



ARDA final status report

Massimo Lamanna
CERN IT-GS

Summary

- “Genesis”
- First years
- Second phase
- Conclusions

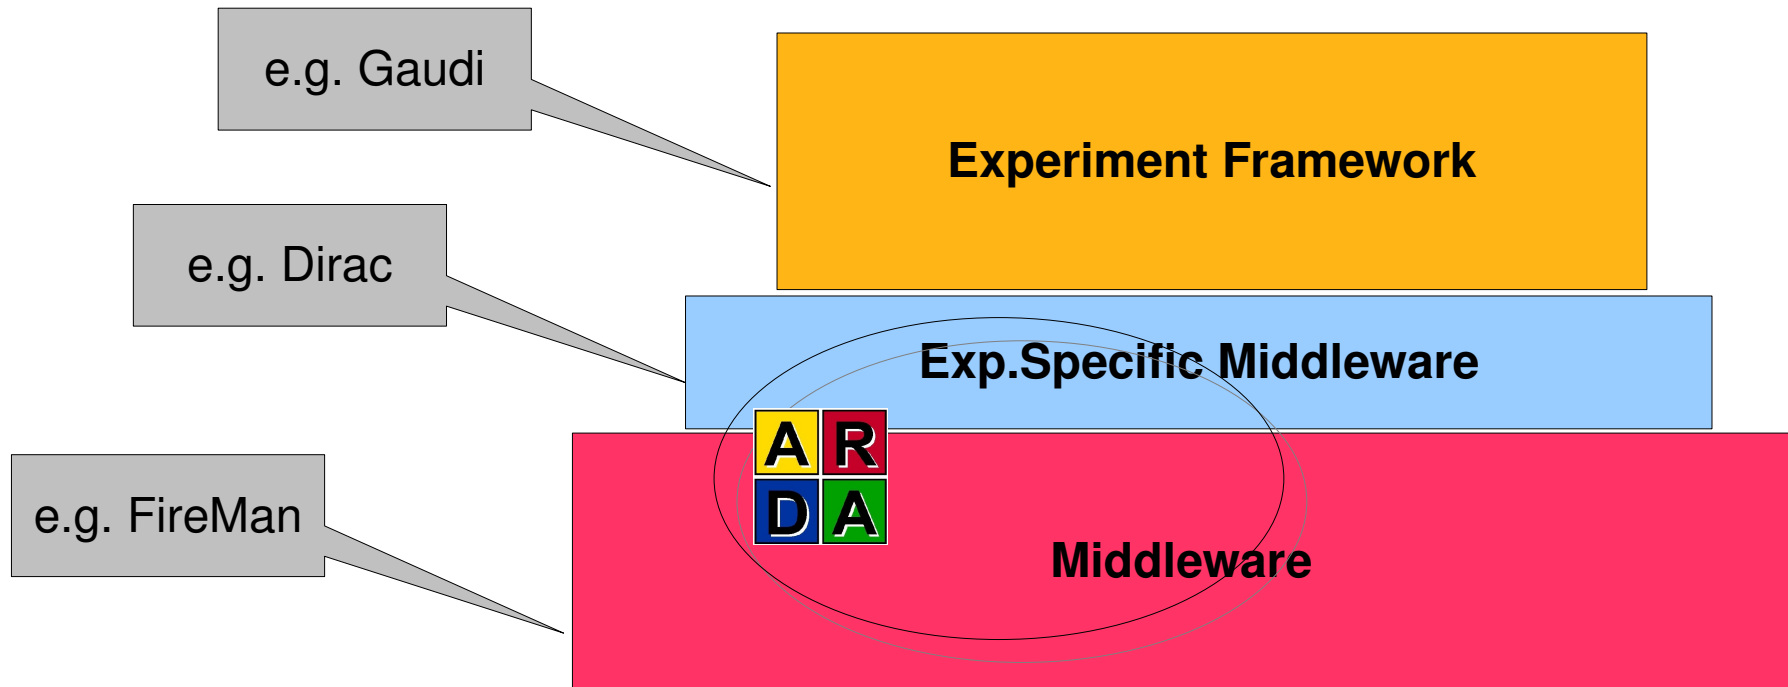
Genesis

- A Roadmap for Distributed Analysis
 - One of the most attended workshops (Miron, Predrag, Torre etc...)
 - Held at the same time of the OGSF -> WSDL announcement (and the premature death of GT3...)
- A Realisation of Distributed Analysis
 - EGEE effort (4 persons) + 4 matching funds from WLCG
 - Initial mantra:
 - Production is understood, Analysis not yet...

First phase

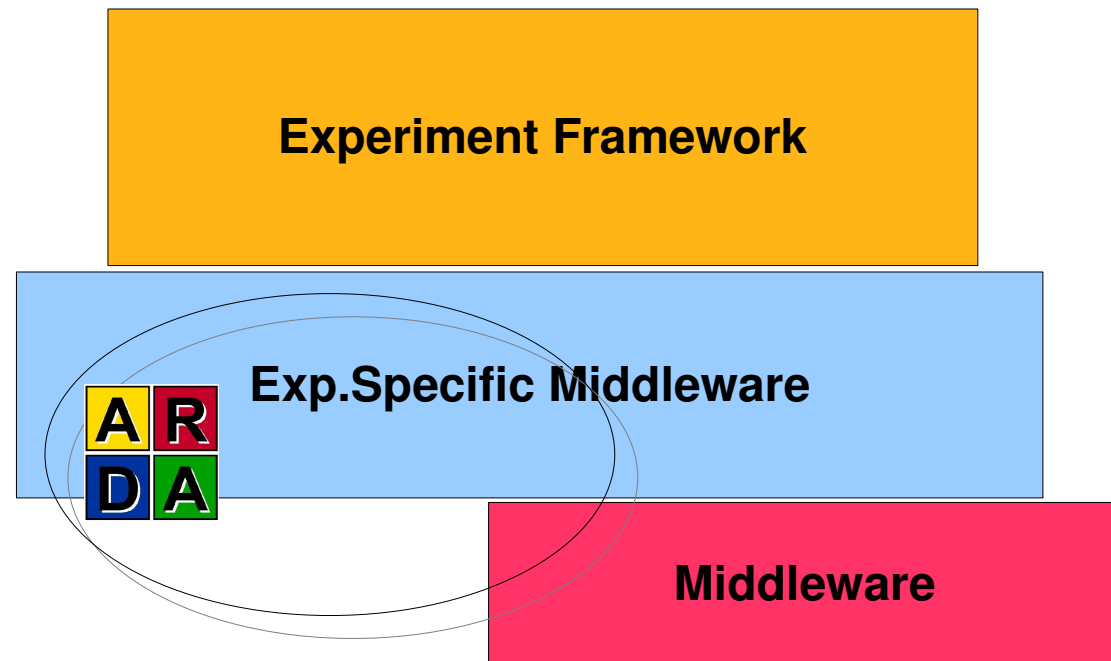
- Contact with the experiments via separate contacts (one per experiments)
 - Agreement on the activity
 - Various levels of integration
 - Sometimes we were a bit side-tracked
 - Exchanges via the ARDA team
 - Exception: Ganga (was already a common project)

Activity (as imagined at the beginning)



Activity

- The relative importance of the middleware diminished (baseline services etc...)
- The scope of the application specific layer increased
- Fortunately some commonality was found
- We attracted a number of collaborators (unfunded) within and without EGEE



Second phase

- Closer collaboration with each experiment
 - Closer means also on “critical path” activities
 - Some prototypes have been stopped
 - e.g. the analysis system ASAP because the official tool CRAB got more and more momentum, but our contributions was reused. In this specific case, in CRAB itself and especially in the monitoring (dashboard)
 - Other activities could be expanded and attract more experiments (Dashboard)

A stroll across ARDA

Largely incomplete...

AMGA



- ARDA Metadata Grid Access
- Metadata catalogue: obvious starting point for ARDA
 - We studied existing systems in the experiments
 - Initially we contributed an I/F (it was outside of the scope of JRA1) and a working prototype (endorsed by GAG)
 - Basis a a few interesting contributions to the fields (master and PhD students)
- Eventually:
 - Part of the gLite distribution since 2006
 - Collaborative effort coordinated by the original developer but all effort coming from outside (Catania, Korea, Clermont-Ferrand,...)
 - Now coordinating the release process, adding new features etc...

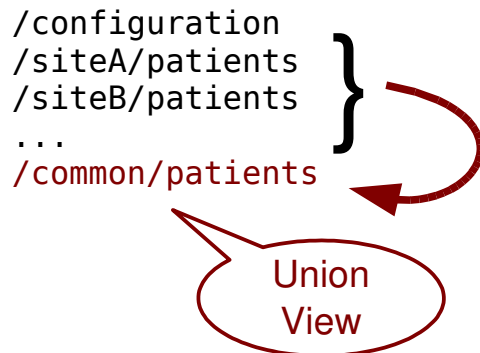
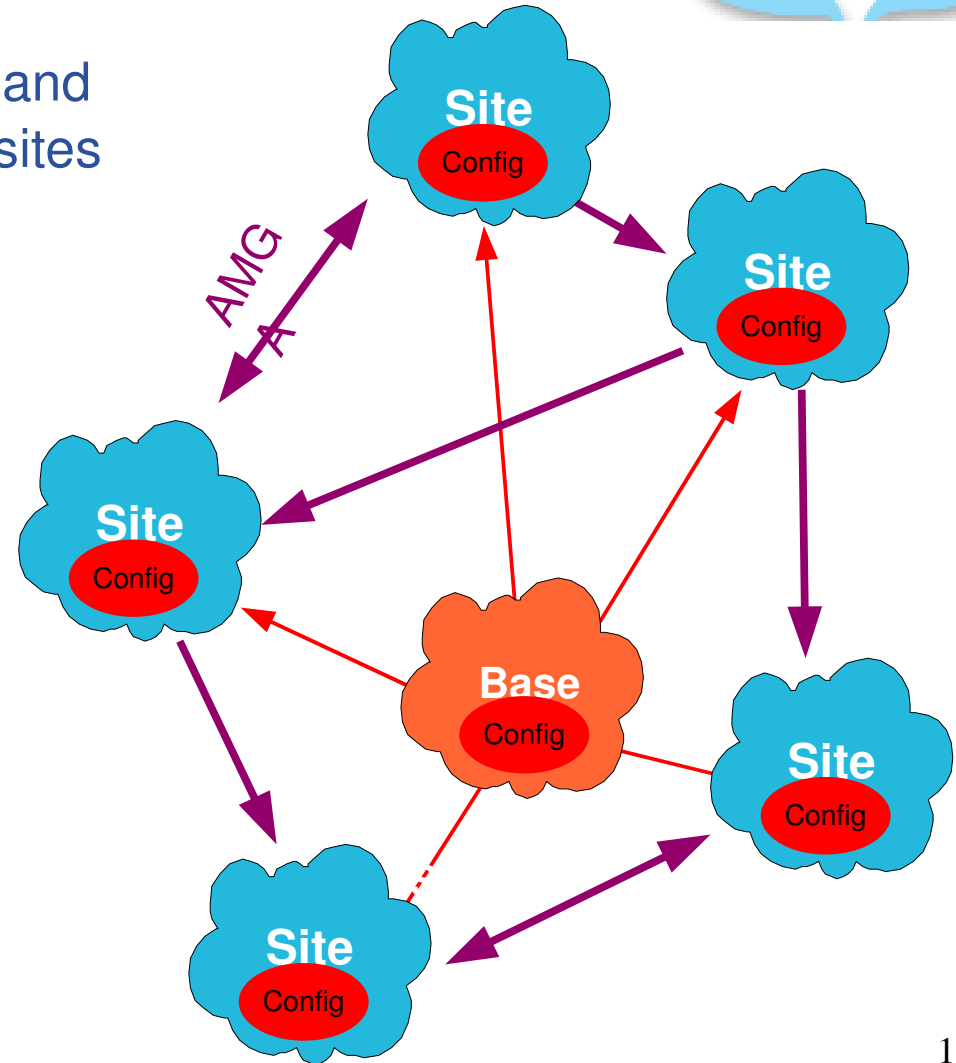
AMGA

- Adopted by LHCb for their Logging and bookkeeping catalogue (used until now; migration taking place now)
- Great success in EGEE. Some examples:
 - Earth sciences:
 - Climatology (Climatology centre in Hamburg DKRZ – also under in DGrid)
 - UNOSAT (Access of satellite images)
 - Biomedical sciences:
 - Wisdom (In-silico drug searches)
 - Health-e-child
 - Digital imaging
 - Non-LHC partners in EGEE3 are using their resources on this subject

Case study: Health-e-Child

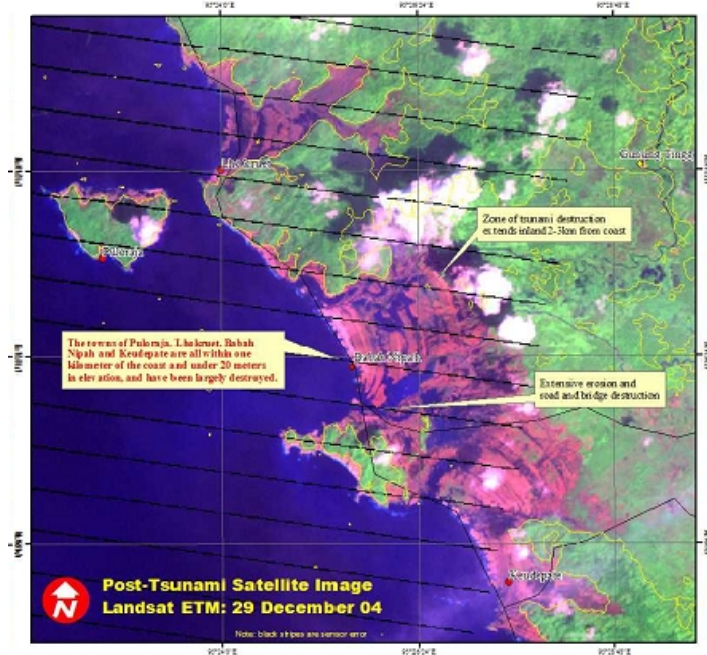


- Several dozens of hospitals providing case-data
- Central server with credentials for participating sites and users (replication mandatory)
- Data replicated from site to site on demand
 - 'Automount' mechanism for joining sites

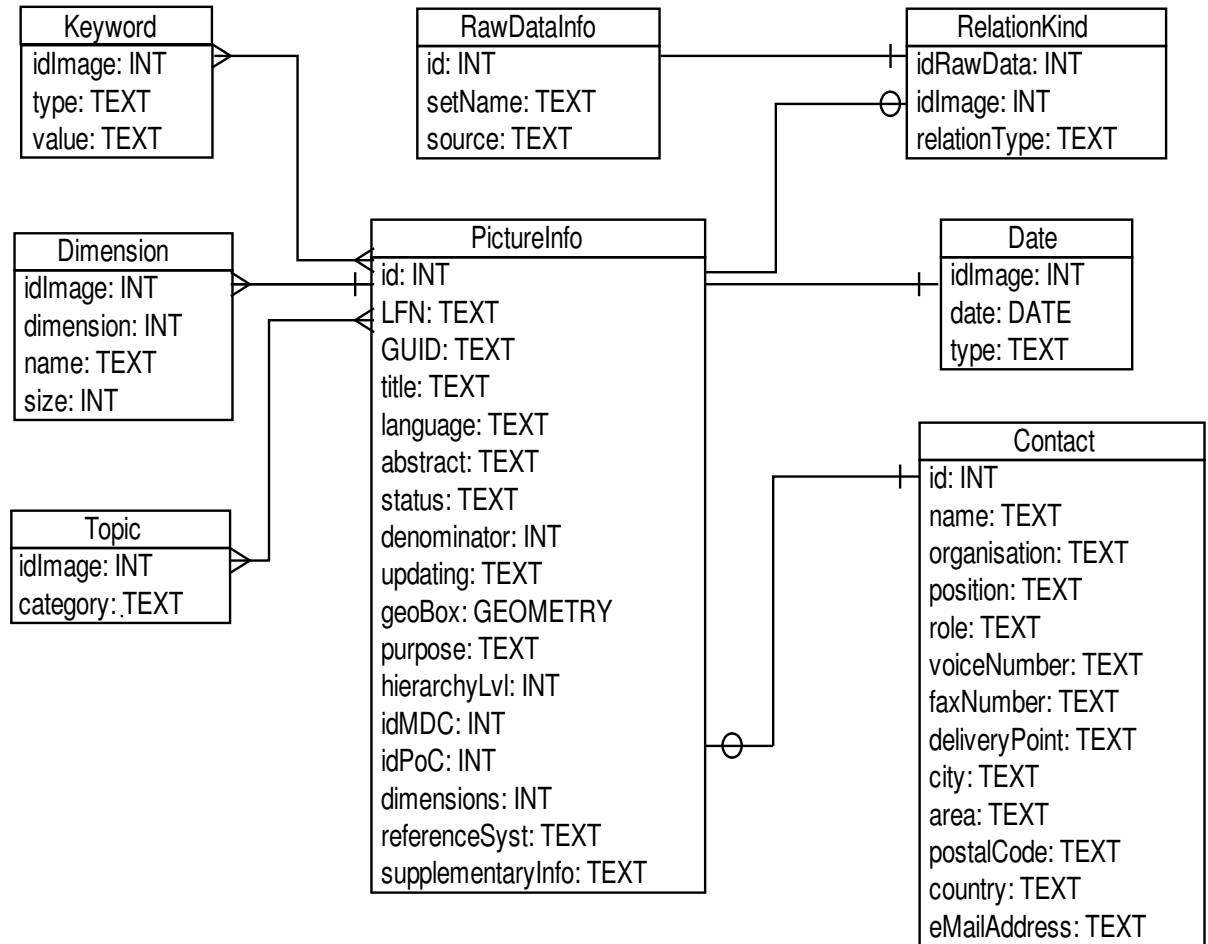


Geographic Metadata

- UnoSat prototype uses AMGA to store GIS (Geographic Information System) Metadata for images
- Accessible also via portable devices (mobile phones)



UNOSAT 
satellite imagery for all




Dashboard

- Initially it was born as CMS Dashboard
 - Reusing components of the CMS analysis prototype
 - Fundamental contribution of MonALISA
- Its scope has progressively grown
 - More experiments!
 - Different VOs (non HEP) interested
 - Other activities (data transfer, site status, middleware errors, ...)

CMS Dashboard - Crab Analysis Jobs

jobs per site



• Summer 2007:

- over 600k jobs/month
- i.e. 20k jobs/day
- 89% grid success rate
- ***Substantial Tier2 contribution***

submitted app-succeeded app-failed app-unknown pending running aborted cancelled



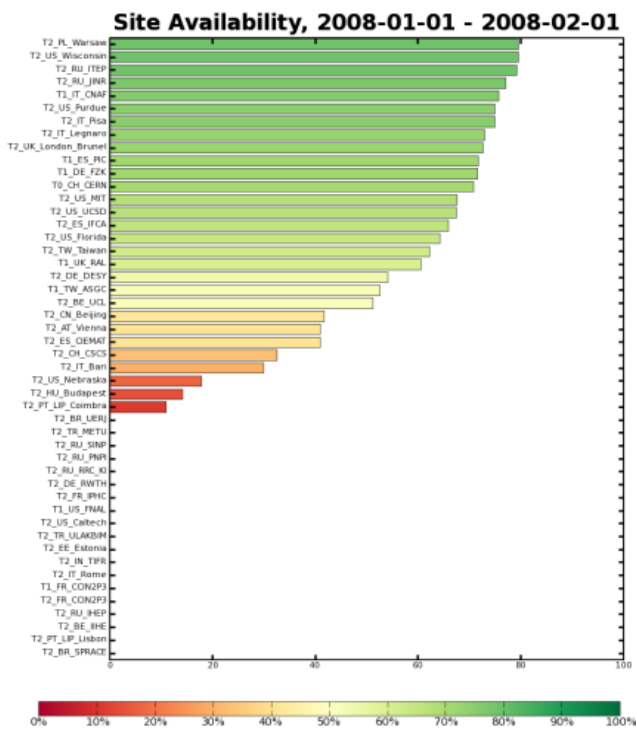
Monitoring improvements of Grid infrastructure

- CMS sites availability monitored with Dashboard application

January 2008

May 2008

September 2008



Example of GridMap – Dashboard integration

CMS GridMap – Visualizing SAM Site Availability

Tier 2

T2 US Florida	T2 FR IPHC	T2 IT Bari	T2 KR KNU	T2 US UCSD	T2 UK London	T2 TR ULAKS	T2 DE DES
T2 UK SGnd B	T2 IN TIFR	T2 IT Pisa	T2 US Nebrask	T2 DE RWTH	T2 PL Warsaw	T2 BE UCL	T2 HU Bud
T2 FR GRIF IR	T2 IT Rome	T2 RU JINR	T2 TW Taiwan	T2 RU IHEP	T2 FR CCIN2	T2 ES IFCA	T2 ES CIE
T2 IT Leonardo	T2 RU ITEP	T2 RU PNPI	T2 PT LIP Col	T2 RU RRC HI	T2 CH CSCS	T2 FI HIP	T2 TR M
T2 US Caltech	T2 EE Estonia	T2 UK London	T2 CN Bellin	T2 UK SGnd R	T2 AT Vienna		
T2 BE IHE	T2 US MIT	T2 BR SPRACET	T2 US Wiscons	T2 US Purdu	T2 RU SINP	T2 PT LIP	T2 BR U

Latest CMS SAM results, Site Availability, 21 Sep 2008 10:00 UTC.
Size of site rectangles is based on tier.

SAM Latest Results - Windows Internet Explorer

http://dashb-cms-sam.cern.ch/dashboard/request.py/latestresultssmy?sites=T2_PT_LIP_Lisbon

SAM VISUALIZATION

Latest Results | Historical View | Feedback | Help

Sites: Tier1s + Tier0
 T0_CH_CERN1
 T1_DE_FZK
 T1_ES_PIC
 T1_FR_CCIN2P3
 T1_IT_CNAF
 T1_TW_ASGC

Service Types: VO critical
 CE
 SRMv2

Test Types: CMS Tests
 Select All
 CE-cms-analysis
 CE-cms-basic
 CE-cms-frontier
 CE-cms-mc
 CE-cms-prod
 CE-cms-squid

Test Exit Status: All Exit Status
 na
 ok
 down
 degraded
 partial
 maint
 error

Show Results

Legend: NA OK MAINTENANCE ERROR WARNING INFO NOTE CRITICAL
 Note: brightest colors: test is 0 - 6 hours old, ... lightest colors: test is more than 24 hours old

Link to the table

SiteName	Service Type	Service Name	frontier	basic	analysis	swinst	jsprod	squid	mc	js	lcp	pfm	from	tfc
T2_PT_LIP_Lisbon	CE	ce02.lip.pt	ok	ok	error	ok	ok	ok	ok	ok	ok	ok	ok	ok
	SRMv2	dcache01.lip.pt										ok	warn	
		srn01.lip.pt										warn	error	

Algorithm for calculating the Site and Service Availability

SAM Service History - Windows Internet Explorer

http://dashb-cms-sam.cern.ch/dashboard/request.py/historicalresultssmy?sites=T2_PT_LIP_Lisbon

Test results for ce02.lip.pt
 48 Hours from 2008-09-19 10:00 to 2008-09-21 10:00 UTC


CE-cms-analysis	ok	ok	ok	ok	ok	ok	ok	ok	ok	ok	ok	ok	ok	ok
CE-cms-basic	ok	ok	ok	ok	ok	ok	ok	ok	ok	ok	ok	ok	ok	ok
CE-cms-dummy	ok	ok	ok	ok	ok	ok	ok	ok	ok	ok	ok	ok	ok	ok
CE-cms-frontier	ok	ok	ok	ok	ok	ok	ok	ok	ok	ok	ok	ok	ok	ok
CE-cms-mc	ok	ok	ok	ok	ok	ok	ok	ok	ok	ok	ok	ok	ok	ok
CE-cms-prod	ok	ok	ok	ok	ok	ok	ok	ok	ok	ok	ok	ok	ok	ok
CE-cms-squid	ok	ok	ok	ok	ok	ok	ok	ok	ok	ok	ok	ok	ok	ok
CE-ath	ok	ok	ok	ok	ok	ok	ok	ok	ok	ok	ok	ok	ok	ok

Click


M4 data taking August 31 2007




Throughput MB/s




Data transferred GB



Completed file transfers

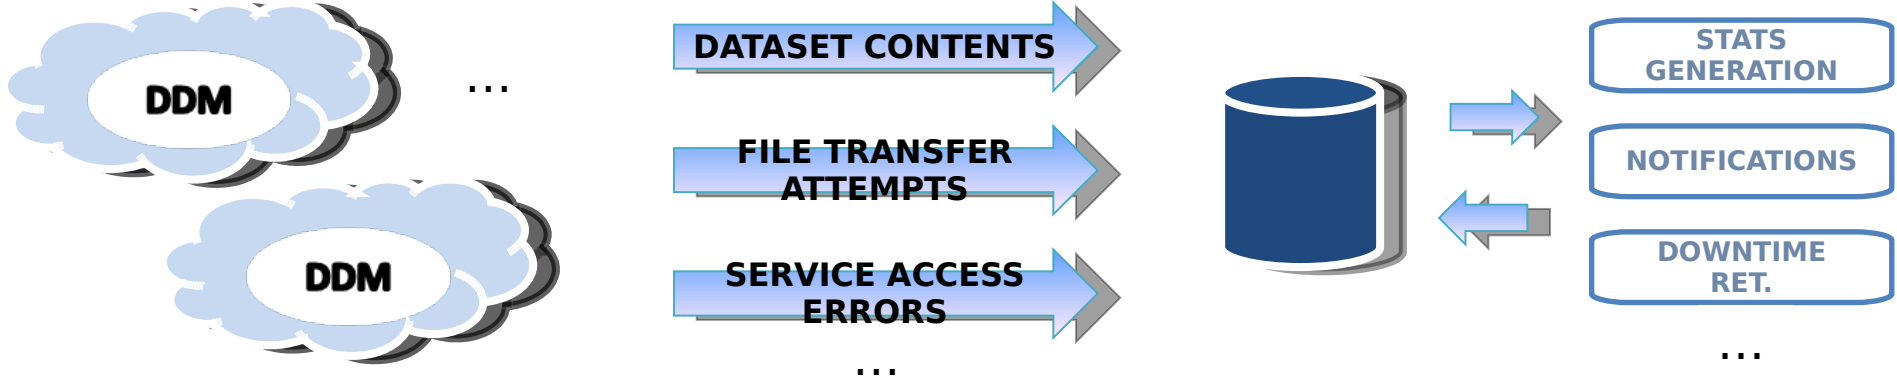


Total number of errors

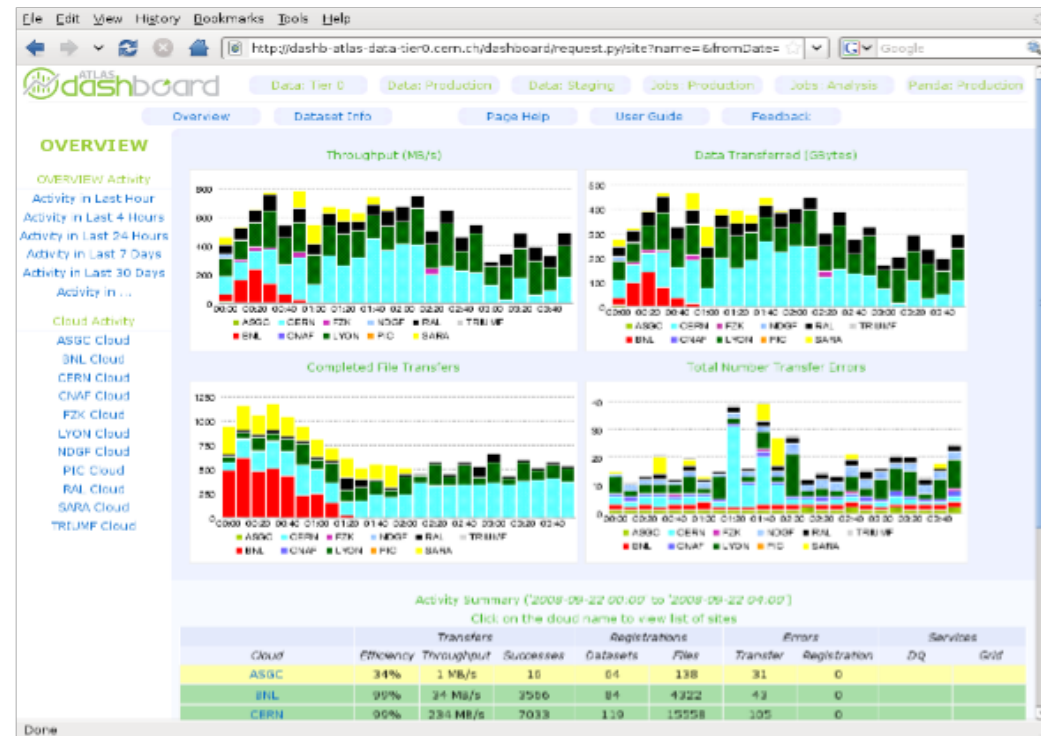


Application Overview

- Monitoring of data movement within clouds and sites



- Topology: clouds, sites, storage space tokens
- Dataset: content, location and completeness
- File: transfer attempt history, location, details on storage (src/dest surl, checksum, ...)
- Resources: SAM results, downtimes
- Statistics: throughput, efficiency, error summaries, transfer attempt number, dataset queued/completion time, ...








find

view

- by grid
- by cloud
- by dest_cloud
- by executortype
- by executor
- by site
- by cluster
- by task

select cloud

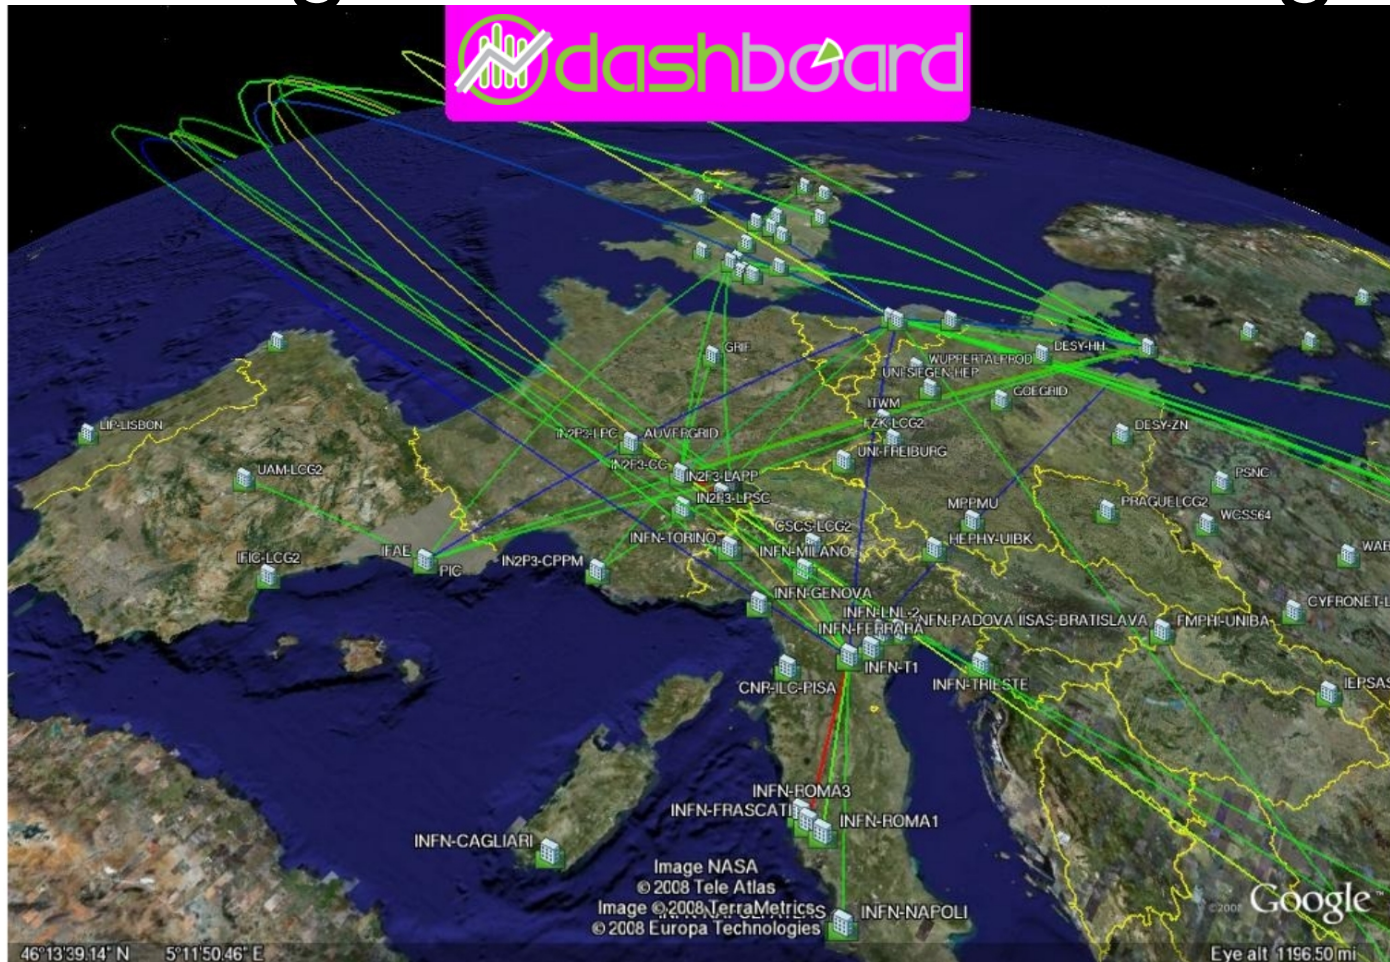
- BNL
- LYON
- CNAF
- None
- FZK
- NDGF
- PIC
- SARA
- TRIUMF
- ASGC
- RAL
- CERN



cloud	defined	assigned	waiting	activated	running	holding	transferring	success	failure	efficiency
BNL	0	192	0	1028	3351	1747	571	6934	310	95.7%
LYON	0	538	0	95	393	636	2030	2027	78	96.3%
FZK	0	358	0	250	303	68	327	915	100	90.1%
PIC	0	8	0	5	36	281	193	687	179	79.3%
CNAF	0	0	0	284	1718	36	196	702	104	87.1%
TRIUMF	0	0	0	0	19	410	5	6	6	50%
RAL	0	0	0	0	16	344	0	4	0	100%
None	165	17	210	1577	12	340	0	0	2	0%
NDGF	0	0	0	0	1656	0	0	0	2	0%
SARA	0	30	0	0	22	427	7	0	0	-
ASGC	0	0	0	0	20	403	0	0	0	-
CERN	0	0	0	0	0	1	0	0	0	-
total	165	1143	210	3239	7546	4693	3329	11275	781	93.5%

CRITICAL
WARNING
NORMAL
GOOD
NO_ACTIVITY

Google Earth Monitoring



- Summaries of activity collected from DDM and ProdSys Dashboards
- Agent generating KML representation of this data
- Automatically refreshed every 10 minutes (*real time like* animation)

APIservice (ALICE)

Late 2004

- Analysis as an interactive service
 - Interface with gLite
- Proposed very early for ALICE
 - Inspiring for other developments in ARDA

eGee ARDA shell + C/C++ API


C++ access library for gLite has been developed by ARDA

- High performance
- Protocol quite proprietary...

Essential for the ALICE prototype


Generic enough for general use

Using this API grid commands have been added seamlessly to the standard shell

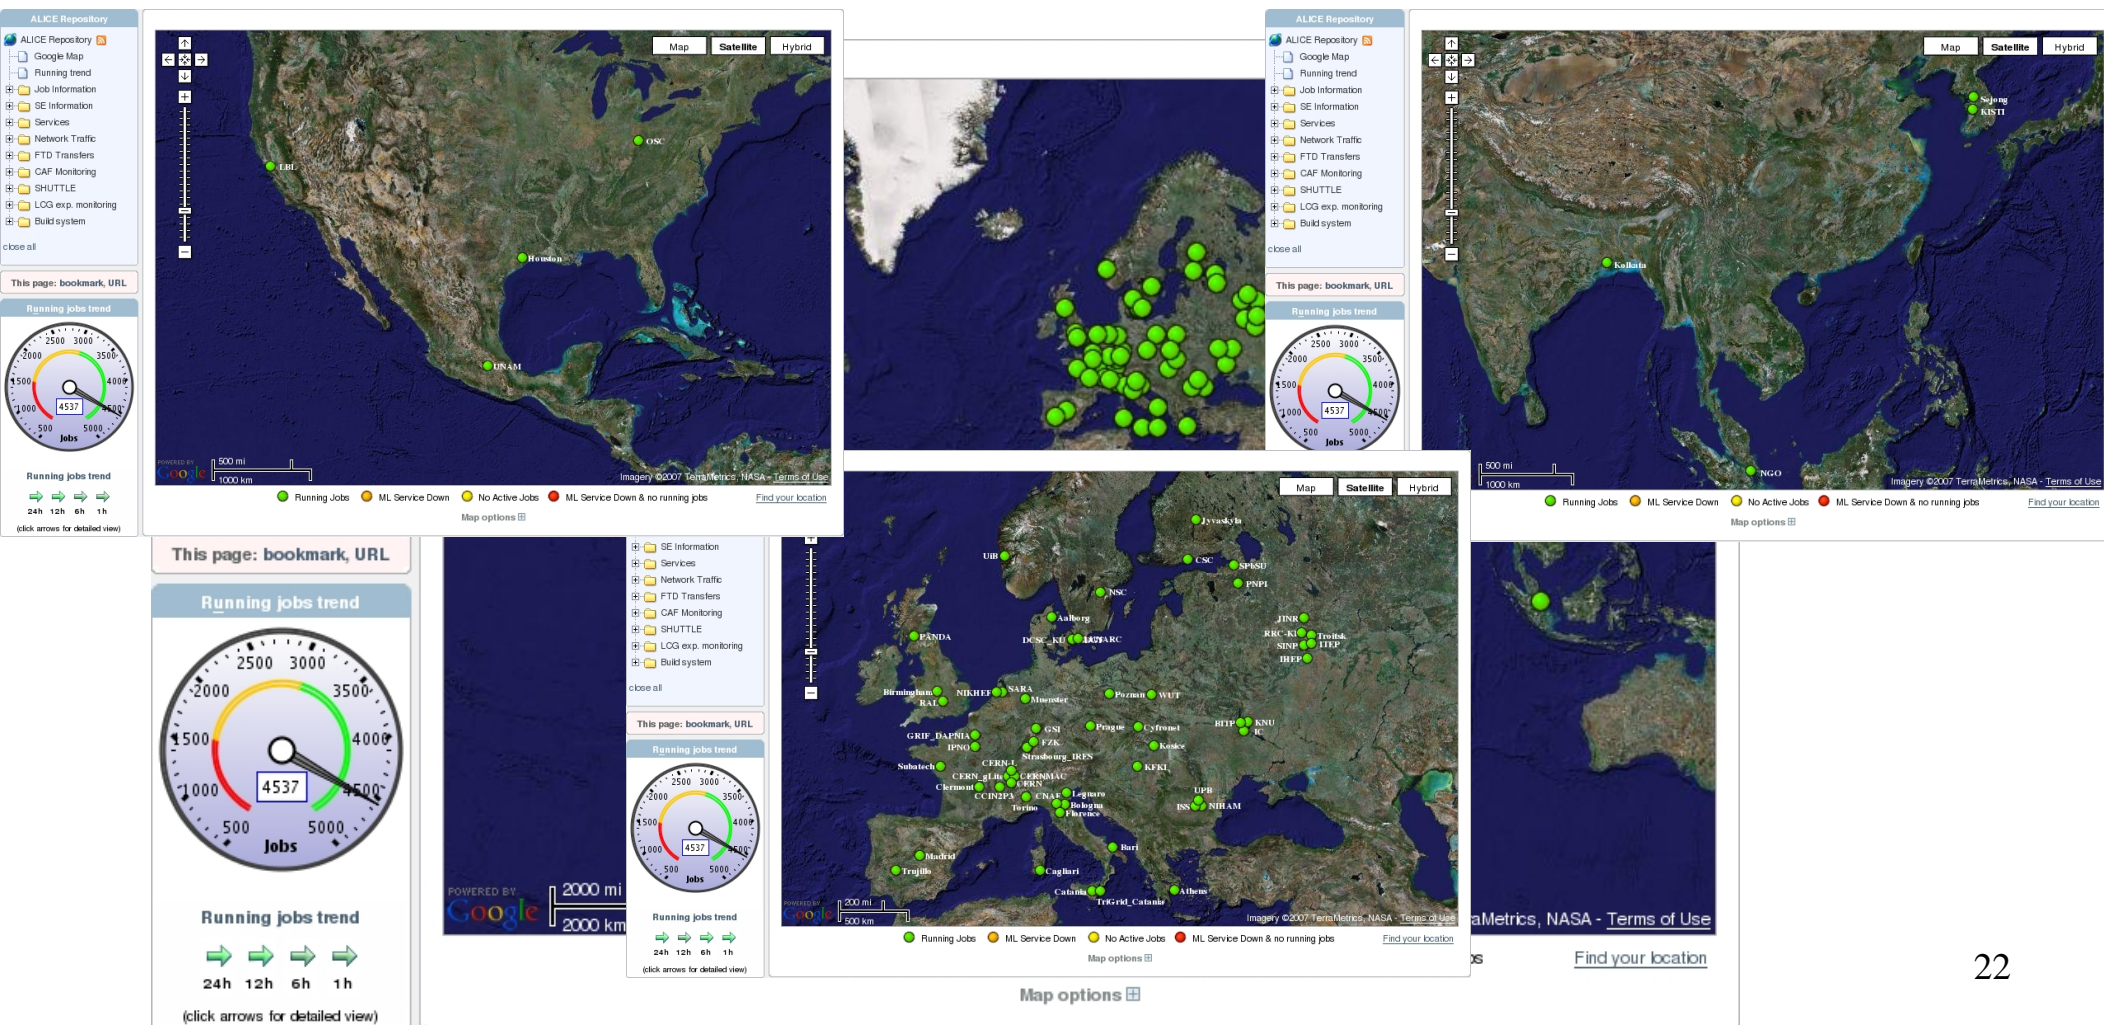


ARDA

Today



PDC 07 <http://pcalimonitor.cern.ch>



Ganga

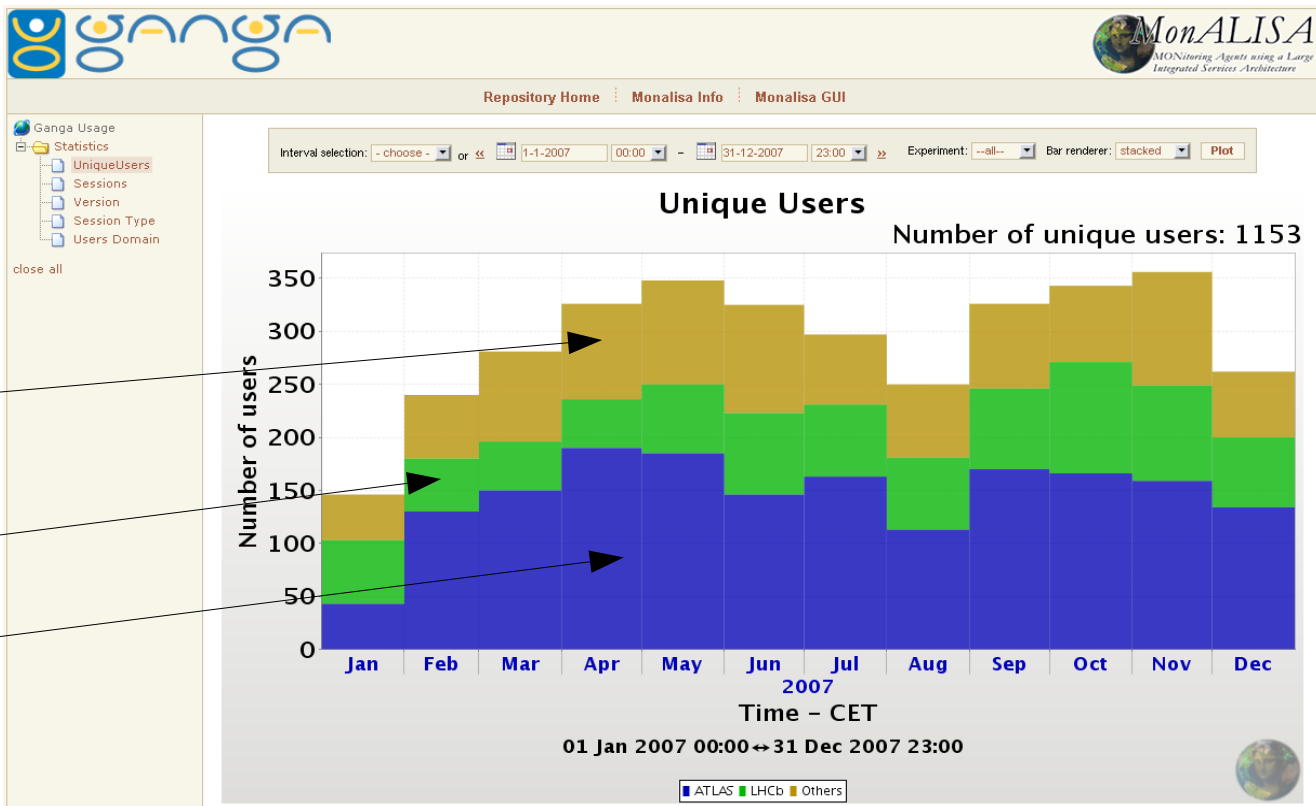


- Common project (ATLAS – LHCb)
 - We started collaborating on LHCb side only
 - After ~18 months, the ARDA/ATLAS contribution joined Ganga

– Due to the EGEE links, considerable interest outside the LHC area

LHCb

ATLAS



Ganga



- Excellent adoption
 - It is the entry point to the ATLAS (PanDA) and LHCb (Dirac) system
 - In the case of ATLAS, complemented by pAthena
- Excellent user feed-back
 - Several tutorials paid off!
 - In addition, we discover of communities discovering and adopting Ganga without our direct involvement (e.g. Minos – discovered by googling... :)

Grid Access for Analysis

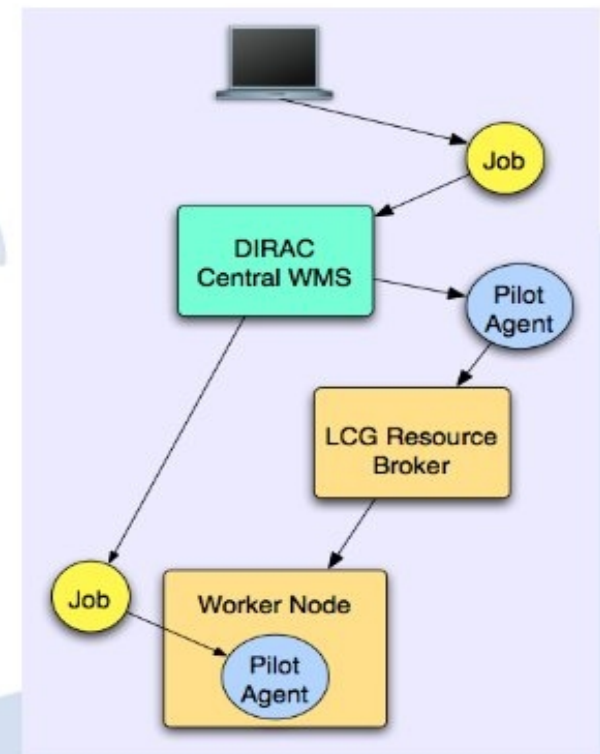
Analysis jobs: No direct submission to LCG

Instead:

Submission to the DIRAC WMS

Advantages:

- Provide transparent access to the LFC file catalogue for reading and writing data
- Allow LHCb to set priorities and or restrictions for analysis jobs
- More see Stuart Paterson's talk



LHCb Analysis Job

Gaudi based applications:

```
In [3]: dv = DaVinci(version='v12r12')
In [4]: print dv
DaVinci {
  version = 'v12r12' ,
  extraopts = None ,
  package = 'Phys' ,
  cmt_user_path = '/afs/cern.ch/user/u/ueqede/cmtuser' ,
  masterpackage = None ,
  optsfile = File {
    name = ''
  }
}
```

Specify extra option file properties appended to the options file




Specify the package you are working on

Specify the options file to be used

... other Ganga "communities" out there!



Geant 4



Academia Sinica
Genomics Research Center

ITU conference (May-Jun 06)


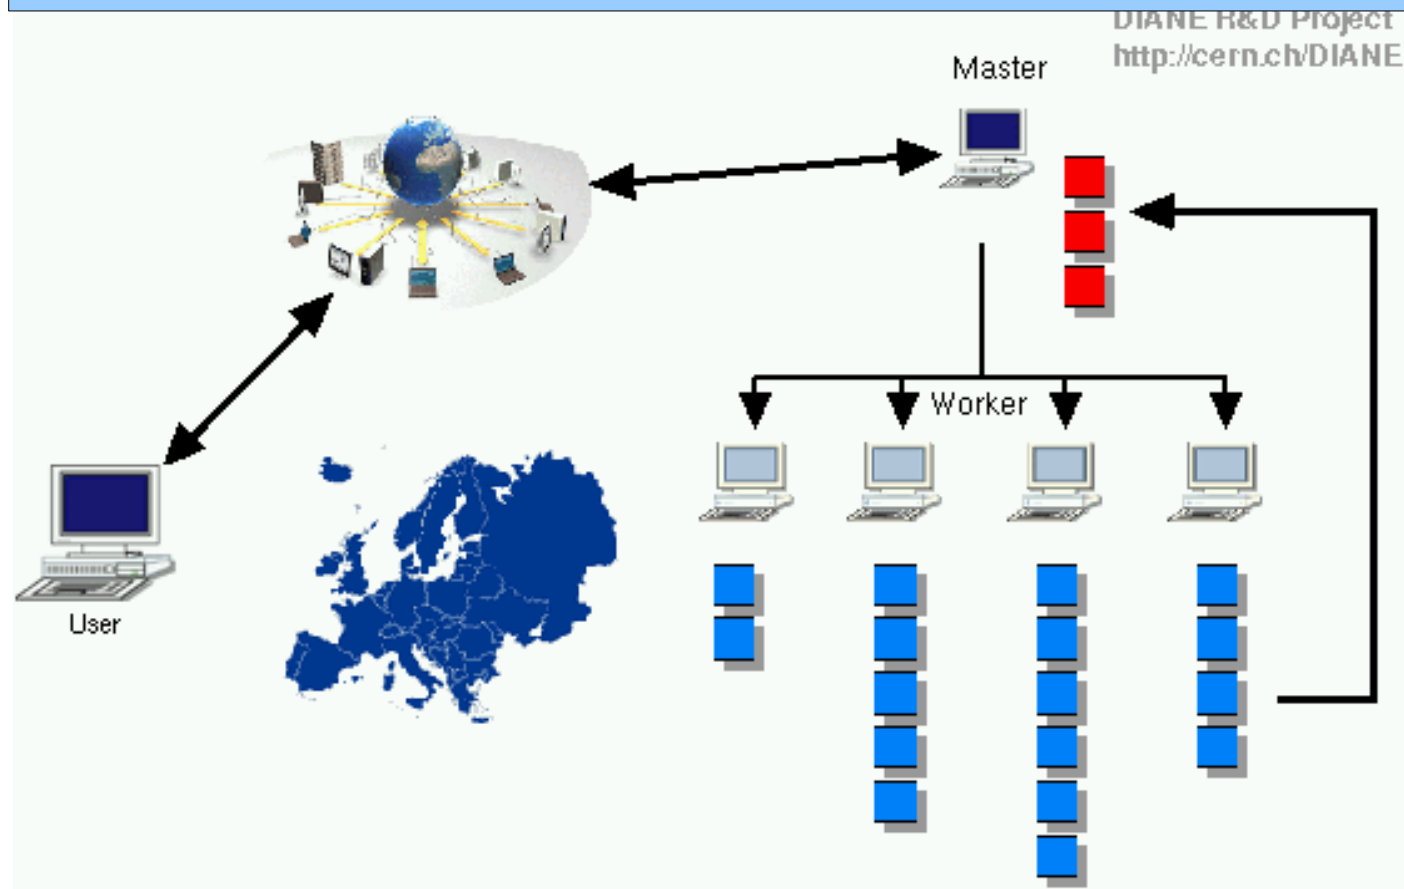


Figure 1
The extent of the planning area for the RRC-06

- Need on **DEPENDABILITY**
 - Large production to be done on a given date within ~24 hours
- Use of Ganga/DIANE
 - client – server worker allowing to execute (grid) jobs in an efficient and dependable way)



Experience with EGEE grid

CERN responsibilities:

- Consulting in GRID technologies
- Online support before and during the Conference
- Preparation of scripts for easy operation of the GRID by ITU personnel



ITU responsibilities:

- Providing datasets
- Operation of the system



A. Manara
(Int. Telecommunication Union - ITU)



Participating institutions

CERN
INFN-CNAF Bologna
(+ sites of GridIT infrastructure)
PIC Barcelona, CNB Madrid
DESY (Hamburg, Zeuthen)
Cyfronet Krakow
Moscow State University



Theoretical physics (Lattice QCD)

- Lattice Conferences in 2007 and 2008
 - Ph. de Forcrand et al.
- Presented at CHEP 2007
- More activity in 2008

A QCD critical point
at small chemical potential:
is it there or not?

Philippe de Forcrand
ETH Zürich and CERN

with

Seyong Kim (U. Sejong) and Owe Philipsen (U. Münster)

ETH

École Polytechnique Fédérale de Zurich
Swiss Federal Institute of Technology Zurich

The legacy

- ARDA smoothly ended with EGEE2
 - CERN plays an important role in EGEE3, but with a slightly smaller effort building also on ARDA experience
- The persons from the ARDA team are in general still contributing into the LHC experiments and WLCG
 - Ganga: centre of the analysis in ATLAS and LHCb
 - Dashboard: more and more used in the experiments and in the infrastructure
- Experience was not lost
 - And “ex-ARDA” people are still providing excellent work in the various area of WLCG

Outlook

- I advocate the approach of close collaboration with the experiments is very positive
 - Experiments benefit from “back-office” collaboration
 - Which leads to true commonality
- Did we fulfil our initial mandate (“*production is understood, analysis not yet*”)?
 - Not yet: analysis is a moving target
 - Analysis (in particular high-performance data access) is still evolving
- An approach à la ARDA might be useful here!