

gLite

Kilian Schwarz
GSI

This is a summary talk composed mainly from existing transparencies.


These are taken from the plenary sessions of the EGEE Den Haag meeting and the last LCG-Review

I am no gLite developer, but rather an ALICE Grid user, so some things may be seen from the ALICE point of view (current interest: use existing prototype for ALICE DC as soon as possible, especially also in the Tier1 centre of Karlsruhe)

Transparencies are taken from

- Frédéric Hemmer
- **Julia Andreeva**
- Leanne Guy
- Fabrizio Gagliardi
- F.Carminati
- Erwin Laure
- Massimo Lamanna
- **Alan Blatecky, Eike Jessen, Thierry Priol, David Snelling**
- Les Robertson
- Ian Bird

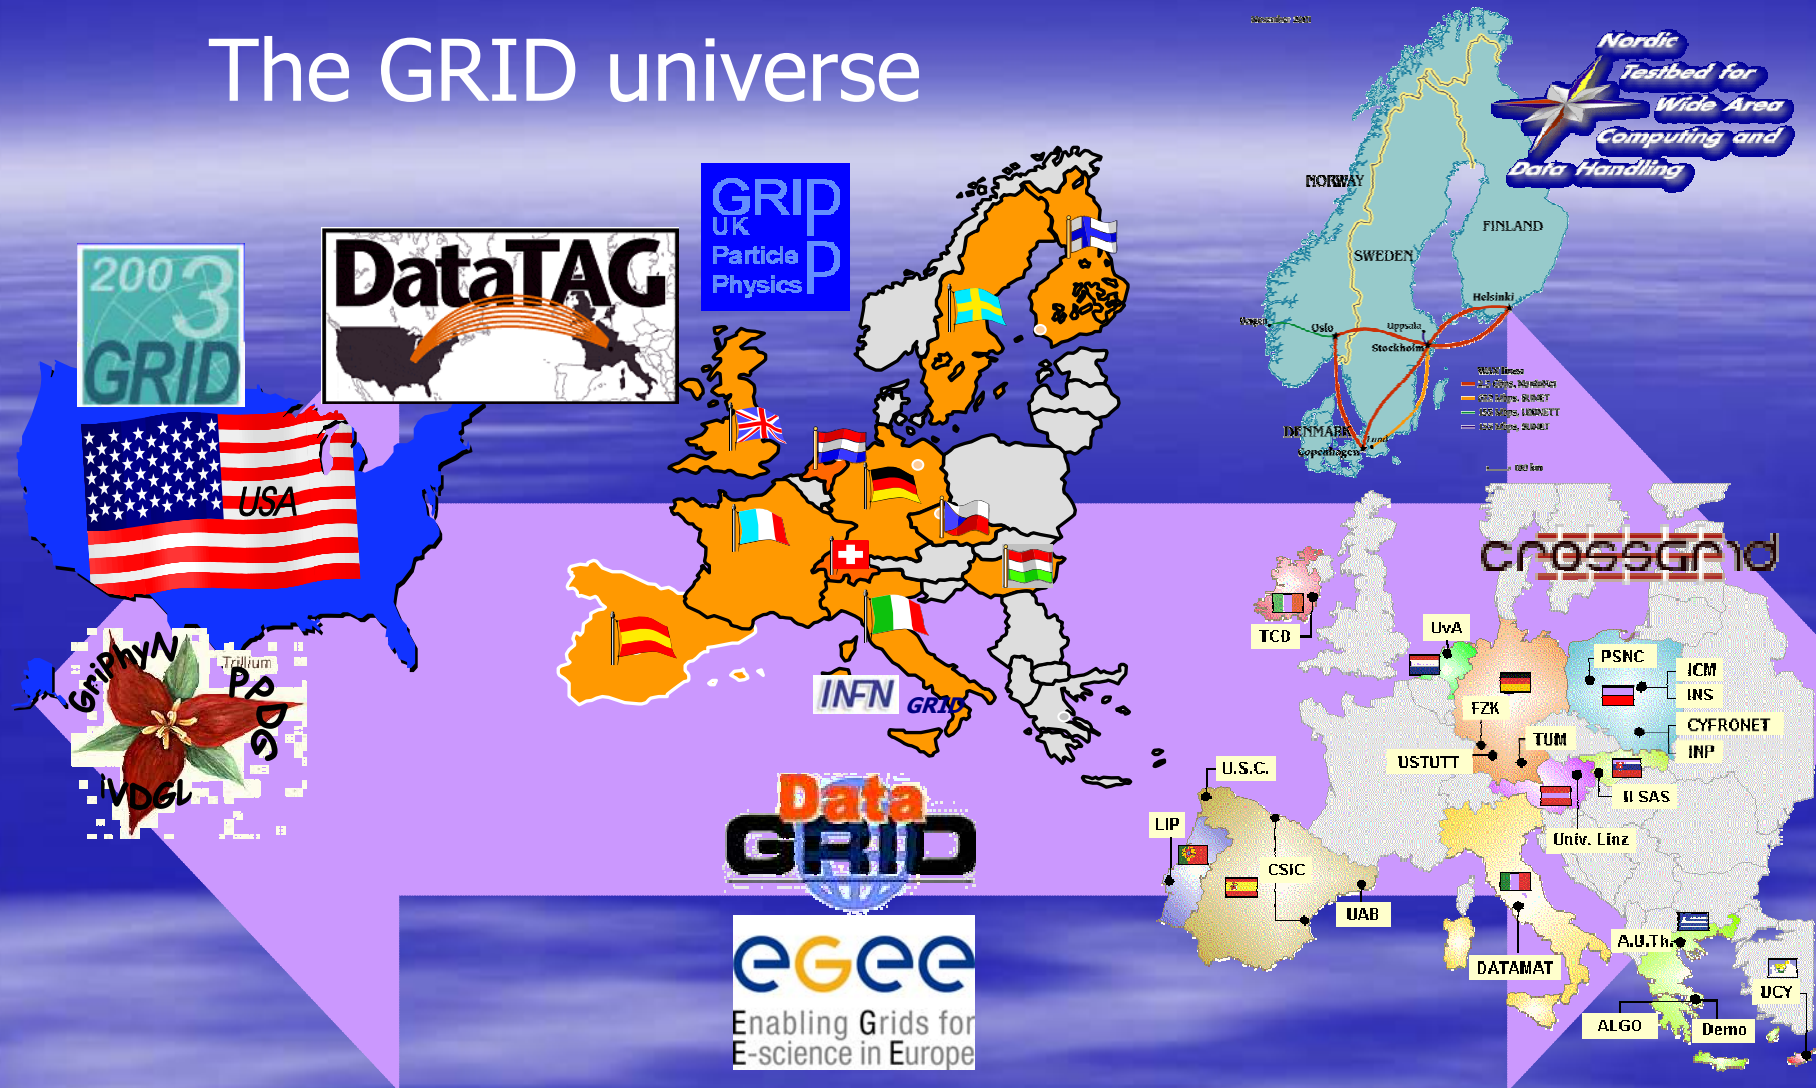
Motivation



Grids are everywhere...

- Europe is flooded by “production” Grid deployment projects
 - [...] more than 6,000 grids have been deployed worldwide (Sun)
 - If by deploying a scheduler on my local network I create a “Cluster Grid”, doesn't my NFS deployment over the same network provide me with a “Storage Grid?” [...] Is there any computer system that isn't a Grid? (Ian Foster)
- Tremendous richness of architectures and products
 - But worrying lack of stable testbeds where to experiment and provide feedback
 - At the moment only friendly and advanced users can use the system
 - Which of course creates a vicious circle...

The GRID universe



- Not the only kid on the block!
- How to avoid divergence!?

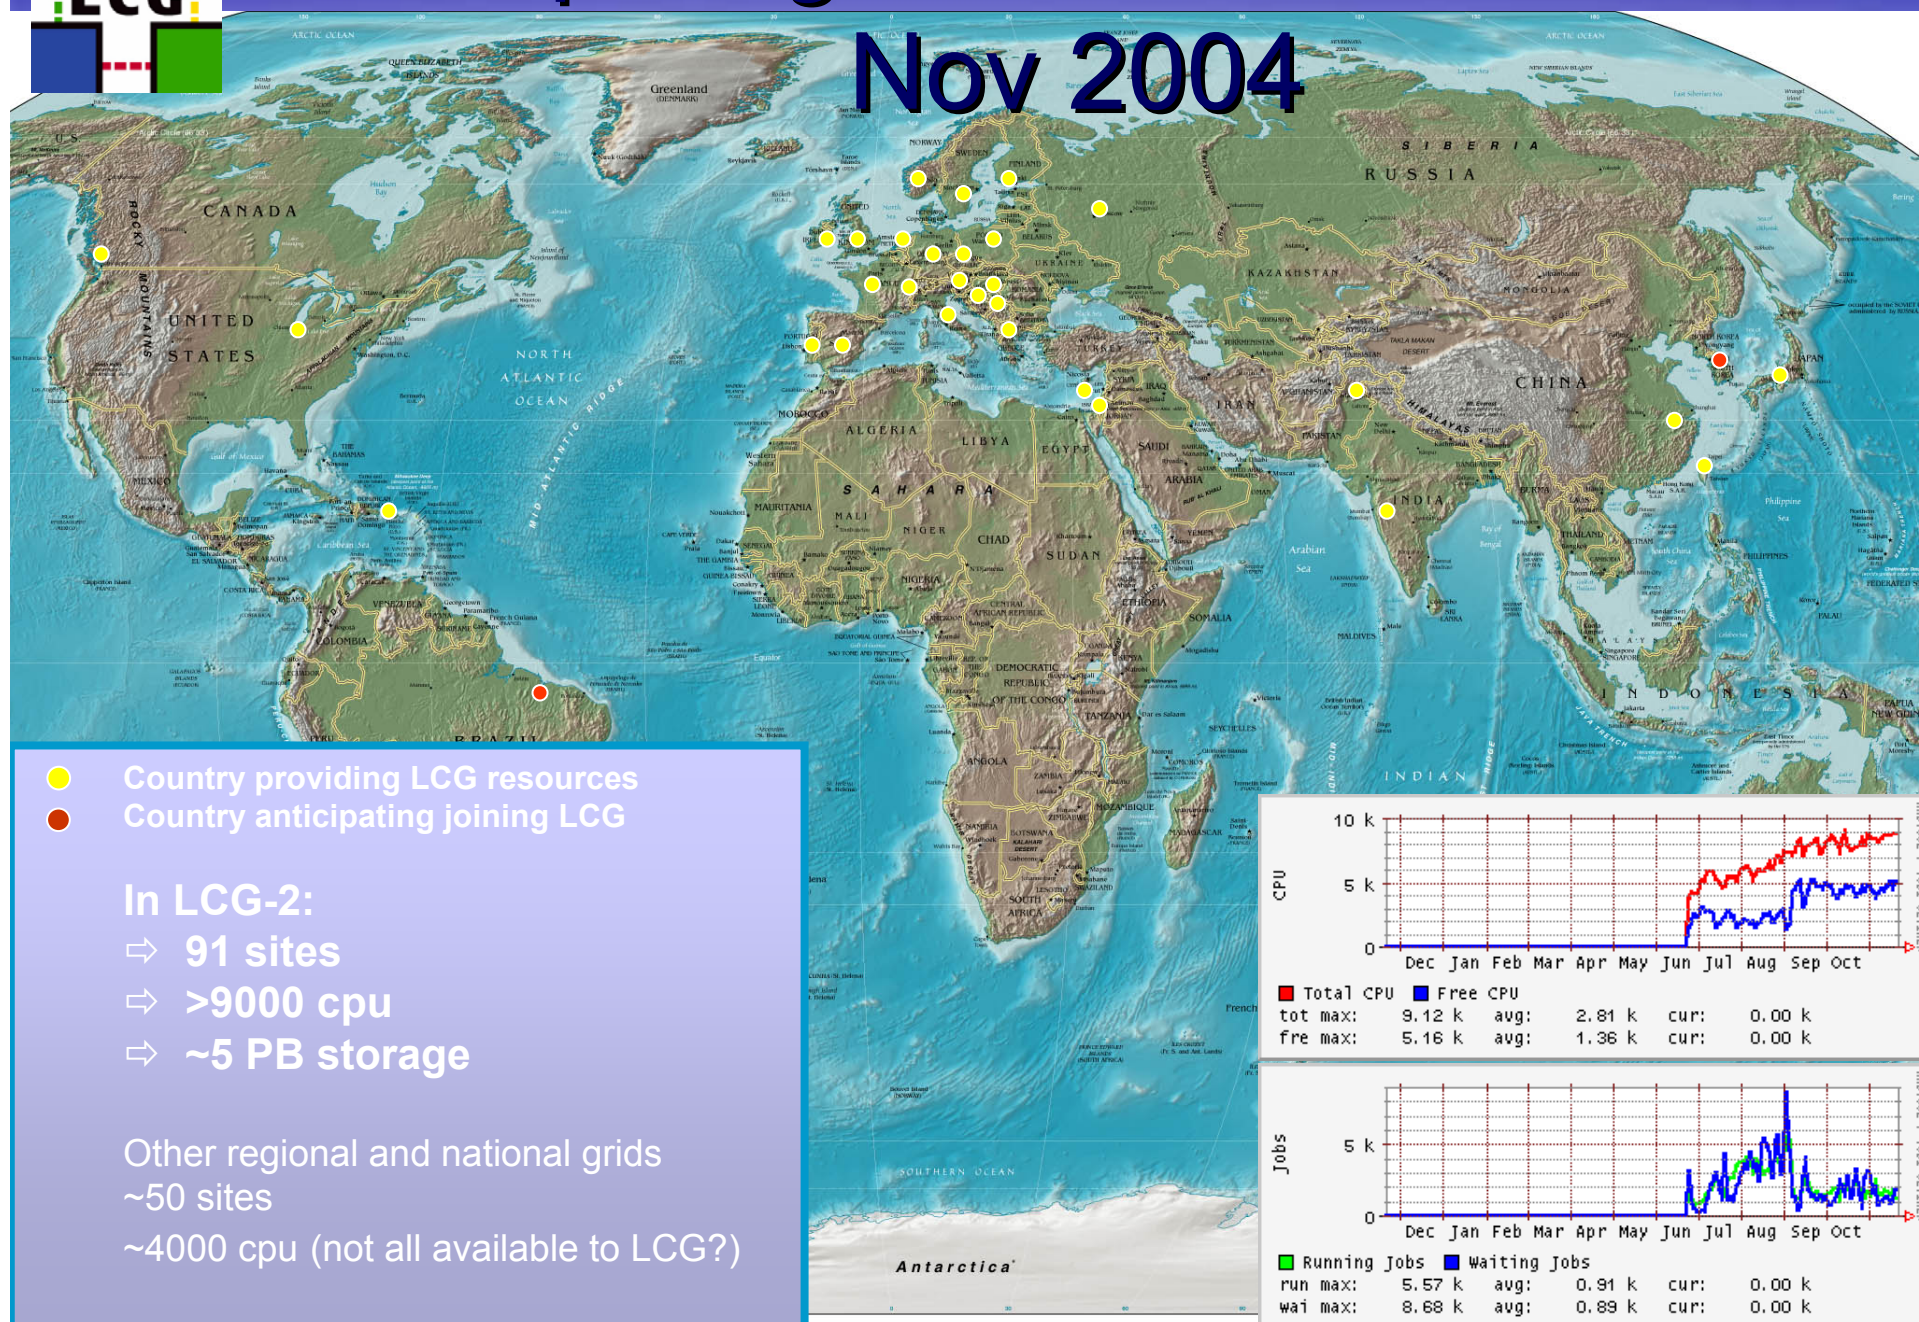
LCG Internal Review 2003

- Federated Grids
 - Currently: also LHC experiments use a number of different Grids
 - Sometimes multiple systems (Grids) are used even within a single experiment
 - Clear that different Grids will coexist (e.g. US Tier2, NorduGrid)
 - First priority should be to show that a single Grid can achieve real production quality.
 - Fortunately, this is the LCG



Computing Resources in LCG –

Nov 2004



Grid Deployment - concerns



- The *basic issues* of middleware reliability and scalability that we were struggling with a year ago have been overcome
BUT - there are many issues of functionality, usability
and performance to be resolved -- **soon**
- Overall job success rate 60-75%
 - Can be *tolerated* for “production” work – submitted by small teams with automatic job generation, bookkeeping systems
 - Unacceptable for end-user data analysis

ARDA and Startup

ARDA (April 2004)

- The ARDA project aims to help experiments prototype analysis systems using grid technology
 - Starting point: existing distributed system in the experiments
 - One prototype for each of the LHC experiments
 - Two people (funded by EGEE and LCG) work closely with each experiment
- Maintain a close relationship with the EGEE middleware team
 - experience with the early versions of the new middleware (gLite)
 - feedback to developers
- ARDA is NOT a middleware or applications development project

ARDA working group recommendations: our starting point

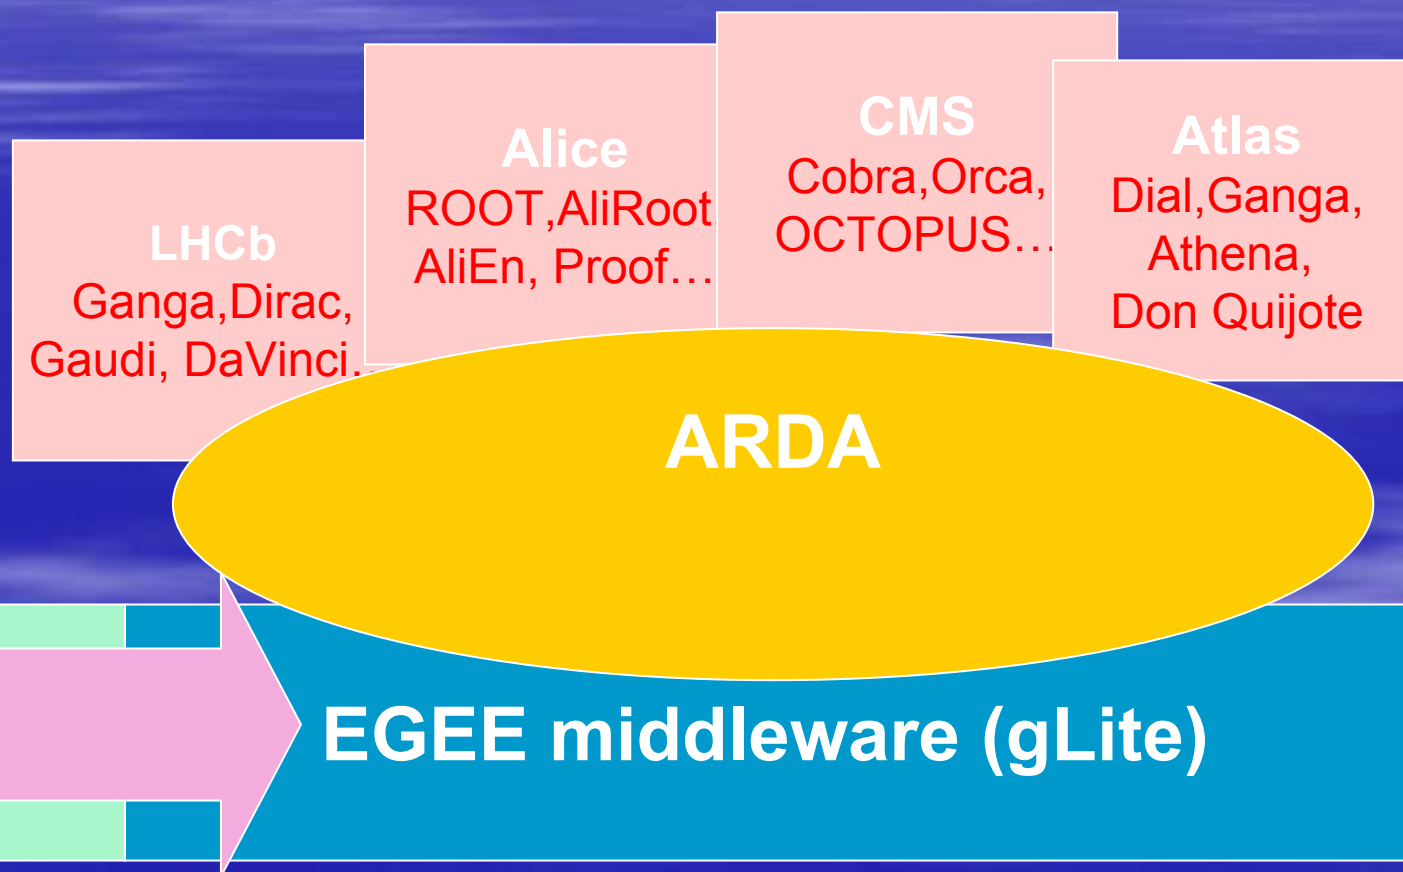
- New service decomposition
 - Strong influence of Alien system
 - the Grid system developed by the ALICE experiments and used by a wide scientific community (not only HEP)
- Role of experience, existing technology...
 - Web service framework

- Interfacing to existing middleware to enable their use in the experiment frameworks
- Early deployment of (a series of) prototypes to ensure functionality and coherence

ARDA and HEP experiments

"Long they laboured in the regions of Eä, which are vast beyond the thought of Elves and Men, until in the time appointed was made Arda..." - J.R.R Tolkien, *Valaquenta*

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



Architecture and Design

- Project started with staffing essentially complete on April 1st, 2004
- Architecture document released in June 2004
- Design Document released in August 2004
- Those documents have been also used by consortiums such as OSG to prepare their blueprint
 - And made available to GGF, GridLab, OMII, etc...

Development Testbed

- A Development Testbed (known as prototype) has been made available as of May 2004
 - To host **prototype middleware** as recommended by the ARDA RTAG
 - Many ideas from the ALICE/AlEn system
 - Started with AliEn, adding additional components from other middleware providers
 - Comprises resources at CERN, University of Wisconsin/Madison and INFN
 - Approximately 60 users registered
 - Being expanded with a second VO in Madison
 - **Will be further expanded as a result of the ARDA Workshop outcome in October 2004**
 - **Used by the ARDA Team to try out new middleware**
 - **Bio-Medical community has been invited to use the development testbed**

– But not very active so far

ARDA @ Regional Centres

- “Deployability” is a key factor of MW success
- On different time scales: gLite prototype and Pre Production Service
 - Understand “Deployability” issues
 - Quick feedback loop
 - Extend the test bed for ARDA users
- Stress and performance tests could be ideally located outside CERN...
 - Pilot sites might enlarge the resources available and give fundamental feedback in terms of “deployability” to complement the EGEE SA1 activity (EGEE/LCG operations; Pre Production Service)
- Running ARDA pilot installations
 - Experiment data available where the experiment prototype is deployed

ALICE & ARDA

- All ALICE grid developers are now hired by EGEE
- So AliEn is some kind of “frozen”
- Only solution: gLite

gLite – the middleware and it's components


Architecture Guiding Principles

- **Lightweight (existing) services**
- **Interoperability**
- **Resilience and Fault Tolerance**
- **Co-existence with deployed infrastructure**
 - Co-existence (and convergence) with LCG-2 and Grids 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**

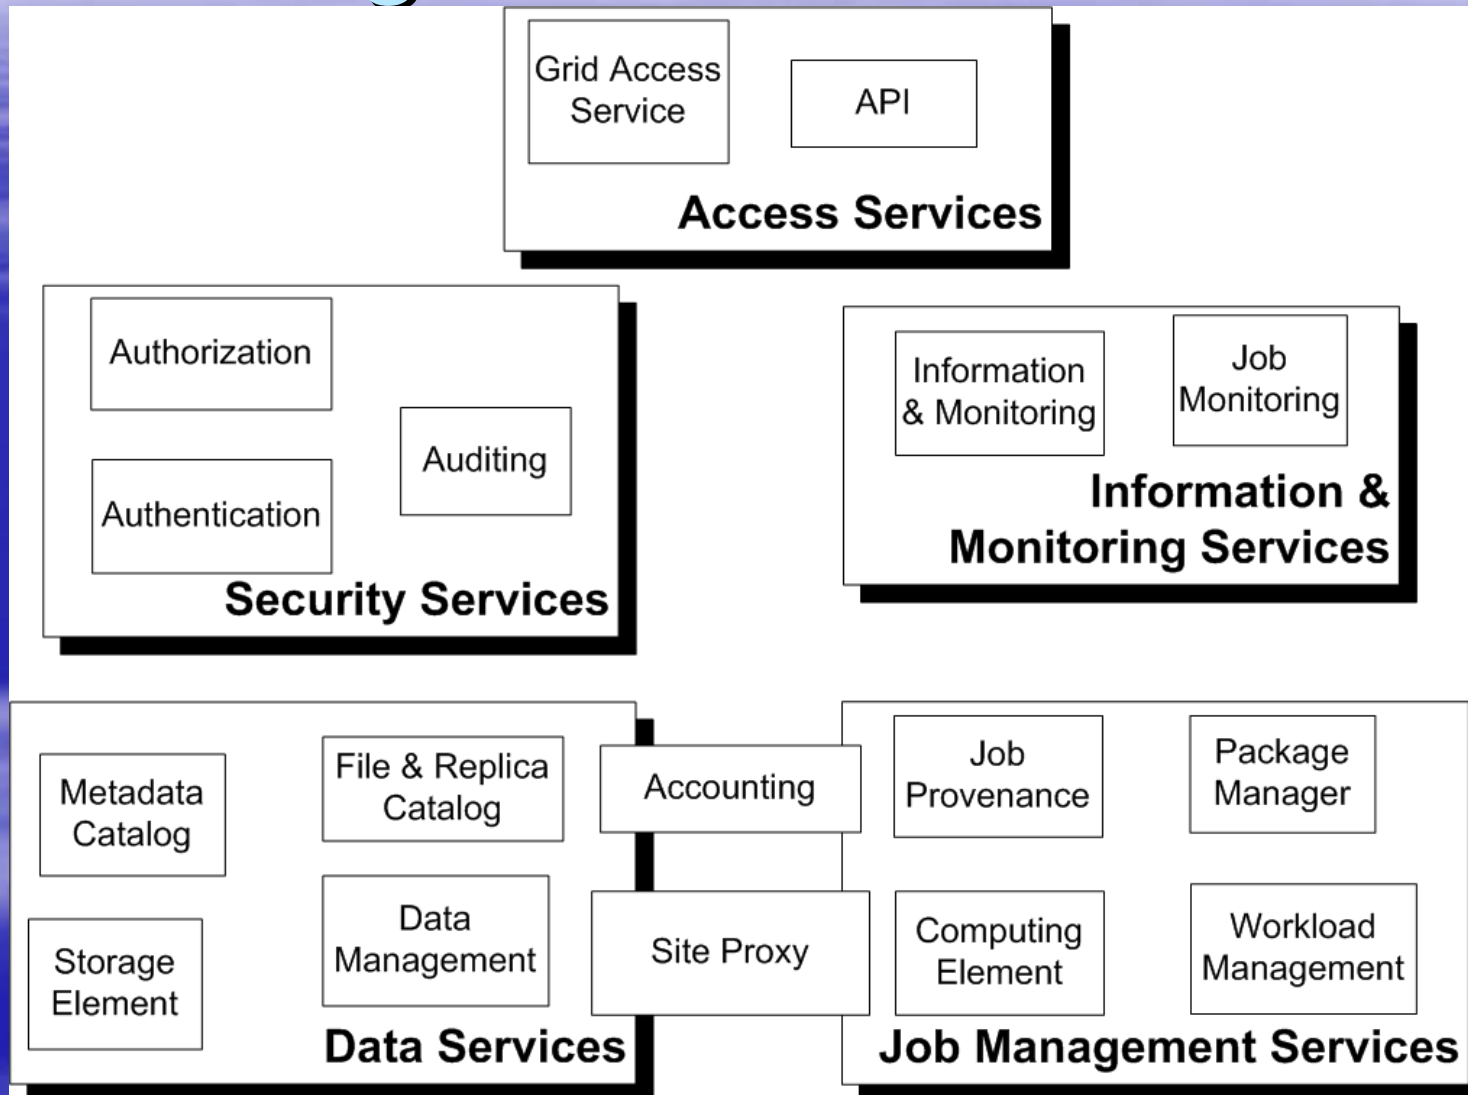


Approach

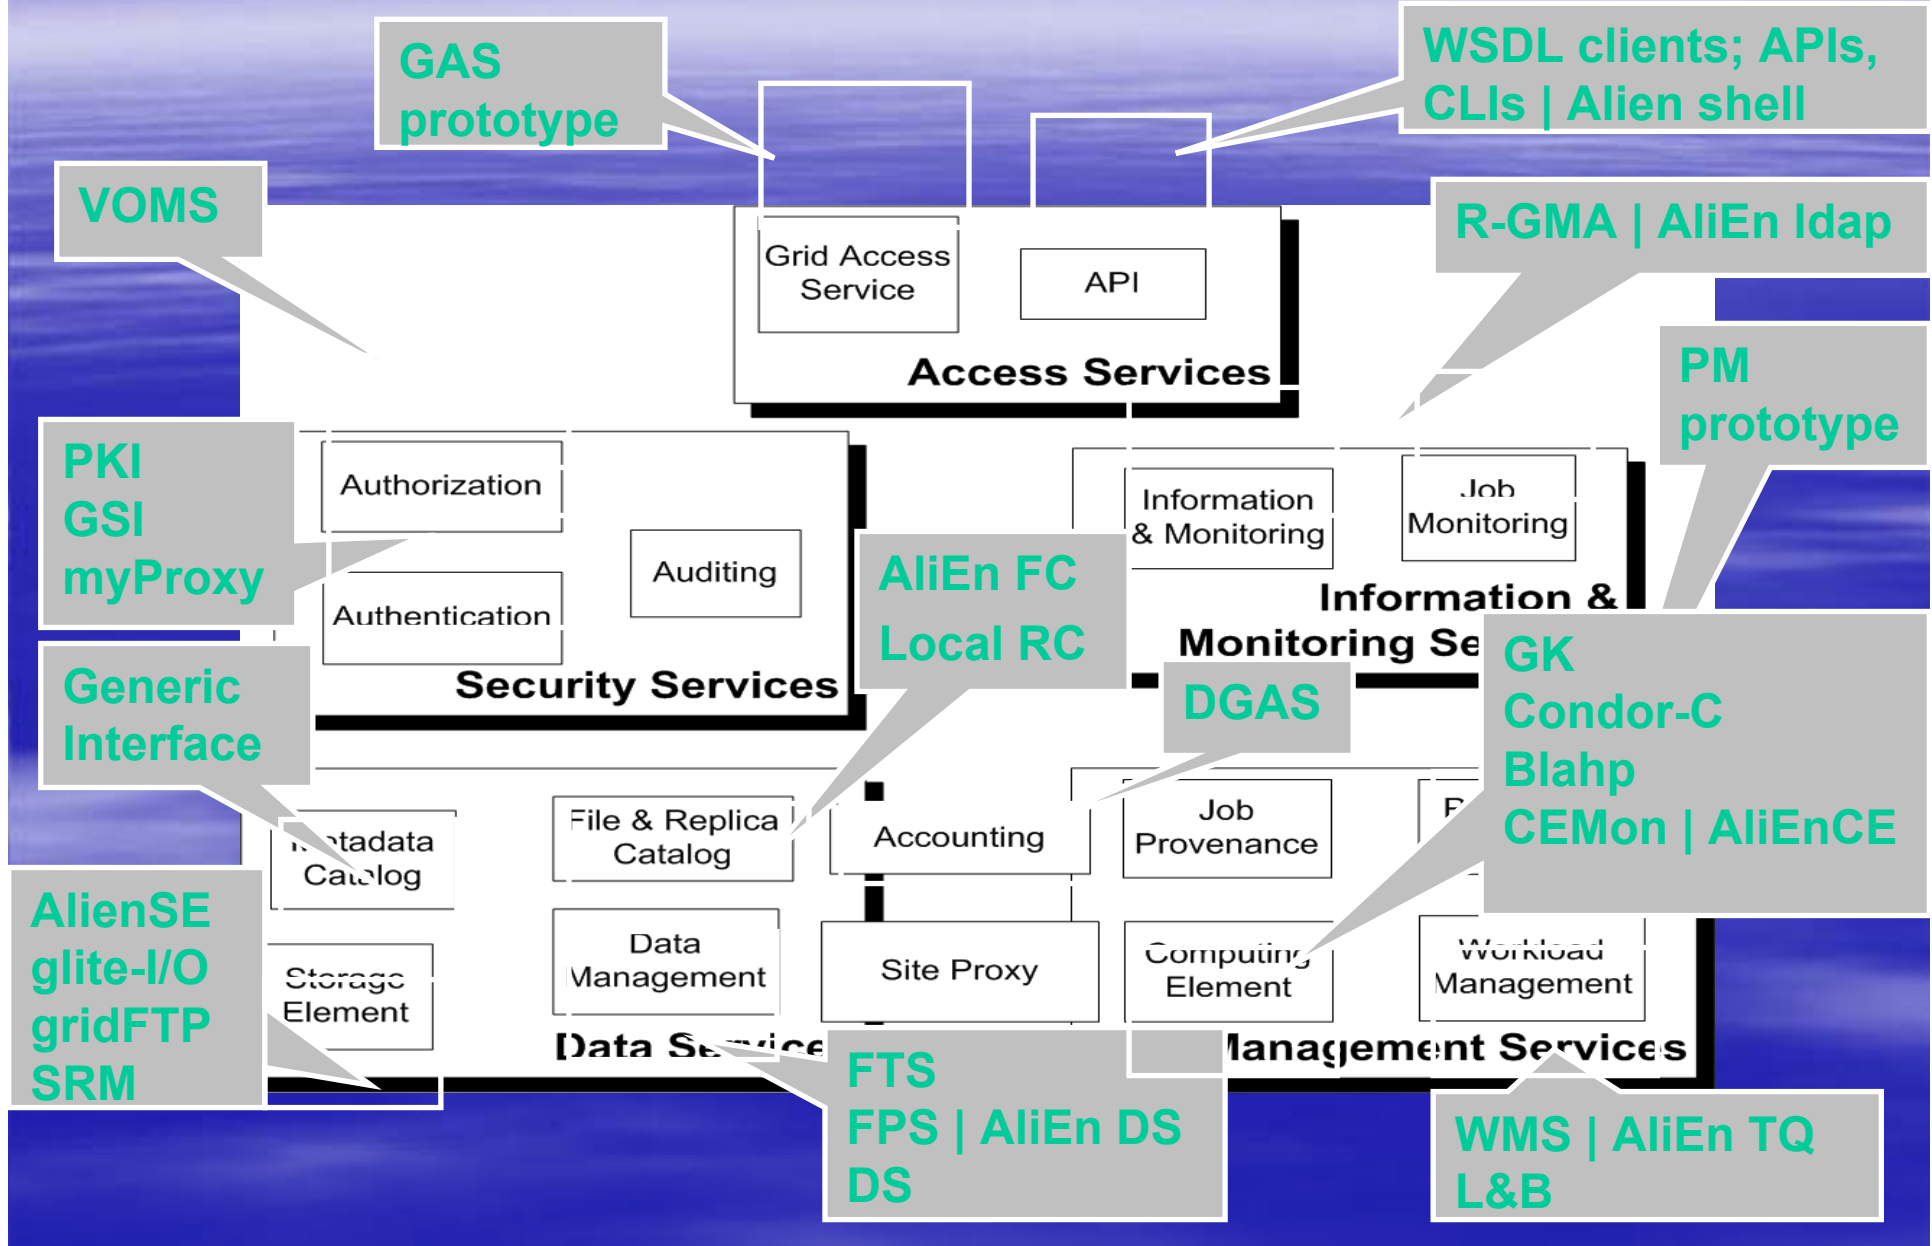
- Exploit experience and components from **existing projects**
 - AliEn, VDT, EDG, LCG, and others
- **Design team** works out architecture and design
 - Architecture: <https://edms.cern.ch/document/476451>
 - Design: <https://edms.cern.ch/document/487871/>
- Components are initially deployed on a **prototype infrastructure**
 - Small scale (CERN & Univ. Wisconsin)
 - Get user feedback on service semantics and interfaces
- After integration and testing components are delivered to SA1



gLite Services



Potential Services for RC1



■ Workload Management

- AliEn TaskQueue
- EDG WMS (plus new TaskQueue and Information Supermarket)
- EDG L&B
- ISM adaptor for CEMon is still missing!
- Query of FC is missing

Blue: deployed on
development
testbed

■ Element


- Globus Gatekeeper + LCAS/LCMAPS
 - **Dynamic accounts (from Globus)**
- CondorC
- Interfaces to LSF/PBS (blahp)
- “Pull components”
 - AliEn CE
 - **gLite CEmon (being configured)**

Red: proposed

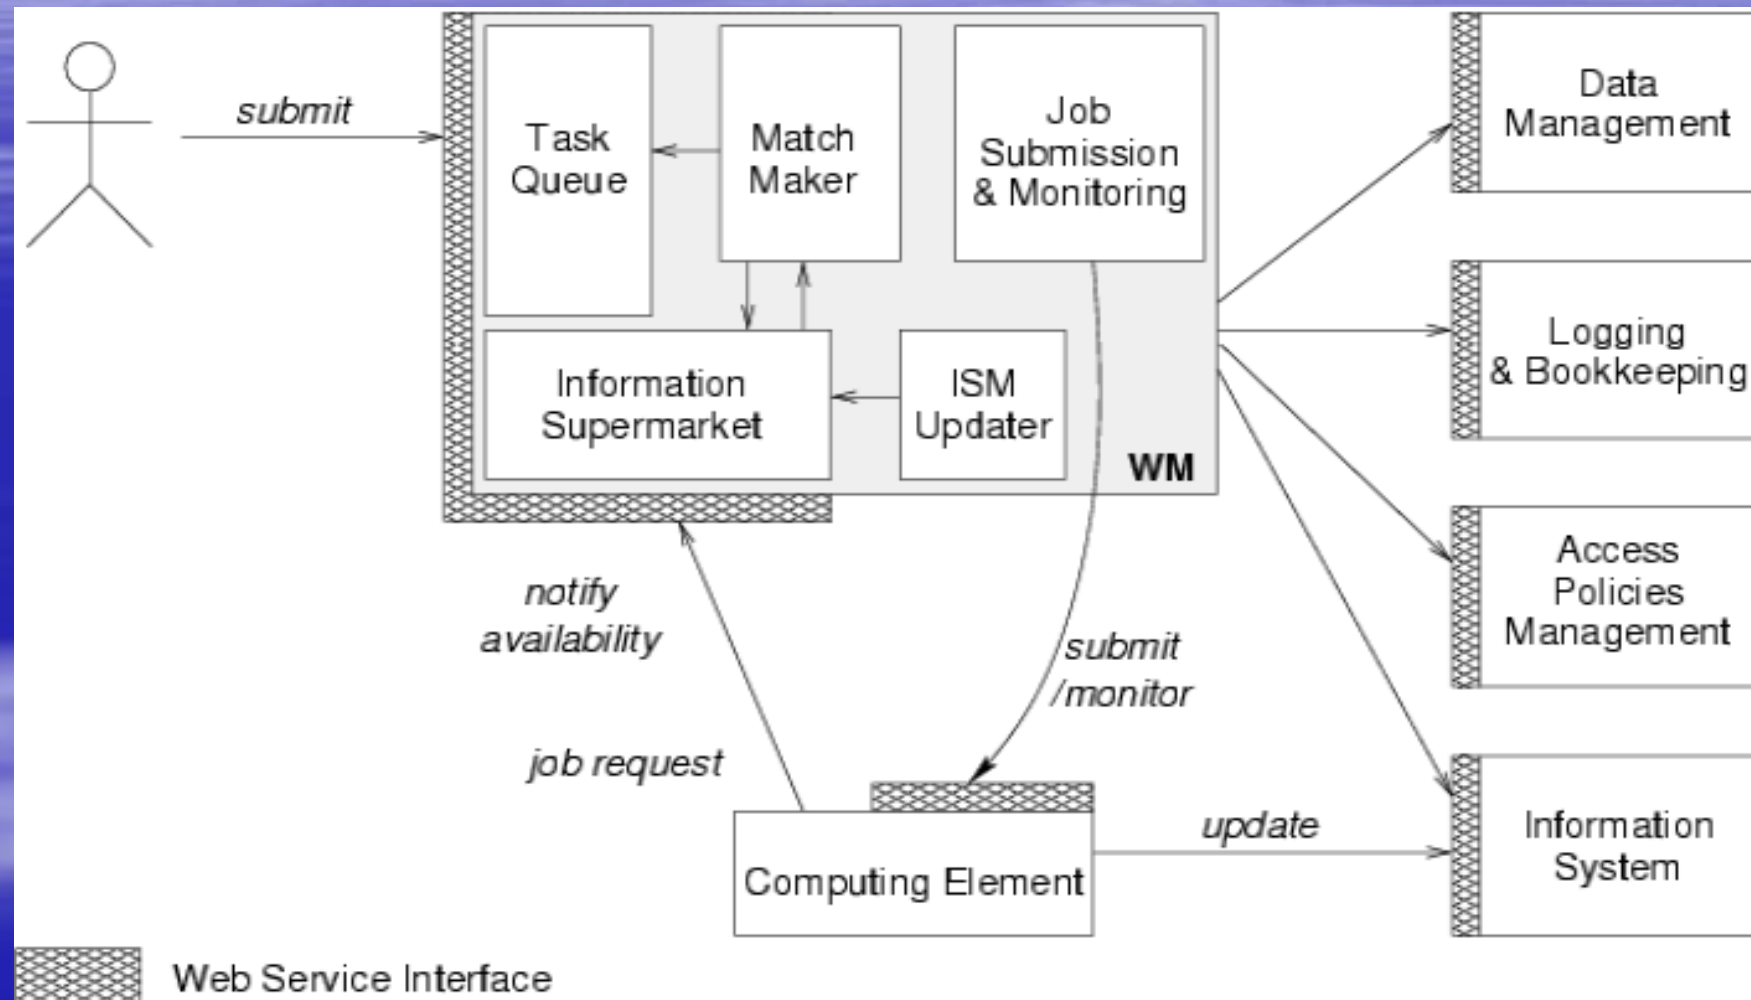
Middleware – RC 1

Accessing gLite

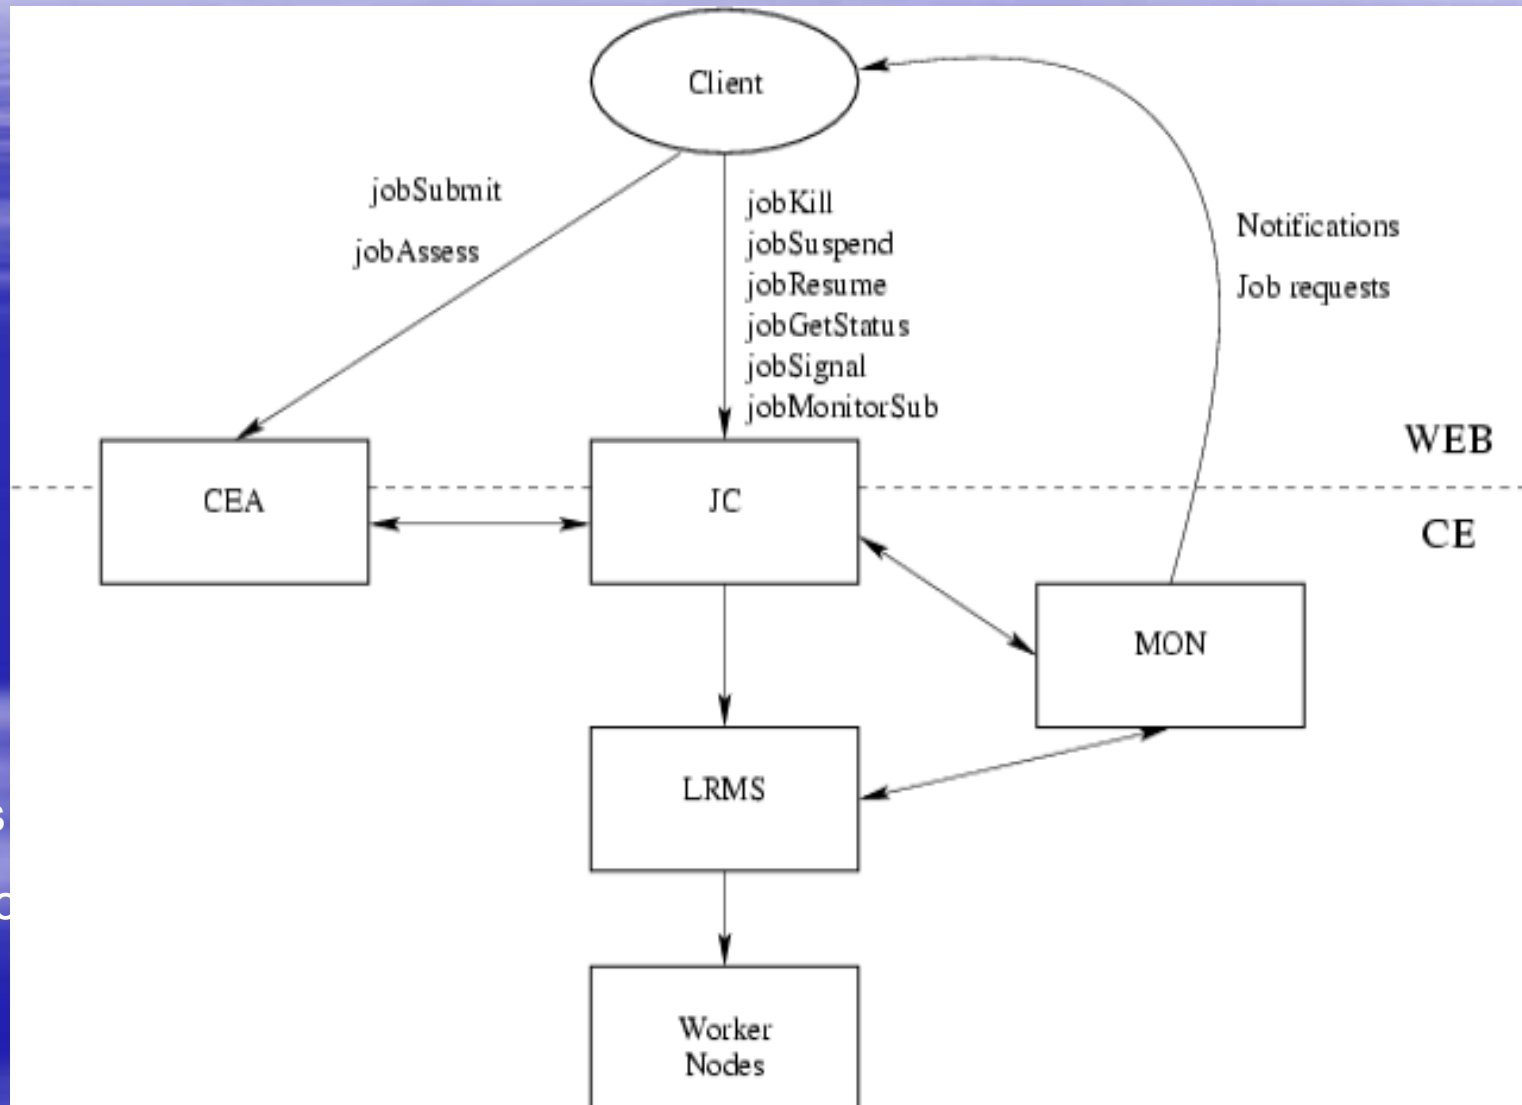
- **Access through gLite shell**
 - User-friendly Shell implemented in Perl
 - Shell provides a set of Unix-like commands and a set of gLite specific commands
- **Perl API**
 - no API to compile against, but Perl-API sufficient for tests, though it is poorly documented (ALICE C++ API with ROOT interface finished)



WMS



CE



- Works
- Site po

- Storage Element

- Existing SRM implementations
 - dCache, Castor, ...
 - **FNAL & LCG DPM**
- gLite-I/O (re-factored AliEn-I/O)

- Catalogs

- AliEn FileCatalog – global catalog
- gLite Replica Catalog – local catalog (Oracle and MySQL)
- **Catalog update (messaging)**
- FiReMan Interface
- **RLS (globus)**

- Data Scheduling

- File Transfer Service (Stork+GridFTP)
- File Placement Service
- gLite I/O
- **Data Scheduler**

- Metadata Catalog

- Simple interface defined (AliEn+BioMed)

- Information & Monitoring

- R-GMA **web service version;**
multi-VO support

Prototype Middleware Status & Plans (II)

Prototype Middleware

- Security
 - VOMS as Attribute Authority and VO mgmt
(show stopper for security)
 - myProxy as proxy store
 - GSI security and VOMS attributes as enforcement
 - fine-grained authorization (e.g. ACLs)
 - globus to provide a set-uid service on CE
- Accounting
 - EDG DGAS (not used yet)
- User Interface
 - AliEn shell
 - CLIs and APIs
 - GAS
 - Catalogs
 - Integrate remaining services
- Package manager
 - Prototype based on AliEn backend
 - evolve to final architecture agreed with ARDA team

Advance Reservation

- Propose to include a 'reservation manager' into WMS for simple and compound agreements
 - Should this reservation mgr become a separate service so it is useful for others (like data mgmt) as well?
- How are the different services acting on a resource represented? Hierarchical?
- Only SRM2 provides space reservation –

VOMS

- Long time to install EGEE version of VOMS server
 - Bug fix integration
 - Manpower problems
 - Scheduled for this week
- VOMS admin still changing rapidly
 - Need to have a more stable version now
- AC format in VOMS core not REC compliant

DM & WMS

- Support for logical data sets (LDS) and generic queries
- WMS will have to deal with multiple catalogs
 - Propose to use a prefix indicating the catalog interface to use for locating the data
 - Other possibility is to nest classads and specify the endpoint
- JDL of WMS and AliEn TQ should (at least for a subset) converge
- DLI interface can use storage index

Some words and numbers about gLite I/O

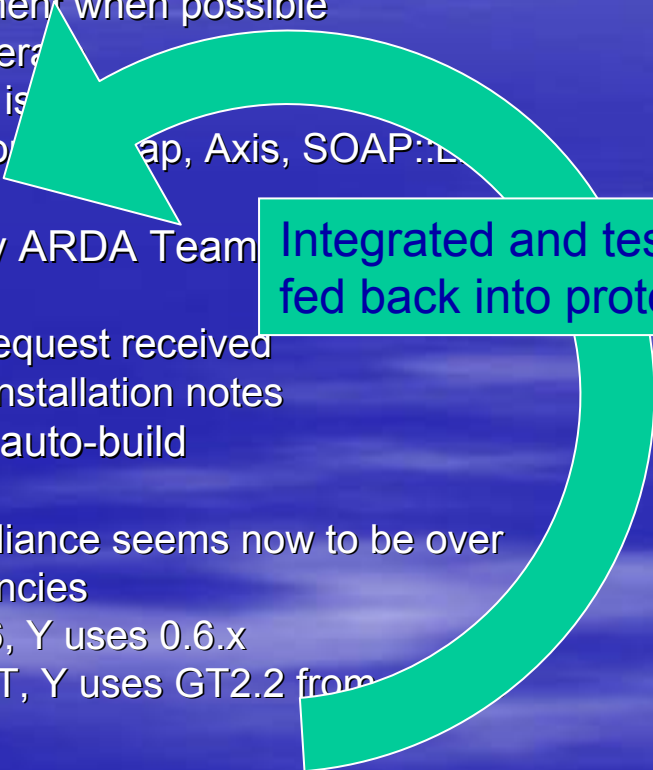
- gLite I/O is not working with the current prototype yet.
- Numbers: using aiiod or xrootd reading from disk server at CNAF (1 Gb) with 100 MB/s
- Castor aiiod-server: 40 Mb/s (problems here is Castor backend, staging, ...)
- CERN Castor SE configured for 150 parallel downloads with 3-8 MB/s each.

Some words and numbers about gLite I/O 2

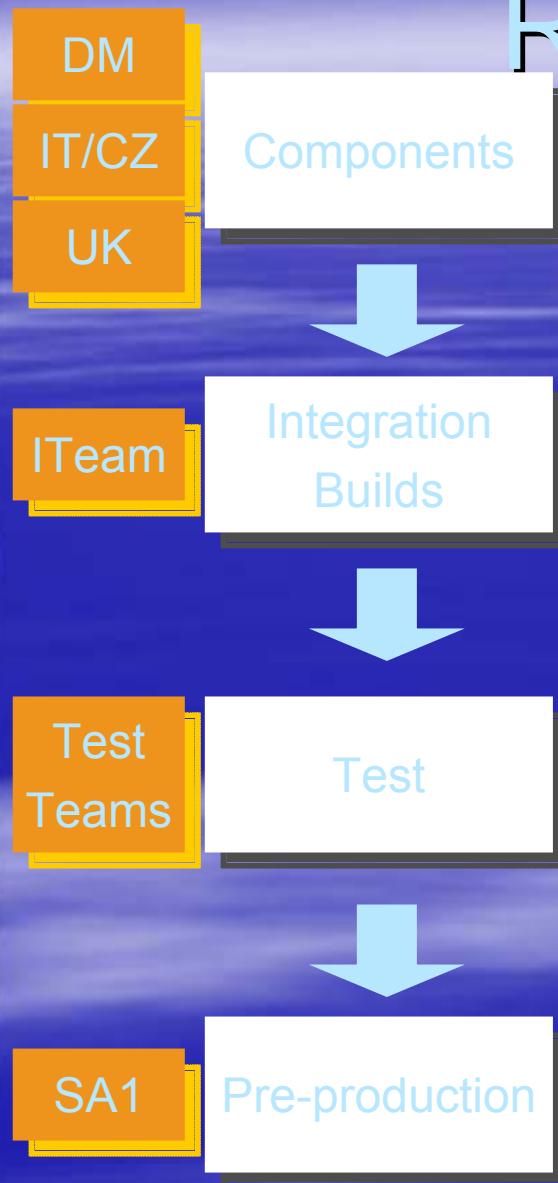
- In ALICE DC Phase 2 for each event 500 MB downloaded from CERN
- Test at CERN using aiode and single tcp connection: 64 MB/s from disk to disk.
- But analysis (gLite was thought to server mainly this purpose) does not need a lot of data transfer anyway (bring the KB to the PB, and not vice versa)

Development Cycle and Release Plans

Development, Integration, Testing processes

- Design discussed, agreed and implemented
 - Developers introduce their components in prototype
 - Interfaces (WDSL usually) published and made available
 - According to Design Document when possible
 - Components modified to interoperate
 - Tricky language dependent issues
 - Immaturity of code generators (e.g. Jax-WS, Axis, SOAP::Lite)
 - Code checked-in in gLite CVS
 - Prototype updated for direct use by ARDA Team
 - Happened 7 times since May
 - Bug reports and enhancement request received
 - Developers required to provide installation notes
 - Code taken over by Integration for auto-build
 - Feedback cycle with developers
 - Early difficulties with SCM compliance seems now to be over
 - Usually difficulties with dependencies
 - E.g. X uses classads 0.9.6, Y uses 0.6.x
 - E.g. X uses GT2.4 from VDT, Y uses GT2.2 from VDT
 - Streamline configuration
 - Contributes to documentation
 - Code taken over by Testing
 - Initial Step is to install Services
 - Usually (especially at the beginning) not suitable docs (installation & configuration)
 - Significant time spent iterating with developers
- 
- In Development**
- In Integration**
- In Testing**
- Available**
- Integrated and tested code fed back into prototype

Release Process



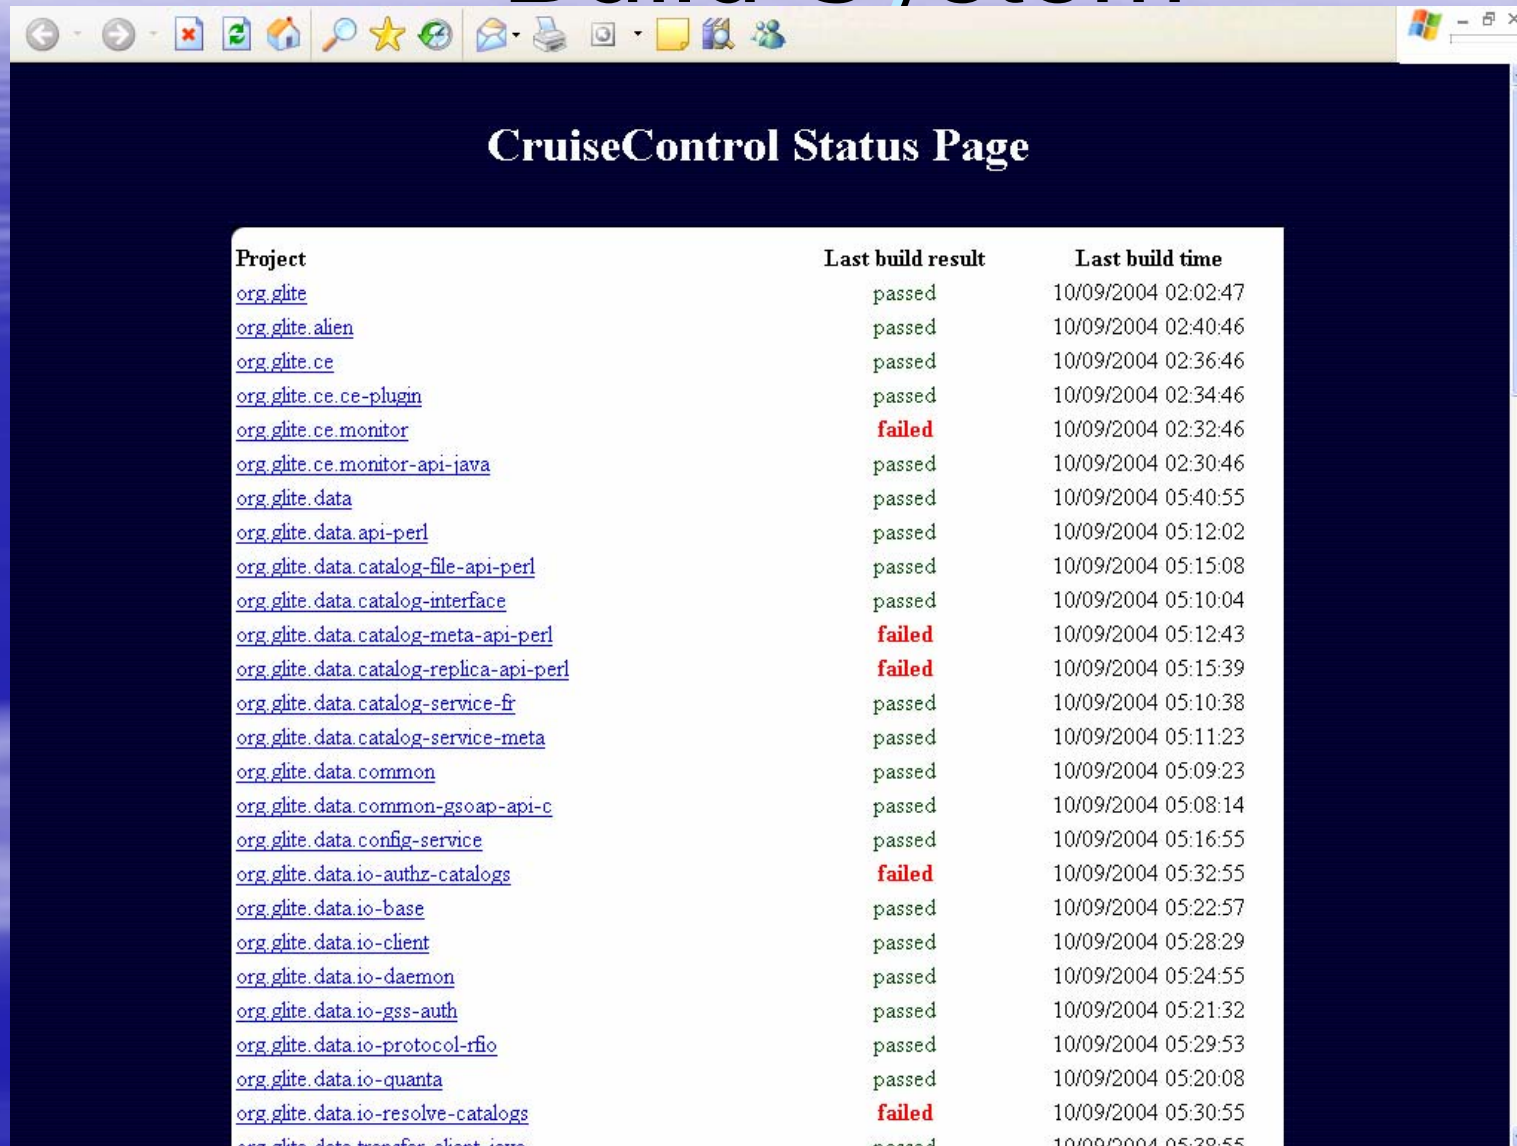
Developers announce every Friday what components are available for a release according to the development plan. Components are tagged, list is sent to Integration Team

ITeam put together all components, verifies consistency and dependencies and add/update the service deployment modules (installation and configuration scripts). The build is tagged as lyyyyMMdd

The integrated build is deployed in the testbeds and validates with functional and regression tests, test suites are updated, packaged and added to the build.

If the build is suitable for release, release notes and installation guides are updated, the build is retagged (RCx, v.0.x.0) and published on the gLite web site for release to SA1

Build System



Project	Last build result	Last build time
org.glite	passed	10/09/2004 02:02:47
org.glite.alien	passed	10/09/2004 02:40:46
org.glite.ce	passed	10/09/2004 02:36:46
org.glite.ce.ce-plugin	passed	10/09/2004 02:34:46
org.glite.ce.monitor	failed	10/09/2004 02:32:46
org.glite.ce.monitor-api-java	passed	10/09/2004 02:30:46
org.glite.data	passed	10/09/2004 05:40:55
org.glite.data.api-perl	passed	10/09/2004 05:12:02
org.glite.data.catalog-file-api-perl	passed	10/09/2004 05:15:08
org.glite.data.catalog-interface	passed	10/09/2004 05:10:04
org.glite.data.catalog-meta-api-perl	failed	10/09/2004 05:12:43
org.glite.data.catalog-replica-api-perl	failed	10/09/2004 05:15:39
org.glite.data.catalog-service-fr	passed	10/09/2004 05:10:38
org.glite.data.catalog-service-meta	passed	10/09/2004 05:11:23
org.glite.data.common	passed	10/09/2004 05:09:23
org.glite.data.common-gsoap-api-c	passed	10/09/2004 05:08:14
org.glite.data.config-service	passed	10/09/2004 05:16:55
org.glite.data.io-authz-catalogs	failed	10/09/2004 05:32:55
org.glite.data.io-base	passed	10/09/2004 05:22:57
org.glite.data.io-client	passed	10/09/2004 05:28:29
org.glite.data.io-daemon	passed	10/09/2004 05:24:55
org.glite.data.io-gss-auth	passed	10/09/2004 05:21:32
org.glite.data.io-protocol-rfio	passed	10/09/2004 05:29:53
org.glite.data.io-quanta	passed	10/09/2004 05:20:08
org.glite.data.io-resolve-catalogs	failed	10/09/2004 05:30:55
org.glite.data.transfer-client-java	passed	10/09/2004 05:29:55

Installation testing reports

Click on each RPM name to see its log file.


RPM Name	Files in RPM	Internal Dep Specified	Libs in \$LOCATION/lib	Bins in \$LOCATION/bin	Correct info in header	Doc Present	Licence Present	RPM Installs	RPM Erases	RPM Relocatable	Outcome
glite-data-catalog-interface-0.1.0-1.noarch.rpm	Files	Yes	N/A	N/A	Yes	Yes	Yes	Yes	Yes	Yes	Pass
glite-data-catalog-service-meta-0.1.0-1.noarch.rpm	Files	Yes	N/A	N/A	Yes	Yes	Yes	Yes	Yes	Yes	Pass
glite-data-common-0.0.1-4.noarch.rpm	Files	Yes	N/A	N/A	Yes	Yes	Yes	Yes	Yes	Yes	Pass
glite-data-config-service-0.1.0-1.i386.rpm	Files	Yes	Yes	Yes	Yes	No	Yes	Yes	Yes	Yes	Fail
glite-data-io-protocol-rfcio-0.1.1-1.i386.rpm	Files	Yes	Yes	N/A	Yes	No	Yes	Yes	Yes	Yes	Fail
glite-data-io-quanta-0.1.1-1.i386.rpm	Files	Yes	Yes	N/A	Yes	No	Yes	Yes	Yes	Yes	Fail
glite-data-io-gss-auth-0.1.1-1.i386.rpm	Files	Yes	Yes	Yes	Yes	No	Yes	Yes	Yes	Yes	Fail
glite-data-io-base-0.2.0-1.i386.rpm	Files	Yes	Yes	N/A	Yes	No	Yes	Yes	Yes	Yes	Fail
glite-data-io-client-0.2.0-1.i386.rpm	Files	Yes	Yes	Yes	Yes	No	Yes	Yes	Yes	Yes	Fail
glite-data-io-daemon-0.2.0-1.i386.rpm	Files	Yes	Yes	Yes	Yes	No	Yes	Yes	Yes	Yes	Fail
glite-data-io-authz-catalogs-	Files	---	NOT	BUILT	---	---	---	---	---	---	Fail
glite-data-io-resolve-catalogs-	Files	---	NOT	BUILT	---	---	---	---	---	---	Fail
glite-data-catalog-service-fr-0.1.0-1.noarch.rpm	Files	Yes	N/A	N/A	Yes	Yes	Yes	Yes	Yes	Yes	Pass
glite-data-transfer-interface-0.1.0-1.noarch.rpm	Files	Yes	N/A	N/A	Yes	Yes	Yes	Yes	Yes	Yes	Pass
glite-data-transfer-client-iava-0.0.0-0.i386.rpm	Files	Yes	N/A	N/A	Warning	No	Yes	Yes	Yes	Yes	Fail
glite-data-transfer-service-0.1.0-1.noarch.rpm	Files	Yes	N/A	N/A	Yes	Yes	Yes	Yes	Yes	Yes	Pass
glite-dcas-ra-client-0.0.0-0.i386.rpm	Files	Yes	Yes	N/A	Warning	No	No	No	Yes	No	Fail
glite-lb-client-interface-0.0.0-0.i386.rpm	Files	Yes	N/A	N/A	Warning	No	No	Yes	Yes	Yes	Fail
glite-lb-common-	Files	---	NOT	BUILT	---	---	---	---	---	---	Fail

WBS	EGEE	Comment	Name	Date MS	MS Done
1.5		—	Middleware		
1.3.4			LCG - EGEE Coordination		
1.3.4.1			EGEE senior management appointed	15-07-03	15-07-03
1.3.4.2			Technical design team established	01-09-03	04-12-03
1.3.4.8		i_jan04	EGEE Middleware people hired	29-02-04	23-03-04
1.3.4.9		i_jan04	EGEE Middleware execution plan available	29-02-04	17-03-04
1.3.4.10		i_jan04	EGEE Contract signed	01-04-04	01-04-04
1.5.2.6	None	i_apr04	First version of prototype available for experiments	16-06-04	16-05-04
1.5.2.7	MJRA1.1	i_apr04	Development and integration tools deployed	30-06-04	30-06-04
1.5.2.8	MJRA1.2	i_apr04	Software cluster development & testing infrastructure in place	30-06-04	30-06-04
1.5.2.9	DJRA1.1	i_apr04	Architecture & Planning Document for release candidate 1	30-06-04	30-06-04
1.5.2.10	MJRA1.3	i_apr04	Integration & Testing infrastructure in place; test plan	31-08-04	31-08-04
1.5.2.11	DJRA1.2	i_apr04	Grid Services design document for release candidate 1	31-08-04	31-08-04
1.5.2.12	None	i_apr04	Second version of prototype available for experiments	31-08-04	31-08-04
1.5.2.13	MJRA1.4	i_apr04	Release Candidate 1	31-12-04	Spring 2005 !?

gLite and LCG-2

Deployment and services

- Current LCG-2 based service continues as production service for batch work
 - Experiments moving to continuous MC production mode
 - Together with work in-hand provides a well-understood baseline service
- Deploy in parallel a pre-production service
 - Deploy LCG-2 components, and
 - gLite components as they are delivered
 - Understand how to migrate from LCG-2 to gLite –
 - Which components can be replaced
 - Which can run in parallel
 - Do new components satisfy requirements – functional and management/deployment
 - Move proven components into the production system
- LCG-2 is also the fallback in case gLite fails




gLite and LCG-2

LCG-2

focus on production,
large-scale data handling

- The service for the 2004 data challenges
- Provides experience on operating and managing a global grid service
- Development programme driven by data challenge experience
 - Data handling
 - Strengthening the infrastructure
 - Operation, VO management
- Evolves to LCG-3 as components progressively replaced with new middleware



gLite

focus on analysis

- Developed by EGEE project in collaboration with VDT (US)
- LHC applications and users closely involved in prototyping & development (ARDA project)
- Short development cycles
- Co-existence with LCG-2
- Profit as far as possible from LCG-2 infrastructure, experience
- Ease deployment – avoid separate hardware
- As far as possible - completed components integrated in LCG-2
- improved testing, easier displacement of LCG-2

Discussion summary – 2

- Coexistence/Migration issues:
 - Workload management – suggested strategy
 - Start with LCG-2
 - Add new (Glite) broker nodes to the LCG-2 based Grid Infrastructure
 - LCG-2 UIs can talk to LCG-2 brokers and to Glite brokers
 - When happy with the Glite broker, update the UIs
 - LCG-2 (Globus based) and Condor based CE can coexist in the same Grid
 - LCG-2 (Globus based) CEs can be used by LCG-2 broker and by Glite broker
 - Condor based CEs can be used only by the Glite Broker
 - So update the CEs (LCG-2 → Condor based CE) when you think it is right

Some Answers

- How can a user
 - find data in either the LCG-2 catalogs or the gLite catalogs using both the LCG-2 and gLite tools?
 - **Explicit mechanism – use both tools**
 - **Or gLite client if LCG-2 catalog exposes a Fireman interface**
 - access data produced already and stored in LCG-2 SE using gLite tools?
 - **Easy if LCG catalogs have a Fireman interface**
 - **Migration is also possible – need to clarify exact semantics**
- How is data
 - on LCG-2 Classic SEs available to gLite clients?
 - **Set up a glite IO server**
 - **Migrate data**
 - on gLite SEs available to LCG-2 clients ?
 - on gLite SEs available to local clients, not using the gLite I/O mechanism?
 - **Update LCG-2 clients**
 - **Without change to LCG-2 clients: Only on SRMs having ACL capability.**

Some Answers

- How can the WMS
 - Do matchmaking against both gLite and LCG-2 catalogs, if the appropriate SEs are available?
 - **LCG catalogs may be also connected into the StorageIndex**
 - **Use the StorageIndex interface**
 - Find available catalogs for LCG-2 and gLite data?
 - **Again easy if LCG catalogs have a Fireman interface and are hooked up through the StorageIndex.**
 - **If not, this is a service discovery issue.**
- Data:
 - Currently :
 - “Classic SE”
 - SRM – dCache, LCG DPM, etc
 - **Advantage :**
 - **Data accessible for gLite**
 - **‘Metadata-only’ migration : no need to move/copy files to go to gLite**

Relationship LCG and EGEE

LCG Project Activity Areas



Applications

Development environment and common libraries, frameworks, tools for the LHC experiments

CERN Fabric

Construction and operation of the central LHC computing facility at CERN

Networking

Planning the availability of the high bandwidth network services to interconnect the major computing centres used for LHC data analysis

Grid Deployment

Coordinate the development, management and operation of the Grids interconnecting resources available for LHC data analysis
Joint activity with EGEE SA1



Middleware

Provision of middleware adapted to LHC requirements – functionality, performance, scale. Identical with EGEE JRA1.



Distributed Analysis

Prototyping LHC experiment distributed analysis systems using a Grid – Component of EGEE
NA4

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

Technical requirements of gLite and Installation

- One frontend machine which hosts the gLite services of the site (CE, SE, gatekeeper/ClusterMonitor)
- This machine has to have access to the local batch system. The installation has to be visible from the WNs (shared HOME)
- For own VO: one machine hosting all server-side services as well as LDAP, MySQL

Technical requirements of gLite and Installation 2

- Installation manual:
<http://people.web.psi.ch/feichtinger/doc/glite-alien-setup.html>
by Derek Feichtinger LCG/ARDA group
- Or download from <http://alien.cern.ch/dist>
AliEn-Base/Client etc. version 1.36-25 or larger
and add current code from CVS
[:pserver:anonymous@jra1mw.cern.ch:/cvs/jra1mw](http://pserver:anonymous@jra1mw.cern.ch:/cvs/jra1mw)
(everything starting with “org.glite.alien” but there
are more than 100 directories and things keep
changing quickly)

Configuration (LDAP) and Firewall

The screenshot shows the GQ configuration window with the following components:

- File Menu:** Search, Browse, Schema, Help
- Left Panel (Tree View):**
 - localhost
 - glite
 - dc=gsi,dc=de
 - o=glite
 - ou=Config** (selected)
 - ou=Packages
 - ou=People
 - ou=Roles
 - ou=Sites
 - ou=Partitions
 - ou=GasModules
- Right Panel (Configuration Table):**

description	
proxyPort	7078
proxyHost	glite001.gsi.de
authPort	7070
authenDatabase	ADMIN
authenHost	glite001.gsi.de:3307
authenDriver	mysql
authHost	glite001.gsi.de
catalogHost	glite001.gsi.de:3307
catalogPort	7072
catalogDatabase	alien_system
catalogDriver	mysql
queuePort	7073
queueHost	glite001.gsi.de
queueDatabase	processes
queueDriver	mysql
queueDbHost	glite001.gsi.de:3307
logPort	7079
logHost	glite001.gsi.de
userDir	/glite/user
isPort	7071
isDatabase	INFORMATIONSERVICE
isDriver	mysql
isDbHost	glite001.gsi.de:3307
isHost	glite001.gsi.de
nwsPort	
nwsDatabase	
nwsDriver	
nwsDbHost	
nwsHost	
clusterMonitorUser	aliproduct
clustermonitorPort	7074
- Buttons:** Apply, Refresh

File

Search Browse Schema

localhost

gLite

dc=gsi,dc=de

o=gLite

ou=Config

ou=Packages

ou=People

ou=Roles

ou=Sites

ou=Partitions

ou=GasModules

clusterMonitorUser	aliproduct
clustermonitorPort	7074
brokerHost	glite001.gsi.de
brokerPort	7080
transferManagerAddress	glite001.gsi.de:7085
transferBrokerAddress	glite001.gsi.de:7086
transferDatabase	glite001.gsi.de:3307/mysql/transfers
transferOptimizerAddress	glite001.gsi.de:7087
jobOptimizerAddress	glite001.gsi.de:7088
jobDatabase	glite001.gsi.de:3307/mysql/processes
ldapmanager	cn=Manager,dc=gsi,dc=de
jobManagerAddress	glite001.gsi.de:7073
jobBrokerAddress	glite001.gsi.de:7080
processPort	7075
	7076
	7077
	7078
	7079
catalogueDatabase	glite001.gsi.de:3307/mysql/alien_system
catalogueOptimizerAddress	glite001.gsi.de:7089
monalisaLUS	
monalisaGroup	
authenSubject	/O=GermanGrid/OU=GSI/CN=host/glite001.gsi.de
proofPort	
proofDatabase	
proofDriver	
proofDbHost	
proofHost	
proofMuxPortRange	
proofCryptPassword	
proofLogin	
proofTimeout	

Apply

Refresh

File

Search Browse Schema

localhost
gLite
dc=gsi,dc=de
o=gLite
ou=Config
ou=Packages
ou=People
uid=alienmaster
uid=vperso
uid=newuser
ou=Roles
uid=admin
uid=aliprod
ou=Sites
ou=GSI
ou=Config
ou=Services
ou=SE
name=File
ou=CE
name=LSF
ou=FTD
ou=Partitions
ou=GasModules
alias=AliEnFileCatalog
alias=AliEnMetaCatalog

dn name=File,ou=SE,ou=Services,ou=GSI,ou=Sites,o=gLite,dc=gsi,dc=de
objectClass AliEnSE
name File
host glite001.gsi.de
mss File
saverdir /glite-data
port 8081
options
tmpdir
certsubject
lvmDatabase mysql/aliendb4.cern.ch:3306/lvm/

File

Search Browse Schema

localhost

- gLite
 - dc=gsi,dc=de
 - o=glite
 - ou=Config
 - ou=Packages
 - ou=People
 - uid=alienmaster
 - uid=vperso
 - uid=newuser
 - ou=Roles
 - uid=admin
 - uid=aliproduct
 - ou=Sites
 - ou=GSI
 - ou=Config
 - ou=Services
 - ou=SE
 - name=File
 - ou=CE
 - name=LSF
 - ou=FTD
 - ou=Partitions
 - ou=GasModules
 - alias=AliEnFileCatalog
 - alias=AliEnMetaCatalog

dn name=LSF,ou=CE,ou=Services,ou=GSI,ou=Sites,o=glite,dc=gsi,dc=de

objectClass AliEnCE

name LSF

type LSF

host glite001.gsi.de

maxqueuedjobs 10

maxjobs 5

submitcmd bsub

submitarg

killcmd

killarg

statuscmd bjobs

statusarg

packages

Outlook

LCG Internal Review 2003

Recommendations (I)

- Evident: the M/W is one of the most important components of LCG.
- While the M/W is not under the exclusive control of the LCG project, its milestones are very important and need to be included in the project overview.
 - They will clearly need to be negotiated between LCG and EGEE

LCG Milestones are aligned with EGEE ones, within the scope of the ARDA project. Milestones and Deliverables are reviewed by both EGEE and LCG projects. EGEE Deliverable reviews reports are available:

DJRA1.1 (Architecture): <http://edms.cern.ch/document/493614/1>

DJRA1.2 (Design):

<http://edms.cern.ch/document/487871/0.8>

- Having the same person in charge of both is clearly good

Outcome of last ARDA Workshop

- Enlarge size and scope of the development testbed
 - Expose more users than just the ARDA team
 - Made enough computing resources available for realistic analysis
 - Madison installation is being expanded with ~60 CPU's
 - FZK and GSI resources will be added
- Deploy current prototype software on ALICE sites
 - To handle Phase III of ALICE Data Challenge 2004 (supported by LCG, EGEE ...?)
 - To provide early feedback to Middleware developers

Enlarged gLite prototype Testbed through FzK and GSI resources

```
alienmaster@glite001:/tmp/gLite/log - Befehlsfenster - Konsole
Sitzung Bearbeiten Ansicht Lesezeichen Einstellungen Hilfe

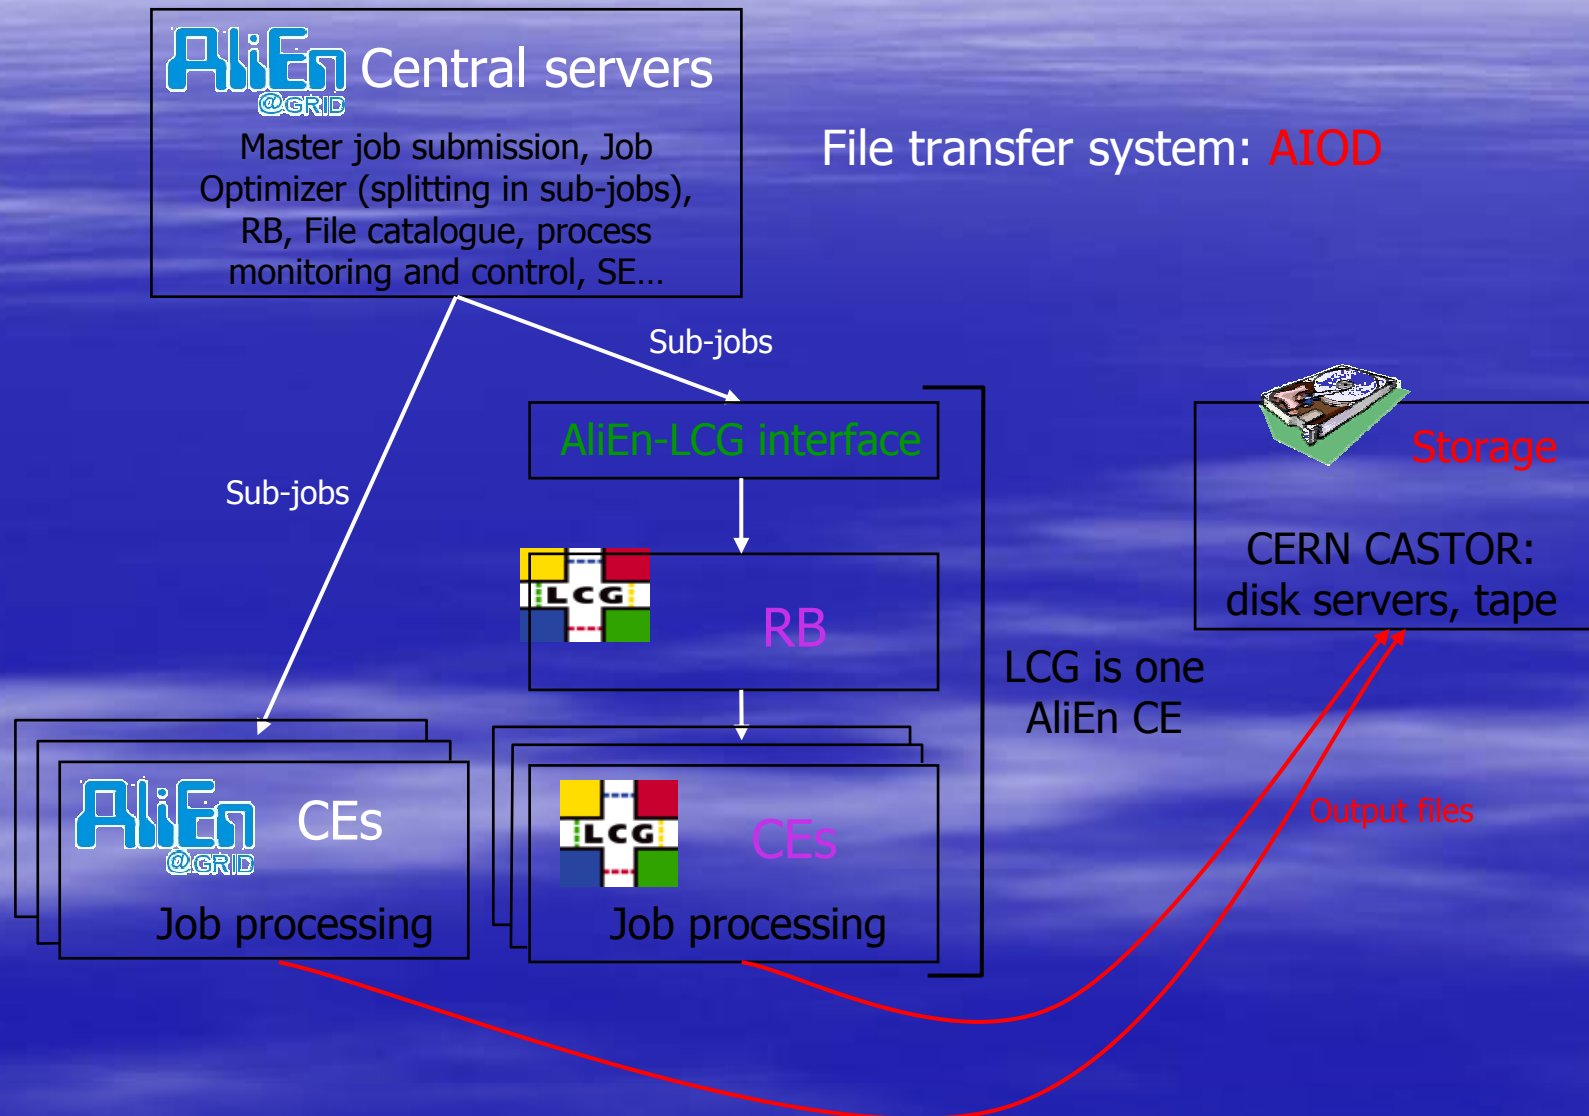
[alienmaster@glite001 log]$
[alienmaster@glite001 log]$
[alienmaster@glite001 log]$
[alienmaster@glite001 log]$
[alienmaster@glite001 log]$
[alienmaster@glite001 log]$
[alienmaster@glite001 log]$
[alienmaster@glite001 log]$
[alienmaster@glite001 log]$
[alienmaster@glite001 log]$
[alienmaster@glite001 log]$
[alienmaster@glite001 log]$ glite login
Warning: No valid proxy. Trying SSH key...
Dec 1 14:19:56 info Error contacting the local SE
Dec 1 14:19:56 notice Starting remotequeue...
Dec 1 14:19:56 info
[
  Requirements = { other.Type == "Job" };
  CloseSE =
  {
    "gLite::GSI::File"
  };
  CE = "gLite::GSI::LSF";
  Host = "glite001.gsi.de";
  LocalDiskSpace = 67364288;
  WHNHost = "glite001.gsi.de";
  Type = "machine";
  Uname = "Linux glite001 2.4.20-30.7.cernsmp #1 SMP Thu Feb 19 12:36:51 CET 2004
i686 unknown"
]
[glite001.gsi.de:3307] /glite/user/a/alienmaster/ > services
== Service == Servicename == Hostname ==
Status ==
-----
Dec 1 14:20:06 error The IS returned an error: No ACTIVE SEs
Dec 1 14:20:06 error The IS returned an error: No ACTIVE CLCs
Dec 1 14:20:06 error The IS returned an error: No ACTIVE CLCs
- [ CluMon ] gLite::GSI::LSF glite001.gsi.de:7074 --
OK
-----
Dec 1 14:20:06 error The IS returned an error: No ACTIVE FTDs
Dec 1 14:20:06 error The IS returned an error: No ACTIVE TopRouters
- [ Core ] IS glite001.gsi.de:7071 --
OK --
- [ Core ] Authn glite001.gsi.de:7070 --
OK --
- [ Core ] Manager/Job glite001.gsi.de:7073 --
OK --
- [ Core ] Logger glite001.gsi.de:7079 --
OK --
- [ Core ] Manager/Transfer glite001.gsi.de:7085 --
OK --
- [ Core ] Broker/Job glite001.gsi.de:7080 --
OK --
- [ Core ] Broker/Transfer glite001.gsi.de:7086 --
OK --
- [ Core ] Optimizer/Transfer glite001.gsi.de:7087 --
OK --
- [ Core ] Optimizer/Catalogue glite001.gsi.de:7089 --
OK
-----
[glite001.gsi.de:3307] /glite/user/a/alienmaster/ >
```

Karlsruhe: comparable status

Both sites (Karlsruhe and GSI) are currently being integrated into the gLite prototype testbed by Pablo Saiz.

ALICE DC and gLite

■ Job structure and production (phase I)



Phase III - Execution Strategy

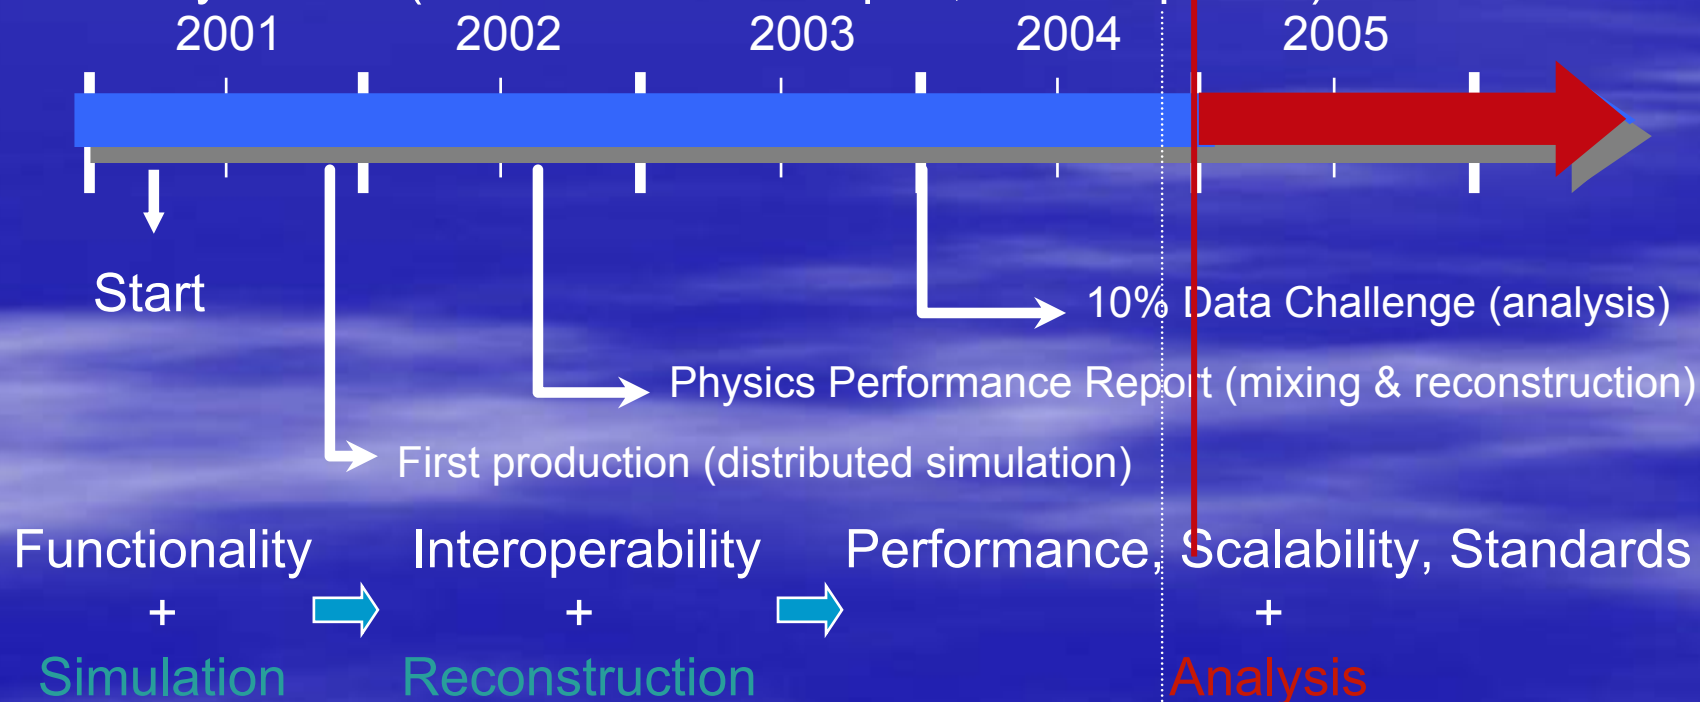
- So... why not doing it with gLite?
- Advantages
 - Uniform configuration: gLite on EGEE/LCG-managed sites & on ALICE-managed sites
 - If we have to go that way, the sooner the better
 - AliEn is anyway “frozen” as all the developers are working on gLite/ARDA
- Disadvantages
 - It may introduce a delay with respect to the use of the present – available – AliEn/LCG configuration
 - But we believe it will pay off in the medium term

Phase III – The Plan


- ALICE is ready to play the guinea-pig for a large scale deployment
 - i.e. on all ALICE resources and on as many existing LCG resources as possible
- We have experience in deploying AliEn on most centres, we can redo the exercise with gLite
 - Even on most LCG centres we have a parallel AliEn installation
 - Many ALICE site-managers are ready to try it
- And we would need little help
 - We need a gLite (beta-) as soon as possible, beginning November
 - Installation and configuration of sites must be as simple as possible
 - I.e. do not require root access
 - We expect help from LCG/EGEE to help us configure and maintain the ALICE gLite server, running common services

The ALICE Grid (AliEn)

- There are millions lines of code in OS dealing with GRID issues
- Why not using them to build the minimal GRID that does the job?
 - Fast development of a prototype, can restart from scratch etc etc
 - Hundreds of users and developers
 - Immediate adoption of emerging standards
- AliEn by ALICE (5% of code developed, 95% imported)



Phase III - Layout



The current plan is, to move to gLite as soon as possible !!!

Nov 9-11: Running Demos



Predrag and Derek giving a demo.

Our Demo

- Demonstrated the feasibility of global distributed parallel interactive data analysis.
- Used 14 sites, each running 4 PROOF workers, i.e. 52 CPU's in parallel.
- Used ALICE MC data that had been produced at these sites during our PDC'04.
- Made a realistic analysis using the ALICE ESD objects.
- Used the AliRoot, ROOT, PROOF, and gLite technologies.

Site A

PROOF SLAVE SERVERS



Site B

PROOF SLAVE SERVERS

LCG

- Proofd
- Rootd
- Forward Proxy
- Forward Proxy



Optional Site Gateway

Only outgoing connectivity

Site <X>

Proofd Startup

Grid Service Interfaces

TGrid UI/Queue UI

Grid Access Control Service

Grid/Root Authentication

Grid File/Metadata Catalogue

Slave
Registration/
Booking- DB

Master Setup

**PROOF
Steer**

**PROOF
Master**

"Standard" Proof Session

Booking Request
with logical file names

Client retrieves list
of logical file (LFN + MSN)

**PROOF
Client**

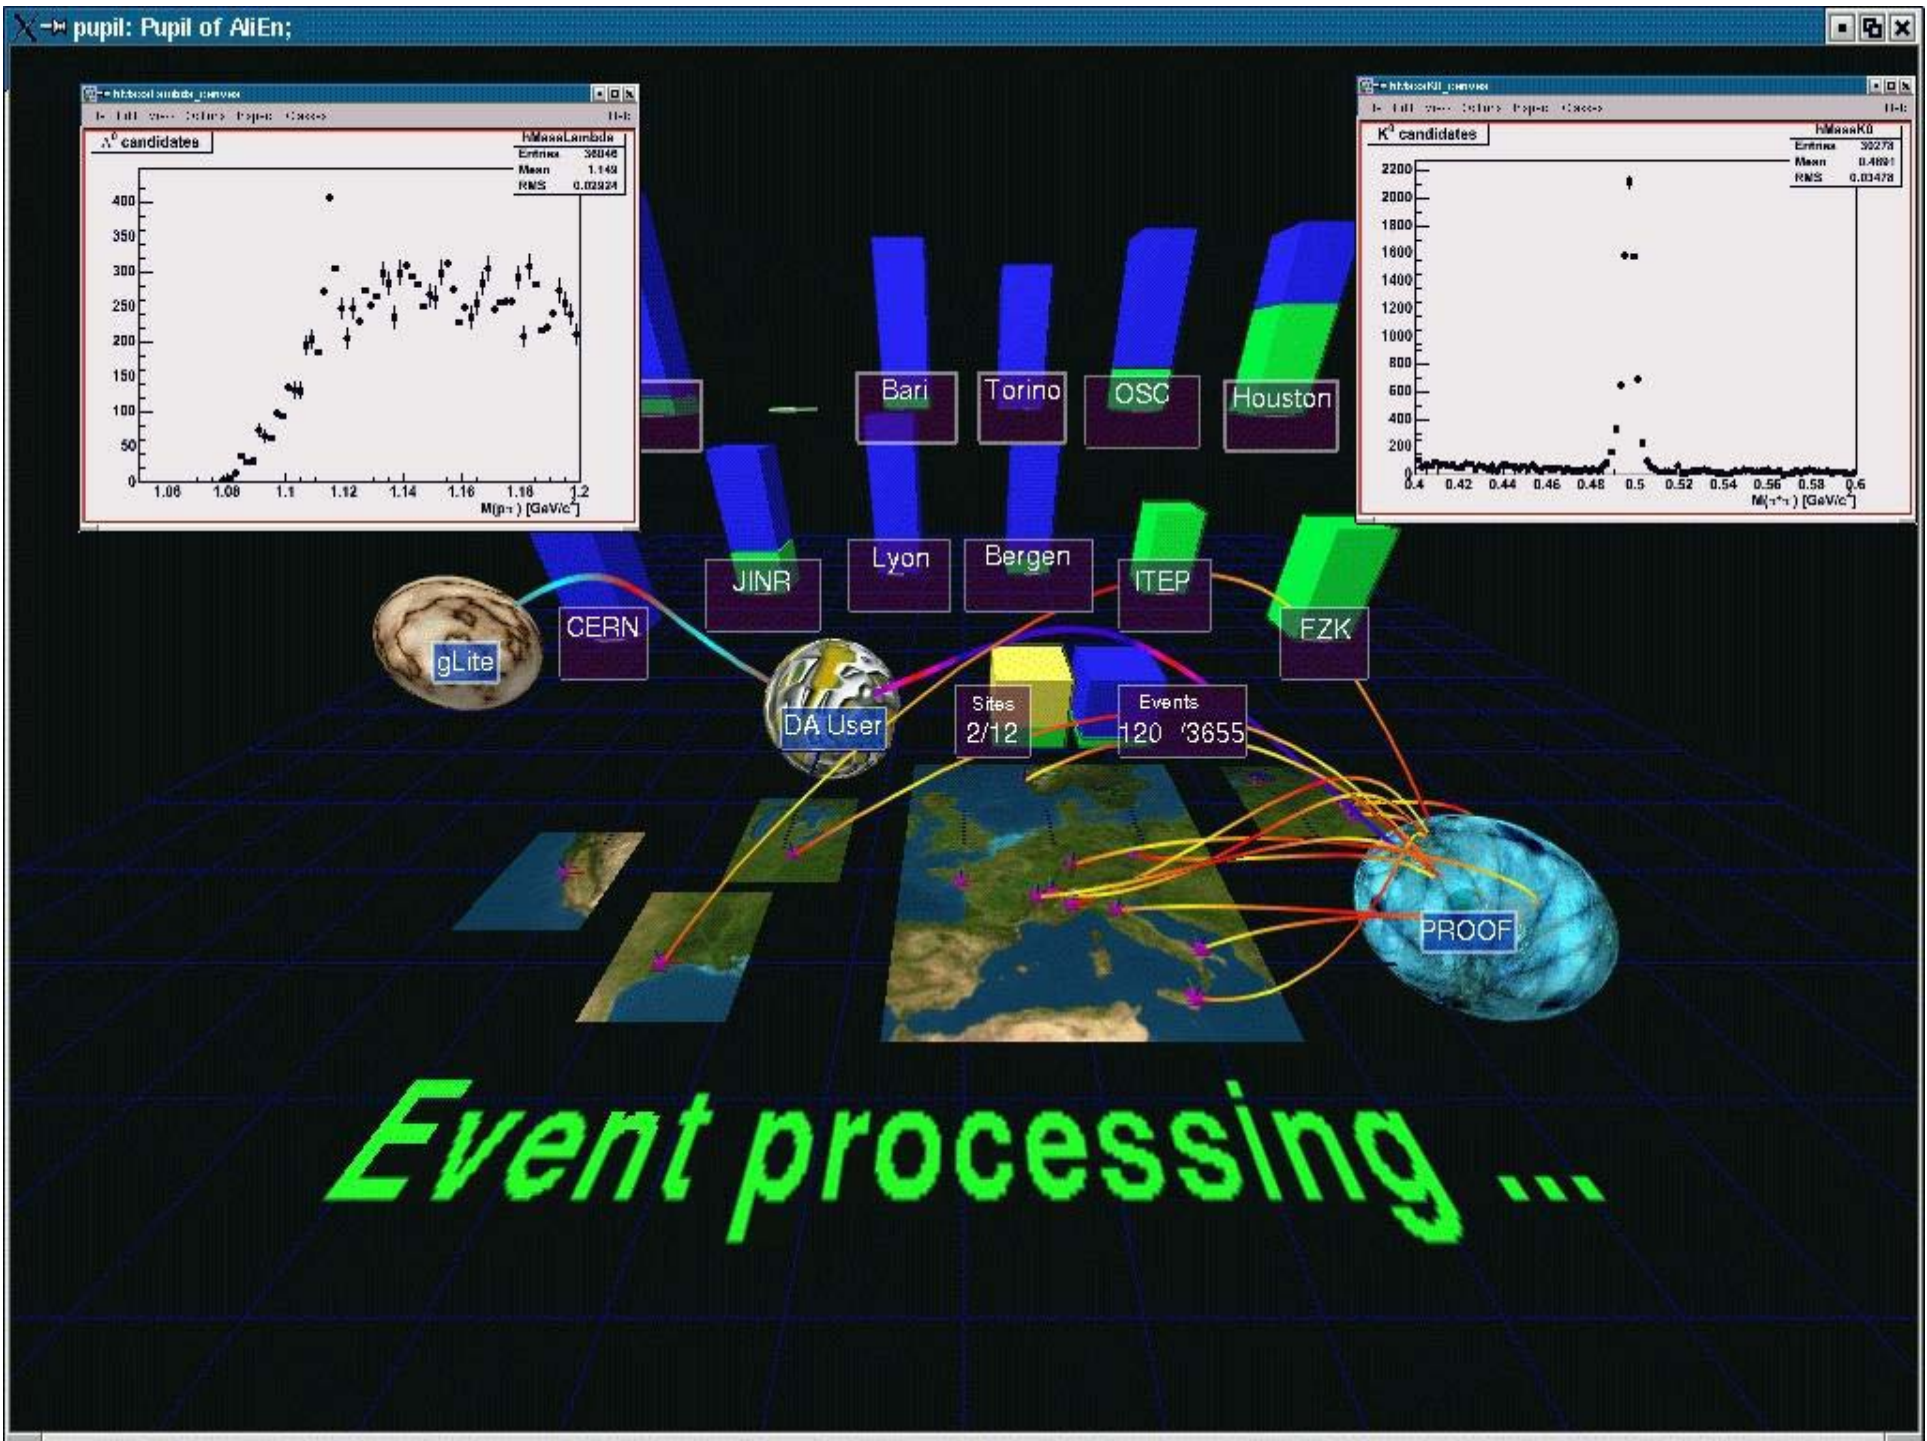
ROOT

Client

Slave ports
mirrored on
Master host

Master

Grid-Middleware independend PROOF Setup



Next Steps (ARDA/gLite)

- Deliver scheduled components to pre-production service
 - Get Operations feedback and adapt as required
- Deploy prototype Middleware to ALICE sites
- Enlarge development testbed and open it to a larger community
- Finalize contents of EU RC1 release
- Deliver and test all components out of integration builds
- Finalize first integrated release as an EU deliverable

Links

- JRA1 homepage
 - <http://egee-jra1.web.cern.ch/egee-jra1/>
- Architecture document
 - <https://edms.cern.ch/document/476451/>
- Release plan
 - <https://edms.cern.ch/document/468699>
- Prototype installation
 - <http://egee-jra1.web.cern.ch/egee-jra1/Prototype/testbed.htm>
- Test plan
 - <https://edms.cern.ch/document/473264/>
- Design document
 - <https://edms.cern.ch/document/487871/>

gLite Web site



The screenshot shows a web browser window displaying the gLite website. The browser's address bar shows 'EGEE > gLite'. The website has a blue header with the 'eGee' logo and the tagline 'Enabling Grids for E-science in Europe'. Below this is the 'gLite' logo and the text 'Lightweight Middleware for Grid Computing'. A yellow banner highlights the 'What is gLite?' section, which describes gLite as the next generation middleware for grid computing, born from the collaborative efforts of more than 80 people in 10 different academic and industrial research centres as part of the EGEE Project. A link to a presentation is provided. Another yellow banner highlights the 'gLite News' section, announcing the new gLite web site unveiled on 13/09/2004. A third yellow banner highlights the 'gLite People' section, listing the organizations involved: CERN, INFN, and Datamat. The left sidebar contains a navigation menu with links to GLITE SUBSYSTEMS (ALIEN, COMPUTING ELEMENT, DATA MANAGEMENT, ACCOUNTING, LOGGING AND BOOKEEPING, MONITORING, SECURITY, WORKLOAD MANAGEMENT), DOWNLOAD (DOCUMENTATION, PACKAGES), ABOUT GLITE (EGEE JRA1, EGEE JRA3, EGEE JRA4), and ABOUT EGEE. The footer includes the page update date (12/09/2004), a link to 'About the website', and a W3C XHTML 1.0 validation logo.

eGee
Enabling Grids for
E-science in Europe

gLite
Lightweight Middleware for Grid Computing

What is gLite?

gLite (pronounced "gee-lite") is the next generation middleware for grid computing. Born from the collaborative efforts of more than 80 people in 10 different academic and industrial research centres as part of the [EGEE Project](#), gLite provides a bleeding-edge, best-of-breed framework for building grid applications tapping into the power of distributed computing and storage resources across the Internet.

Want to know more about gLite? Read the following [presentation](#).

gLite News

New gLite web site unveiled (13/09/2004)
The new gLite web site has officially gone online on Monday 13 September. The web site offers a single point of access to public documentation, installation packages and guides and loads of other useful information. The web site has been developed by the gLite [Integration Team](#) with the collaboration of all project members using original web templates from [TERENA](#).

gLite People

The gLite software is produced as part of the EU EGEE Project funded by the European Communities. The following academic and industrial research centres are collaborating to the development of the software organized in three different Activities: [JRA1](#) (data management, workload management, monitoring, accounting, computing element, logging and bookkeeping), [JRA3](#) (security) and [JRA4](#) (network monitoring and provisioning).

 The European Organization for Nuclear Research (CERN)

 Istituto Nazionale di Fisica Nucleare (INFN), Italy

 Datamat Spa, Italy

Page updated: 12/09/2004
[About the website](#)

W3C XHTML 1.0

Conclusions and outlook

- Larger infrastructure needed to
 - Attract real users
 - Continue the validation problem on a credible scale
 - Incremental process on the prototype (functionality) and its extension (scale)
- ARDA created multiple channels of communication
 - The most important being experiments \leftrightarrow gLite
 - Assume some natural selection/bootstrap will happen
 - Continue with the ARDA workshops + regular meetings (every fortnight) to start (recommendation of the last workshop)
 - Other opportunities will be exploited
- ARDA produced a large feedback from the experiments to gLite

Conclusions and outlook 2

- ARDA uses all components made available on the gLite prototype
 - Experience and feedback
- First version of analysis systems are being demonstrated
 - We look forward to have users!

Major issues

- Documentation and installation procedures
- Applications: we need real users other than our enthusiastic HEP colleagues using our infrastructure in production and fast (if we want EGEE-II...)
- Training and education are also means to build new EGEE user communities

Major issues

- The Project Mgmt Board unanimously supported the plan to adhere to the project work-plan (Annex 1) and ensure a release of gLite is ready for deployment in March 2005
- ALL effort (funded or unfunded, full-time or part-time) in JRA1 will be concentrated on bringing a selected set of high priority components to production-ready status
- Any groups that wish to take earlier versions of gLite are welcome to do so but the support of these deployments is not the responsibility of JRA1

Foreseen risk

- There exists a high risk that the project may not meet its objective due to conflicting requirements and interests in the development of the gLite middleware
- The project is facing a difficulty in the development of gLite with two possible scenarios
 - Focus JRA1 integration and testing on AliEn components
 - High-energy physics application will take benefit of such a scenario
 - Continue delivery to pre-production service as planned
 - Most of the applications will benefit of such a scenario
- Such situation must be addressed urgently by the Project Director having in mind the objective of the project
 - “Enabling Grids for e-Science in Europe”
 - We recommend thus to follow the second scenario

To conclude

Everyone has to take ...

“EGEE : Enabling Grid for e-Science in Europe”

... as a day to day Mantra !