http://cern.ch/arda

LCG PEB, 1 June 2004

"The ARDA project: status report"

Massimo Lamanna

Enabling Grids for E-science in Europe

www.eu-egee.org

cern.ch/lcg

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

Contents

- ARDA Project status
  - Installation
  - Planning of activity
  - Activity (so far and plans)
    - Experiment software
    - Experiment prototypes
  - Forum activities

People

- Massimo Lamanna
- Birger Koblitz

- Andrey Demichev
- Viktor Pose

Russia

- Derek Feichtinger
- Andreas Peters

ALICE

- Wei-Long Ueng
- Tao-Sheng Chen

Taiwan

- Dietrich Liko
- Frederik Orellana

ATLAS

- Julia Andreeva
- Juha Herrala

CMS

- Andrew Maier
- Kuba Moscicki

LHCb

Experiment interfaces

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

Logistics and installation in bd. 510

- Technicalities and preliminary installation solved
  - Erwin Mosselmans
  - John Harvey
  - John Fergusson
- Final installation more or less completed
  - Not easy…
  - End of April, all people had a desk close to bd.510
- Probably a bit more space would be necessary
  - PhD students (over 2…)
  - F. Harris
  - Room for visitors (more coming?)
  - At least 2/3 phone conferences a week

Preliminary activities

- Existing system as starting point
  - Every experiment has different implementations of the standard services
  - Used mainly in production environments
    - Few expert users
    - Coordinated update and read actions
  - ARDA
    - Interface with the EGEE middleware
    - Verify (help to evolve to) such components to analysis environments
      - Many users
        - » Robustness
      - Concurrent "read" actions
        - » Performance
- One prototype per experiment
  - A Common Application Layer might emerge in future
    - ARDA emphasis is to enable each of the experiment to do its job

Glite "disclosed" May 18th ☺

Since the beginning

Time consuming (see next section)

Milestone	Date	Description
1.6.2-5	May 2004	E2E x prototype definition agreed with the experiment
1.6.6-9	September 2004	E2E x prototype using basic EGEE middleware
1.6.14-17	November 2004	E2E x prototype [illegible]
1.6.18	December 2004	E2E prototype for experiment x, capable of analysis
1.6.19	December 2005	E2E prototype for experiment x, capable of analysis and production

All ARDA milestones

LHCb

- The LHCb system within ARDA uses GANGA as principal component.
- The LHCb/GANGA plans:
  - enable physicists (via GANGA) to analyse the data being produced during 2004 for their studies
  - It naturally matches the ARDA mandate
  - Have the prototype where the LHCb data will be the key (CERN, RAL, …)
- At the beginning, the emphasis will be to validate the tool focusing on usability, validation of the splitting and merging functionality for users jobs
- The DIRAC system (LHCb grid system, used mainly in production so far) could be a useful playground to understand the detailed behaviour of some components, like the file catalog. Convergence between DIRAC and GANGA foreseen.

ARDA contribution to Ganga

- Integration with EGEE middleware
  - Waiting for the EGEE middleware, we developed an interface to Condor
  - Use of Condor DAGMAN for splitting/merging and error recovery capability
- Design and Development
  - Command Line Interface
  - Future evolution of Ganga
- Release management
  - Software process and integration
    - Testing, tagging policies etc.
  - Infrastructure
    - Installation, packaging etc.
  - It looks to be effective!

LHCb Metadata catalog

- Used in production (for large productions)
- Web Service layer being developed (main developers in the UK)
- Oracle backend
- ARDA contributes a “testing” focused on the analysis usage
  - Robustness
  - Performances under high concurrency (read mode)

Measured network rate
vs no. of concurrent clients

CERN/Taiwan tests

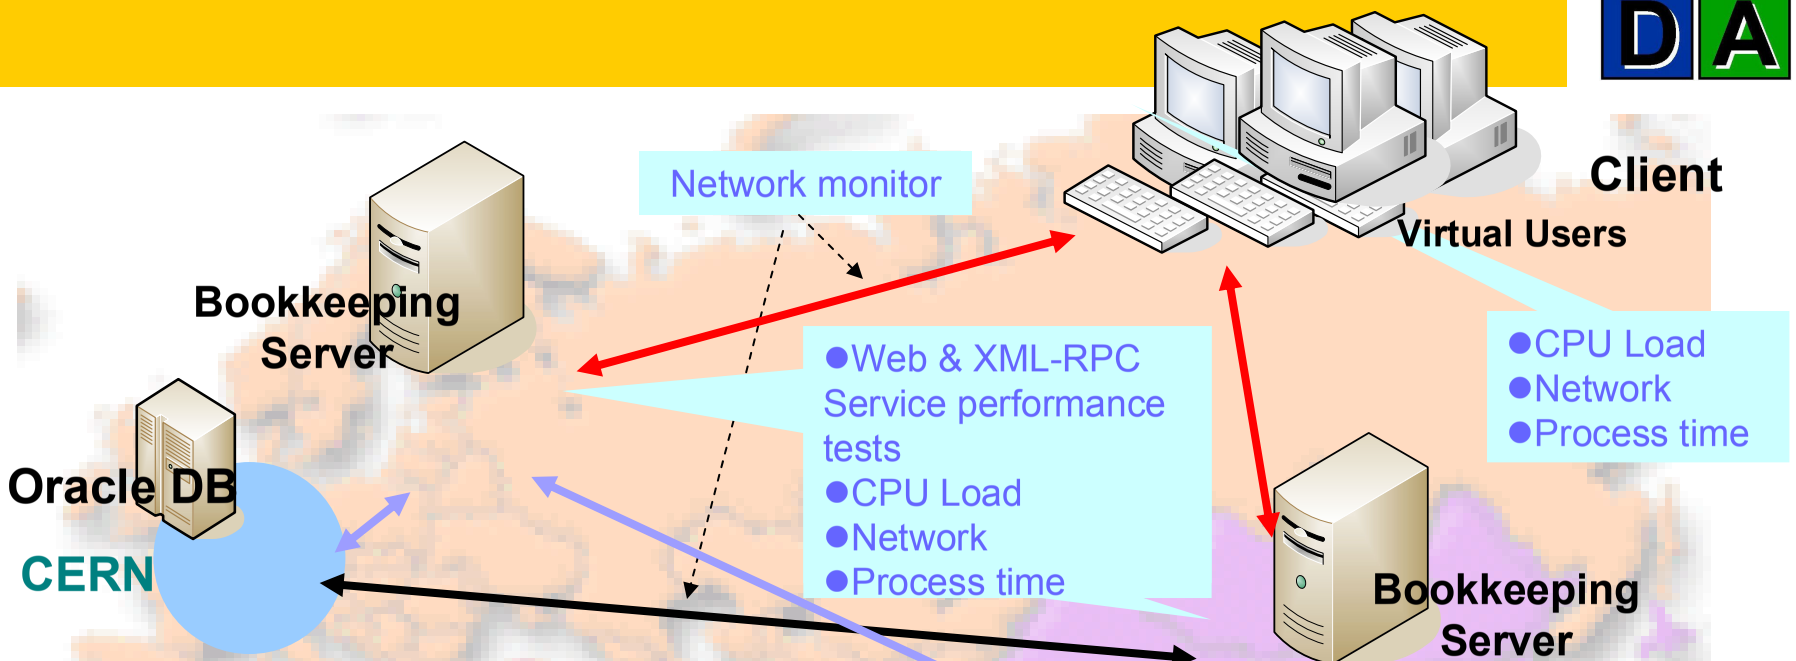

- Clone Bookkeeping DB in Taiwan
- Install the WS layer
- Performance Tests
  - Database I/O Sensor
  - Bookkeeping Server performance tests
    - Taiwan/CERN Bookkeeping Server DB
    - XML-RPC Service performance tests
    - CPU Load, Network send/receive sensor, Process time
  - Client Host performance tests
    - CPU Load, Network send/receive sensor, Process time

ALICE

- Strategy:
  - The ALICE/ARDA will evolve the analysis system presented by ALICE at SuperComputing 2003
- Where to improve:
  - Heavily connected with the middleware services
  - “Inflexible” configuration
  - No chance to use PROOF on federated grids like LCG in AliEn
  - User libraries distribution
- Activity on PROOF
  - Robustness
  - Error recovery

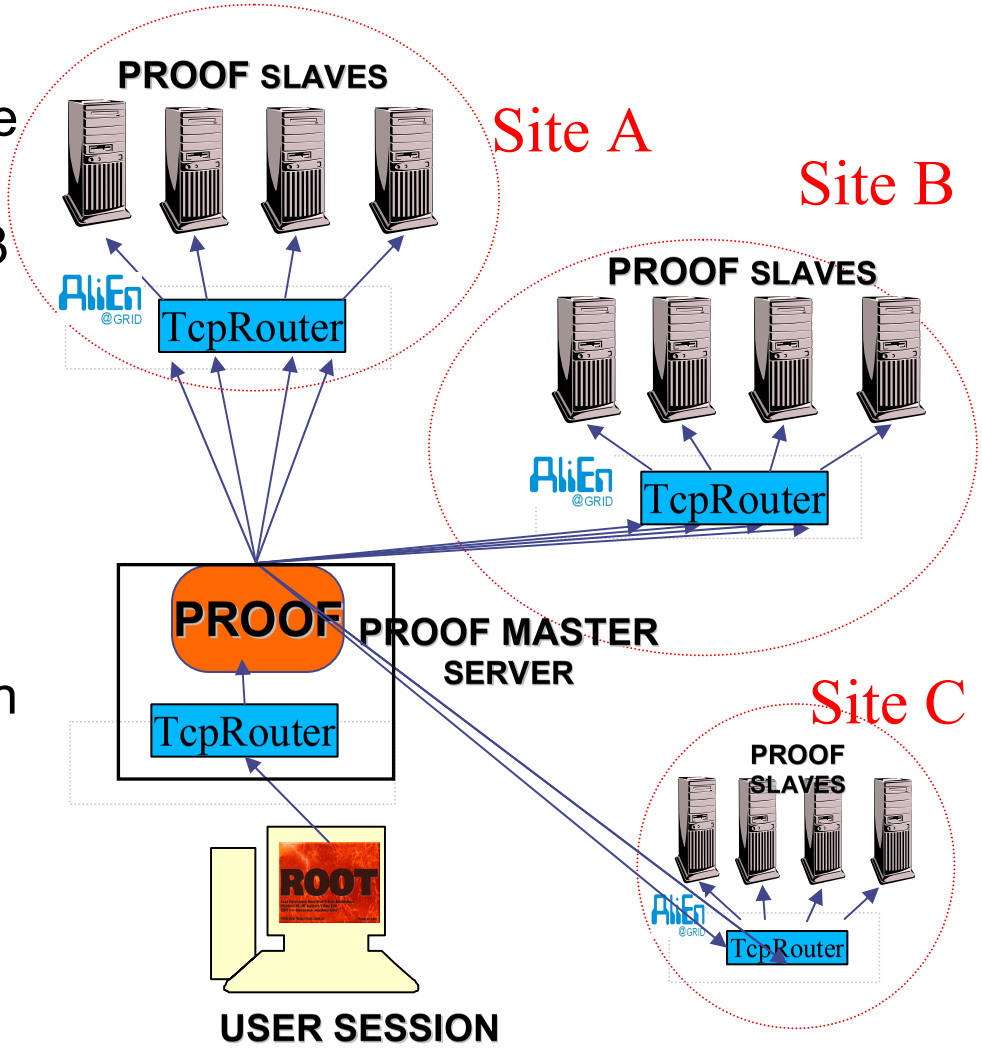

ALICE-ARDA improved system

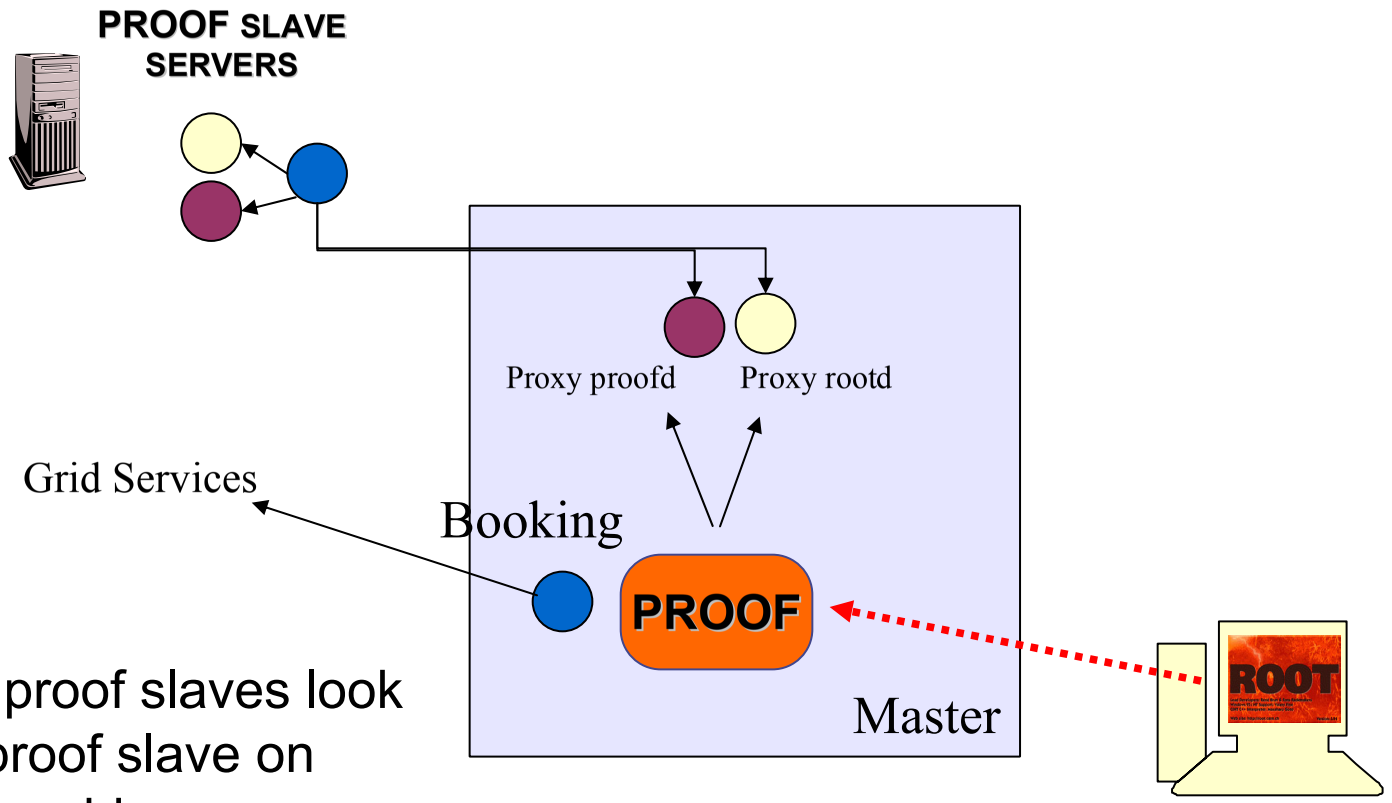

- The remote proof slaves look like a local proof slave on the master machine
- Booking service is usable also on local clusters

ATLAS

- The ATLAS system within ARDA has been agreed
  - ATLAS has a complex strategy for distributed analysis, addressing different area with specific projects (Fast response, user-driven analysis, massive production, etc…: see http://www.usatlas.bnl.gov/ADA/)
  - Starting point is the DIAL analysis model system
- The AMI metadata catalog is a key component
  - mySQL as a back end
  - Genuine Web Server implementation
  - Robustness and performance tests from ARDA
- In the start up phase, ARDA provided some help in developing ATLAS production tools
  - Being finalised

AMI studies in ARDA

- Atlas Metadata- Catalogue, contains File Metadata:
  - Simulation/Reconstruction-Version
  - Does not contain physical filenames
- Many problems still open:
  - Large network traffic overhead due to schema independent tables
  - SOAP proxy supposed to provide DB access
    - Note that Web Services are “stateless” (not automatic handles to have the concept of session, transaction, etc…): 1 query = 1 (full) response
  - Large queries might crashed server
    - Shall proxy re-implement all database functionality?
- Good collaboration in place with ATLAS-Grenoble
- N.B. This has to be considered a preparation work in addition to the agreed prototype (no milestone associated)

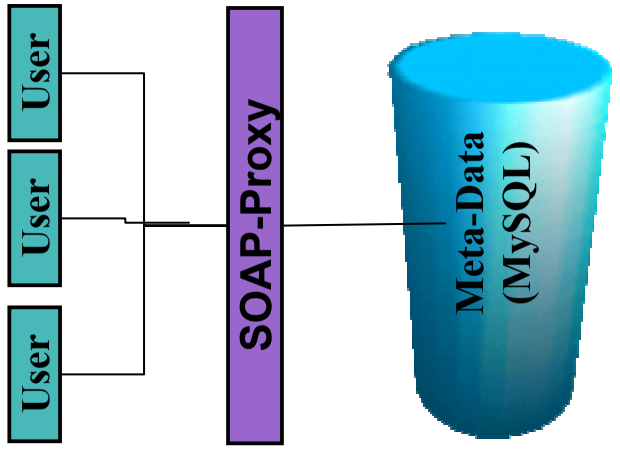

Studied behaviour using many concurrent clients:

	Rows in Response					
Clients	5	10	20	50	100	150
1	0.22	0.27	0.35	0.87	2.49	5.26
5	0.40	0.48	0.74	2.94	10.99	27.98
10	0.67	0.75	1.74	4.77	21.99	56.17
20	1.02	1.34	2.46	9.51	41.79	timeout
30	1.42	2.36	3.10	14.21	66.61	timeout
40	1.80	2.33	4.84	19.94	timeout	timeout
50	2.32	6.43	5.02	21.43	timeout	timeout
100	9.94	9.82	SOAP-Err	SOAP-Err		
150	16.51	SOAP-Err				

CMS

- The CMS system within ARDA is still under discussion
- Provide easy access (and possibly sharing) of data for the CMS users is a key issue
  - RefDB is the bookkeeping engine to plan and steer the production across different phases (simulation, reconstruction, to some degree into the analysis phase)
  - It contained all necessary information except file physical location (RLS) and info related to the transfer management system (TMDB)
  - The actual mechanism to provide these data to analysis users is under discussion
  - Measuring performances underway (similar philosophy as for the LHCb Metadata catalog measurements)

RefDB in CMS DC04

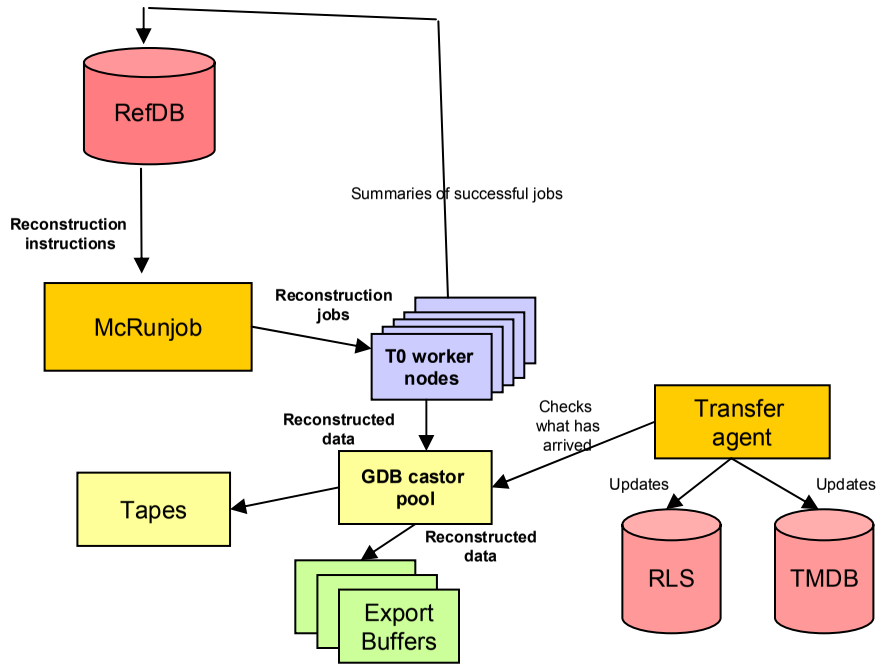

CMS refDB tests

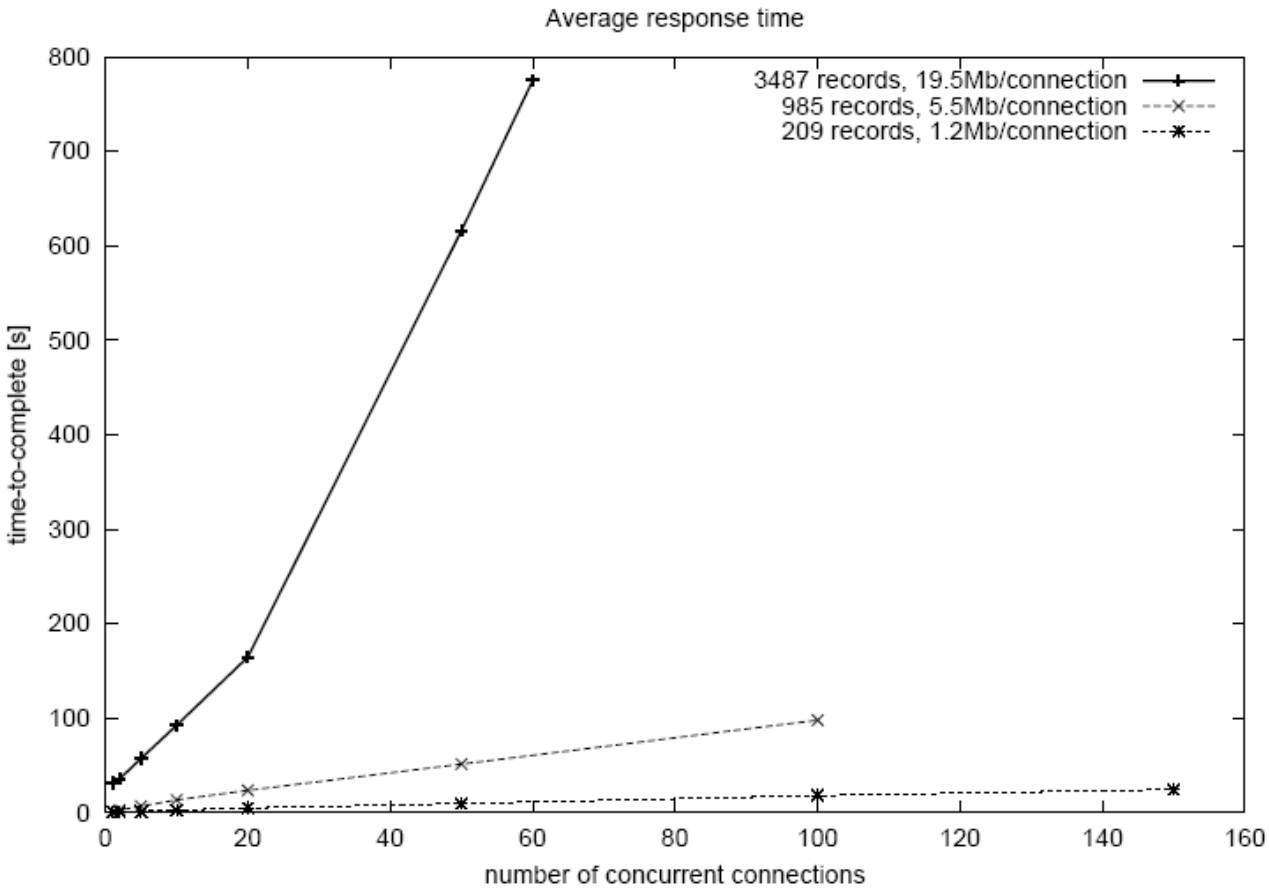

Figure 1: The average time-to-complete of a request presented as a function of the number of concurrent queries. The collections were discribed in the Table 1.

LHCb status

- Easy to agree on the prototype
  - Naturally aligned with the GANGA plans
  - Just started to play with Glite
- Other contributions
  - GANGA technical contribution
  - LHCb metadata catalogue measurements
    - Taiwan (ARDA + local DB know-how on Oracle)
  - DIRAC
    - Coherent evolution with Ganga
    - Expose DIRAC experience in the ARDA workshop

ALICE status

- Easy to agree on the prototype
  - Evolution of SC2003
  - Just started to play with Glite
- Other contributions
  - Investigate/survey Data Transfer protocols (comparison with RFIO, gridFTP… emphasis on robustness, error recovery and security)
    - PROOF (starting)
    - ROOTD (feedback loop closed)
    - AIOD (to be done)
    - XROOTD (to be done)
  - AliEn testing (activity started before ARDA, now completed; info handed over also to EGEE JRA1)

ATLAS status

- Difficult to agree on the prototype:
  - ATLAS complex strategy to be made coherent with the ARDA prototype spirit
  - Major role of the DIAL model agreed ☺
    - Minimal system as a starting point (run ATHENA jobs on a local cluster)
- Other contributions
  - Production system (activity started before ARDA; finishing)
  - ATLAS metadata catalogue measurements
    - Mainly at CERN (on the ARDA side)
    - Nice collaboration (feedback) with ATLAS Grenoble (S. Albrand et al.)
  - DIAL
    - Exercise with the old DIAL version

CMS status

- Difficult to agree on the prototype
  - CMS complex strategy to be made coherent with the ARDA prototype spirit
  - Major role of the catalogues
    - RefDB (metadata)
    - RLS (replica location)
    - POOL catalogues
  - Agreement before the ARDA workshop!
- Other contributions
  - Production system (new usage of COBRA metadata in RefDB)
  - RefDB catalogue measurements
    - Mainly at CERN (on the ARDA side)
    - Nice collaboration with many CMS people: exploratory work (share/agree on tests etc…)

HEP/Grid and ARDA

- LCG GAG
  - Massimo invited to be in the GAG: 1 per month
  - GAG has the key role to keep the HEP requirements/use cases
    - No duplication
    - ARDA contribution is complementary
    - NA4 LHC representative sit in GAG (Piergiorgio, Laura, Claudio, Andrey)
- Many invitations
  - HEP
    - DESY
    - gridPP (RAL and CERN)
  - GGF in Honolulu… (postponed to GGF Brussels if useful)

Difficult message: NA4 HEP mandate is to support the LHC experiments in using the Grid. A loosely coupled collaboration is possible on specific subjects, like metadata.

“The first 30 days of the EGEE middleware” ARDA workshop

- CERN: 21-23 of June 2004

- Monday, June 21
  - ARDA team / JRA1 team
  - ATLAS (Metadata database services for HEP experiments)
- Tuesday, June 22
  - LHCb (Experience in building Web Services for the Grid)
  - CMS (Data management)
- Wednesday, June 23
  - ALICE (Interactivity on the Grid)
  - Close out

"The first 30 days of the EGEE middleware" ARDA workshop

- Effectively, this is the 2nd workshop (January '04 workshop)
- Given the new situation:
  - Glite middleware becoming available
  - LCG ARDA project started
  - Experience + need of technical discussions
- New format:
  - "Small" (30 participants vs 150 in January)
  - To have it small, by invitation only…
  - ARDA team + experiments interfaces
  - EGEE Glite team (selected persons)
  - Experiments technical key persons (2-3 times 4)
  - Technology experts (Dirk, Fons, Iosif, Rene)
  - NA4/EGEE links (4 persons, Cal Loomis included)
- Info on the web:
  - http://lcg.web.cern.ch/LCG/peb/arda/LCG_ARDA_Workshops.htm

Workshop activity

- 1st ARDA workshop (January 2004 at CERN; open)
- 2nd ARDA workshop (June 21-23 at CERN; by invitation)
  - "The first 30 days of EGEE middleware"
- NA4 meeting mid July
  - NA4/JRA1 and NA4/SA1 sessions organised by M. Lamanna and F. Harris
- 3rd ARDA workshop (September 2004?; open)

- "Forum activities" are fundamental (see LCG ARDA project definition), on the other hand there are no milestones proposed for this (removed a proposed one)

EGEE and ARDA

- Strong links already established between EDG and LCG. It will continue in the scope of EGEE
- The core infrastructure of the LCG and EGEE grids will be operated as a single service, and will grow out of LCG service
  - LCG includes many US and Asia partners
  - EGEE includes other sciences
  - Substantial part of infrastructure common to both
- Parallel production lines as well
  - LCG-2
    - 2004 data challenges
  - Pre production prototype
    - EGEE MW
    - ARDA playground for the LHC experiments

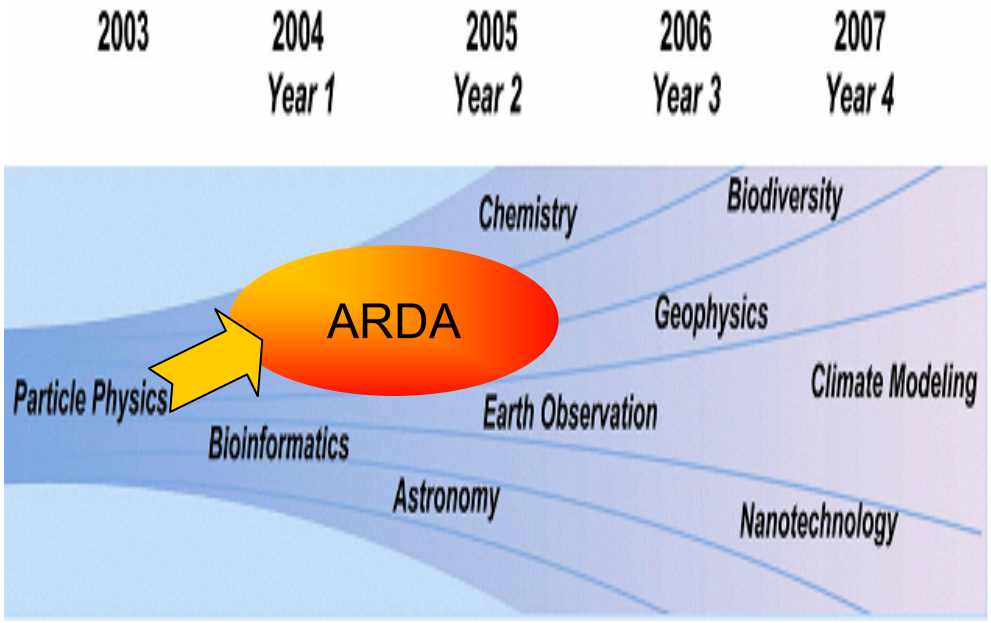

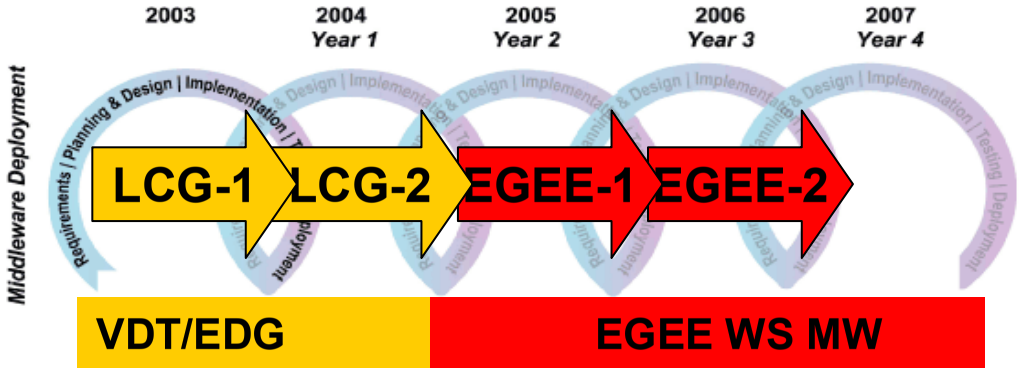

EGEE and ARDA

- EGEE/LCG effort to ARDA
  - 4 FTE from EGEE
  - NOW: 6 LCG + 4 persons from regional centres
- EGEE Conferences: 2 per year. April = Cork; Ireland
  - Not too efficient (could not attend to the full conference; maybe next one more interesting)
  - Opportunity to meet people outside the LCG circle
- EGEE “All Activity Meetings”: several a year?
  - First one (for me) June 18th
- NA4 AWG (Application working group): 1 meeting per week (Massimo and Frank)
  - NA4 steering body
  - Nice atmosphere, clearly the goals are not always the same for all (not a surprise)
  - ARDA should be there
- EGEE PEB: 1 meeting per week (Frank)
  - Excellent relation with Frank
- EGEE PTF (Project Technical Forum): 1 per month (Massimo and Jeff Templon)
  - New body. First meeting June 17th. Close to the Architecture Team (alias?). It reports to the Technical Director. Convener: Cal Loomis
  - ARDA should be there. I hope concrete technical issues will dominate the discussion
- NA4 meeting in Catania (mid July):
  - JRA1/NA4: organised by Massimo
    - I would like to use it as another ARDA workshop (benefiting from non-LHC ideas and experience)
  - SA1/NA4: organised by Frank
    - Very important (preproduction service)

ARDA @ Regional Centres

- "Deployability" is a key factor of MW success
- A few Regional Centres will have the responsibility to provide early installation for ARDA to supplement the LCG preproduction service
- Stress and performance tests could be ideally located outside CERN
  - This is for experiment-specific components (e.g. a Meta Data catalogue)
  - Leverage on Regional Centre local know how
    - Data base technologies
    - Web services
    - …
- Ease the interaction with "the rest of" HEP?
  - DESY
  - Non LHC experiments?
- Running ARDA pilot installations
  - Experiment data available where the experiment prototype is deployed
  - CERN, RAL, all Tier1s… The strategy is not clear yet
- As for the "Forum activities", no milestones proposed for these activities

Status

- Prototype definition
  - 3 out of 4 OK (1 milestone late)
- Prototype status
  - ALICE and LHCb OK
  - ATLAS/DIAL “starting point” not yet available
- EGEE Middleware
  - GLite software available ☺
- EGEE
  - Useful contacts
  - Sizable but manageable overhead so far