The South African National Compute Grid

07/25/08

Bruce Becker,
INFN Cagliari & UCT-CERN Research Centre

Bruce Becker :

- Ph.D. 2007, University of Cape Town, Physics Department – experimental ultrarelativistic nuclear physics
  - Work on **ALICE experiment at LHC** since 2003 – dimuon spectrometer high-level trigger (dHLT)
  - Specialised in high-performance computing for physics applications
- 2005 : Visiting researcher, Commissariat à l'Energie Atomique, Saclay (Paris)
- late 2006 – present :
  - **Postdoctoral Fellow** at Istituto Nazionale di Fisica Nucleare, Sezione di Cagliari
  - ALICE Physics Analysis in Working Group 3 (heavy probes and quarkonia)
  - Physics performance of dHLT (muons from heavy probes)
  - ALICE commissioning (muon spectrometer and HLT)
  - massively distributed computing (production and analysis on the grid)
  - Site operations of ALICE grid Tier 2 (CyberSar) in Cagliari
- Since April 2008, co-ordinating ***South African National Compute Grid*** project, together with several partners.
  - Project management, co-ordination, deployment
  - Organising site admin, user training, application porting sessions, with GILDA, Catania

A brief digression : Our experiment : ALICE @ LHC

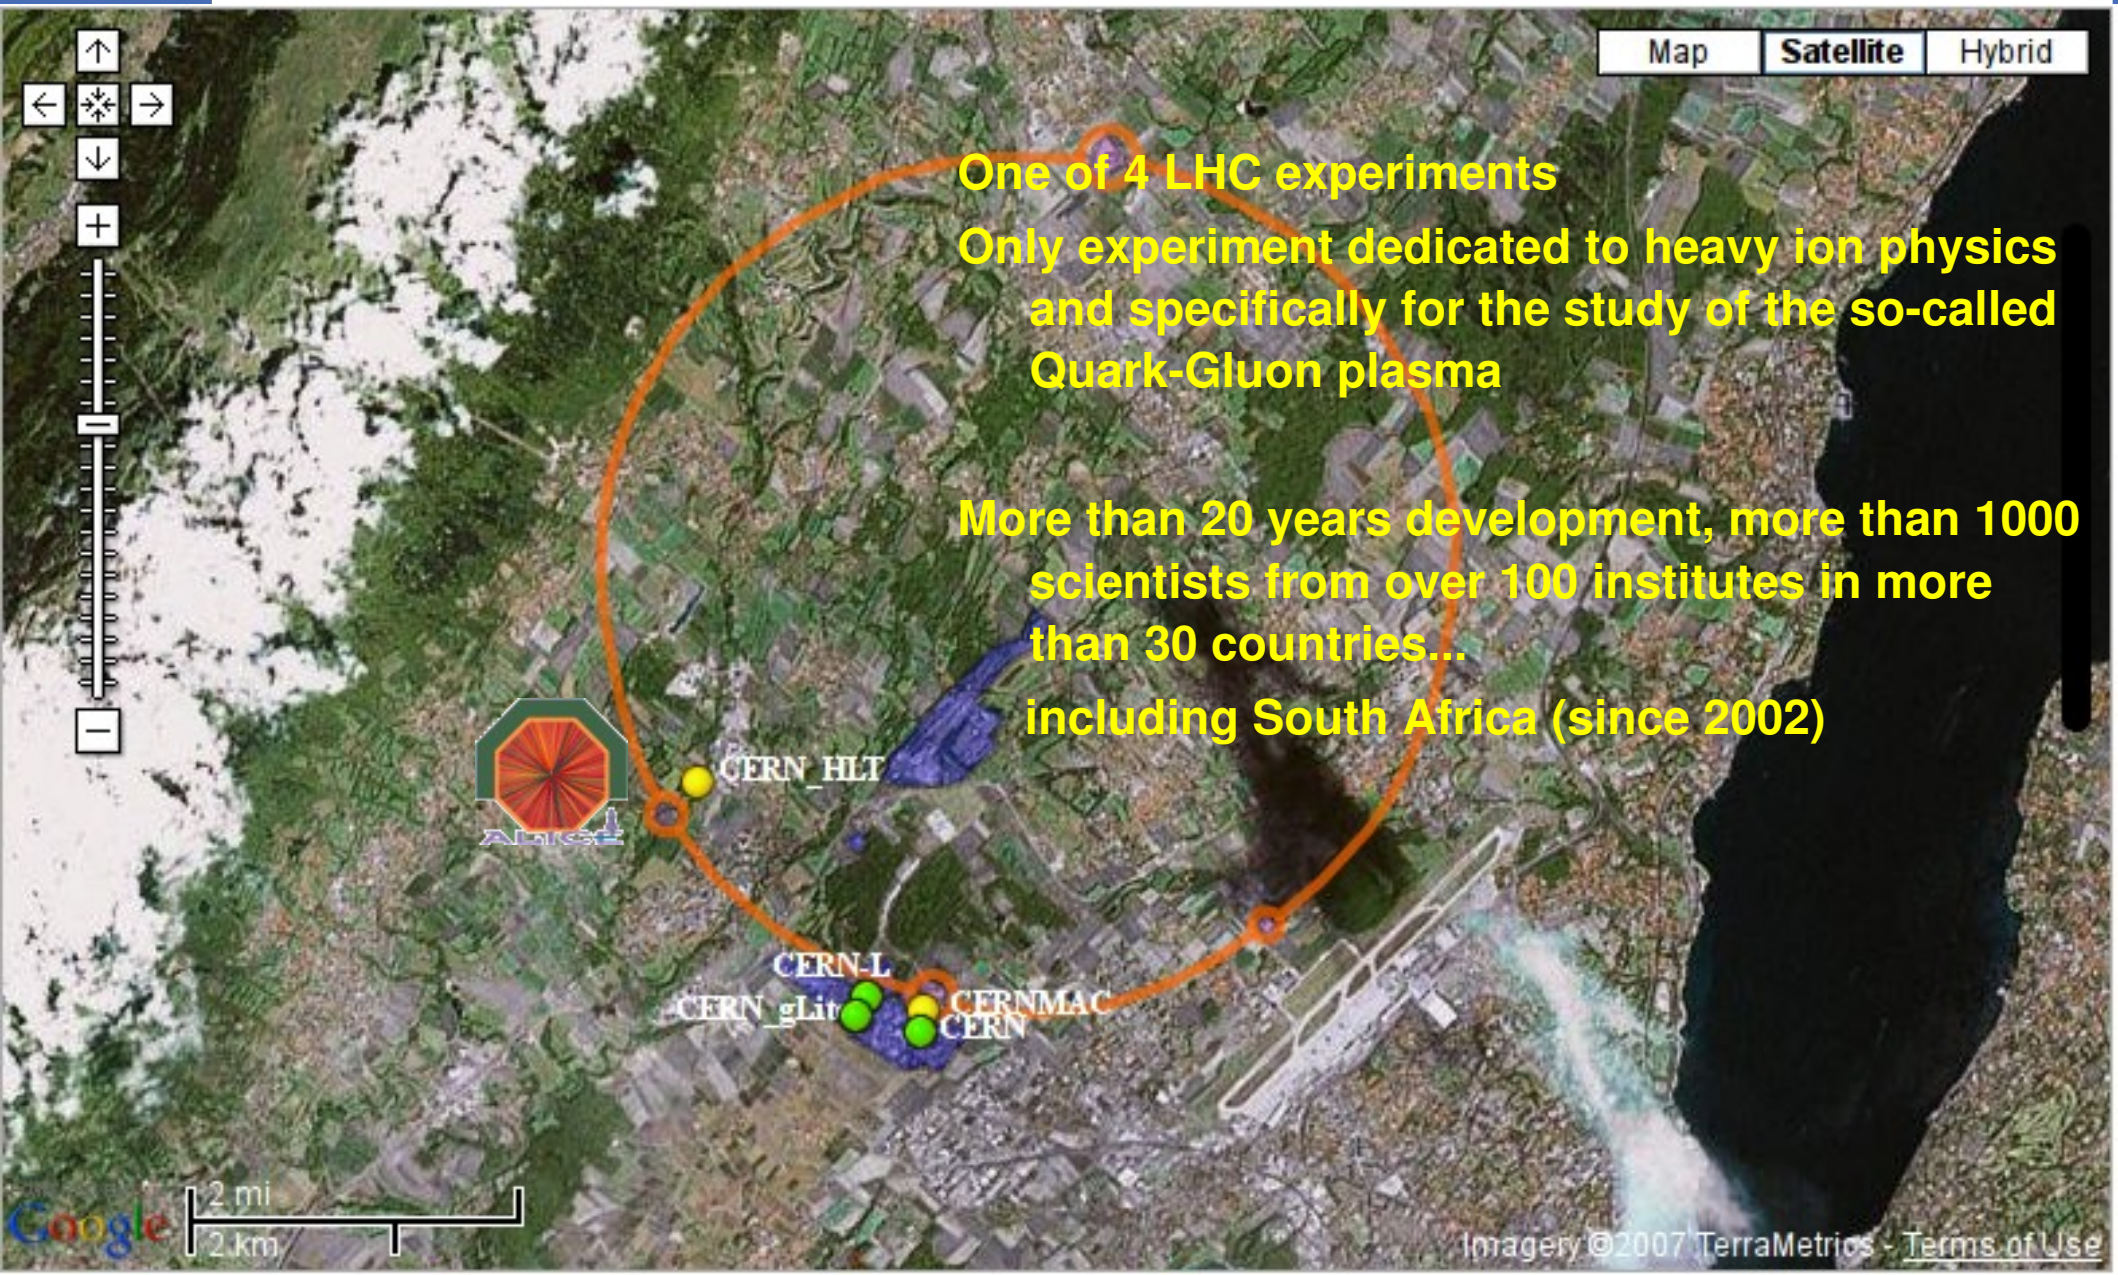

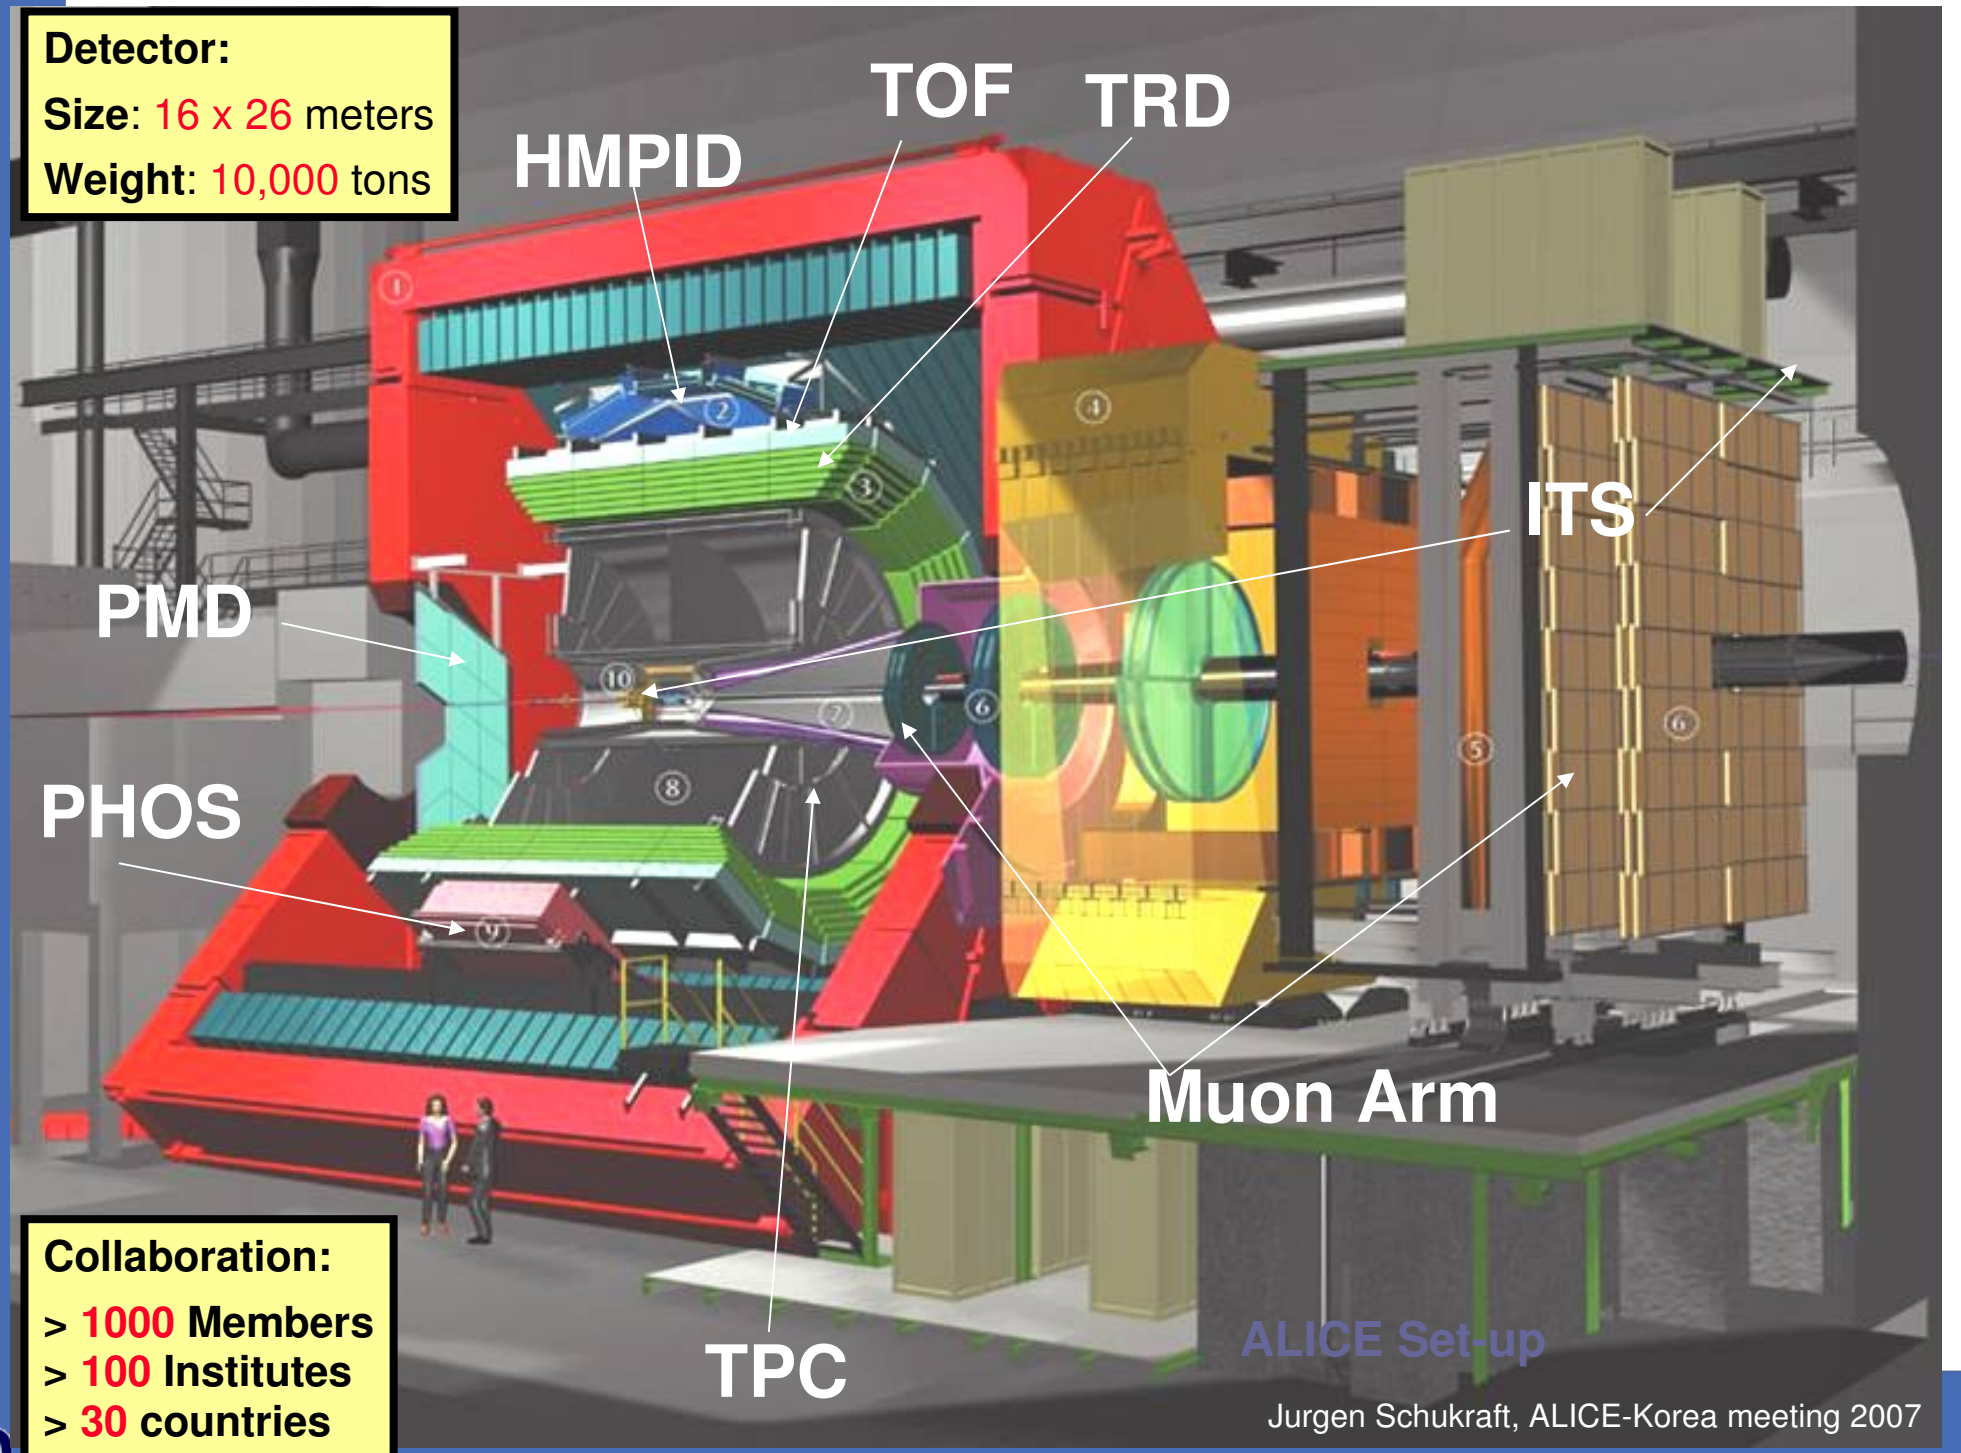
Detector:
Size: 16 x 26 meters
Weight: 10,000 tons
HMPID
TOF
TRD
ITS
PMD
PHOS
Muon Arm
TPC
Collaboration:
> 1000 Members
> 100 Institutes
> 30 countries
ALICE Set-up
Jurgen Schukraft, ALICE-Korea meeting 2007

The Large Hadron Collider : the coldest tunnel in the world

Not the Alps... the magnets !

- 27 kilometers of cryogenic tunnel
- currently being cooled to a few mK....

CERN Open Day – March 2008
The last time you, the public,
will see these experiments for a few years

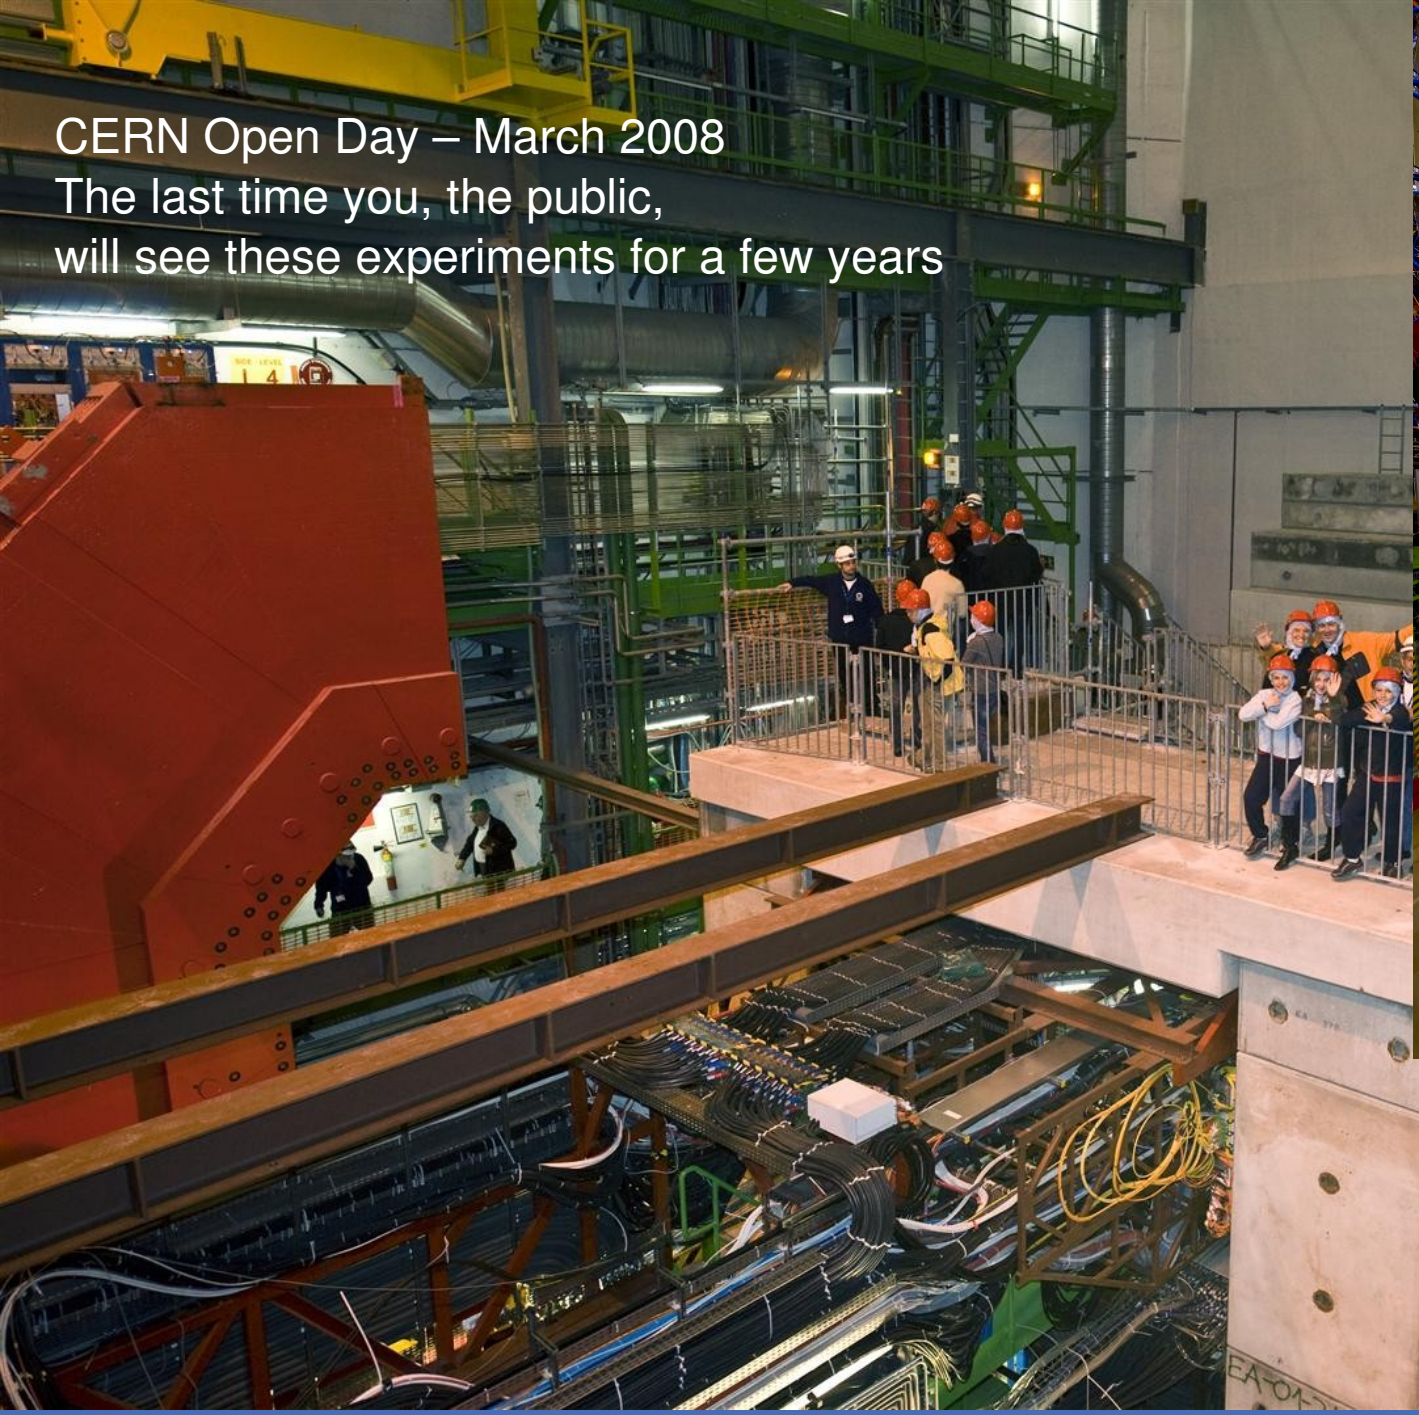


HLT III

LCG has been in SA since some years... but only specifically for ALICE

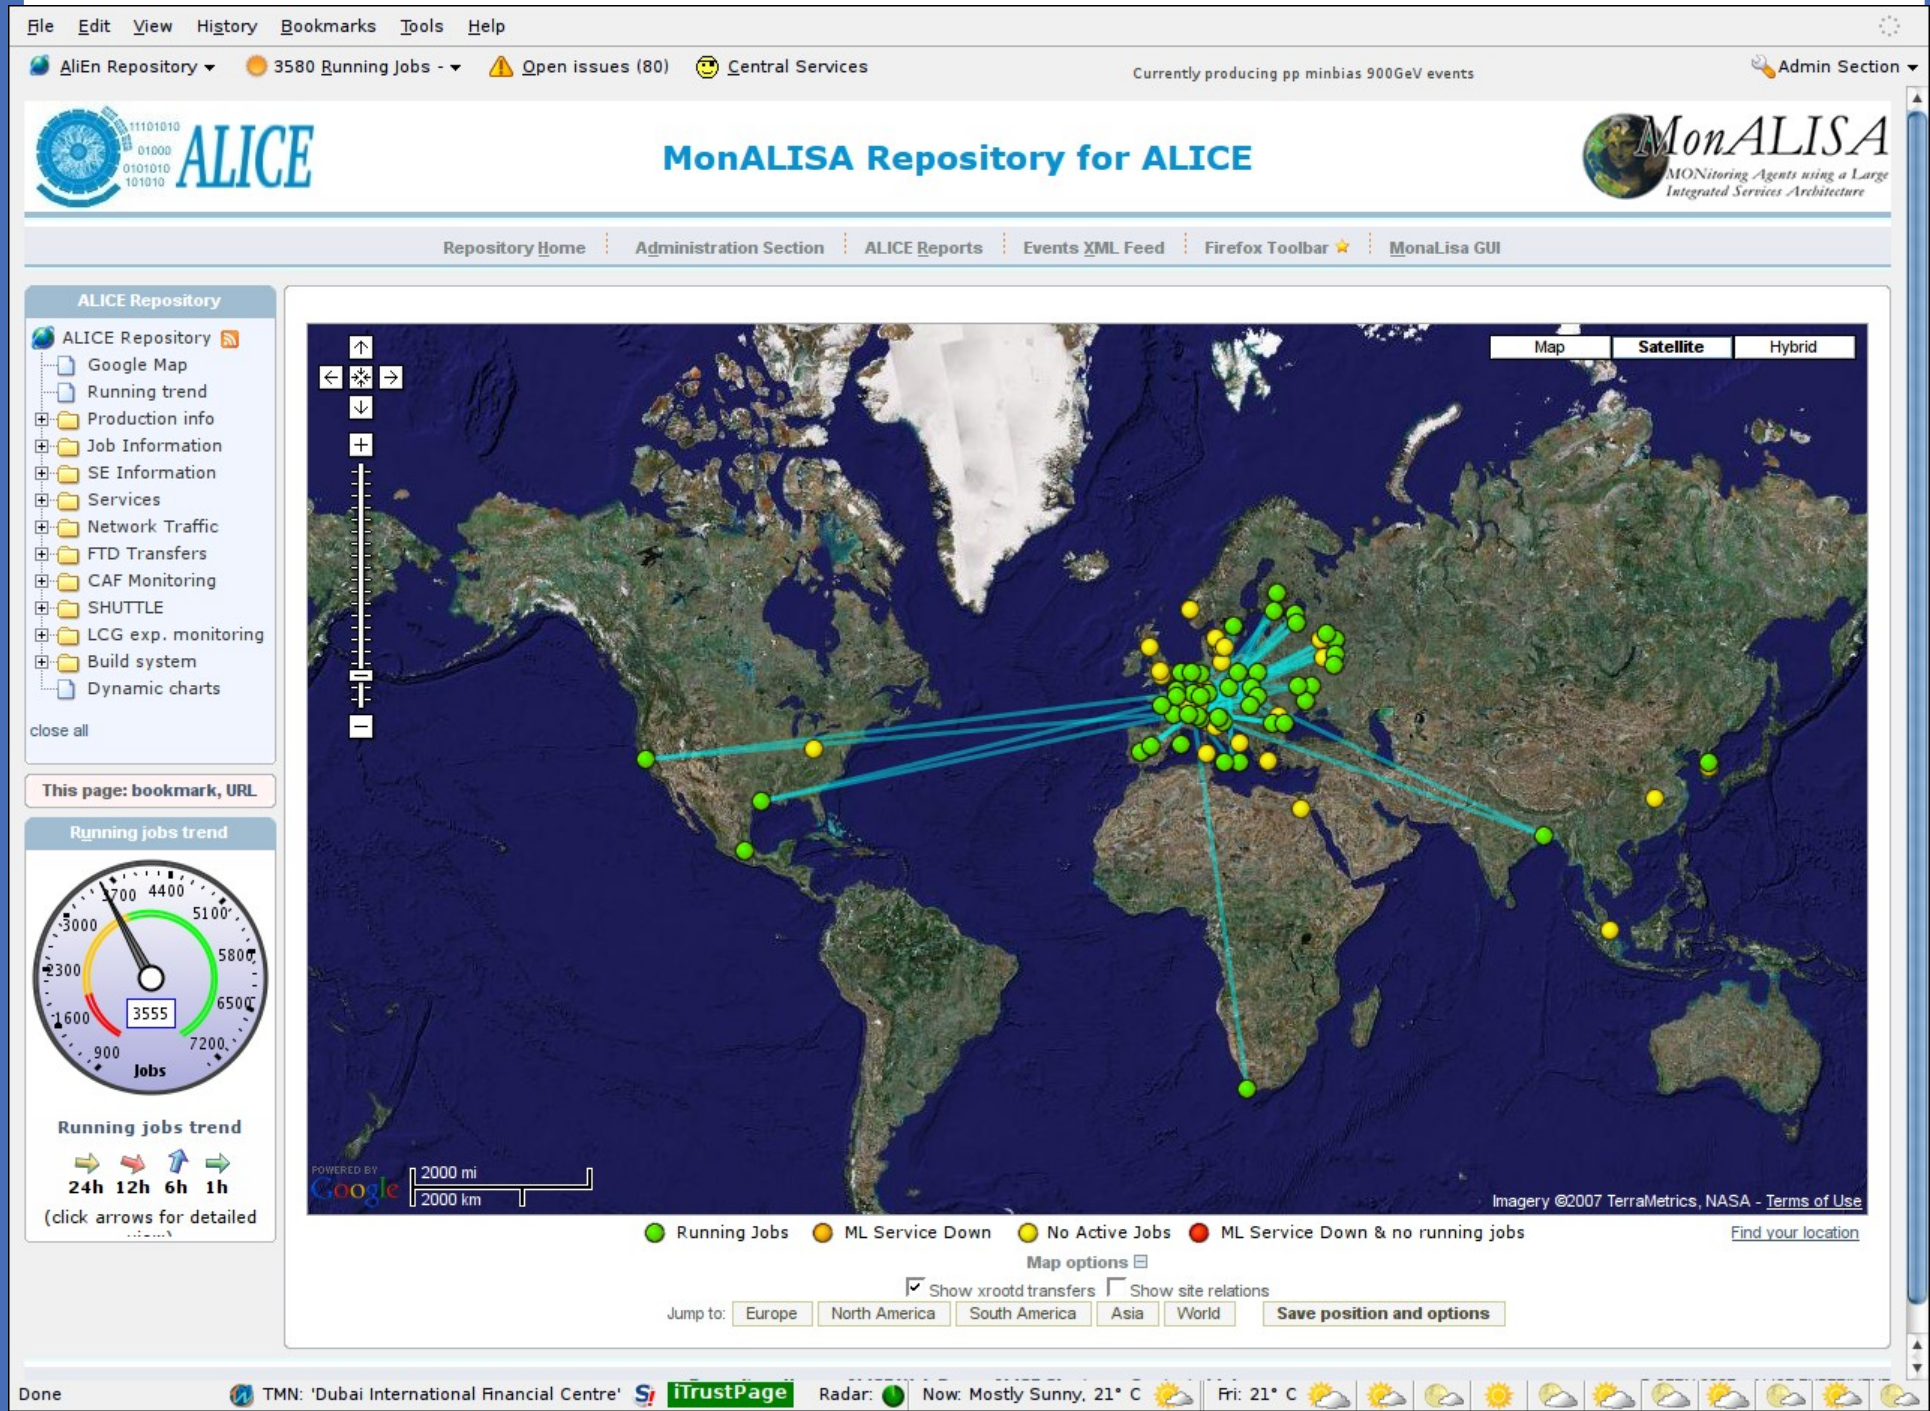

But I digress .... Back to Grid Computing : The state of play in South Africa : December 2007

- Decision to build the CHPC made serious HPC a priority; several disciplines relying more and more on HPC, domains several different applications
- But
  - still concentrated in one place
  - almost exclusively for “flagship” projects.
  - Not “open” for smaller groups, which still relied on setting up departmental and group-level clusters, on various campuses
- Missing :
  - No security infrastructure
  - No proper collaboration model in SA (but several ad-hoc collaborations) and internationally.
  - No enabler for smaller groups
  - No model for fast-turnaround computing
  - No model for industrial/commercial/academic collaboration in computing
- All of which and more can be provided by a well-funded and widely supported national grid initiative

What was happening in the meantime

- “Gigantic” transnational infrastructure may have started at CERN, but it's now grown far beyond that.
  - LHC experiments still take up a very large part of computing requirements provided by grid in EU/US
  - Data challenges prove that the infrastructure can work, in the hierarchical model of the experiments
  - LHC experiment Commissioning currently under way... first beam towards the end of the year, a very exciting time in physics, entirely impossible to do without grid computing.
- However, EU projects in several fields of science now rely on grid computing - Considered a national imperative in most EU countries.
- Several regional efforts (e.g. CyberSar in Sardinia, and Consorzio COMETA in Sicily...)
- Middleware development
  - Ease of use
  - Standardised – user base covers most of the world
  - Extended functionality – supports mpich jobs natively, handles distributed licences of proprietary software, etc etc...

EGEE and gLite – some information

- **EGEE** : **E**nabling **G**rids for **E**scienc**E**
- Supported by a large consortium of labs and organisations which promote the use of large-scale grid computing (not just physics)
- Used by all 4 LHC experiments, and several other VO's, on hundreds of sites worldwide (in co-operation with several other middelware stacks)
- Have released a middleware stack with which to do this : gLite 3.1
  - Globus functionality where needed (e.g. security infrastructure)
- gLite is a lightweight set of services needed to build your grid
- Easily configurable and very modular - can extend functionality and various interfaces

The LHC experiments are a very big deal, but ... well, lots of people couldn't care less !

- It's not just physics... we want to support all research :
- The “traditional” hard sciences :
  - Physics
    - Theoretical, high-energy, solid-state, fluid dynamics, medical, etc
  - astronomy/astrophysics
    - SKA, SALT, sky surveys, etc
  - Engineering
    - Space, mechanical, chemical, nuclear
  - Earth observation, climate studies, etc
  - Pure computer science-related research
- But also, and very importantly, life sciences :
  - Bioinformatics, drug discovery, DNA analysis, flora and fauna studies,
  - HIV studies
  - Sociology

And eventually, be open for business

- Financial services
- Commercial/industrial “production” research
- Film, media, entertainment sectors
- Government

- This is hard to get right and may never take off, but there are several reasons to plan for it seriously
  - The grid provides computing services, the problem is fundamentally one of contracts, not a technical one.
  - The grid has to be guaranteed financial and operational support, come what may – it is a good idea to diversify your funding sources
  - We are starting from scratch, we can be as imaginative as we like in our funding models, but it will take willpower and co-operation from site administrators.

- We had to start somewhere – and one of the most important choices was **which middleware ?**

The short answer : gLite from EGEE.
The long answer....

- EGEE : Enabling Grids for EsciencE
- Supported by a large consortium of labs and organisations which promote the use of large-scale grid computing (not just physics)
- Used by all 4 LHC experiments, and several other VO's, on hundreds of sites worldwide.
- Have released a middleware stack with which to do this : gLite
  - evolution of previous stacks :
    - EDG (eu-datagrid)
    - LHC computing grid (LCG)
    - Globus functionality where needed (e.g. security infrastructure)
- gLite is a lightweight set of services needed to build your grid
- Easily configurable and very modular
- Well-defined support structure, if you're part of the grid

How CERN got into grids

Enabling Grids for E-sciencE

- Partially decentralized model
  - replicate the event data at about five regional centres
  - data transfer via network or movable media

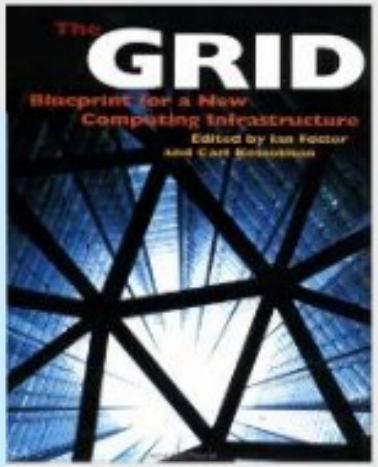

GriPhyN, iVDGL, PPDG → GRID 3 → OSG → WLCG

EU DataGrid → EGEE 1 → EGEE 2 → EGEE 3

LCG 1 → LCG 2

1994 1995 1996 1997 1998 1999 2000 2001 2002 2003 2004 2005 2006 2007 2008

Service Challenges

Cosmics

First physics

Data Challenges

The accelerator's timeline has slipped... the experiments not really, and the grid development only slightly

From: Fabrizio Gagliardi – then director of EDG, WLCG (2006)

EGEE Training Infrastructure : GILDA

- EGEE has a dedicated training infrastructure : GILDA – Grid INFN Laboratory for Dissemination Activities
- Organisation based in Catania, Italy, responsible for user training, application support, and deployment projects, such as ours
- A virtual laboratory for testing applications and new uses for the grid.
- Heavy use of virtualisation (quite similar to EC2 in that sense, but specifically for training)
- All based on gLite, with some specific extensions for training purposes.
- Has been used for several national grid infrastructure projects :

And many many more

Why gLite in South Africa

- One of the most stable and mature instruments to implement grids on a large scale.
- Several parts are based on Globus, but there are modular and fairly easy to use interfaces to these services
- Integrated with several data- and job-management tools used in back-end clusters
- The most general tool to satisfy the needs of the various applications present in the various research centres, as well as present “friendly” interfaces
- Possibility to integrate with other national and international grids currently implemented or under construction.
- **And remember** : At the time of the initiation of this project, there was an almost complete vacuum of information on various middleware options, and lots of confusion regarding which could do what. Glite grids had been working since many years, and middleware was not only stable but several extensions were being developed.
- **The right choice at the right time.**

What can gLite do ? Overview of the services provided.

gLite Services : General Services

The basic gLite middleware consists of:

- **Authentication** and **Authorization System**
- **Information System**
- **Workload Management System**
- **Data Management System**

some sites run :

- **Various Monitoring Services**

and a lot of sites run :

- **Computing Element** Grid interface to local computing cluster (LRM)
- **Worker Nodes** the local computing cluster nodes
- **Storage Element** the grid interface to the local storage solution

gLite Services : Security and Information Systems

- The **Authentication** and **Authorization System**:
  - Contains the list of all the people authorized to use gLite
    - divided by VO
    - downloaded to all machines running Grid services
    - map the gLite users to the local users of the machine
- The **Information System**:
  - provides information about gLite resources and their statuses.
  - Information published by the individual resources and copied into central databases.
  - Used by:
    - WMS/RB: match resources against job requirements and to rank them
    - DMS: storage resources and file catalog
    - monitoring systems

gLite : Description of some general services – WMS, DMS

- The **Workload Management System**:
  - manages jobs submitted by users
    - matches the job requirements to the available resources
    - schedules the job for execution on an appropriate computing cluster
    - tracks the job status
    - allows the user to retrieve the job output when ready
- The **Data Management System**:
  - Allows users to
    - move files in and out of the Grid
    - replicate files among different locations
    - locate files.
  - This is achieved by:
    - transferring data via a number of protocols (GridFTP is the most commonly used)
    - interacting with a central file catalog

gLite site services : Compute Services

- **The CE runs a gatekeeper**
  - Accepts jobs from the WMS
  - Creates a **job manager** ***(JM)*** per job
    - Generic interface to the batch system
  - The JM *only* submits or cancel a job
  - The **grid monitor** queries the status of the jobs. One instance per CE per user
- **The local batch system**
  - Last element of the chain
  - Often a server runs on the CE node
- **The WNs are the host executing the job**
  - A set of WNs managed by a CE by **local batch system** constitutes a compute cluster
  - The gLite WN does NOT run any service, but requires a minimal amount of Grid middleware, in order to interact with the other services (mostly monitoring and storage services)
  - The WN runs a job wrapper around the user executable and transports the input/output sandbox from/to the RB

Authentication and Authorization System

- User authentication based on central databases
  - one database per VO
  - database contains the certificate subjects of all gLite users.

- Databases accessed by RBs (a part of WMS), CEs and SEs
  - locally build a list of authorized users (/etc/grid-security/grid-mapfile)
  - The list maps user certificate subjects to local “pool” accounts

```
. . .
"/C=IT/O=INFN/OU=Personal Certificate/L=COSMOLAB/CN=Francesca Mocci" .cybersar
"/C=IT/O=INFN/OU=Personal Certificate/L=COSMOLAB/CN=Giuseppe Saba" .cybersar
"/C=IT/O=INFN/OU=Personal Certificate/L=COSMOLAB/CN=Ignazio Pillai" .cybersar
. . .
```

  - Users with a ‘.’ in are illegal in UNIX.
    - Signal to the globus libraries to allocate a pool account, eg cybersar005

Information System

- Based on the **Globus Metadata Directory Service** (MDS)
  - Modified by gLite to improve scalability and robustness
- The information is organized following the **Glue Schema**
  - a common data model for Grid resources monitoring and discovery, consisting of three main components:
    - Attributes of Computing Elements
    - Attributes of Storage Elements
    - Binding information for Computing and Storage Elements
- The Information System uses **OpenLDAP databases**
  - Caches and publishes information
  - Hierarchically structured

GRIS, and GIIS

- Each Grid resource is served by a **Grid Resource Information Service** (GRIS)
  - USUALLY runs on the same machine.
- At the back of the GRIS is the **Information Provider**
  - Basically a script that collect required information
    - Static Info (i.e. GlueCEUniqueID)
    - Dynamic info (i.e. GlueCEStateWaitingJobs)
- The information is written in openLDAP database
- An openLDAP server returns the result of eventual queries
  - Runs on port 2135
  - Runs in insecure mode

- Local GRISes register to the **Grid Information Index Service** (GIIS of BDII site), normaly hosted by the CE. One GIIS per site running in insecure mode (port 2135 or 2170). When queried, returns info (site name, software, mpich support, disk space, vo supported etc etc *see config file for site*) from each registered GRIS, in aggregated mode

The BDII

- The top of the Information System the **Berkeley Database Information Index** (BDII)
  - Introduced by NIKHEF during EDG
  - avoid the problem of site GIISes hanging an entire in hierarchy of GIISes within MDS.
- BDII scripts query a list of GIISes (and GRISes)
  - and populate a standard OpenLDAP tree of the data.
- The BDII uses two databases to improve scalability
  - One read-only db and one write-only db
    - Switched when an update is completed
- NO soft state registration from site to higher level
  - A configuration file references all the resources to check for
  - Each VO can configure its BDII to query only relevant sites
    - Helps in excluding “problematic” sites

The Workload Management System

- Inherits several elements of the Globus Toolkit
  - Grid Security Infrastructure (GSI)
  - Globus Resource Allocation Manager (GRAM)
  - Global Access to Secondary Storage
  - (…)
- WMS was developed by EDG and VDT
  - Several modifications by LCG
  - It also relies on regular batch systems to manage worker nodes
    - LSF, PBS (in several flavors), Condor …
- The LCG-2 WMS is deployed on five kinds of machines:
  - **User Interface** (UI)
  - **Resource Broker** (RB)
  - **Computing Element** (CE)
  - **Worker Node** (WN)
  - **Proxy Server** (PS)

The Data Management system

- The DMS relies on two kind of services:
  - The File Catalog (FC)
  - The Storage element (SE)
- The **File Catalog**
  - Needed to maintain mappings between
    - Global Unique Identifiers of files (GUID)
    - Logical file names (LFN)
    - Phisical File Names i.e. Locations (PFN)
  - Is an application server frontend to a database backend
    - Currently ORACLE, but support also for MySQL
- The **Storage Element**
  - Provide uniform access to storage space
  - SE can manage several kinds of back-ends...

Grid Topology and Services

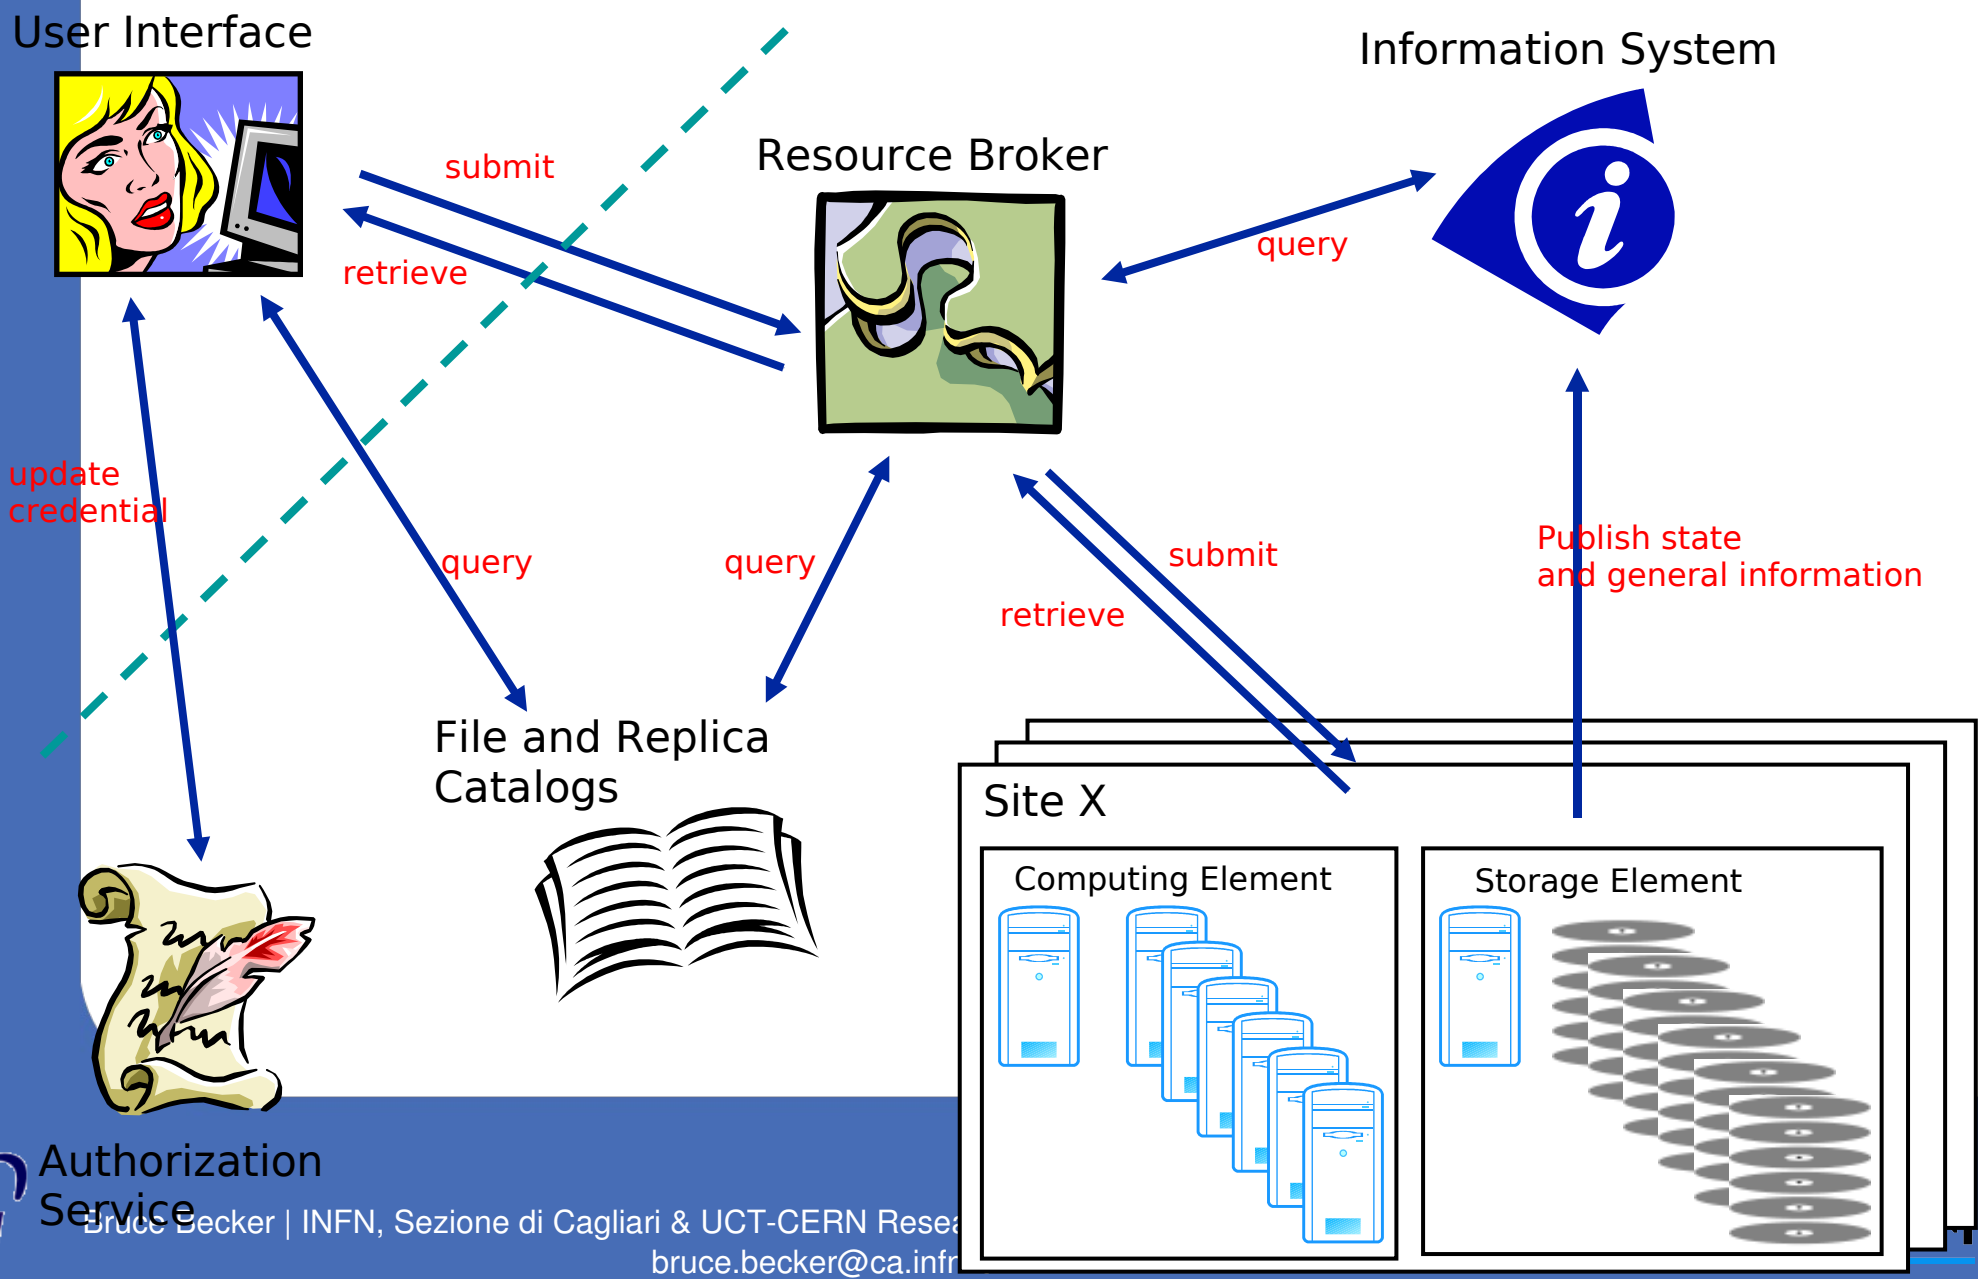

Job Preparation – Job Description Language

- Information to be specified when a job has to be submitted:
  - Job characteristics
  - Job requirements and preferences on the computing resources
    - Also including software dependencies
  - Job data requirements

- Information specified using a Job Description Language (JDL)
  - Based upon Condor's *CLASSified ADvertisement language (ClassAd)*
    - Fully extensible language
    - A ClassAd
      - Constructed with the classad construction operator
      - It is a sequence of attributes separated by semi-colon (;).

- So, the JDL allows definition of a set of attribute, the WMS takes into account when making its scheduling decision

Having made a basic technology choice, we can begin.

SA National grid : Partners (not a closed list, but not a free-for-all)

- A national grid initiative doesn't happen in isolation, collaboration is required both between national institutes and international partners
- Our collaborators
  - South African Universities and National Labs
  - Istituto Nazionale di Fisica Nucleare, Sezione di Cagliari and Catania
  - The EGEE training infrastructure GILDA (Grid INFN Laboratory for Dissemination Activities). Also supported by EU FP7 Project 'EPIKH'
  - Centre National de Recherche Scientifique (CNRS)
  - UNESCO/HPLabs anti brain-drain project with African countries
  - Staff exchange and training funding from EU FP7 project “EPIKH” - Exchange Programme to advance e-Infrastructure Know-How. (8 Meuros for 4 years)
- (Current) Sponsors
  - UCT-CERN Research Centre Consortium project, via CHPC
  - Institutes hosting phase-1 sites
  - Microsoft South Africa
  - Sun Microsystems EMEA
  - HP South Africa

The South African National Compute Grid : Where are we ?

- **Our Task : Implement a national production grid infrastructure for all researchers who want to use it.**
- **First** : implement the security infrastructure
  - SA had no Certificate Authority
  - We have now identified the Meraka Institute as the host for the CA
  - Working very closely with the INFN CA, in terms of technical deployment and liaison with IGTF, EUGridPMA.
  - This process takes several months though :
    - In the meantime, we have assigned two Registration Authorities (from Meraka and iThemba LABS)
    - These can issue certificates to users and hosts, which are certified by the INFN CA.
- **Next** : deploy sites...

South African National Compute Grid : Timelines for Deployment

- Organise deployment into an initial prototype deployment with 3 phases :
  - Phase 1 : May - July. First site admin training, install first sites, make simple tests
  - Phase 2 : July – September. Second site admin training, user training, application porting tutorials. Install second round of sites. Make advanced tests of the infrastructure
  - Phase 3 : October – December. Full runs of pilot applications, application debugging, preparation for production; “Opening Day” in coincidence with CHPC National Conference
- Concurrently to all phases :
  - dissemination of information about the grid
  - polling and informing user communities,
  - testing infrastructure
  - Implementation of security infrastructure

South African National Cyberinfrastructure

- The grid will not run on fresh air – in fact, it will mostly run on already existing resources :
- The “big iron” : CHPC cluster in CT, Blue Gene-P, CHPC phase 2... etc
- The national research network : SANReN
- The institutional clusters : several universities and labs are implementing O(100) core clusters :
  - UFS, UNW, US, UKZN, UP
- The Departmental clusters : very many smaller (O(10)) clusters on very many campuses around the country.... too many to mention.
- The National Facilities
  - Telescopes, observatories
  - Acelerators
  - Large Databases
  - etc...

We want to integrate this...

CERN Computer Centre

But also this...

- Little clusters
- Departmental clusters
- Transient sites
- etc...

Into this :

We need
- More yellow dots
- Define how to include the blue dots
- To fill in the empty parts of the country...

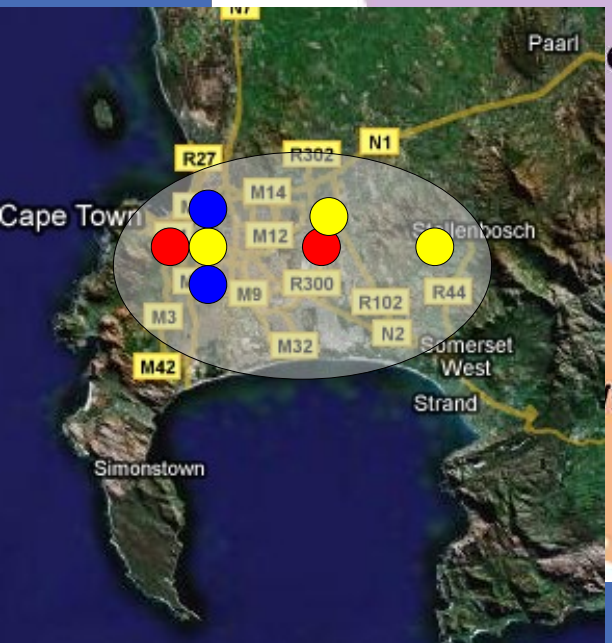

LIMPOPO PROVINCE
Pietersburg
MPUMALANGA
Mmabatho
NORTHWEST
GAUTENG
?
FREE STATE
Kimberley
KWAZULU/ NATAL
Maseru
Lesotho
?
Pietermaritzburg
ORTHERN CAPE
SALT/SAAO
EASTERN CAPE
Bisho (King William's Town)
ESTERN CAPE
?

Phase 1 sites (currently operational)37

Phase 2 sites

Big Fish

South African Institutes Involved So Far

- Phase 1 :
  - UCT-CERN Research Centre
  - University of the Free State
  - Northwest University
  - iThemba LABS, Cape Town
  - University of the Witwatersrand and University of Johannesburg
  - CSIR Cluster Computing Centre (C4)
- All phase 1 sites have been already installed and configuration is now being tested.
- Full co-operation of GILDA t-infrastructure

The SA Grid data model

- It depends a lot on who is in the first 3 phases and what support we can expect from government
- A strictly hierarchical model such as DataGrid appears not to be well-adapted to our needs
  - No one single overwhelming source of data
  - Several “small” research groups (compared to LHC experiments), more equally distributed resources
  - No massive financial backing
- We need something that will be USEFUL and not just bragging rights.
  - Initial model is one of two regions of central services, north and south, perhaps 3 including a coastal region, with strong co-ordination between these sites
  - Other sites relate publish availability first to the nearest central service, but are visible nation wide
- Main factor contributing to data model at this stage is the **bandwidth**, but it will not **always** be like that. We need to think in wider terms and **design something that we want, not around obstacles.**
- **We get to decide on this...**

Picture of LHC (big data) model

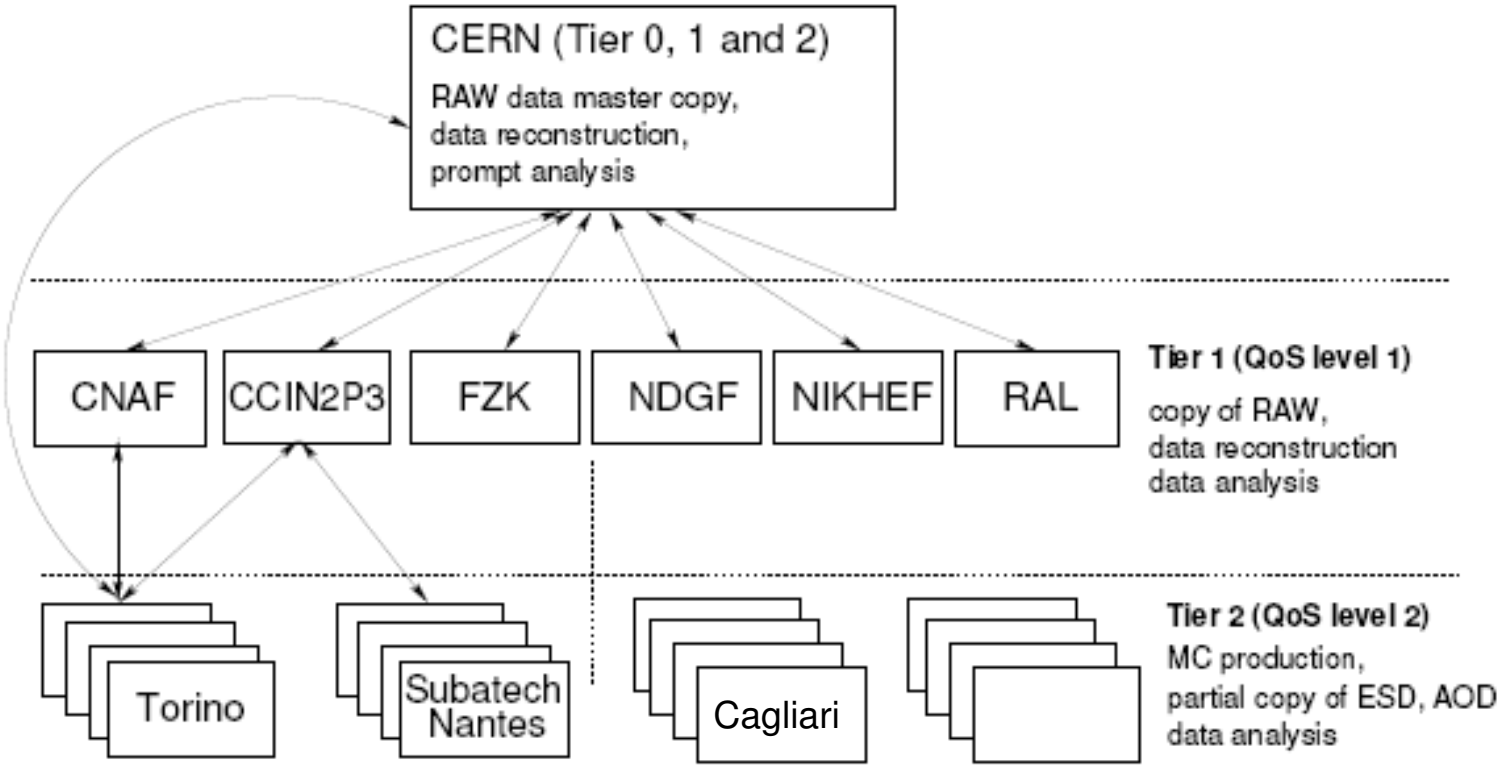

We probably want something more like

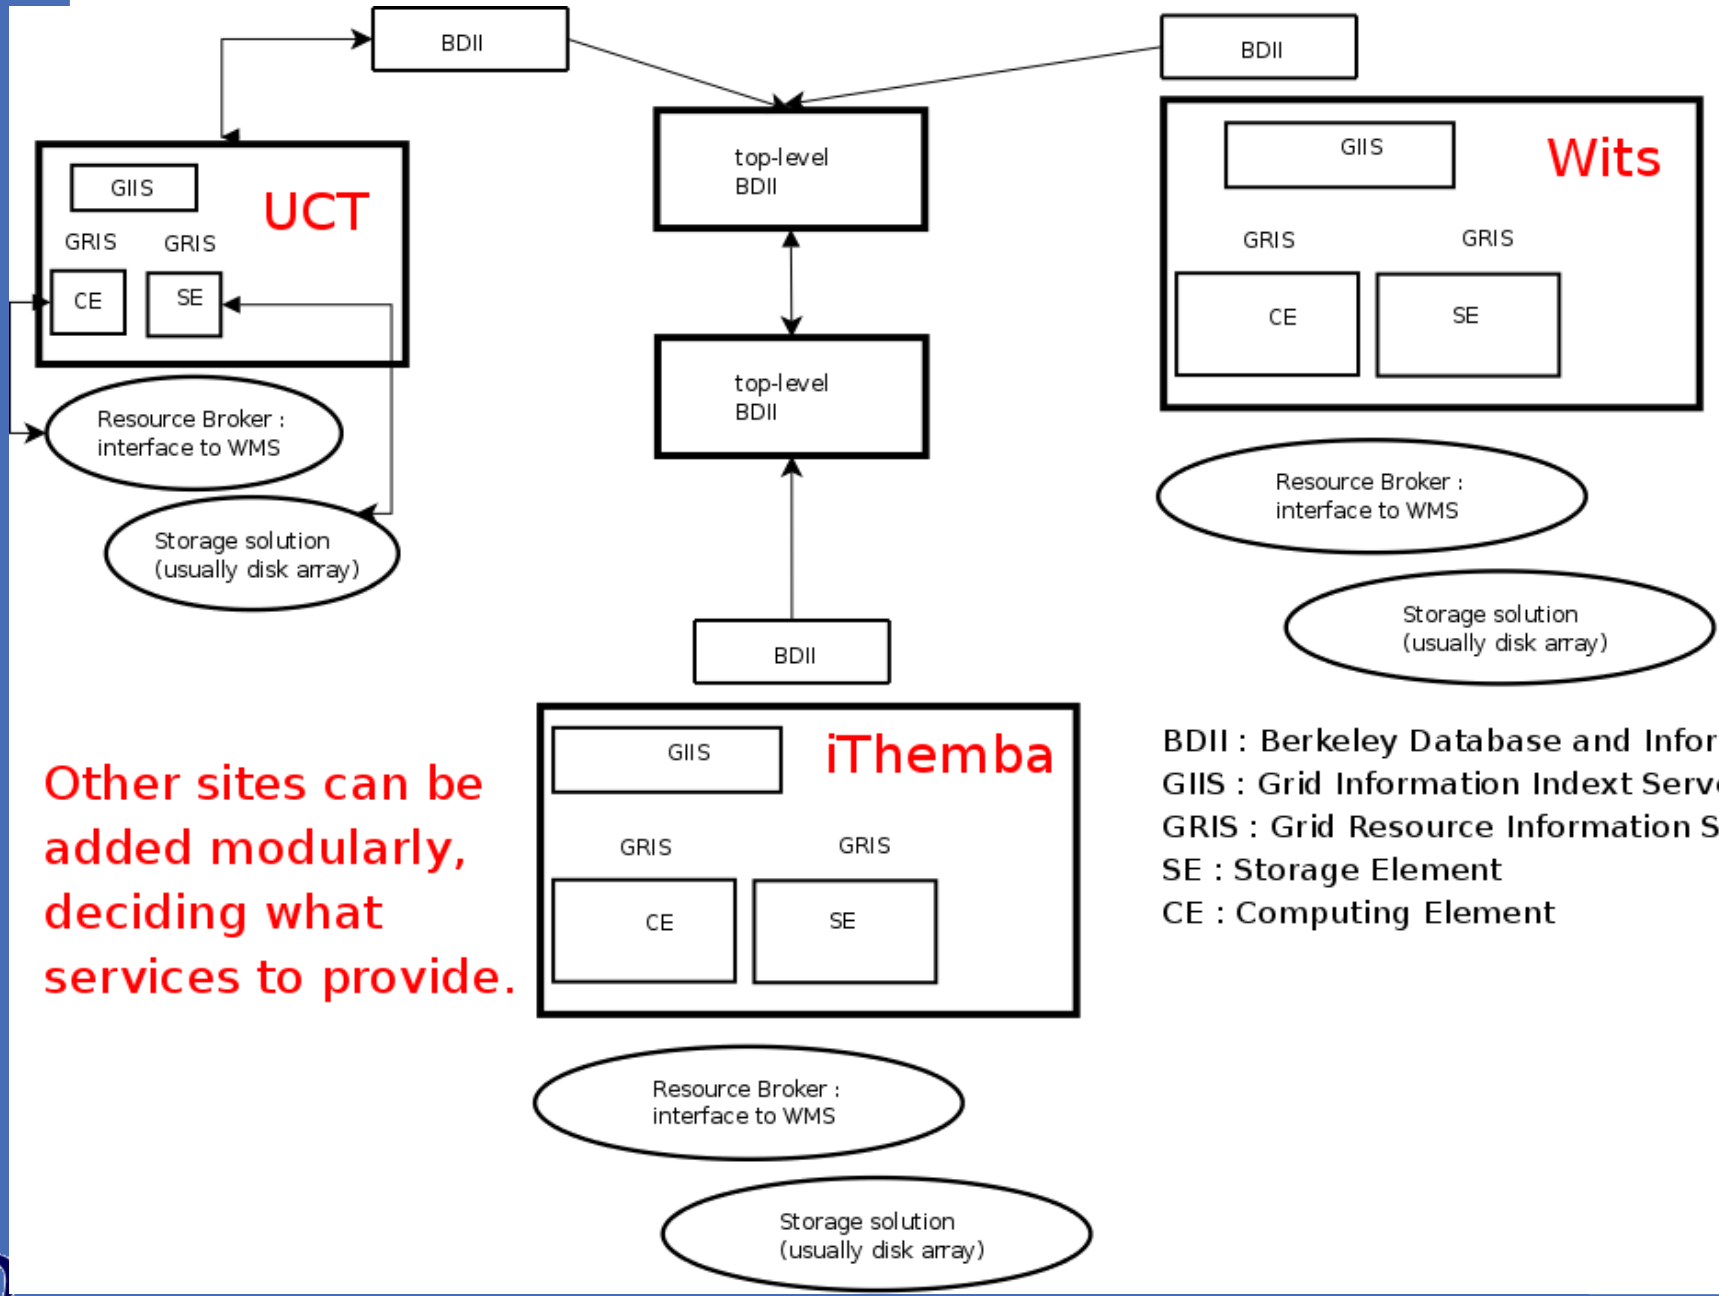

Other sites can be added modularly, deciding what services to provide.

BDII : Berkeley Database and Information Index
GIIS : Grid Information Indext Server
GRIS : Grid Resource Information Service
SE : Storage Element
CE : Computing Element

Applications that will be run on the grid

Initially, we are porting applications for the following fields -

- Physics (a lot of physics...)
- Astrophysics
- Bioinformatics (a lot of bioinformatics) + WISDOM grid
- Chemical Engineering
- Most of these are already CHPC Flagship or Consortium projects, and are well-developed applications.

But....

If you don't see your application here, we haven't yet heard from you.

- Experience so far with decent national research networks has shown that almost any (but not strictly every) application is suitable for the grid
- We are providing dedicated training and are eager to profile lots of different applications to help you port them
  - Application questionnaire
  - Hands-on training
  - Support from GILDA after the training
- We specifically want to hear from the groups who are skeptical about using the grid.
  - We will want you to tell us why...
  - And figure out if you're a really special case, or if specific support and middleware extensions would allow you to port the applications

Fitting in with the rest of the world

- Once we have our house in order, we want to join the party in the rest of the world :
  - Plug into international grids
  - Support several large science experiments (not just CERN !)
  - Offer more freedom to sites to configure themselves as they like
- We are learning
  - PI2S2 and COMETA consortia in Sicily, several industrial and commercial applications.
  - CyberSar in Sardinia, and several implementations of “our” scale.
  - GILDA working on multi-middleware, gLite extensions
  - and the WLCG/EGEE/OSG experience in integration.
- It's not easy, but...
  - We need to focus on building something functional now
  - These problems are being solved in the rest of the world right now, we will benefit from the experience in the near future.
  - We have very strong, long-term, dedicated support

Multi-middleware

- For specific cases, OSG and gLite production grids are already since a long time integrated at the site level.
- Lots of work has been done on providing a unified access to these grids by, amongst others, the ICEAGE project :
  - http://www.iceage-eu.org/v2/index.cfm
  - The International Collaboration to Extend and Advance Grid Education
- https://grid.ct.infn.it/twiki/bin/view/ICEAGE/MultiMiddleware

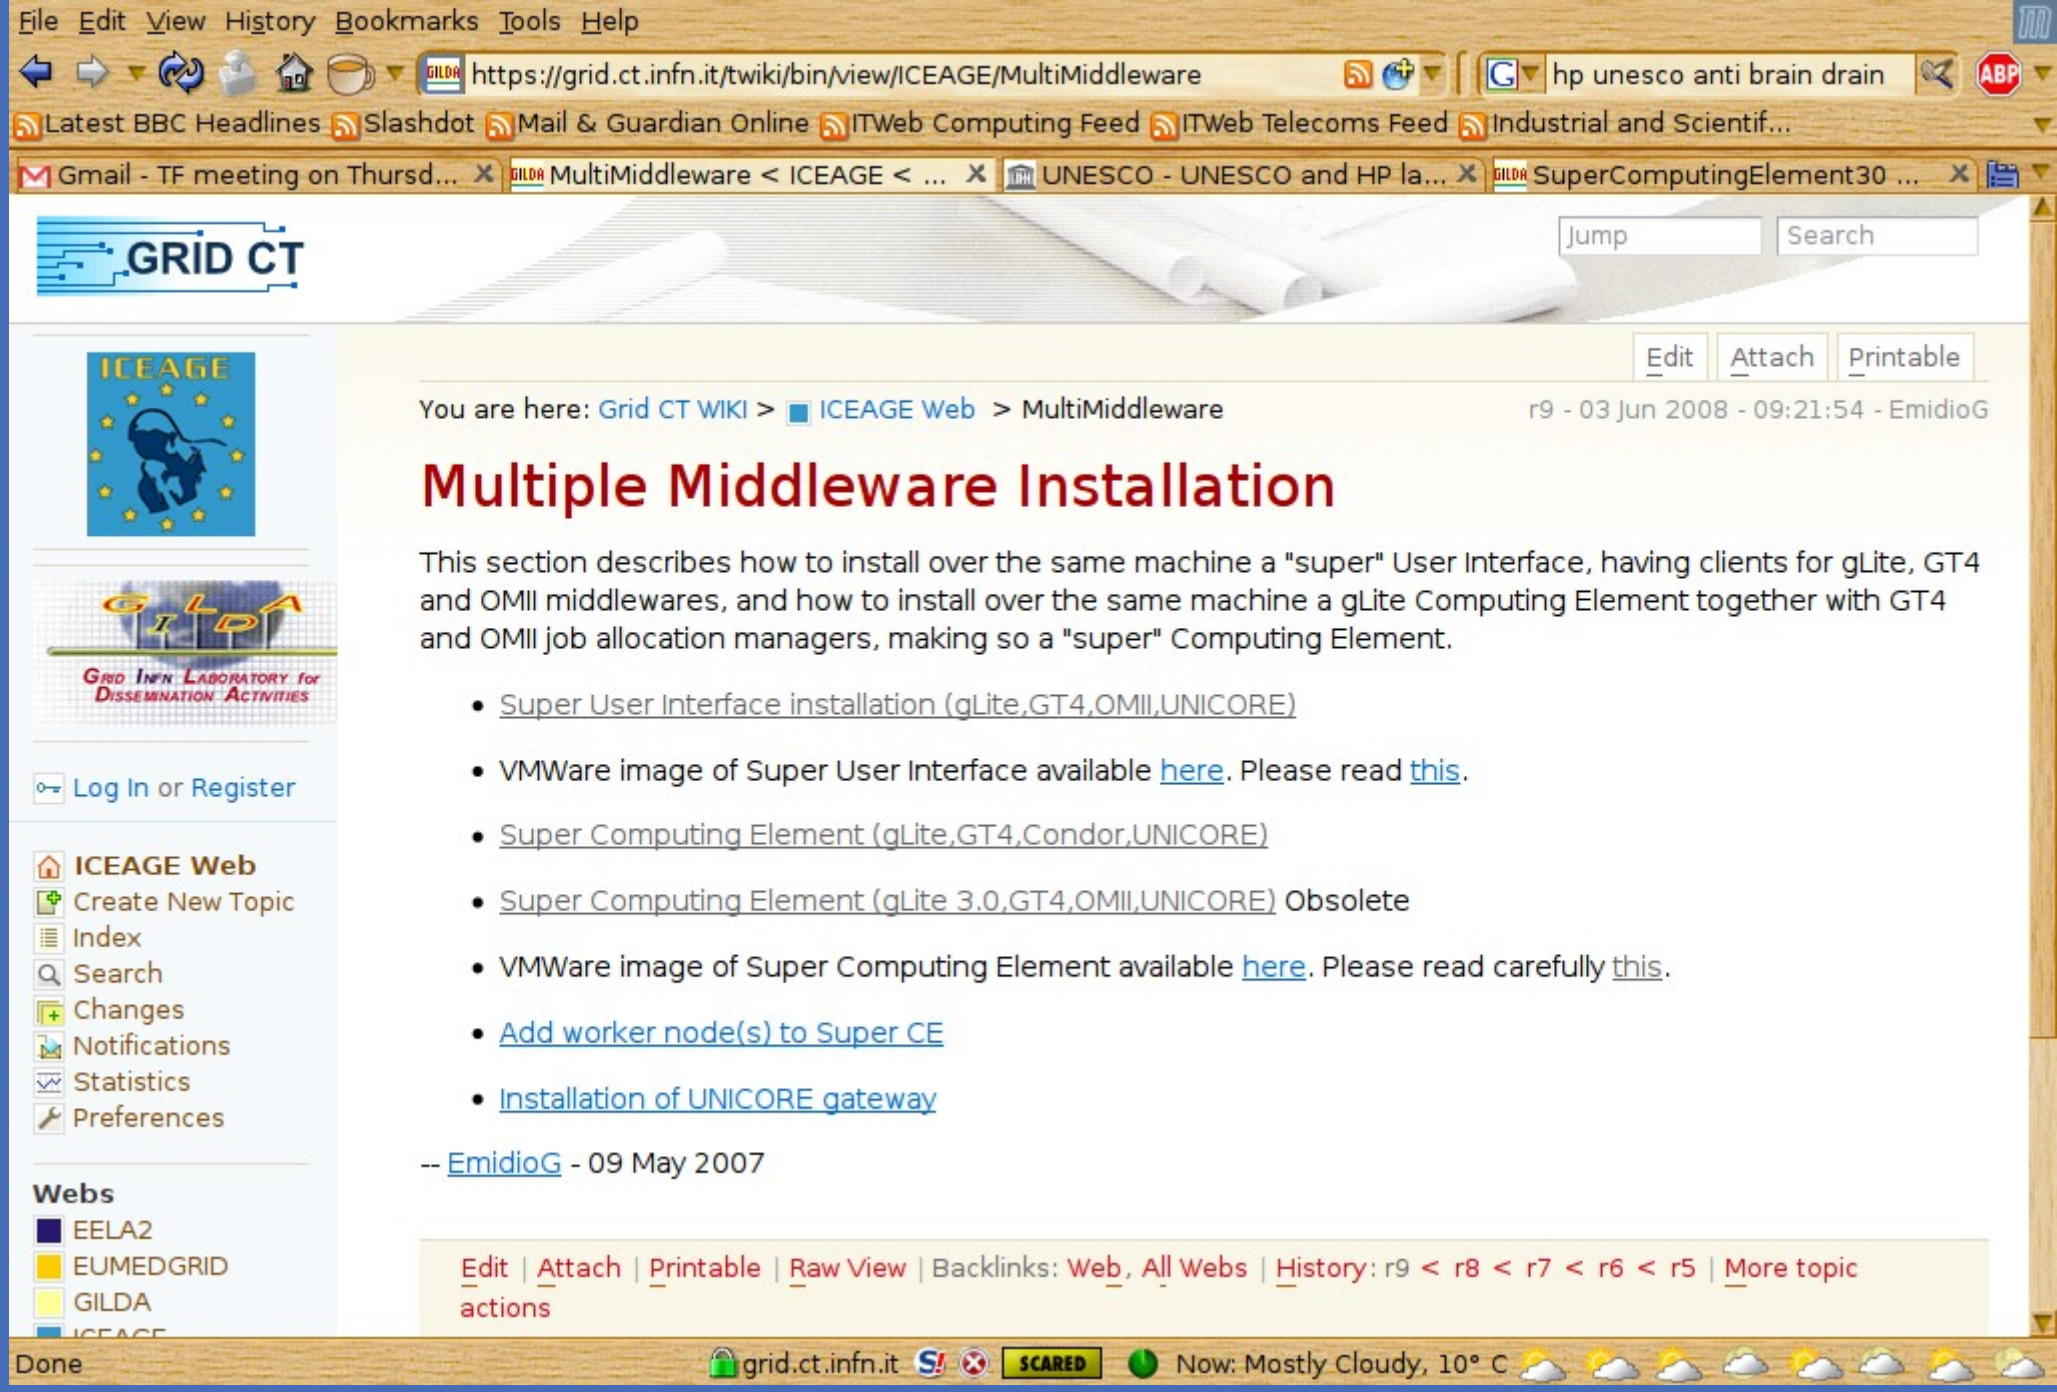
File Edit View History Bookmarks Tools Help
https://grid.ct.infn.it/twiki/bin/view/ICEAGE/MultiMiddleware
hp unesco anti brain drain
Latest BBC Headlines
Slashdot
Mail & Guardian Online
ITWeb Computing Feed
ITWeb Telecoms Feed
Industrial and Scientif...
Gmail - TF meeting on Thursd...
MultiMiddleware < ICEAGE < ...
UNESCO - UNESCO and HP la...
SuperComputingElement30 ...
GRID CT
Jump
Search
ICEAGE
GILDA
Grid Infn Laboratory for Dissemination Activities
Log In or Register
ICEAGE Web
Create New Topic
Index
Search
Changes
Notifications
Statistics
Preferences
Webs
EELA2
EUMEDGRID
GILDA
Edit
Attach
Printable
You are here: Grid CT WIKI > ICEAGE Web > MultiMiddleware
r9 - 03 Jun 2008 - 09:21:54 - EmidioG
Multiple Middleware Installation
This section describes how to install over the same machine a "super" User Interface, having clients for gLite, GT4 and OMII middlewares, and how to install over the same machine a gLite Computing Element together with GT4 and OMII job allocation managers, making so a "super" Computing Element.
Super User Interface installation (gLite,GT4,OMII,UNICORE)
VMWare image of Super User Interface available here. Please read this.
Super Computing Element (gLite,GT4,Condor,UNICORE)
Super Computing Element (gLite 3.0,GT4,OMII,UNICORE) Obsolete
VMWare image of Super Computing Element available here. Please read carefully this.
Add worker node(s) to Super CE
Installation of UNICORE gateway
-- EmidioG - 09 May 2007
Edit | Attach | Printable | Raw View | Backlinks: Web, All Webs | History: r9 < r8 < r7 < r6 < r5 | More topic actions
Done
grid.ct.infn.it
SCARED
Now: Mostly Cloudy, 10° C

Let's not forget the rest of Africa...

- We are collaborating with a UNESCO/HP project to alleviate brain drain in Africa, using high-performance computing with a grid aspect.
- We work together at this point only on training aspects
  - http://tinyurl.com/6mp8b2
  - Currently involving Zimbabwe, Nigeria, Egypt, Ghana.
- However, we would likeSouth Africa to become a support and operations centre for sub-saharan Africa, like has already been done in the Mahgreb by GILDA.

Our timeline until the end of this year.

- June (Catania):
  - site admin + phase 1 sites
  - RA of INFN CA
- July (Cagliari):
  - site admin + user + phase 2 sites
  - Open day
  - Deployment meeting
  - Application porting
- August
  - Standalone prototype in SA, apart from security infrastructure
  - Run first jobs, obtain first user experience and feedback
  - Start CA implementation, talks with EUGridPMA, IGTF

Our timeline 2

- September (Lyon)
  - site admin
  - Application porting,
  - VO setup
  - User and site admin feedback
- October (TBD)
  - User training
- November (Cape Town)
  - Phase 3 sites + user + site admin training
  - User + application porting
  - Fine tuning user and site admin experience
- December (Durban)
  - Open Day
  - Inauguration... from prototype to production
- And then the really hard work will start...

You're interested ? What to do...

- Come to Cape Town next week
  - Look at the agenda http://indico.ct.infn.it/conferenceDisplay.py?confId=67
  - And register...
- Contact me
  - bruce.becker@ca.infn.it
- Tell your friends – and your bosses !

basta.

Information System Hierarchy

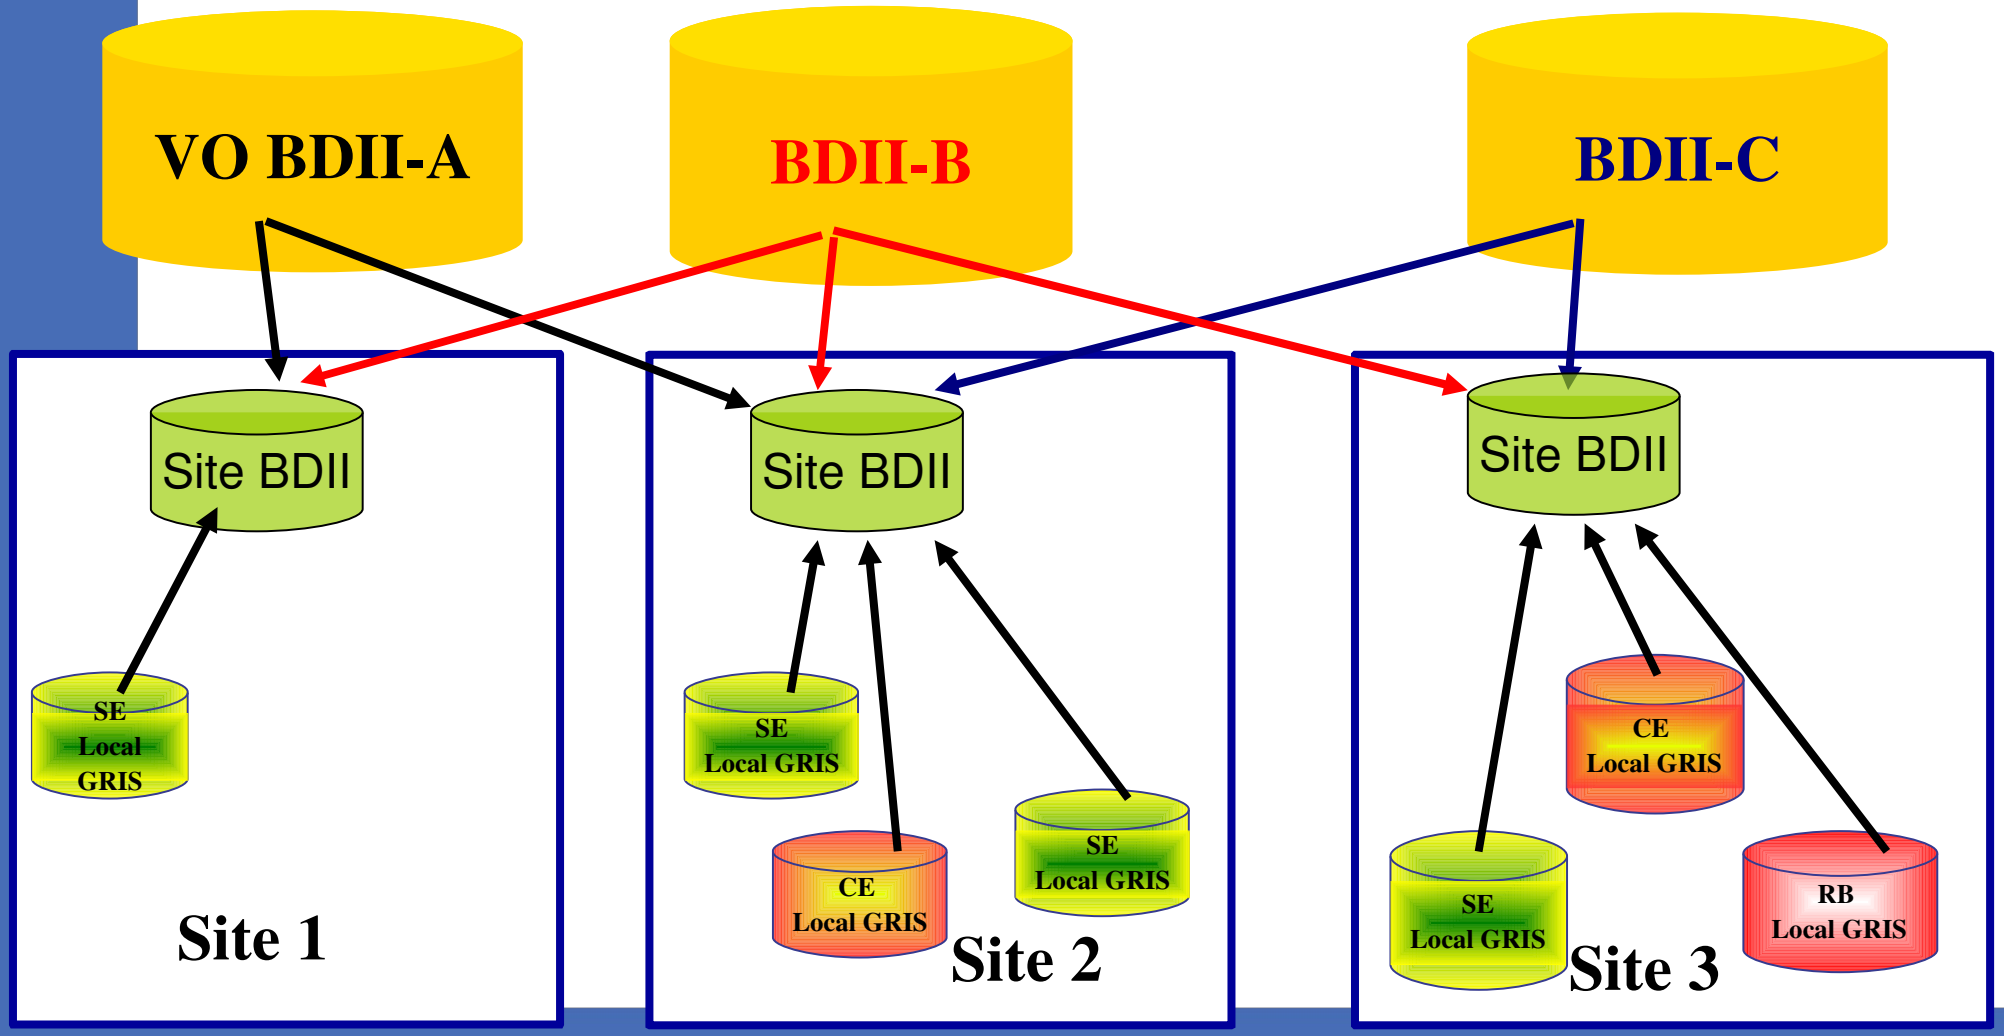

JDL relevant attributes

- **JobType**
  - *Normal* (simple, sequential job), *Interactive*, *MPICH*, *Checkpointable*
  - Or combination of them
- **Executable** (mandatory)
  - The command name
- **Arguments** (optional)
  - Job command line arguments
- **StdInput**, **StdOutput**, **StdError** (optional)
  - Standard input/output/error of the job
- **Environment (optional)**
  - List of environment settings
- **InputSandbox** (optional)
  - List of files on the UI local disk needed by the job for running
  - The listed files will automatically staged to the remote resource
- **OutputSandbox** (optional)
  - List of files, generated by the job, which have to be retrieved
- **VirtualOrganisation** (optional)
  - A different way to specify the VO of the user

Essential JDL

- **At least one has to specify the following attributes:**
  - the name of the executable
  - the files where to write the standard output and standard error of the job
  - the arguments to the executable, if needed
  - the files that must be transferred from UI to WN and viceversa

```
[
Executable = “ls -al”;
StdError = “stderr.log”;
StdOutput = “stdout.log”;
OutputSandbox = {“stderr.log”, “stdout.log”};
]
```


Bruce Becker | INFN, Sezione di Cagliari & UCT-CERN Research Centre | OSG School Wits | July 2008
bruce.becker@ca.infn.it

Job Preparation

- Information to be specified when a job has to be submitted:
  - Job characteristics
  - Job requirements and preferences on the computing resources
    - Also including software dependencies
  - Job data requirements
- Information specified using a Job Description Language (JDL)
  - Based upon Condor's *CLASSified ADvertisement language (ClassAd)*
    - Fully extensible language
    - A ClassAd
      - Constructed with the classad construction operator []
      - It is a sequence of attributes separated by semi-colon (;).
- So, the JDL allows definition of a set of attribute, the WMS takes into account when making its scheduling decision

General Jdl

```
Type = "Job";
JobType = "Normal";
Executable = "/bin/sh";
StdOutput = "std.out";
StdError = "std.err";
InputSandbox = {“script.sh",“program.exe"};
OutputSandbox = {"std.out","std.err"} ;
Arguments = “script.sh";
Requirements = other.GlueCEUniqueId == "ce-
  cybr.ca.infn.it:2119/jobmanager-lcglsf-infinite"
```