



Enabling Grids for
E-science in Europe

www.eu-egee.org

LCG SC3 Workshop – CERN June 13, 2005

ALICE Data Challenges on the Grid: Past experience

Stefano Bagnasco
I.N.F.N. Torino



EGEE is a project funded by the European Union under contract IST-2003-508833

(Very) few details on AliEn

- The Workload Management is “pull-model”: a server holds a master queue of jobs and it is up to the CE that provides the CPU cycles to call it and ask for a job. There is no IS.
- The system is integrated with a large-scale job submission and bookkeeping system “tuned” for Data Challenge-like productions, with job splitting, statistics, pie charts, automatic resubmissions, etc.
- The Job Monitoring model requires no “sensors” installed on the WN. It is the jobwrapper itself that talks to the server.

Production on different grids

- Several Grid infrastructures are becoming available: LCG, INFN GRID, possibly others, maybe in the U.S.
- Lots of resources but, in principle, different middlewares
- Pull-model is well-suited for implementing higher-level submission systems, since it does not require knowledge about the periphery, that may be very complex:

“A Grid is a system that [...] coordinates resources that **are not subject to centralized control** [...] using standard, open, general-purpose protocols and interfaces [...] to deliver nontrivial qualities of service.”

I. Foster
“What is the Grid? A three Point Checklist”
Grid Today (2001)

Production on different grids

Design strategy:

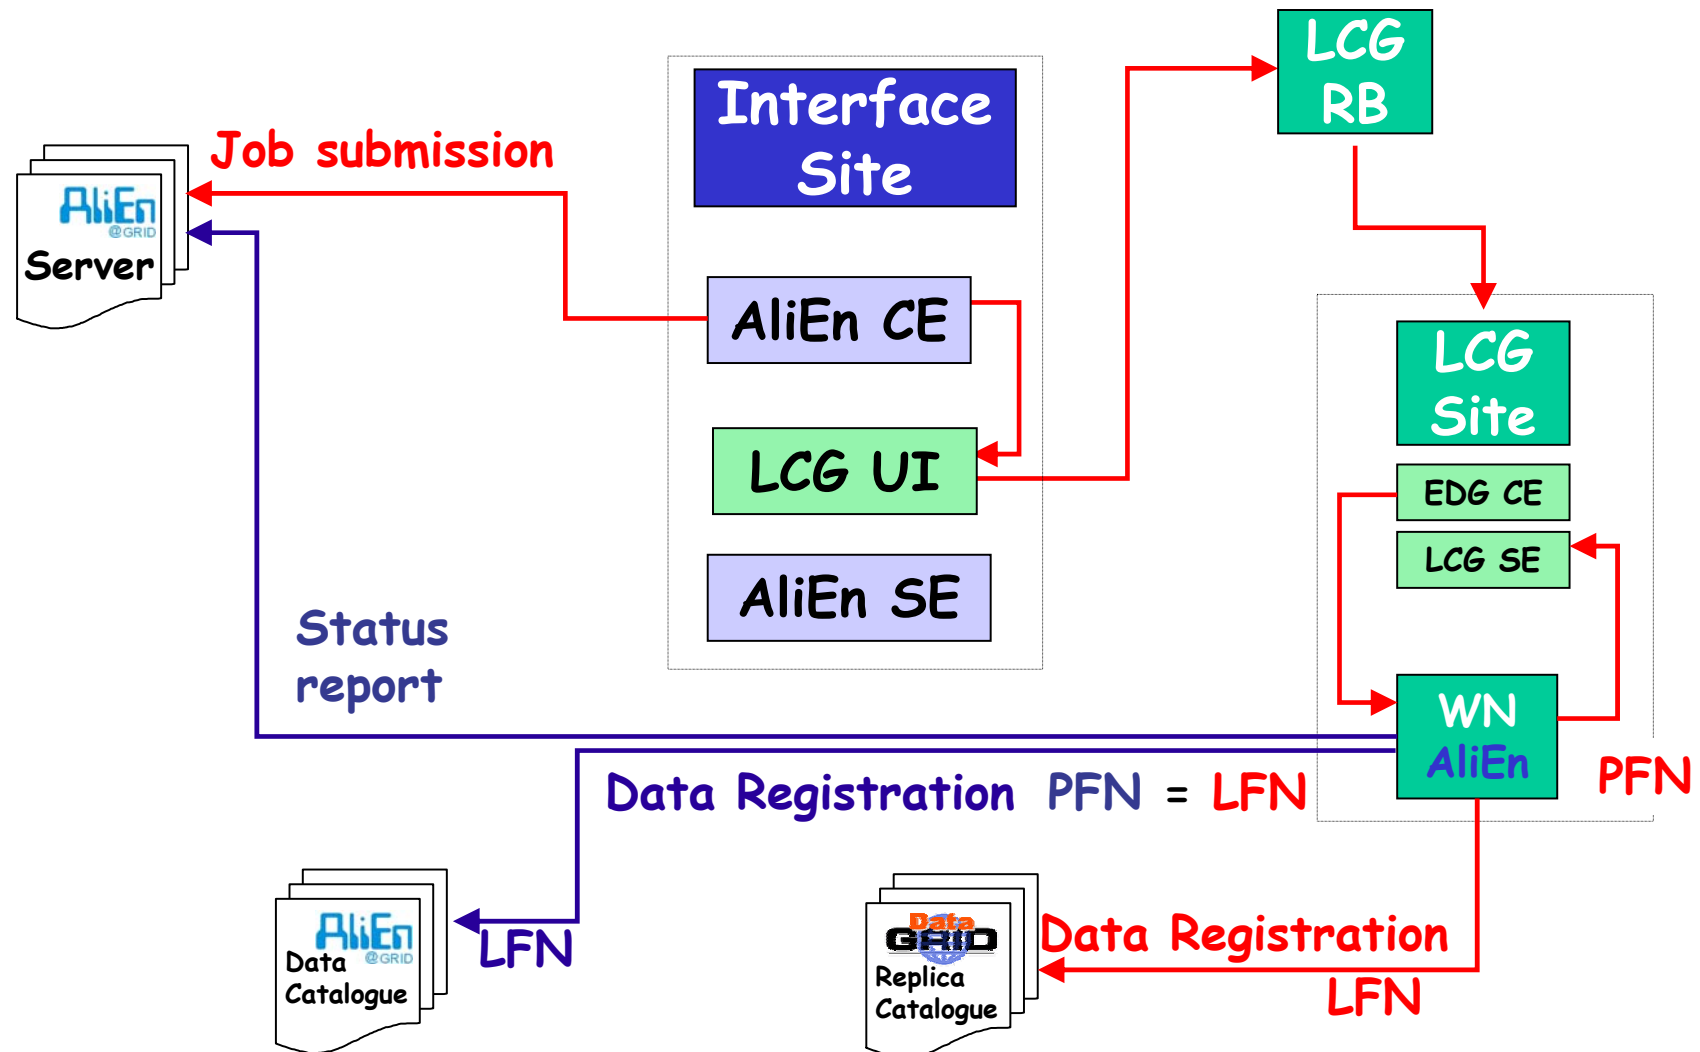
- **Use AliEn as a general front-end**
 - Owned and shared resource are exploited transparently
- **Minimize points of contact between the systems**
 - No need to reimplement services etc.
 - No special services required to run on remote CE/WNs
- **Make full use of provided services:** Data Catalogues, scheduling, monitoring...
 - Let the Grids do their jobs (they should know how)
- **Use high-level tools and APIs** to access Grid resources
 - Developers put a lot of abstraction effort into hiding the complexity and shielding the user from implementation changes

Available resources

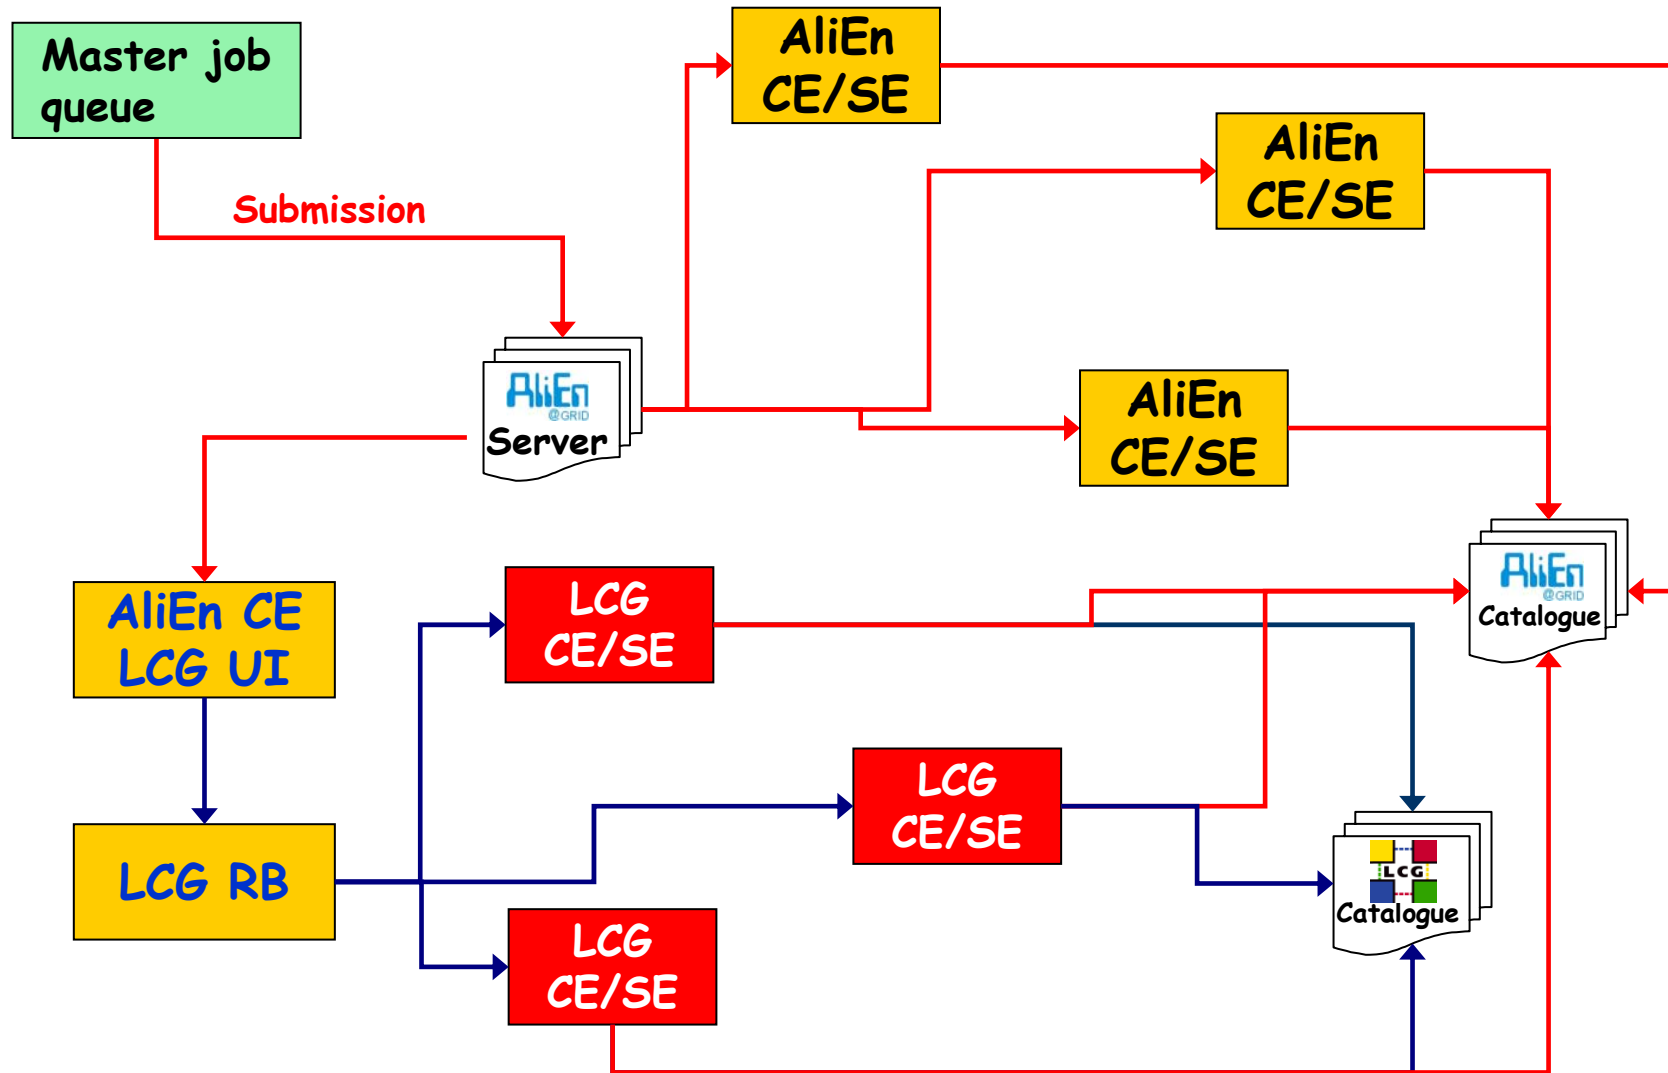
- **Several AliEn “native” sites** (some rather large)
 - CERN, CNAF, Catania, Cyfronet, FZK, JINR, LBL, Lyon, OSC, Prague, Torino
- **LCG-2 core sites**
 - CERN, CNAF, FZK, NIKHEF, RAL, Taiwan (more than 1000 CPUs)
- **INFN GRID sites**
 - LNL.INFN, PD.INFN and several smaller ones (about 400 CPUs not including CNAF)
- **Implementation:** manage LCG resources through a “gateway”: an AliEn client (CE+SE) sitting on top of an LCG User Interface

The whole of LCG computing is seen as a single, large AliEn CE associated with a single, large SE

Interfacing AliEn and LCG

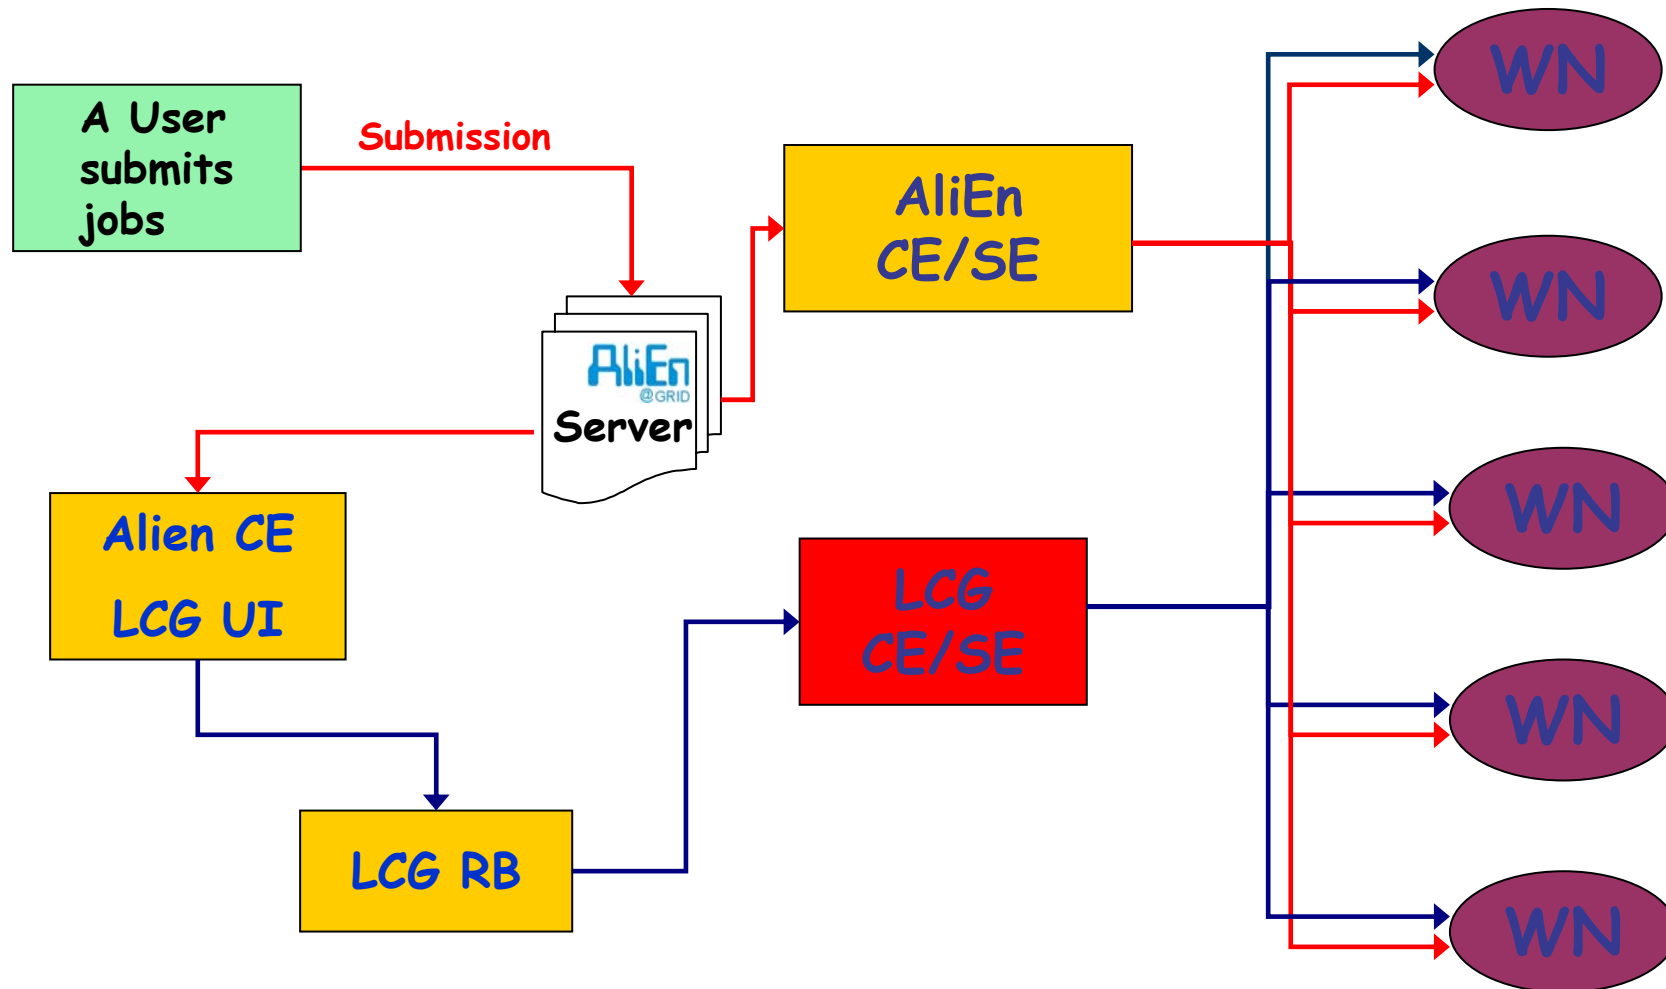


Production on two grids




Two grids, same resources!

- “Double access” for selected sites (CNAF and CT.INFN)

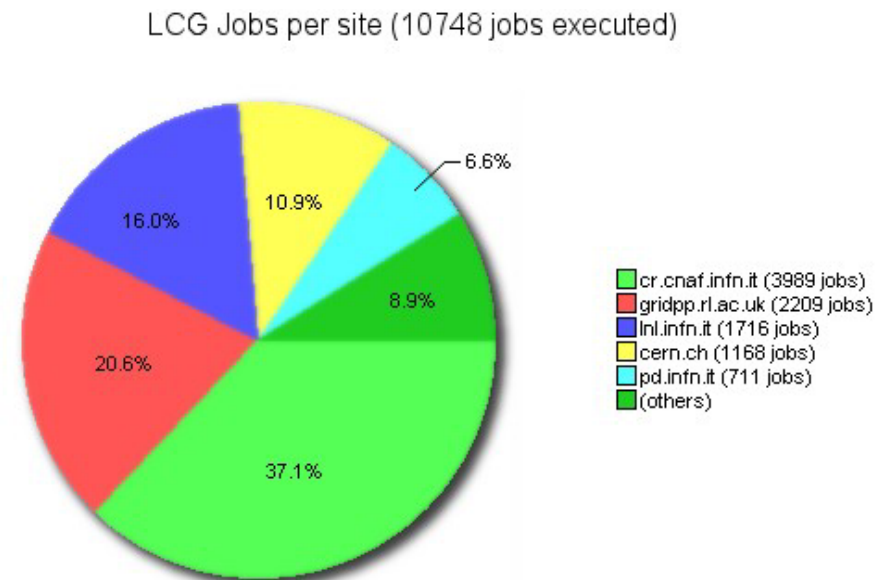
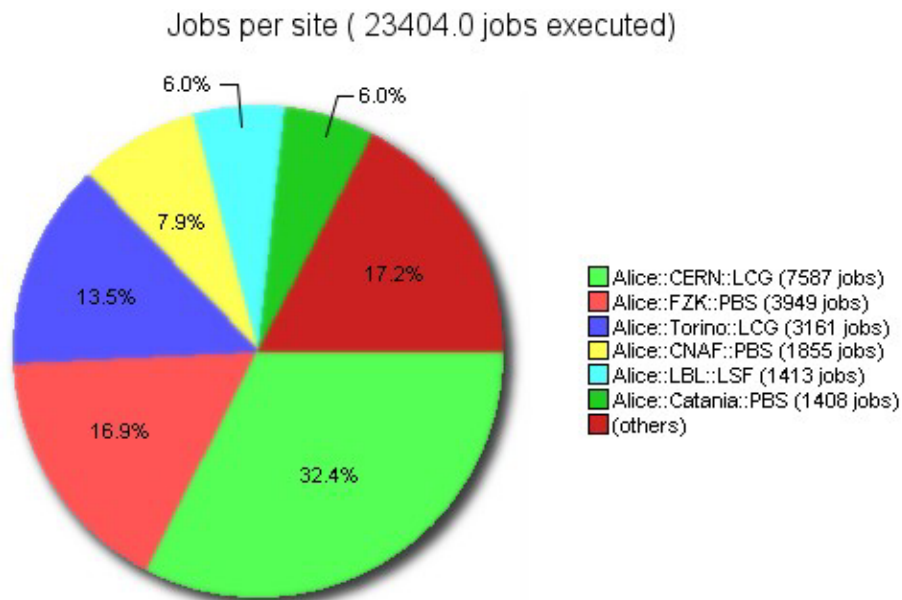


- **Phase 1: Production of RAW + Shipment to CERN (Mar-May 2004)**
 - Large output files (up to 1GB/event in ~25 files)
 - 1a: Central events (long jobs, large files)
 - 1b: Peripheral events (short jobs, smaller files)



PDC2004 - Status


- **Statistics after phase 1 (ended April 4, 2004):**
 - Alice::CERN::LCG is the interface to **LCG-2**
 - Alice::Torino::LCG is the interface to **GRID.IT**



- **~ 1.3 million files, 26 TB data volume**

Jobs history

- **Up to 1800 CPU simultaneously**
 - 1400 running jobs + 400 saving
 - Two interface sites deployed (to LCG-2@CERN, to GRID.it@TO.INFN)
 - About half “native AliEn”, half LCG-2+GRID.IT



Issues from Phase I


- **Small sites performed consistently better than large ones**
 - Often could not run for a long time on some large site trying to solve some strange idiosyncrasy
 - Some solutions not really performant – e.g. LRMS at CNAF
 - Support from LCG (EIS) very good and helpful, support from sites looks a bit “best effort” (sometimes good and prompt, sometimes less...)
- **Local site configuration and running the major source of problems**
 - Misconfigured experiment software area prevents installation or upgrade of the software
 - Strange WN configuration (std libraries...) in larger sites
 - The nasty “NFS Black Hole” (which still haunts the Grid, BTW)

Issues from Phase I cont'd

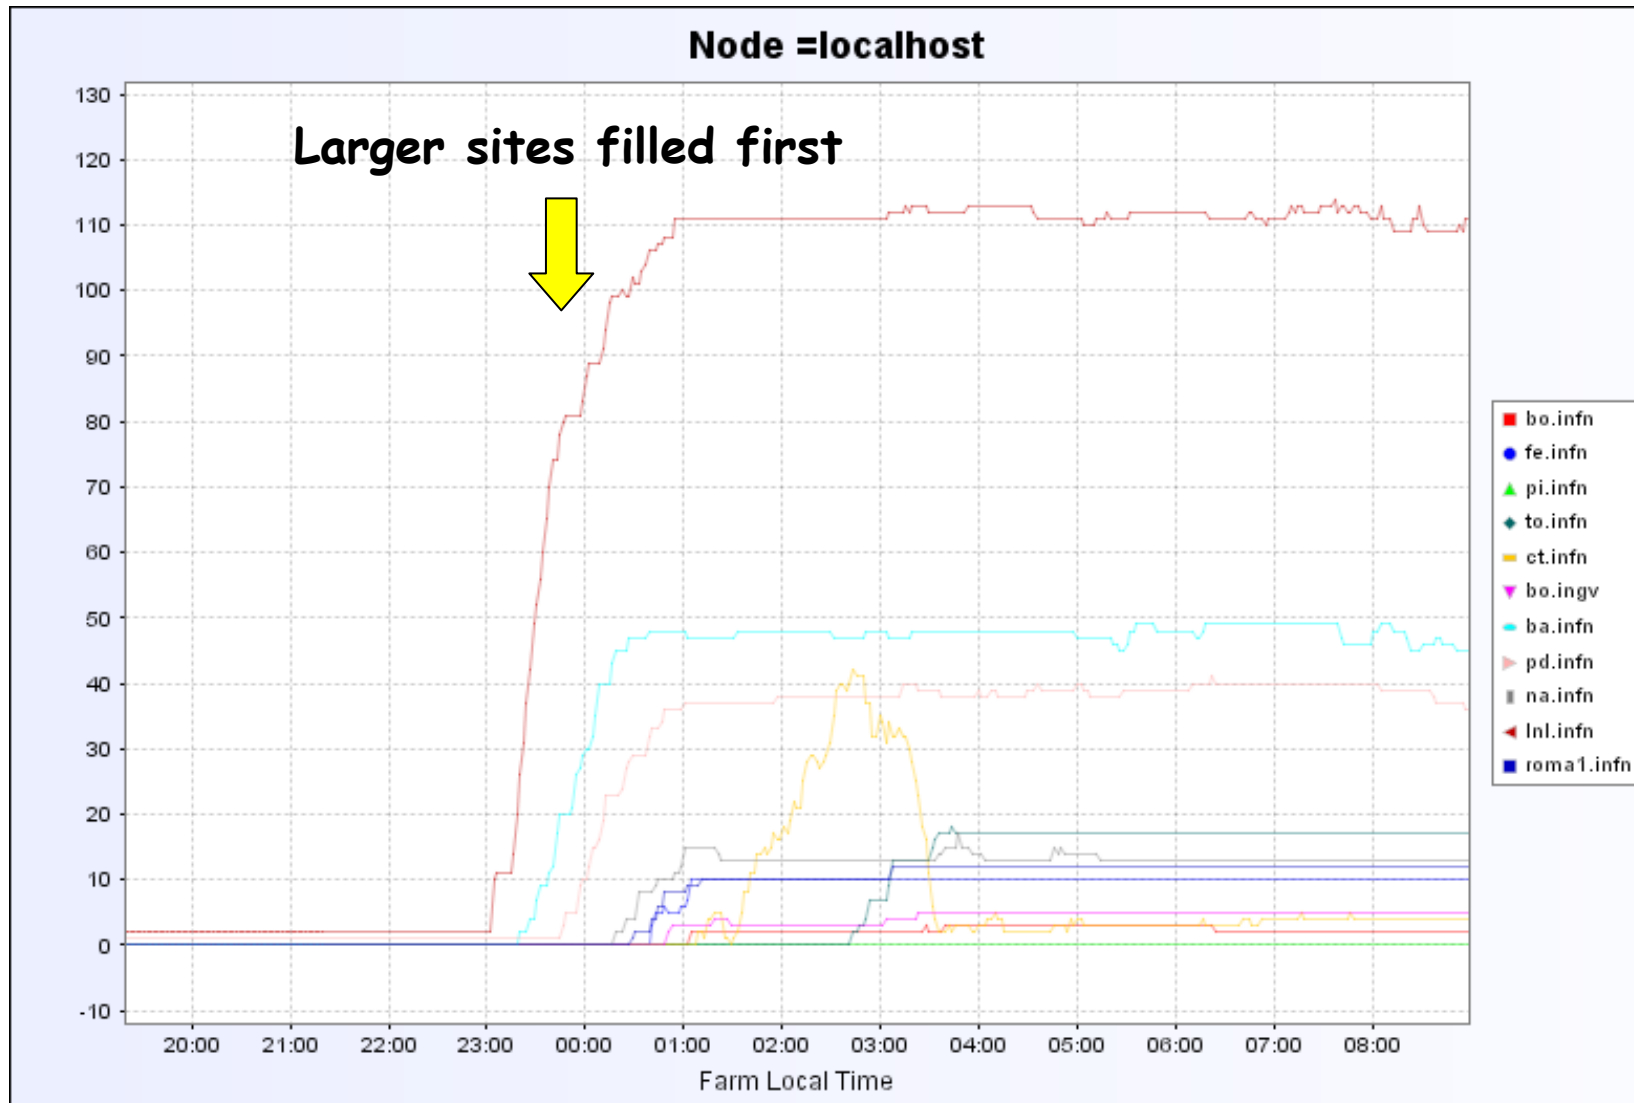
- **Some limitations to our “keyhole” approach**
 - “Keyhole” means you see all the complexity as a single site – all metrics integrated!
 - Difficult to debug: e.g. AliEn and LCG State Machines did not map onto each other
 - Different philosophies: unexecuted jobs sit (mostly) in the Master Queue in AliEn, in the local queues in LCG
- **Difficult to saturate the available resources**
 - Ranking criteria not always working as expected
 - Slow submission a severe limit for shorter jobs
 - You can favour quickly filling large sites or spreading jobs evenly, but not both

AliEn Vs. AliEn+LCG


- **LCG-2 jobs seen through AliEn MonaLisa monitoring**
 - Ramp-up slopes show small performance degradation



INFN GRID starting up



- **Phase 1:** Production of RAW + Shipment to CERN (Mar-May 2004)
 - Large output files (up to 1GB/event in ~25 files)
 - 1a: Central events (long jobs, large files)
 - 1b: Peripheral events (short jobs, smaller files)
- **Phase 2:** Merging + Reconstruction in all T1's (Jun-Sep 2004)
 - Events are redistributed to remote sites before merging and reconstruction
 - Smaller merged output (~100MB/event)




- **Phase 1: Production of RAW + Shipment to CERN (Mar-May 2004)**
 - Large output files (up to 1GB/event in ~25 files)
 - 1a: Central events (long jobs, large files)
 - 1b: Peripheral events (short jobs, smaller files)
- **Phase 2: Merging + Reconstruction in all T1's (Jun-Sep 2004)**
 - Events are redistributed to remote sites before merging and reconstruction
 - Smaller merged output (~100MB/event)
- **Phase 3: Distributed analysis**
 - Postponed...

Phase II results

- **Number of jobs:**
 - 180K, 5M events;
 - Jobs running: 430 average, 1150 max.
 - average duration: 4.5 hours
 - total CPU work: 780 MSi2K hours
- **Number of files:**
 - AliEn file catalogue: 5.4M (+3.8M from Phase I, no degradation in performance observed)
 - About 0.5M files, generated on LCG, also registered in the LCG RLS
- **File size:**
 - Total: 5.5 TB, all on remote SEs
- **Resources provided by 17 AliEn sites + 12 LCG sites**
 - Total: 780 MSI-2K hours
 - LCG: 80 MSI-2K hours (10%)

Phase II results



Issues with Phase 2

- **More complex use case**
 - Generations of a large number (10^6) of (relatively small: ~7MB) files
 - Use of local storage – even double-hopping to LCG SEs
 - However, no large or complex “data management” operations
 - the LCG fraction dropped from potentially 50% to 10%
- **Large number of files posed problems**
 - Basically a limitation in CASTOR
 - Was solved by tar'ring files together
- **Basic storage problems**
 - The space available on the some LCG SEs not always scaled to site's CPU power
 - Actually available storage space not always matched pledged from the beginning

Issues with Phase 2 – cont'd

- **Even more difficult to saturate resources**
 - Thousands of LHCb “disposable” joblets all over the place!
- **Hard to get really up-to-date information on storage**
 - Will the SE still have room for our files when, in the end, they get to run?
 - Our output files are large; no information available on the local WN disk
- **Small sites on the average performed better than large ones**
 - We were essentially never able to run smoothly and continuously on Tier-1s

Conclusions (ARDA workshop 2004)

- LCG support & performance is good
- Local site configuration issues are the largest source of problems
- LCG-2 shows some instabilities, but we are really pushing it!
- Heavy testing of the LCG-SE may bring more surprises, as large parts of it are new
- LCG-2 really looks and feels as a first-generation Grid
 - Very good for getting experience
 - Needing a new fresh start to build a production-grade system
- **Hope to test the first ARDA prototype in late summer/autumn**