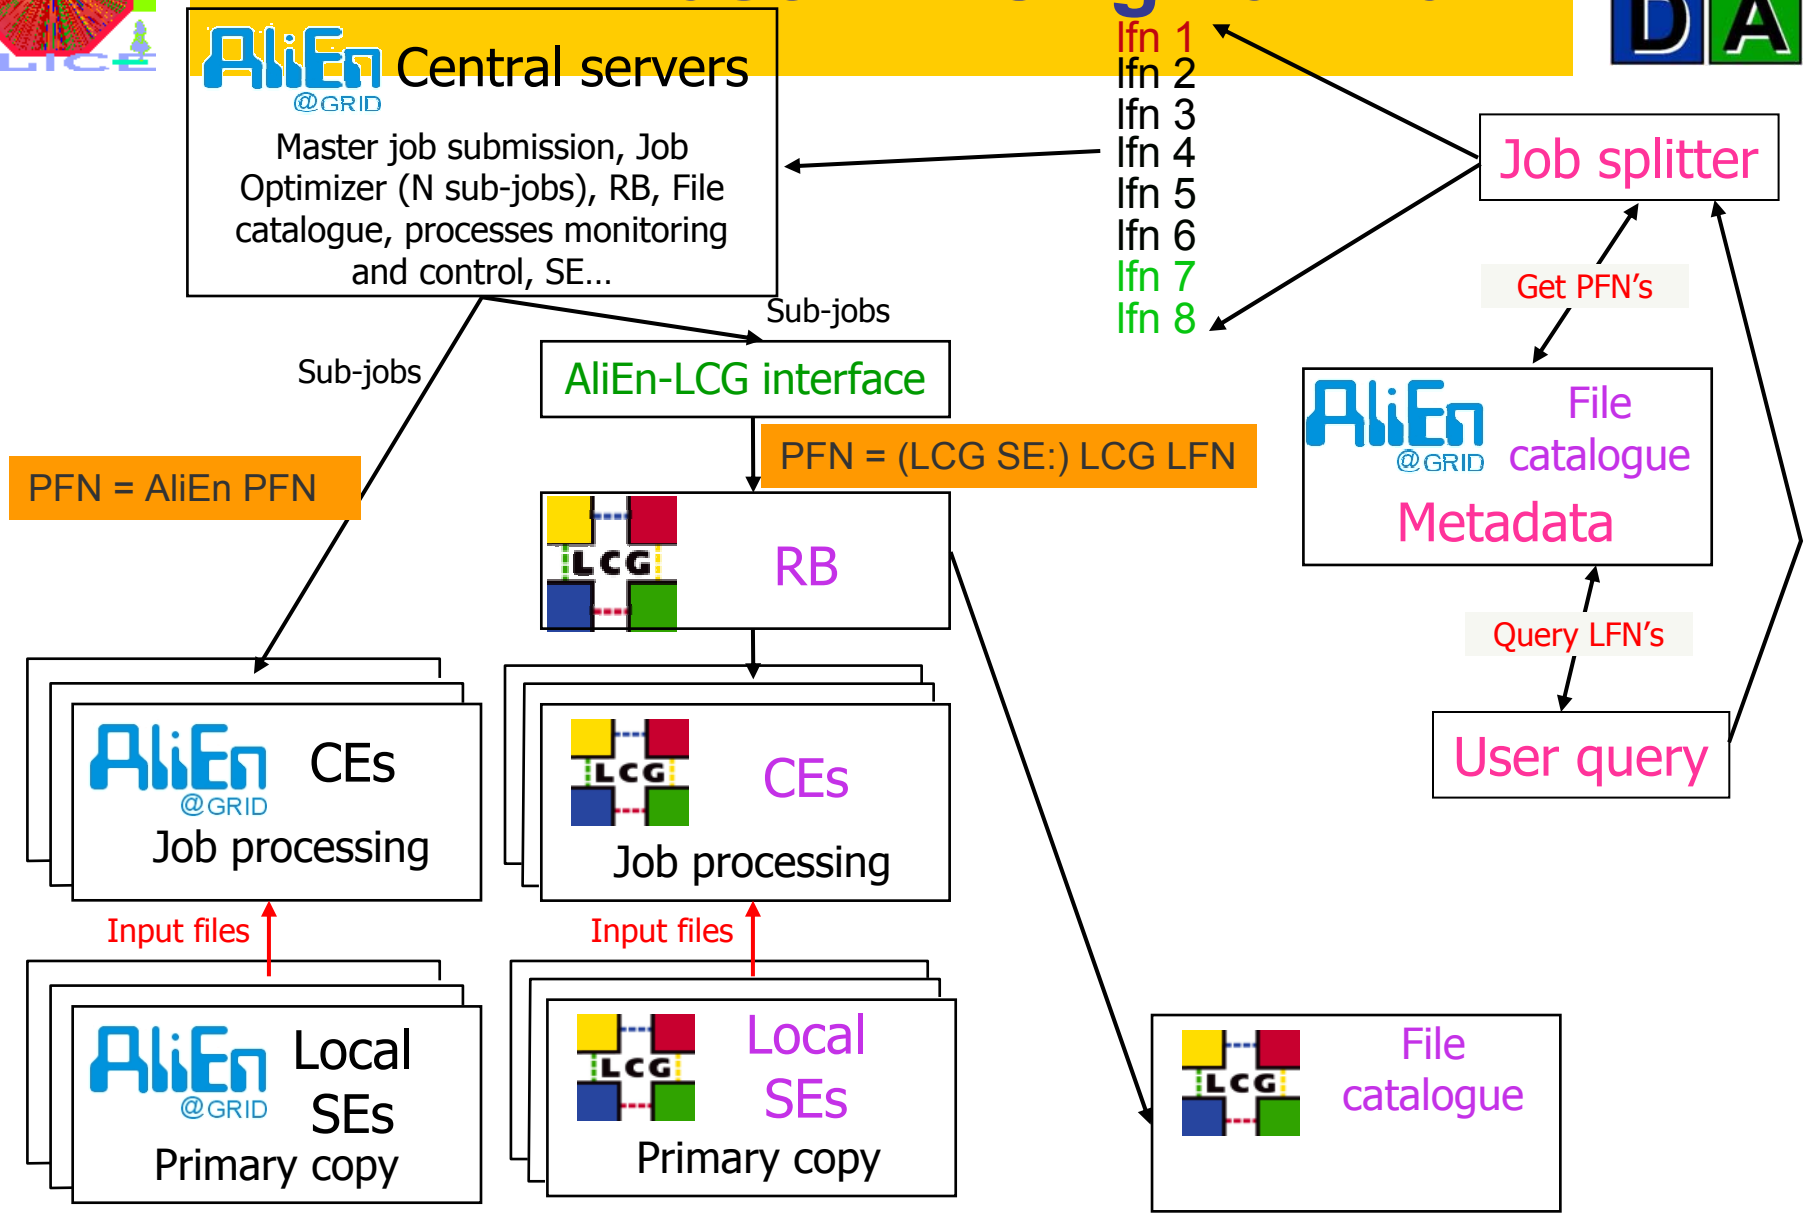

Phase III - Execution Strategy

- Very labour intensive
  - The status of LCG DMS is not brilliant
- Does not “leverage” the (excellent!) work done in ARDA
- So… why not doing it with gLite?
- Advantages
  - Uniform configuration: gLite on EGEE/LCG-managed sites & on ALICE-managed sites
  - If we have to go that way, the sooner the better
  - AliEn is anyway “frozen” as all the developers are working on gLite/ARDA
- Disadvantages
  - It may introduce a delay with respect to the use of the present – available – AliEn/LCG configuration
  - But we believe it will pay off in the medium term

Phase III - Layout

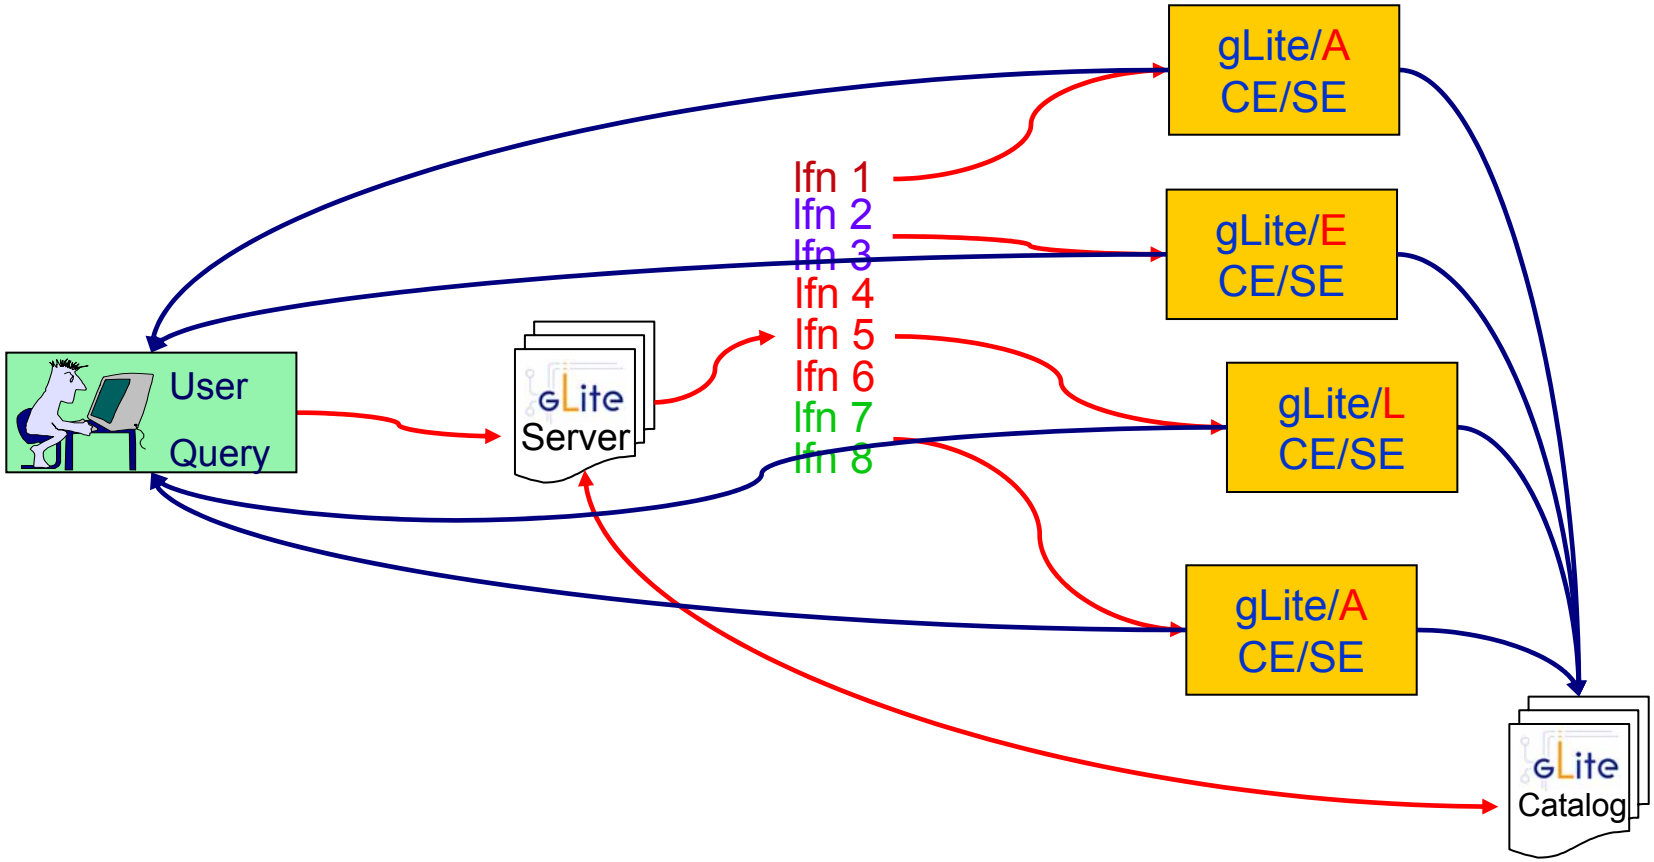

Phase III – The Plan

- ALICE is ready to play the guinea-pig for a large scale deployment
  - i.e. on all ALICE resources and on as many existing LCG resources as possible
- We have experience in deploying AliEn on most centres, we can redo the exercise with gLite
  - Even on most LCG centres we have a parallel AliEn installation
  - Many ALICE site-managers are ready to try it
- And we would need little help
  - We need a gLite (beta-) as soon as possible, beginning November
  - Installation and configuration of sites must be as simple as possible
    - I.e. do not require root access
  - We expect help from LCG/EGEE to help us configure and maintain the ALICE gLite server, running common services

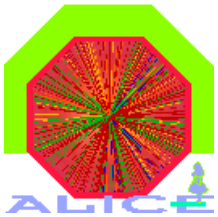

Phase III – Steps to get started

- Data Management Services
- Input Data
  - They are already registered in the AliEn & LCG Data Catalogues and stored on AliEn & LCG Storage Elements
  - Access the Alien & LCG Catalogues from gLite
  - … or Translate the AliEn & LCG Catalogues to a gLite Catalogue instance
  - Input Data is scattered on about 30 sites: we need ALL of them to become gLite sites
  - … or we may do some data movement
- Output Data
  - Must be made available at the User's site for merging and optional registration to the gLite Data Catalogue

Phase III – Steps to get started

- Workload Management Services
  - No special requirement, as the job distribution is intrinsically defined by the input location, but:
    - We need the functionality to split a job into sub-jobs according to the input distribution -- ARDA has it!
  - Jobs submitted by users must be registered to a Master Queue which keeps the record of ALL the ALICE-VO tasks

Conclusion

- Phase I and II of our Data Challenge were completed
  - LCG support and resources were very important
  - LCG middleware has limited the usability of the resources
- Phase III
  - With gLite on the verge to be released we think it would be absurd not to “bite the bullet” and use it now
  - We will provide an experience “complementary” to the component by component strategy of LCG
  - We feel we would gain many months and acquire a precious experience with no special additional load on the deployment team
  - This will help bootstrapping a process which we feel is much too slow and timid
  - And we have already done it with AliEn
- ALICE intends to be gLite-only in the shortest possible time