Middleware and Interaction with EDG/EGEE and other Developments in Middleware

LCG Comprehensive Review

22-23 November 2004

F.Forti, *for the referees*

Outline

- Introduction
- Middleware, development model and EGEE
- LCG-2, gLite, ARDA
- LCG-2 achievements and limitations
- gLite status and plans
- ARDA status and plans
- Concerns and recommendations

As always, we concentrate on problems.

Introduction

- The referees congratulate the LCG project for the large progress shown in all fields.
- A large effort was put in the review preparation: two days of in-depth and informative talks and useful discussion.
  - Schedule conflict with EGEE workshop prevented us from having more discussion with people. To avoid in the future.
- Compared to the 2003 C.R., the tone of presentations is more concrete and specific: a welcome shift, keep going !
  - We would welcome more contact with developers, Tier1 managers, etc.
- Many new referees and committee members: significant effort to understand and remember all TLAs. Thanks for your patience in explaining the obvious.
- Still, there are some areas where a level of fuzziness remains in our understanding, and we haven't been able to penetrate the technical details

Overview of the review

- Middleware and Interaction with EDG/EGEE and other Developments in Middleware (30')
  F. Forti
- LCG Fabric and Wide Area Networking (30')
  V. Guelzow
- Grid Deployment and Regional Centres (30')
  D. Boutigny
- LCG Application Areas (30')
  C. Niebuhr
- LCG Status Update, Management and Planning (30')
  T. Wyatt

- + P. McBride chief referee (getting away without a talk)

- We haven't had time to polish our talks and remove all overlaps

Functional Overview of a Grid

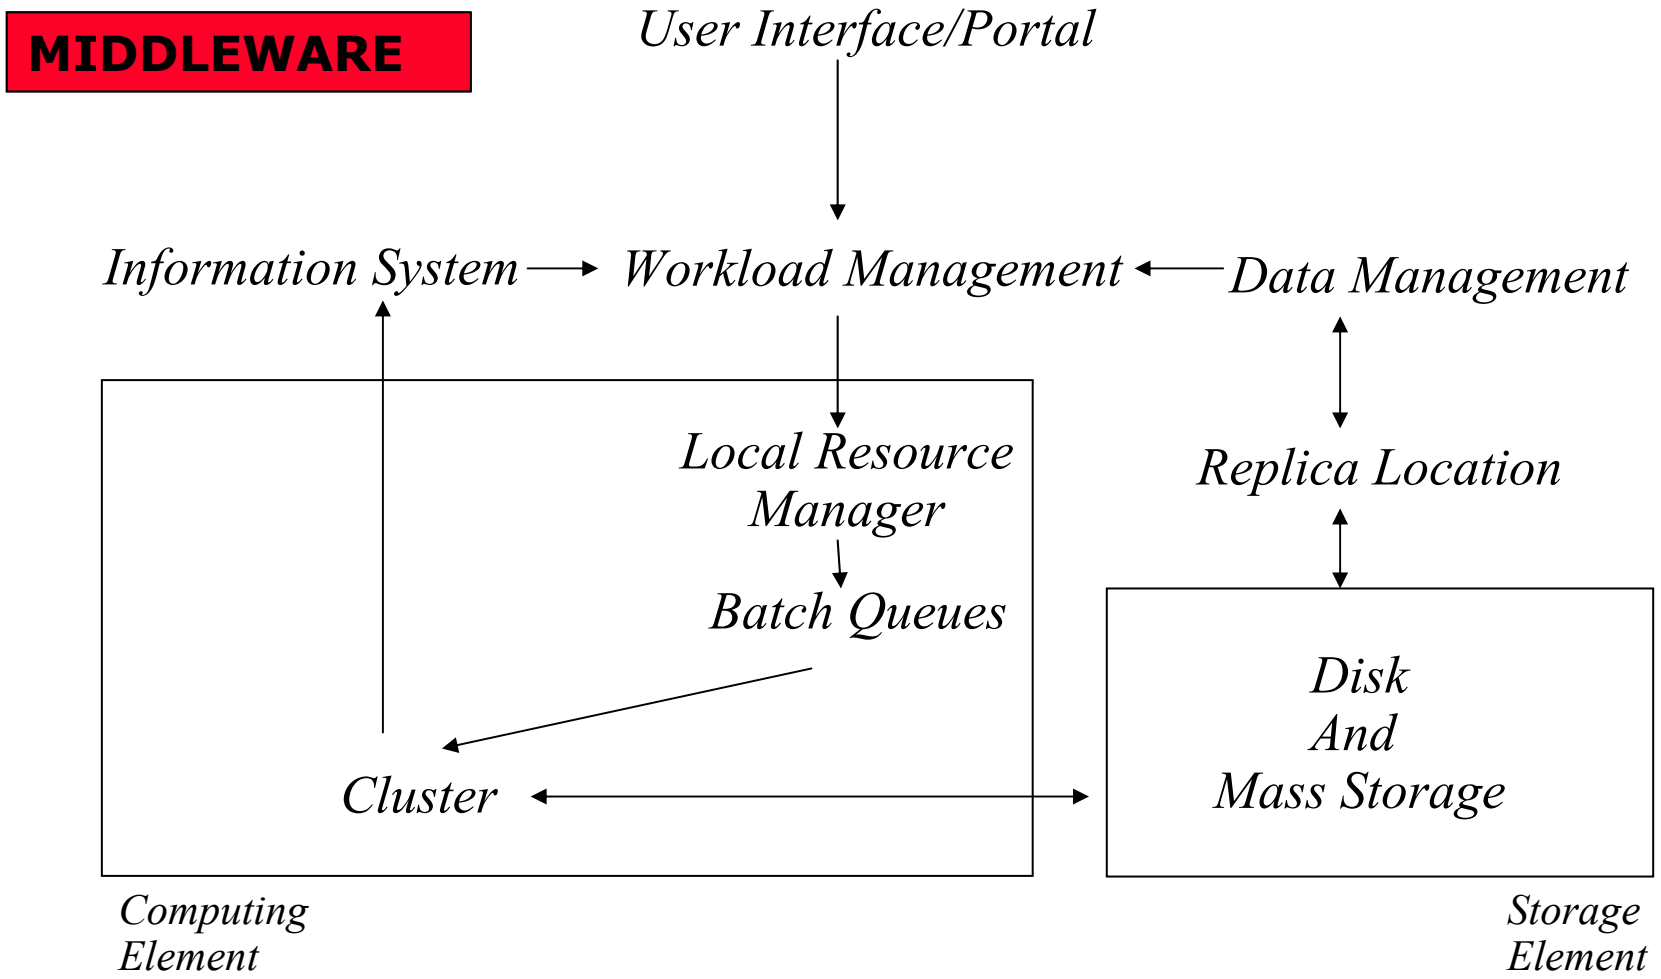

Middleware development

- Middleware is not HEP-specific
  - Developed outside LCG direct control, but in close collaboration and contact:
    - Previously EDG, now EGEE.
    - Pieces and ideas coming from many sources
      - Globus, VDT, Condor, Alien, EDG, ……
- First middleware, LCG-2, has been deployed and used in data challenges 2004
  - A success, albeit qualified, see Dominique's talk.
- New middleware development launched with EGEE:
  - Code name is gLite (Lightweight Middleware for Grid Computing)
  - Has the LCG Big Brother ARDA (A Realization of Distributed Analysis for LHC)

ARDA, gLite, and HEP experiments

ARDA is an LCG project whose main task is to enable LHC analysis on the GRID

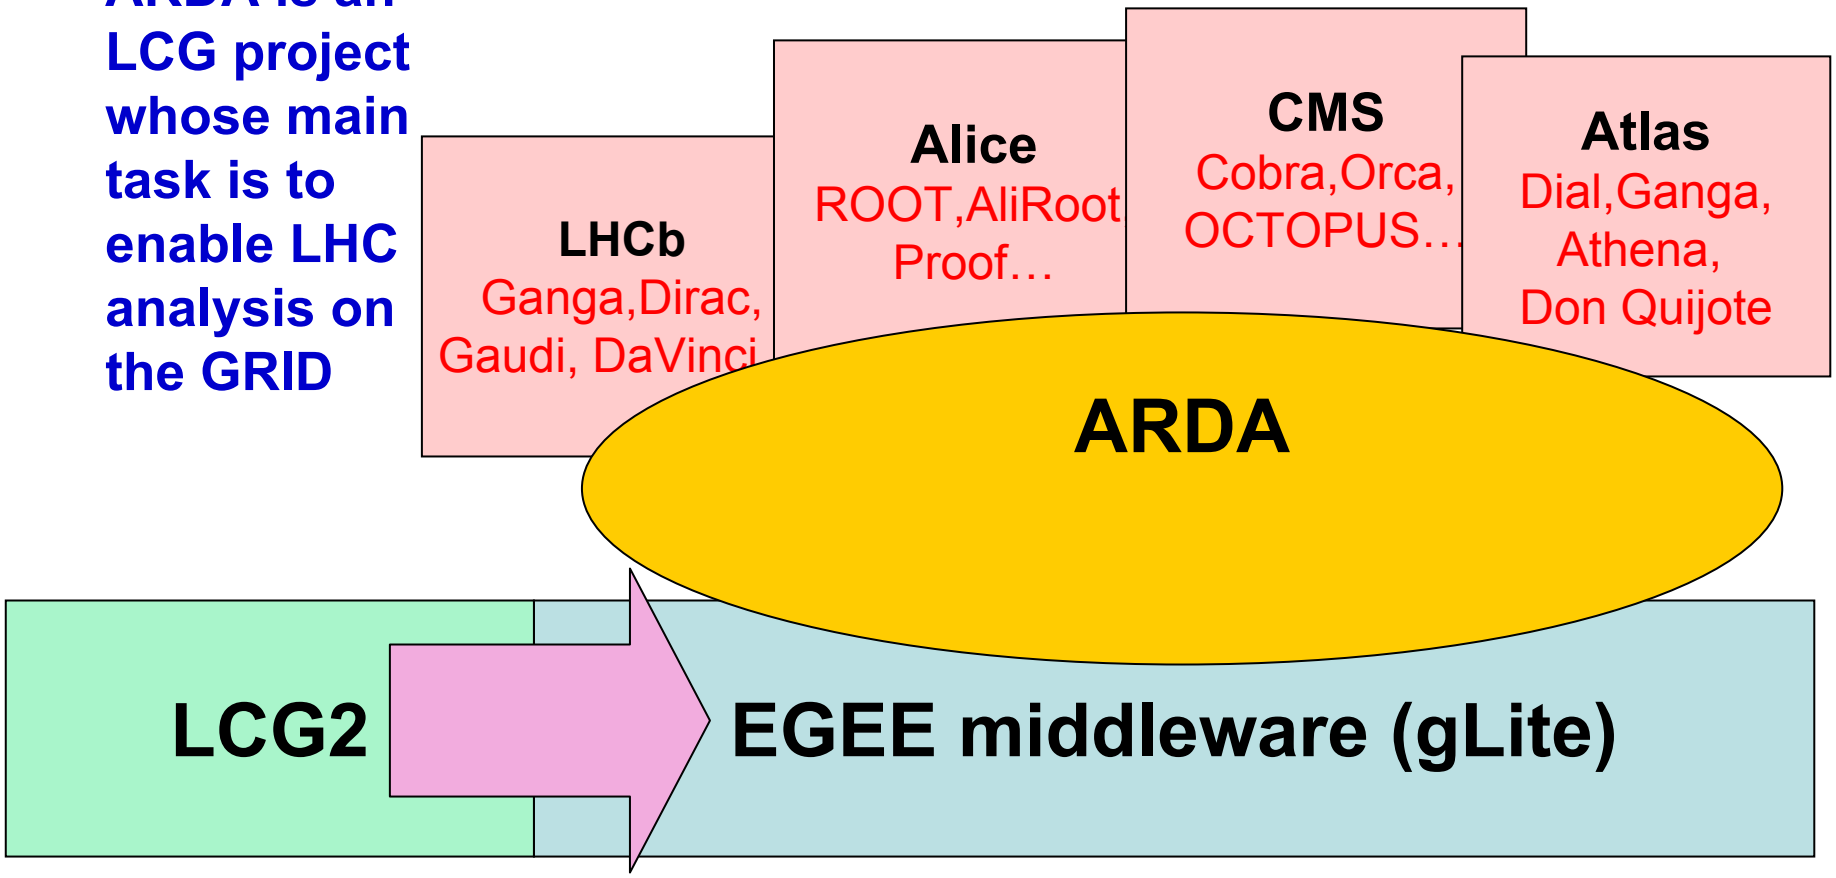

Relation of LCG and EGEE

- LCG has been setting up the operation of a grid infrastructure for HEP experiments
  - combining national and regional grids
  - middleware certification, integration, distribution, operation management, ..
  - agreed middleware package – currently “LCG-2”
- EGEE has a goal of extending this into a general grid infrastructure for science in Europe and beyond
- EGEE brings new manpower –
  - 50% of its budget is for operation
  - 70 partners – almost all involved in HEP
- EGEE is funded for 2 years, with a *probable* extension for at least another 2 years
- LHC comes fully online just as EGEE ends

LCG & EGEE

- Goals of LCG in working closely with EGEE
  - Maximise the benefits of EGEE funding for LHC
  - Generalise the HEP grid to other sciences
  - Be well positioned to benefit from a future science grid infrastructure (cf GEANT research network backbone in Europe)
  - BUT - ensure that LCG can continue if there is no EGEE successor
- LCG-2 and EGEE share the core operation
  - LCG Grid Deployment Manager acts as EGEE Operations Manager
- EGEE middleware development activity has a strong focus on HEP requirements
  - Hardening, re-engineering of functionality needed by LCG (***gLite***)
  - Very close collaboration between EGEE middleware and the ARDA project of LCG (early access to development prototypes, regular technical meetings)
  - Commitment to deliver rapidly to LCG/EGEE deployment – - taking over, as quickly as practicable, from the LCG-2 middleware
  - EGEE middleware manager sits in the PEB
- EGEE provides funding (4 FTEs) for ARDA
  - Which in turn acts as the HEP pilot application for EGEE
- Cross-representation on management boards (PEB, POB)

LCG2

- LCG2 has been deployed to 90-odd sites for the data challenges.
- Largely successful, but:
  - Overall job success rate not satisfactory (60%)
    - Due to composition on many different problems, not necesarily related to middleware
    - Installation and configuration of MW and experiment SW are crucial and part of the problem
  - Workload management system insufficient
    - Information about remote queue situation unreliable
    - Slow response and lack of bulk submit limit the throughput
  - Data management system insufficient
    - Some tools just came too late (lcg-utils replacing edg)
    - In many case a “perfect network” is assumed

Other middleware

- In the meantime other flavours of middleware have emerged
  - Grid3 → Open Science Grid (OGS)
    - Based on the same tools and schema, with relatively small incompatibilities
    - Deployed in the US network
  - NorduGrid
    - Based on significant modifications of the Globus toolkit, and much more incompatible
    - Deployed on nordic countries (including sites in Switzerland and Germany)
    - Atlas-only
  - AliEN
    - Based on different concepts (pull instead of push)
    - Developed by Alice collaboration, largely based on ROOT experience
    - Deployed only on Alice computers
- Interoperability becomes mandatory

The genesis

Now the world of computers was formless and empty, darkness was over the surface of the storage element, and the Spirit of Grid was hovering over the workload manager.

And Grid said……

“Let there be gLite”

- The plan, broadly:
  - Start from the existing middleware
    - EDG, Alien, Globus, VDT, ……
  - Factor in feedback
    - “ARDA Blueprint” redefining functional decomposition
    - GAG (Grid Application Group) reports
  - Produce new, lightweight, refactored and repackaged middleware
    - Feedback from development testbed is part of the cycle
    - Focus on usability and matching user’s requirements.

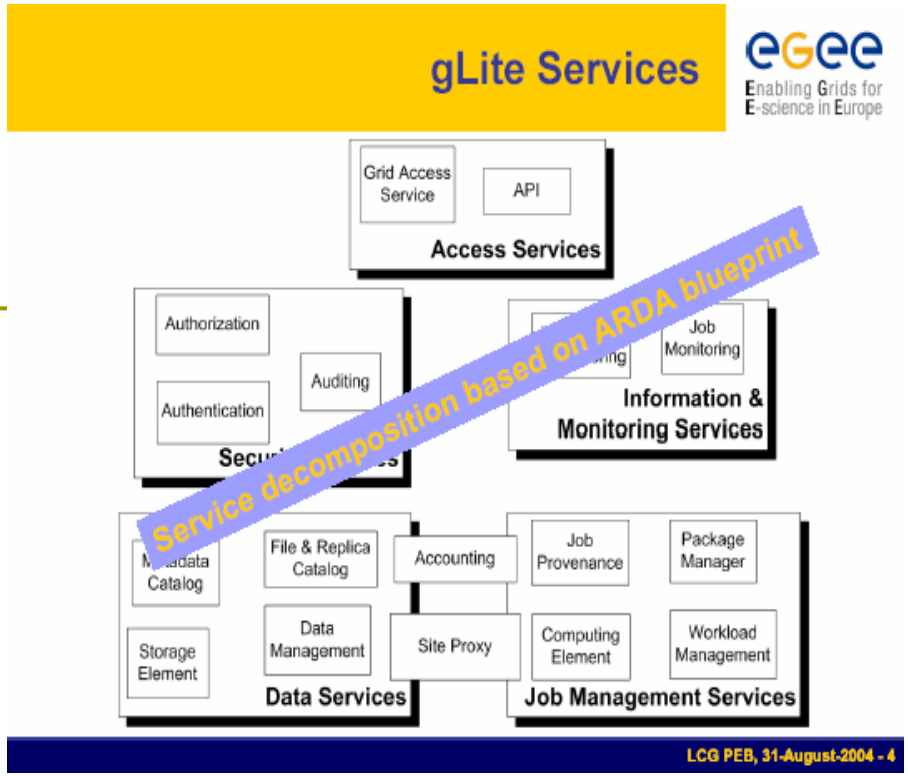

Architecture Guiding Principles

- Lightweight (existing) services
  - Easily and quickly deployable
  - Use existing services where possible as basis for re-engineering
- Interoperability
  - Allow for multiple implementations
- Resilience and Fault Tolerance
- Co-existence with deployed infrastructure
  - Run as an application (e.g. on LCG-2; Grid3)
  - Reduce requirements on site components
    - Basically globus and SRM
  - Co-existence (and convergence) with LCG-2 and Grid3 are essential for the EGEE Grid service
- Service oriented approach
  - WSRF still being standardized
  - No mature WSRF implementations exist to date, no clear picture about the impact of WSRF hence: start with plain WS
    - WSRF compliance is not an immediate goal, but we follow the WSRF evolution
    - WS-I compliance is important

All GOOD THINGS

Turning on gLite.

- What has been done so far
  - Staffing completed in April 2004 – about 80 people
    - We didn't see the details of that
  - Design Team
    - Architecture and design document June 2004
    - Design document August 2004
    - Used also by OSG (hope for interoperability)
  - Development Testbed (a.k.a. prototype)
    - Available as of May 2004
    - Collection of many existing tools (based on AliEn)
    - Used by **ARDA** to develop physics analysis prototypes
    - Deployed to CERN + Wisconsin
  - Software Development, Integration and Testing Processes
    - Important for successful deployment
    - Testing clusters at CERN, Nikhef, RAL
  - … and even some Software

… and there was gLite.

Schedule for pre-production service

• gLite I/O	–	Available
• Logging & Bookkeeping, WMS, CE, WN	–	In testing – end November
• R-GMA	–	In integration/testing – mid December
• CE-Notification	–	In integration – mid/end November
• Replica, File, Combined Catalog	–	In development – December
• File Transfer Service	–	In integration/testing – December
• File Placement Service	–	In integration/testing – December
• VOMS	–	In integration/testing – December
• UI	–	In integration – December
• AliEn Task Queue & CE	–	In integration/testing – To be deployed on ALICE sites
• Package Manager	–	Discussions w/experiments, deployment – prototype exists
• Grid Access	–	Prototype exists
• Accounting (DGAS)	–	In integration – Prototype exists
• Job Provenance	–	Proof of concept exists

…not quite as bright as it should be.

- Many gLite components are late and eagerly expected
  - Nervousness expressed about what exactly it will look like when released
  - Risk of being overly ambitious and wanting to include too much functionality
  - "Hope it's not hopeless" to have the developers support the deployment instead of working on the next functionality.
- Experiments need a more or less stable service for their TDRs
  - This is LCG-2, which therefore needs short-term support.
  - LCG-2 is also the fallback in case gLite fails.
  - This may in turn delay the usage of gLite
- Still, if it works as advertised it will represent a big step forward.

gLite next steps

- Enlarge development testbed and open it to a larger community
  - Although only a prototype, it must be used to a wide community to provide feedback
- Deliver scheduled components to pre-production service
  - we aren't sure exactly what the pre-production service is
  - As they become available, even if in bits and pieces, to ensure early exposure to community feedback
  - Understand how to migrate from LCG-2 to gLite
- Deploy prototype Middleware to ALICE sites for DC
  - Needed resource still in discussion
  - Is this really the right strategy ?
- Further steps
  - Finalize contents of EU RC1 release
  - Deliver and Test all components out of integration builds
  - Finalize first integrated release as an EU deliverable

gLite comments

- gLite is clearly the way forward, that capitalizes on previous developments and re-engineers things that need it
- The delay in the delivery of gLite is a real concern, and all effort should be made to conlude this development rapidly, delivering every piece to the development testbed as it becomes available
  - Nevertheless integration and certification should be carried out rigorously, especially for the deployment on the pre-production service
- Connection with EGEE introduces level of fuzziness and poses difficult issues including
  - Unclear whether manager has real authority over people working in the project
  - Will the developers support gLite once it is delivered and during deployment ?

ARDA, gLite, and HEP experiments

ARDA is an LCG project whose main task is to enable LHC analysis on the GRID

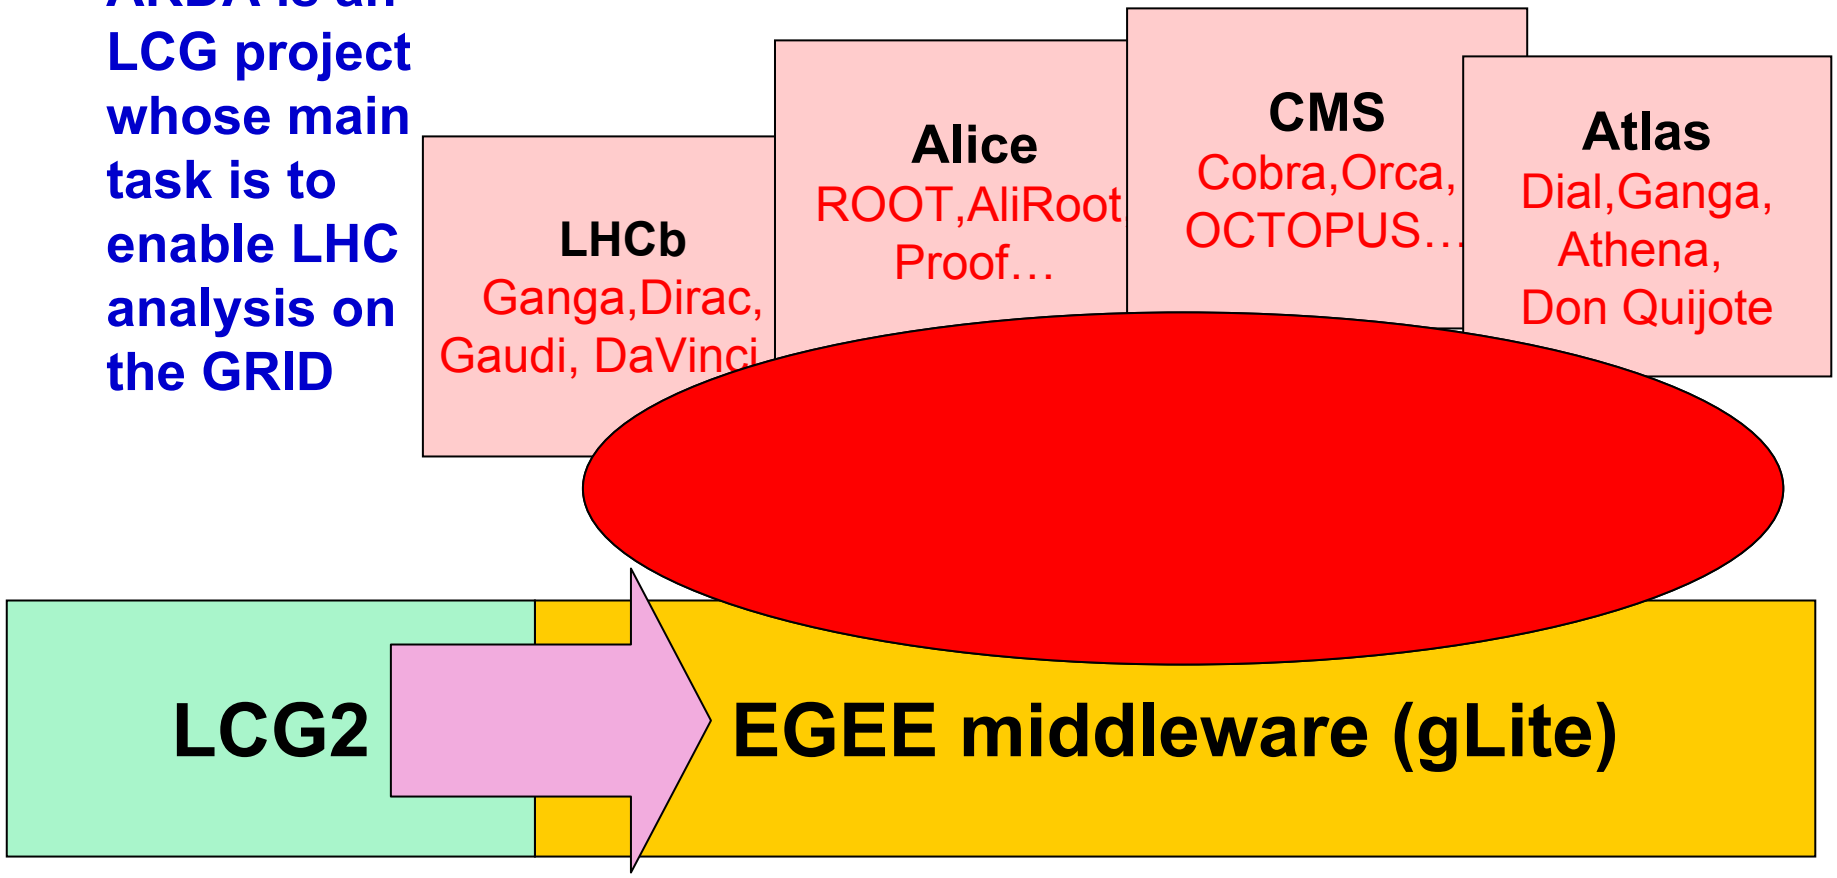

ARDA People

- Massimo Lamanna
- (EGEE NA4 Frank Harris)
- Birger Koblitz

- Dietrich Liko
- Frederik Orellana

ALICE

- Derek Feichtinger
- Andreas Peters

ATLAS

- Julia Andreeva
- Juha Herrala

CMS

- Andrew Maier
- Kuba Moscicki

LHCb

- Andrey Demichev
- Viktor Pose
- Alex Kryukov

Russia

- Wei-Long Ueng
- Tao-Sheng Chen

Taiwan

Visitors

- 2 PhD students (just starting)
- Many students requests

Experiment interfaces

Piergiorgio Cerello (ALICE)
David Adams (ATLAS)
Lucia Silvestris (CMS)
Ulrik Egede (LHCb)

ARDA Mission

- Interface to existing middleware to enable their use in the experiment frameworks
  - Test and compare different m/w pieces
- Early deployment of (a series of) analysis prototypes to ensure functionality and coherence
  - Provide physicists something to try
- Provide link between experiments and middleware

ARDA Activity

- Used components available in the gLite prototype testbed
  - Workload management systems
    - AliEn Taskqueue and EDG WMS
    - Job submission: job splitting policies
  - Data Management Systems
    - gLite file catalog and FireMan (just delivered)
  - Package Management
    - Dynamic vs. static experiment software installation
  - Metadata studies
    - Need to be able to efficiently select data with a query on metadata (eg “I want the dimuon sample”)
    - Feedback to gLite on this requirement

ARDA Prototypes overview

LHC Experiment	Main focus	Basic prototype component	Experiment analysis application framework	Middleware prototype
LHCb	**GUI to Grid**	GANGA	**DaVinci**	gLite
ALICE	**Interactive analysis**	PROOF ROOT	**AliROOT**	gLite
ATLAS	**High level service**	DIAL	**Athena**	gLite
CMS	**Use of maximum native gLite functionality**	Aligned with the APROM activity	**ORCA**	gLite

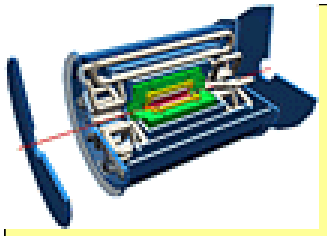

ATLAS Combined Test Beam

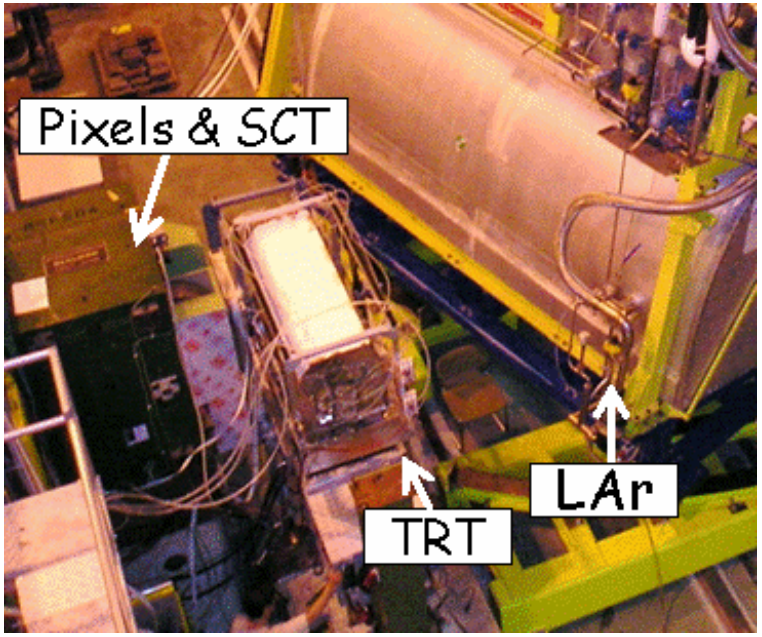

Real data processed at gLite

Standard RecExTB

Data from CASTOR

Processed on gLite worker node

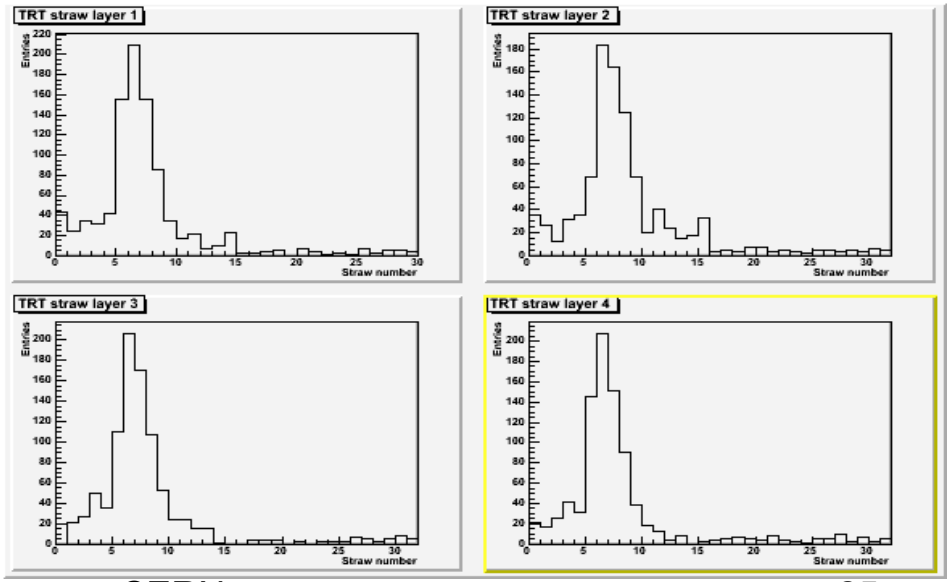

Example:

ATLAS TRT data analysis done by PNPI St Petersburg

Number of straw hits per layer

ARDA other activities

- Exploring the possibility to expand the gLite prototype testbed
  - Good perspective with FZK. Should explore CNAF as well
  - Installing prototype gLite on ALICE sites welcome but:
    - Need to ensure no divergence between the middleware on the infrastructure used by ALICE for their DC04 and the regular prototype and deployment chain.
- Help gLite team on integration and operational issues
  - Extra people on board soon
  - Not clear why gLite large team can't handle this
- Participate in EGEE NA4
  - Networking activity: Application Identification and Support
  - Contact with other sciences and be in the loop for future gLite development
  - Get 4 extra FTEs from EGEE

Summary and Comments

- The current LCG-2 m/w is working, albeit with limitations.
- gLite is the way forward and must be pursued vigorously and rapidly
- Connection between gLite, LCG and the experiments seems to be too weak, and there is the risk that gLite, when delivered, will not satisfy the experiments needs.
- In particular ARDA should renew its effort to involve experiment people in running analysis on the gLite prototype and provide feedback
- We feel the effort should be focused on what software is absolutely needed for the first year of data taking. Avoid duplication of effort
- Interoperability of different grid flavor is clearly become an absolute need, since there will be more than one grid.
- Is it possible to define a common interface protocol which grids participating in LHC must follow ? Is the Grid Deployment Board the right body to take this on ?

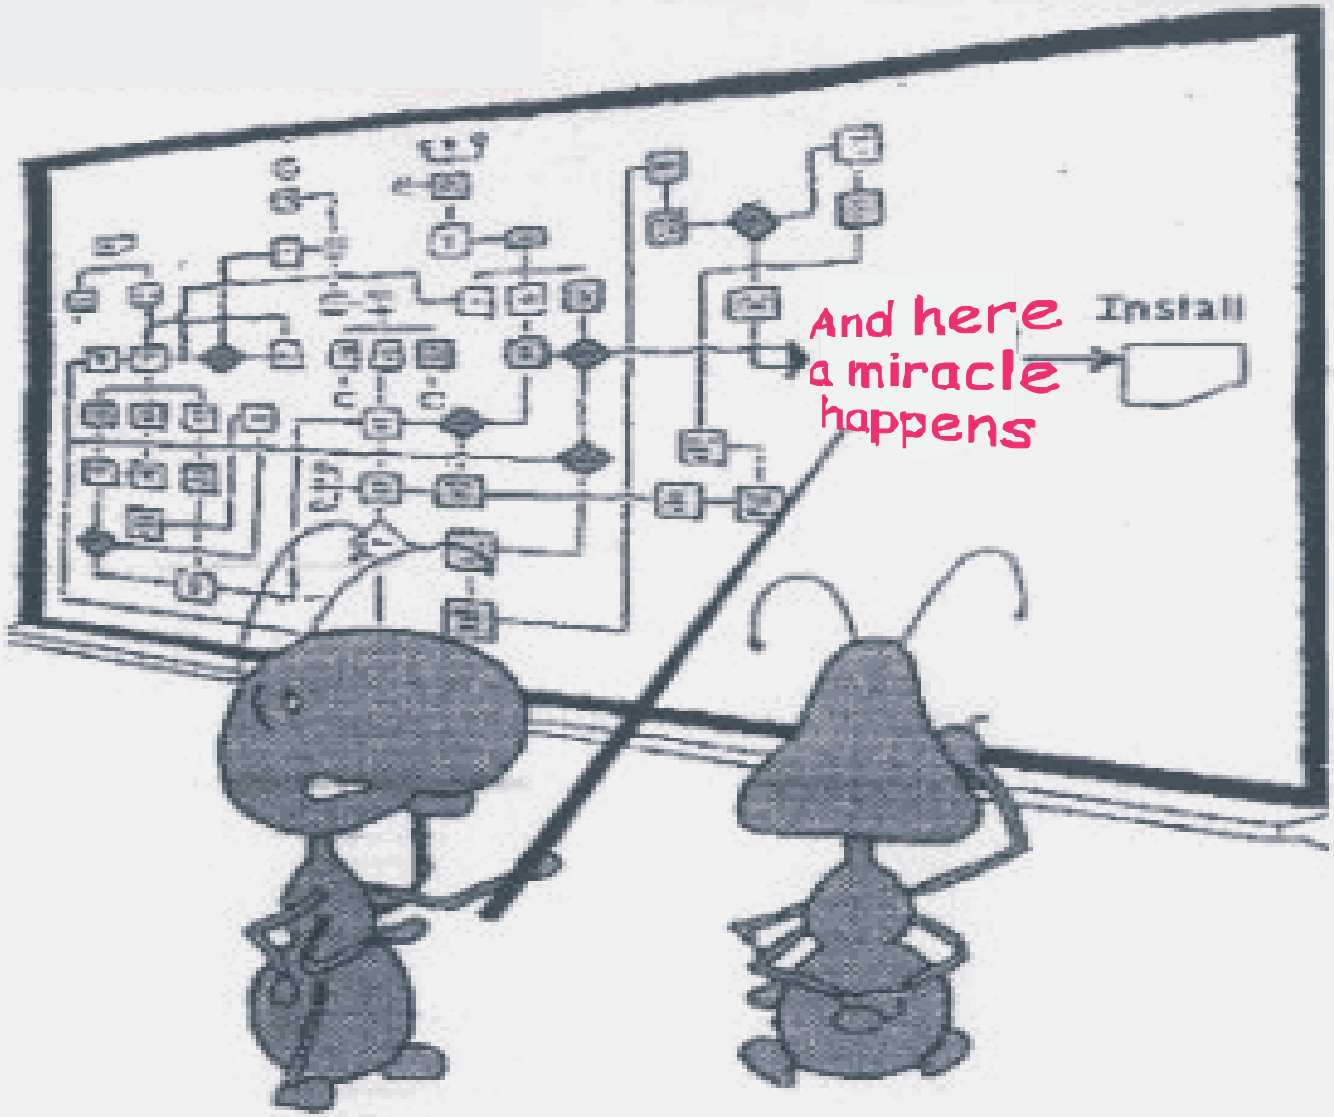

Very nice work!
Dont you think we should perhaps get a little bit more detailed here...?