Grid Computing in HEP and other Fields

Dirk Duellmann, CERN

CERN School Thailand 2010

12 October 2010

Chulalongkorn University, Bangkok





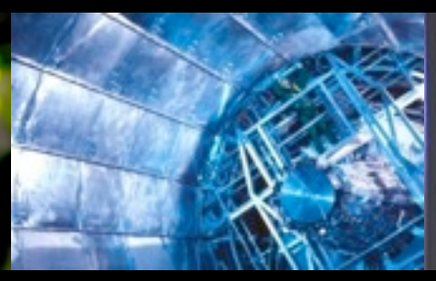


WLCG Worldwide LHC Computing Grid

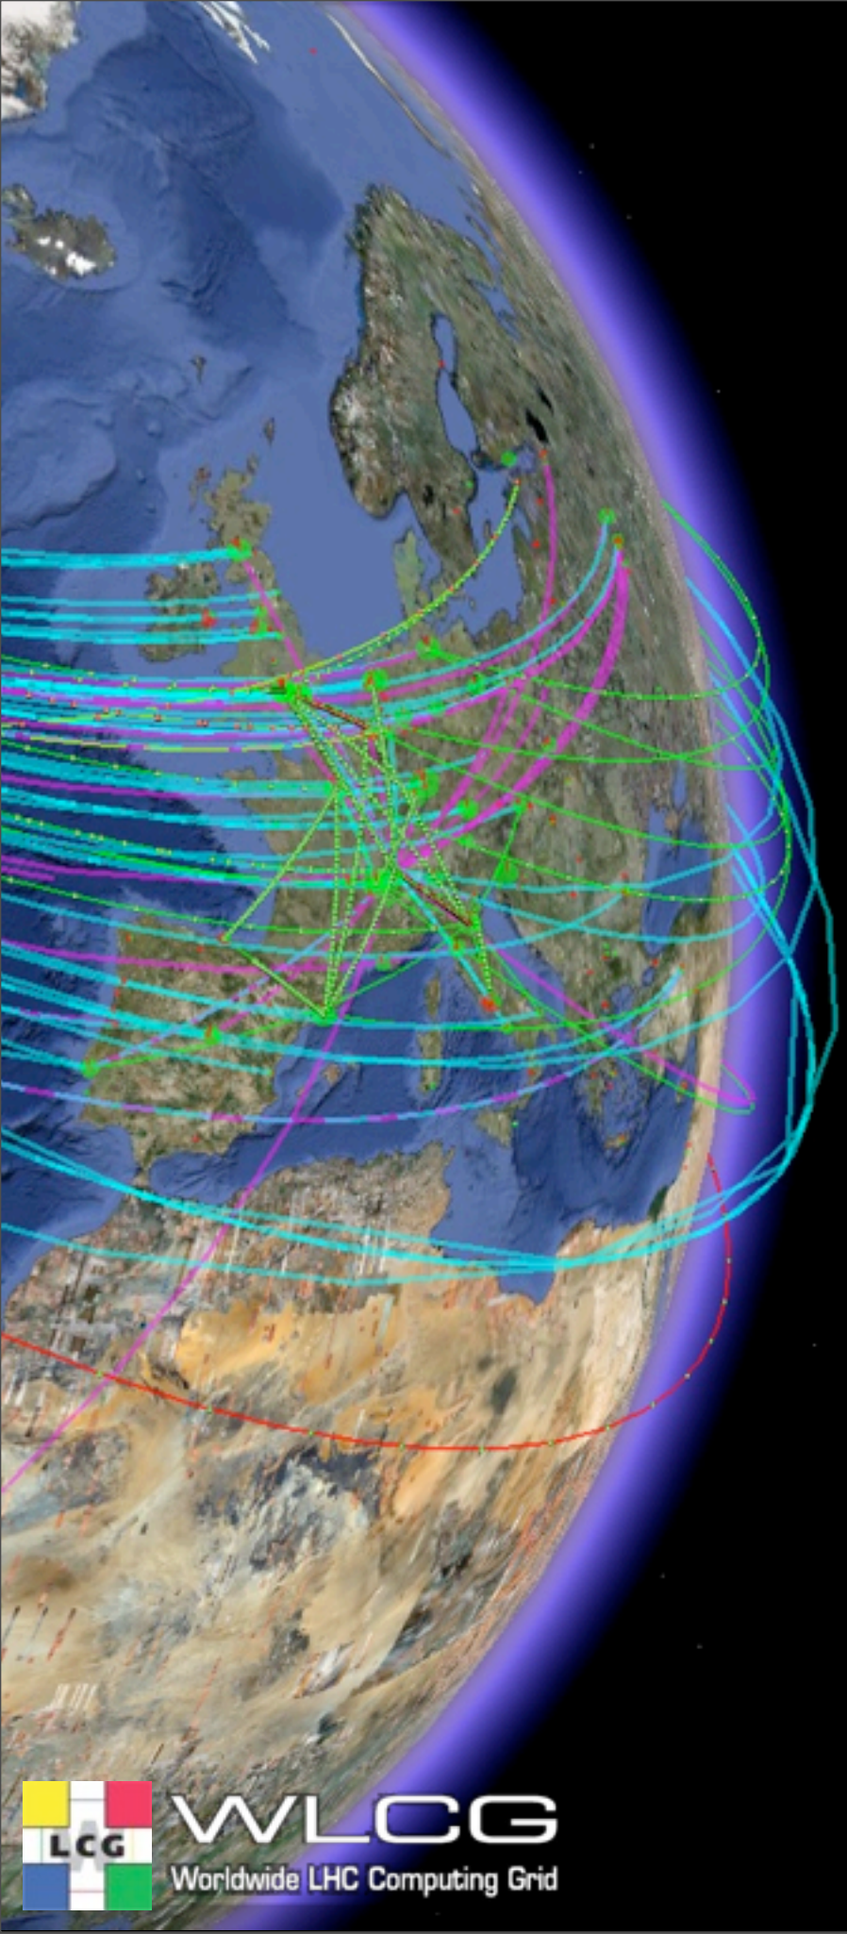

Outline

- Computing requirements of large scale projects
  - Eg LHC data processing
  - Hardware and Software base in HEP
- What is grid computing?
  - Why is it needed?
- Which problems are solved and how?
  - Concepts introduced by the grid
  - Components of today's production grids
  - How are large scale grids organised/coordinated?
  - How are they operated?
- What has be achieved up until today?
  - Which problems are still being addressed?
- Which future changes are expected?
  - Technology, market and user changes

WLCG
Worldwide LHC Computing Grid

Production Grids

- WLCG relies on a *production quality* infrastructure
  - Requires standards of:
    - Availability/reliability
    - Performance
    - Manageability
  - Is used 365 days a year ... (has been for several years!)
  - Tier 1s must store the data for at least the lifetime of the LHC - ~20 years
    - Not passive – requires active migration to newer media
- Vital that we build a fault-tolerant and reliable system
  - That can deal with individual sites being down and recover

CERN

Amsterdam/NIKHEF-SARA

Taipei/ASGC

Bologna/CAF

TRIUMF

BNL

NGDF

FNAL

Lyon/CCIN2P3

Barcelona/PIC

RAL

CERN

KIT

Bologna/CAF

Taipei/ASGC

BNL

TRIUMF

FNAL

Lyon/CCIN2P3

NGDF

FZK

Barcelona/PIC

RAL

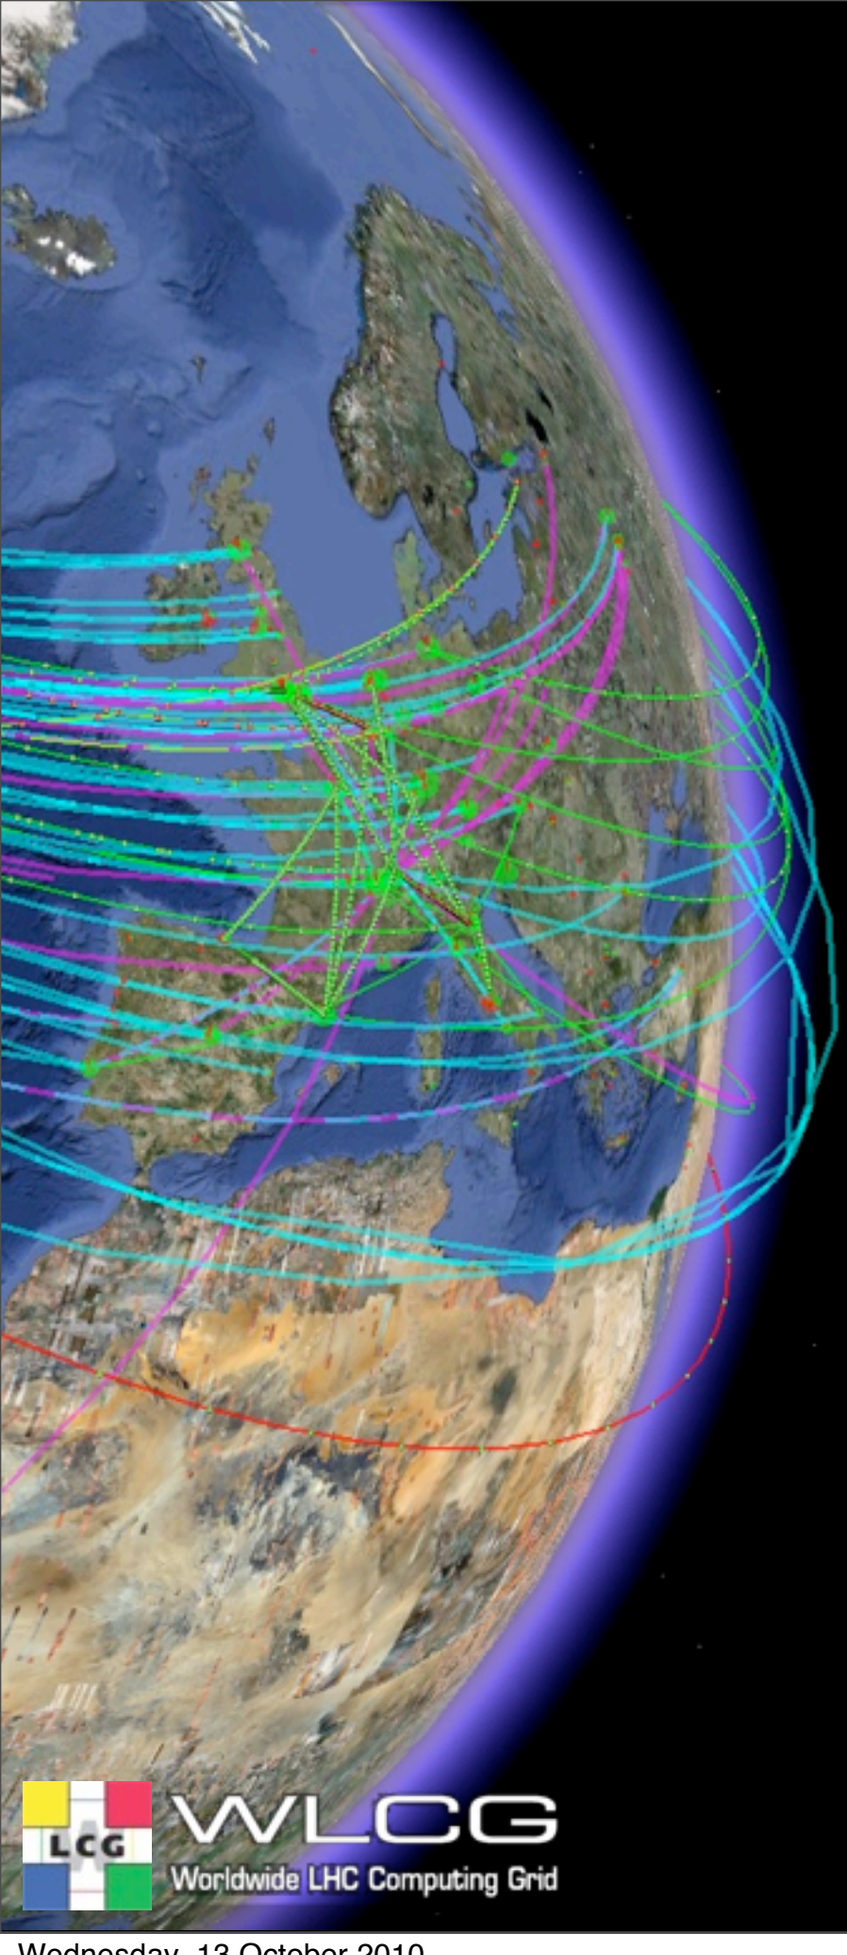

What is needed to make it work?

- Apart from Middleware
- Apart from Computer Centers

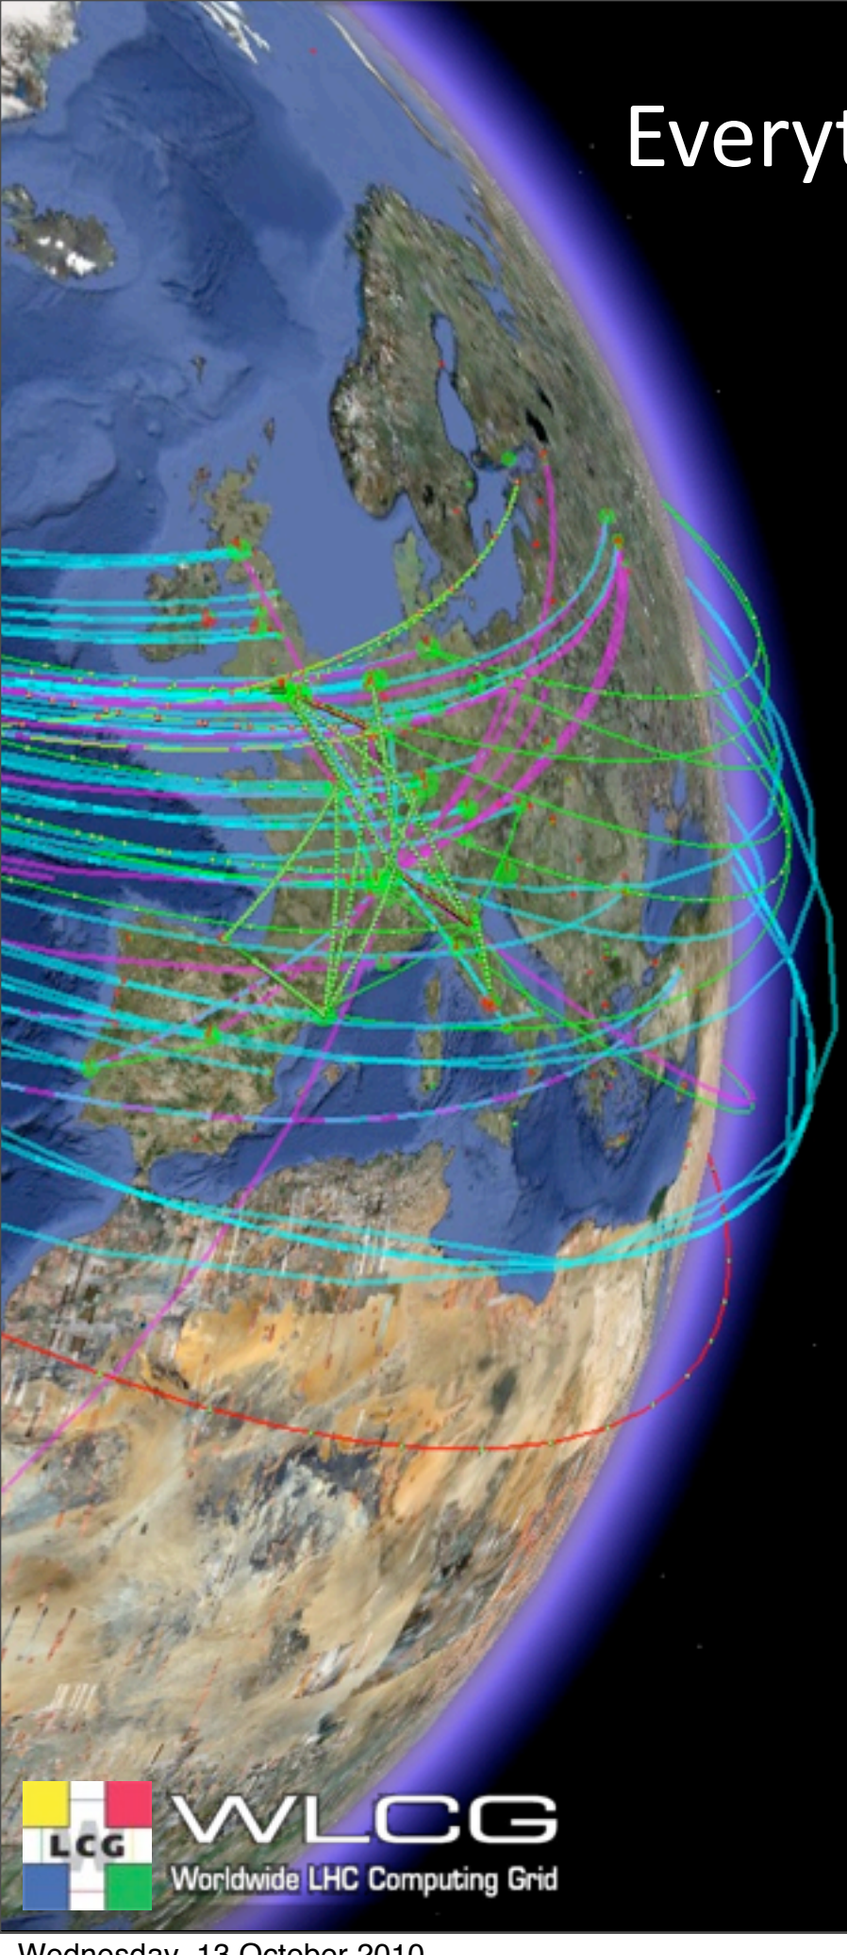

Everything you need in a Computer Center!

- Management
- Fabric
- Networking
- Security
- Monitoring
- User Support
- Problem Tracking
- Accounting
- Service support
- SLAs…..

- But now on a global scale!
  - Respecting the sites' independence
  - Linking the different infrastructures
    - NDGF, EGEE (EGI), OSG

Markus Schulz

Worldwide LCG Organisation

LHC Committee – LHCC
Scientific Review

Computing Resources Review Board – C-RRB
Funding Agencies

Resource Scrutiny Group – C-RSG

Collaboration Board – CB
Experiments and Regional Centres

Overview Board - OB

Management Board
Management of the Project

Architects Forum
Coordination of Common Applications

Grid Deployment Board
Coordination of Grid Operations

Activity Areas

- **Physics Applications Software**
- **Service & Support**
- **Grid Deployment**
- **Computing Fabric**

Worldwide LCG Organisation

LHC Committee – LHCC
Scientific Review

Computing Resources Review Board – C-RRB
Funding Agencies

Resource Scrutiny Group – C-RSG

Collaboration Board – CB
Experiments and Regional Centres

Overview Board - OB

EGEE, OSG representation

Management Board
Management of the Project

Architects Forum
Coordination of Common Applications

Grid Deployment Board
Coordination of Grid Operations

Activity Areas

- **Physics Applications Software**
- **Service & Support**
- **Grid Deployment**
- **Computing Fabric**

WLCG Coordination

- WLCG Operations Meetings
  - Experiments, Tier 1 and Tier 0
  - 30 minutes Mon-Fri
    - Follow up on current problems
- WLCG T1 Service Coordination meeting
  - Every two weeks
    - Longer term scope
    - Operational Planning
  - Incidents followup
    - Service Incident Reports
- Detailed monitoring of the SLAs.

Global Grid User Support

- GGUS: Web based portal
  - –about 1000 tickets per months
  - –Grid security aware
  - –Interfaces to regional/national support structures

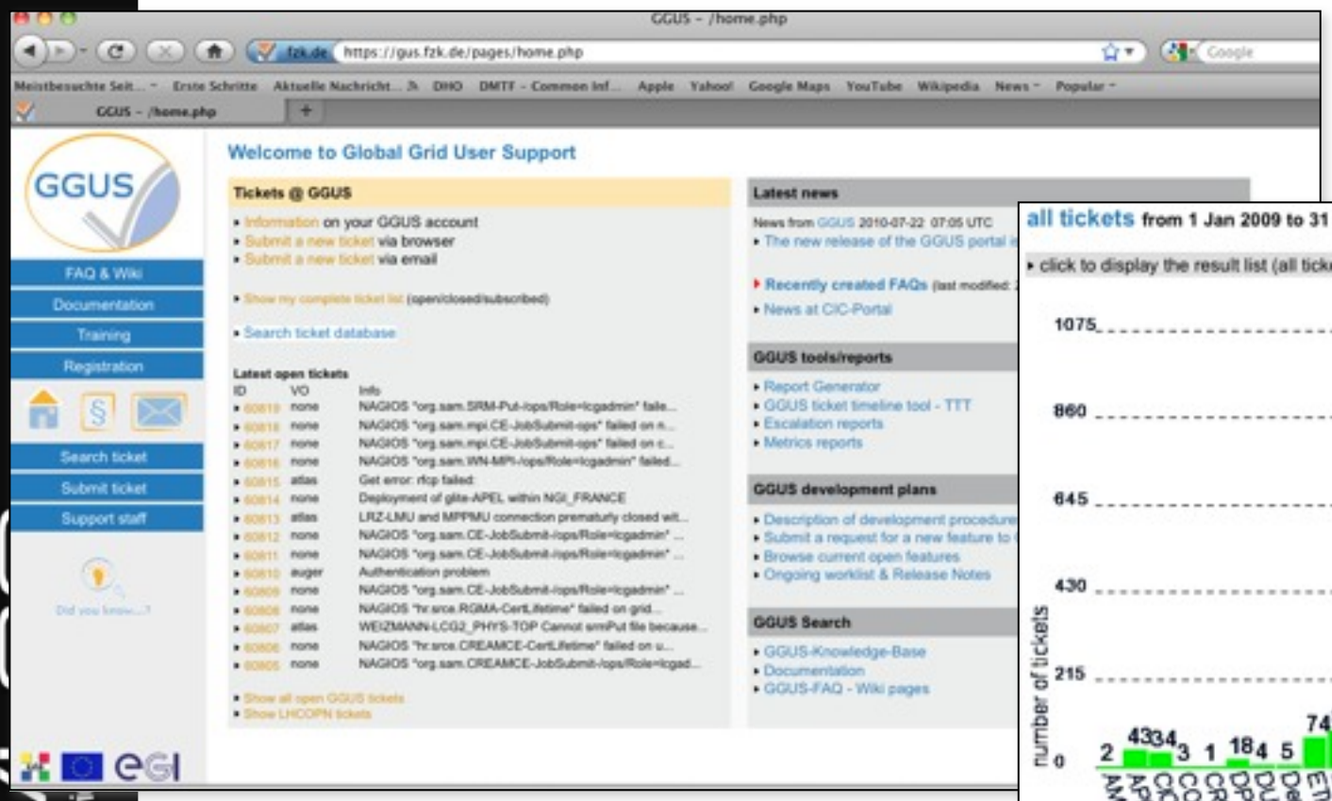

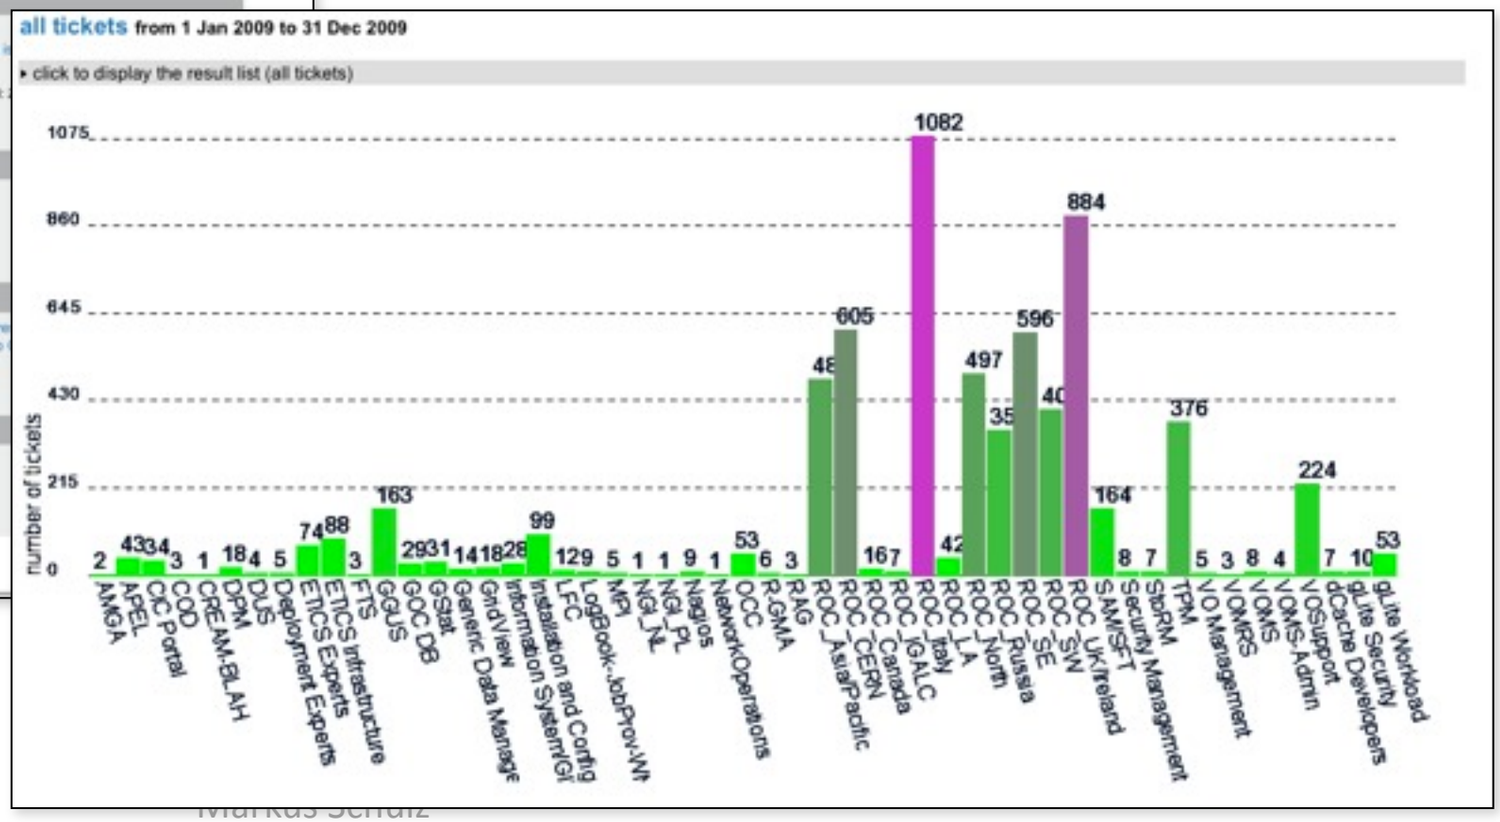

Reliabilities

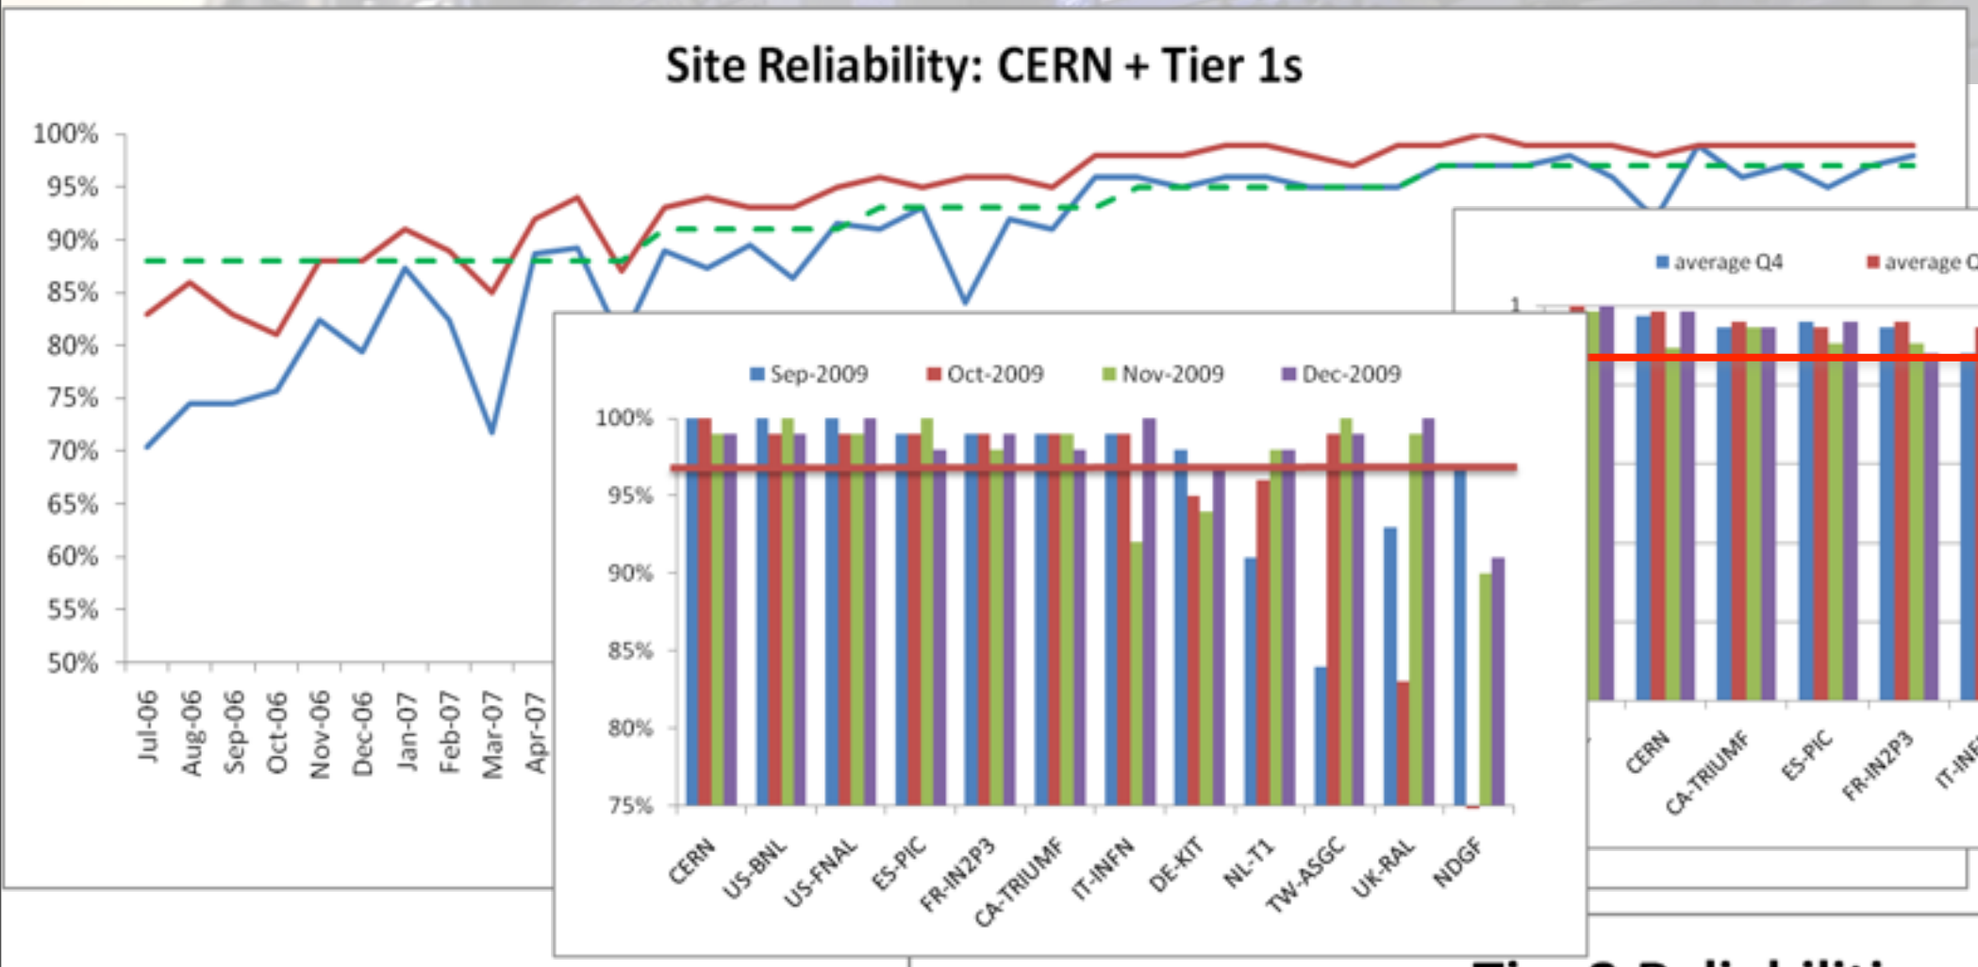

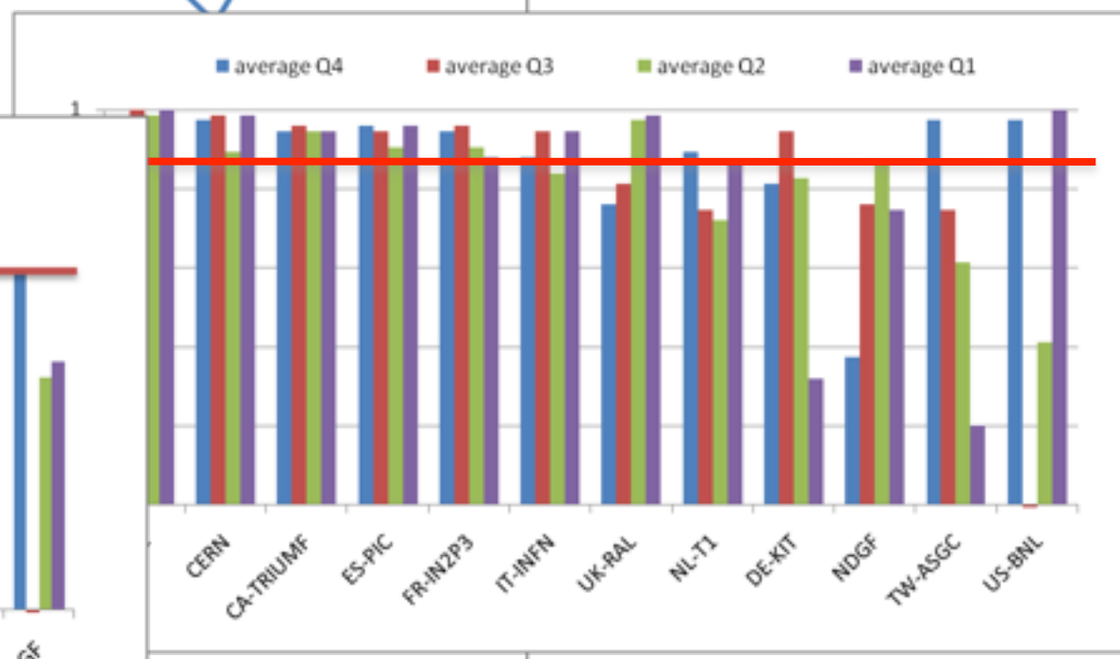

- Experiment-specific measures give a complementary view
- Need to be used together with some understanding of underlying issues

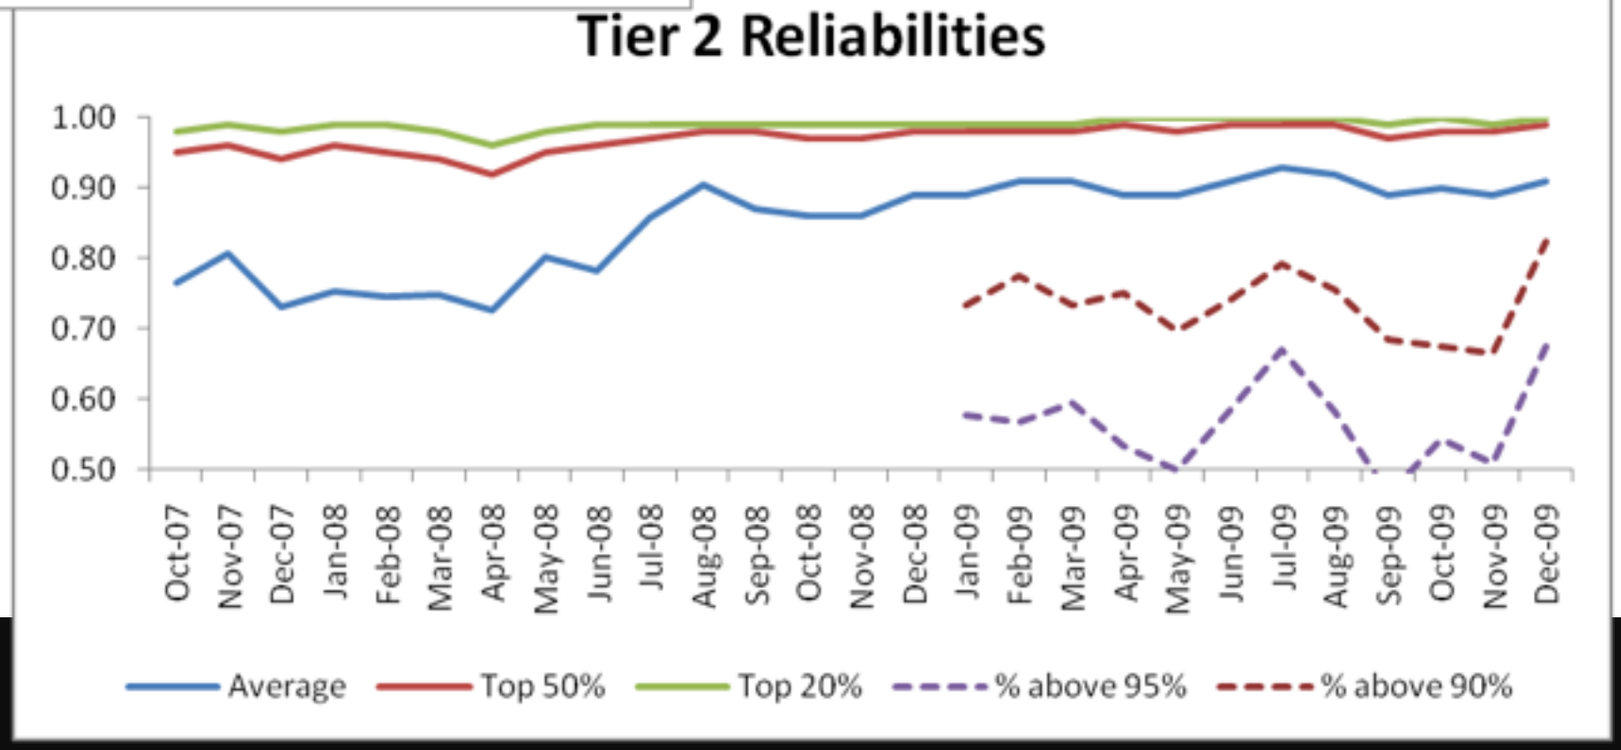

Monitoring to Improve Reliability

- Monitoring
- Metrics
- Workshops
- Data challenges
- Experience
- Systematic problem analysis
- Priority from software developers

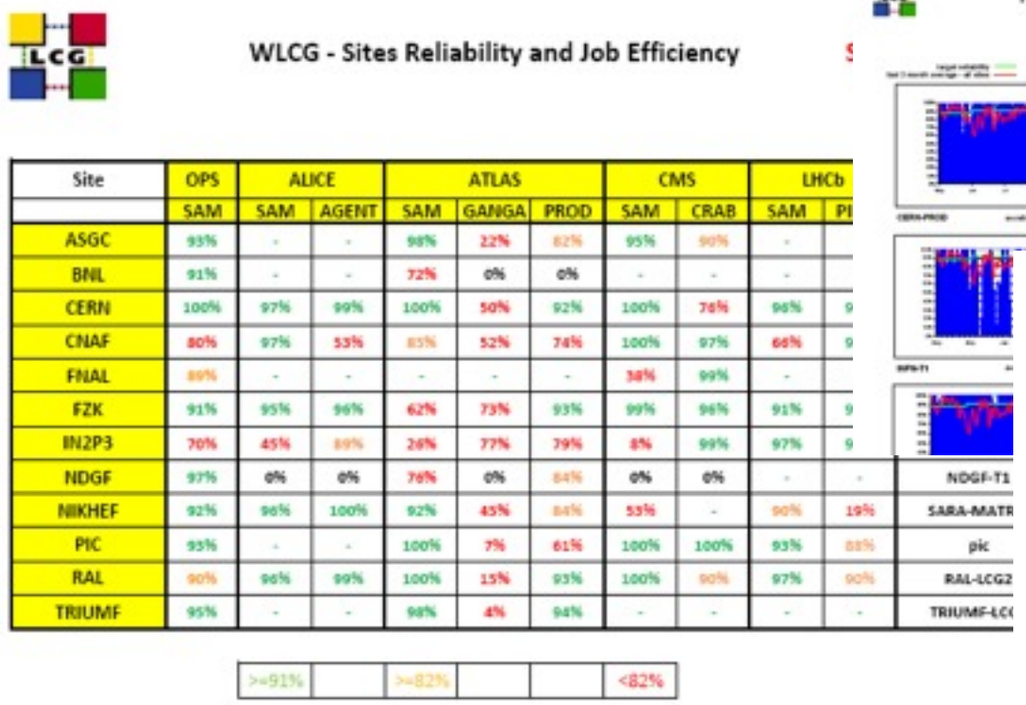

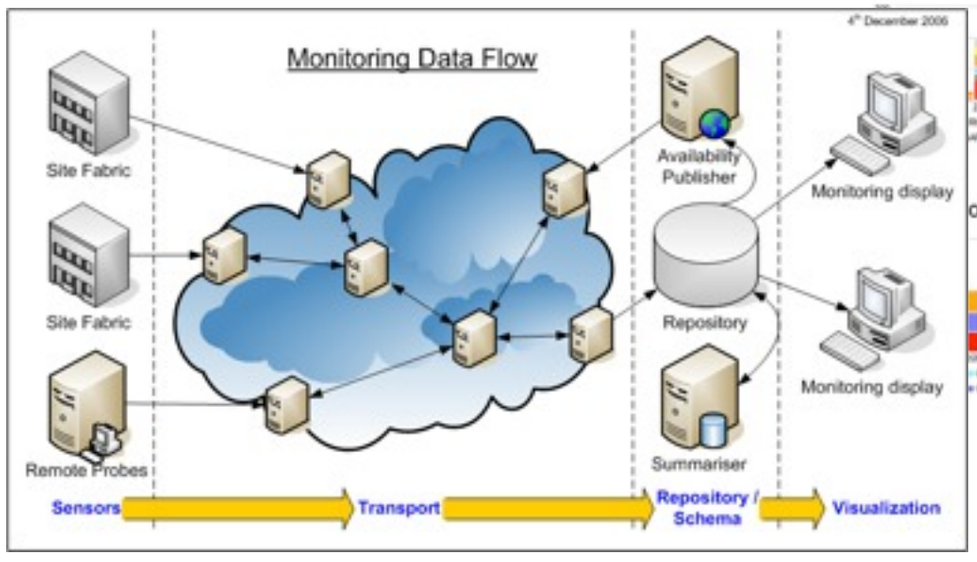

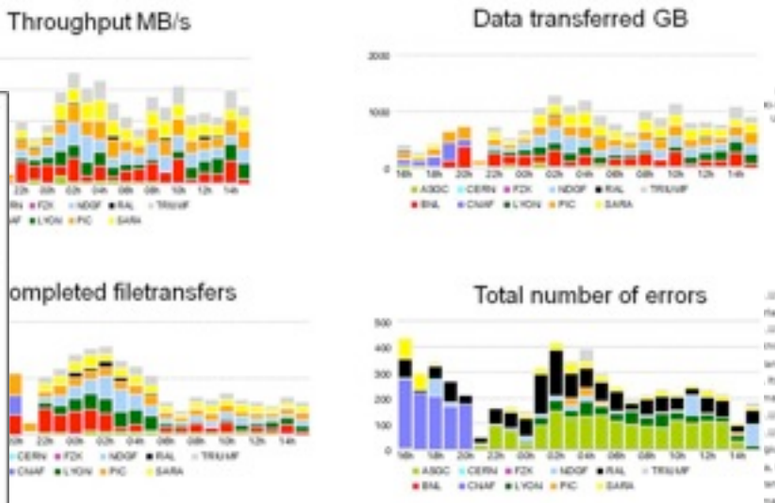

ATLAS M4 Data Monitoring – August 31

GridMap for Visualization

Drilldown into region by clicking on the title

Link: http://gridmap.cern.ch

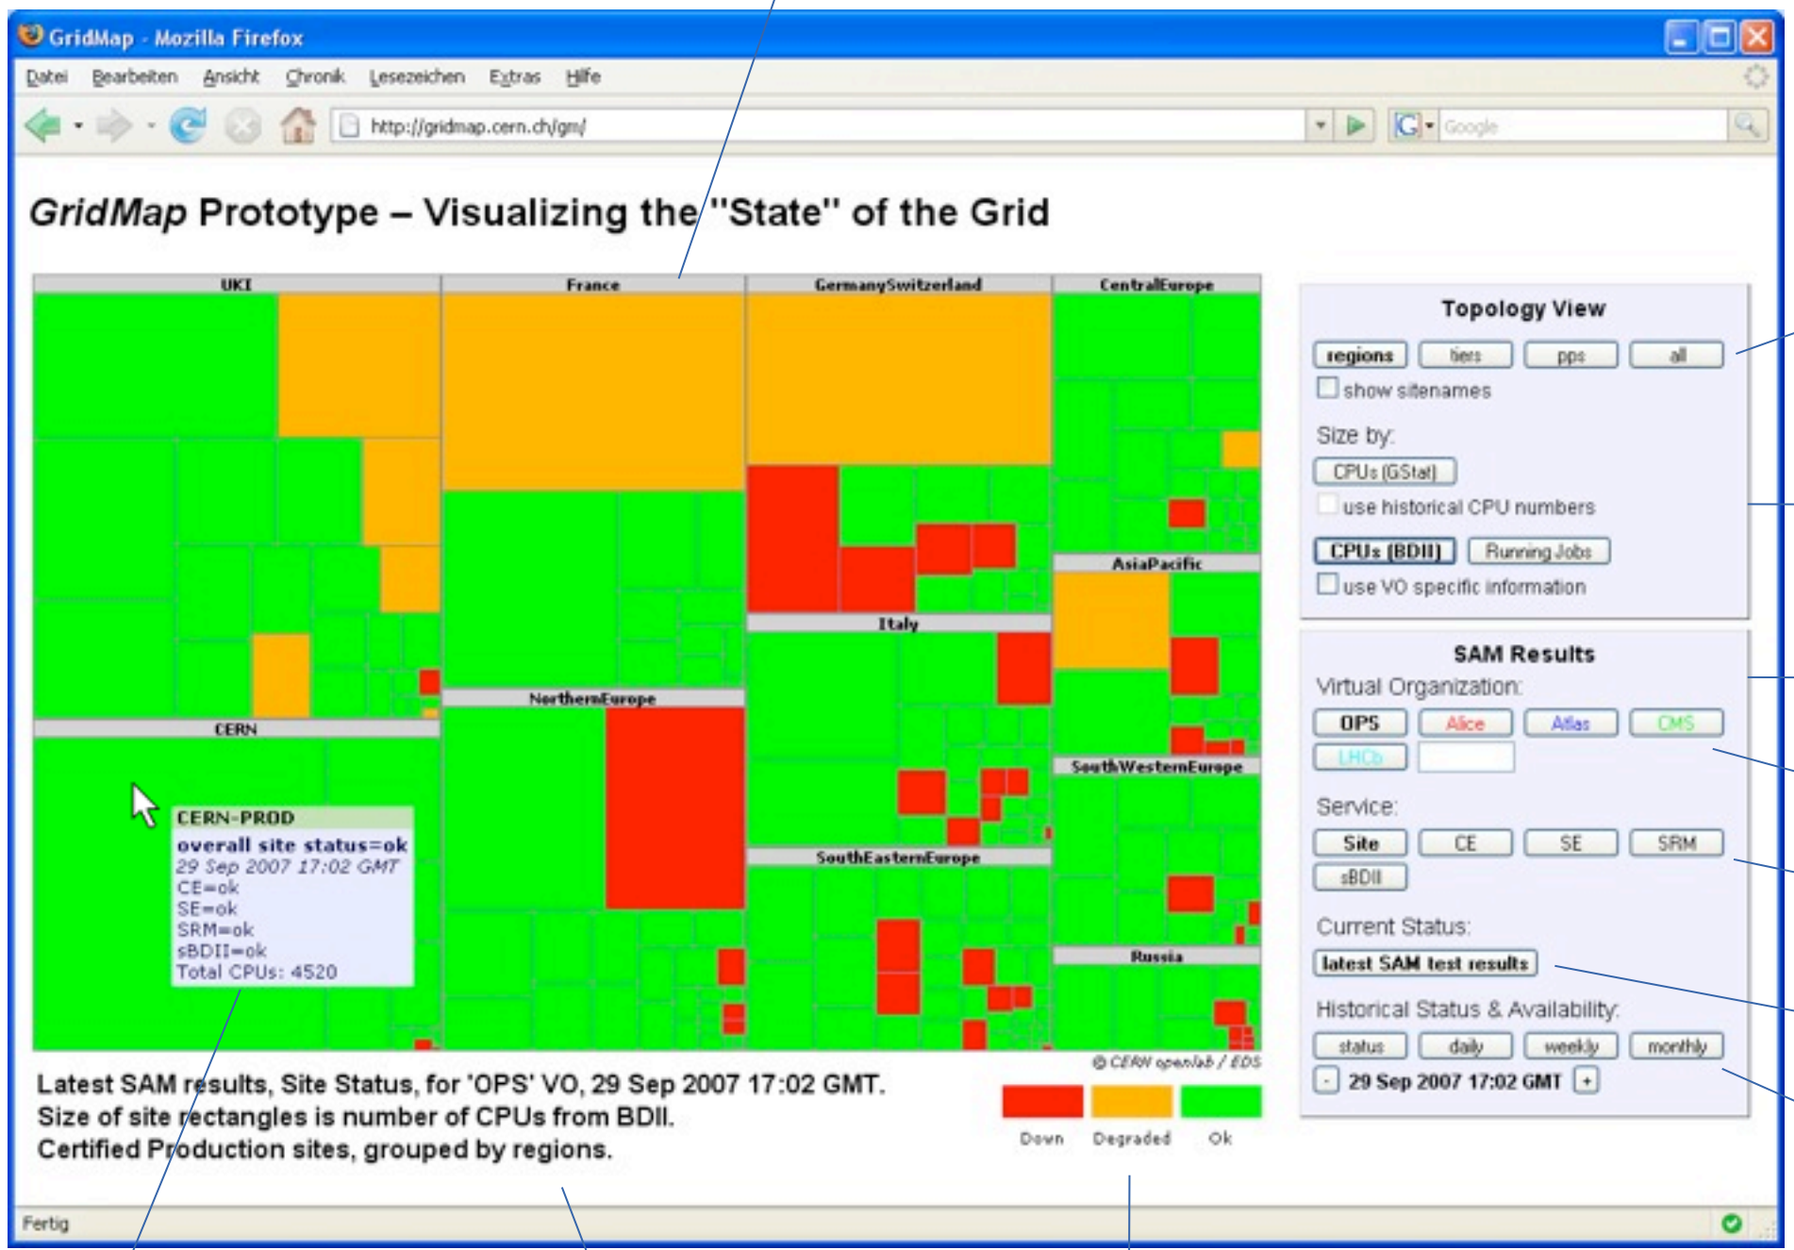

Grid topology view (grouping)

Metric selection for **size** of rectangles

Metric selection for **colour** of rectangles

VO selection

Overall Site or *Site Service* selection

Show SAM status

Show GridView availability data

Context sensitive information

Description of current view

Colour Key

Monitoring

ActiveMQ

- Availability/Reliability monitoring
  - SAM tests and infrastructure
    - recently migrated to NAGIOS based system, decentralized
  - Visualization: GridView, GridMap, dashboards……
  - Solid foundation: Monitoring Infrastructure

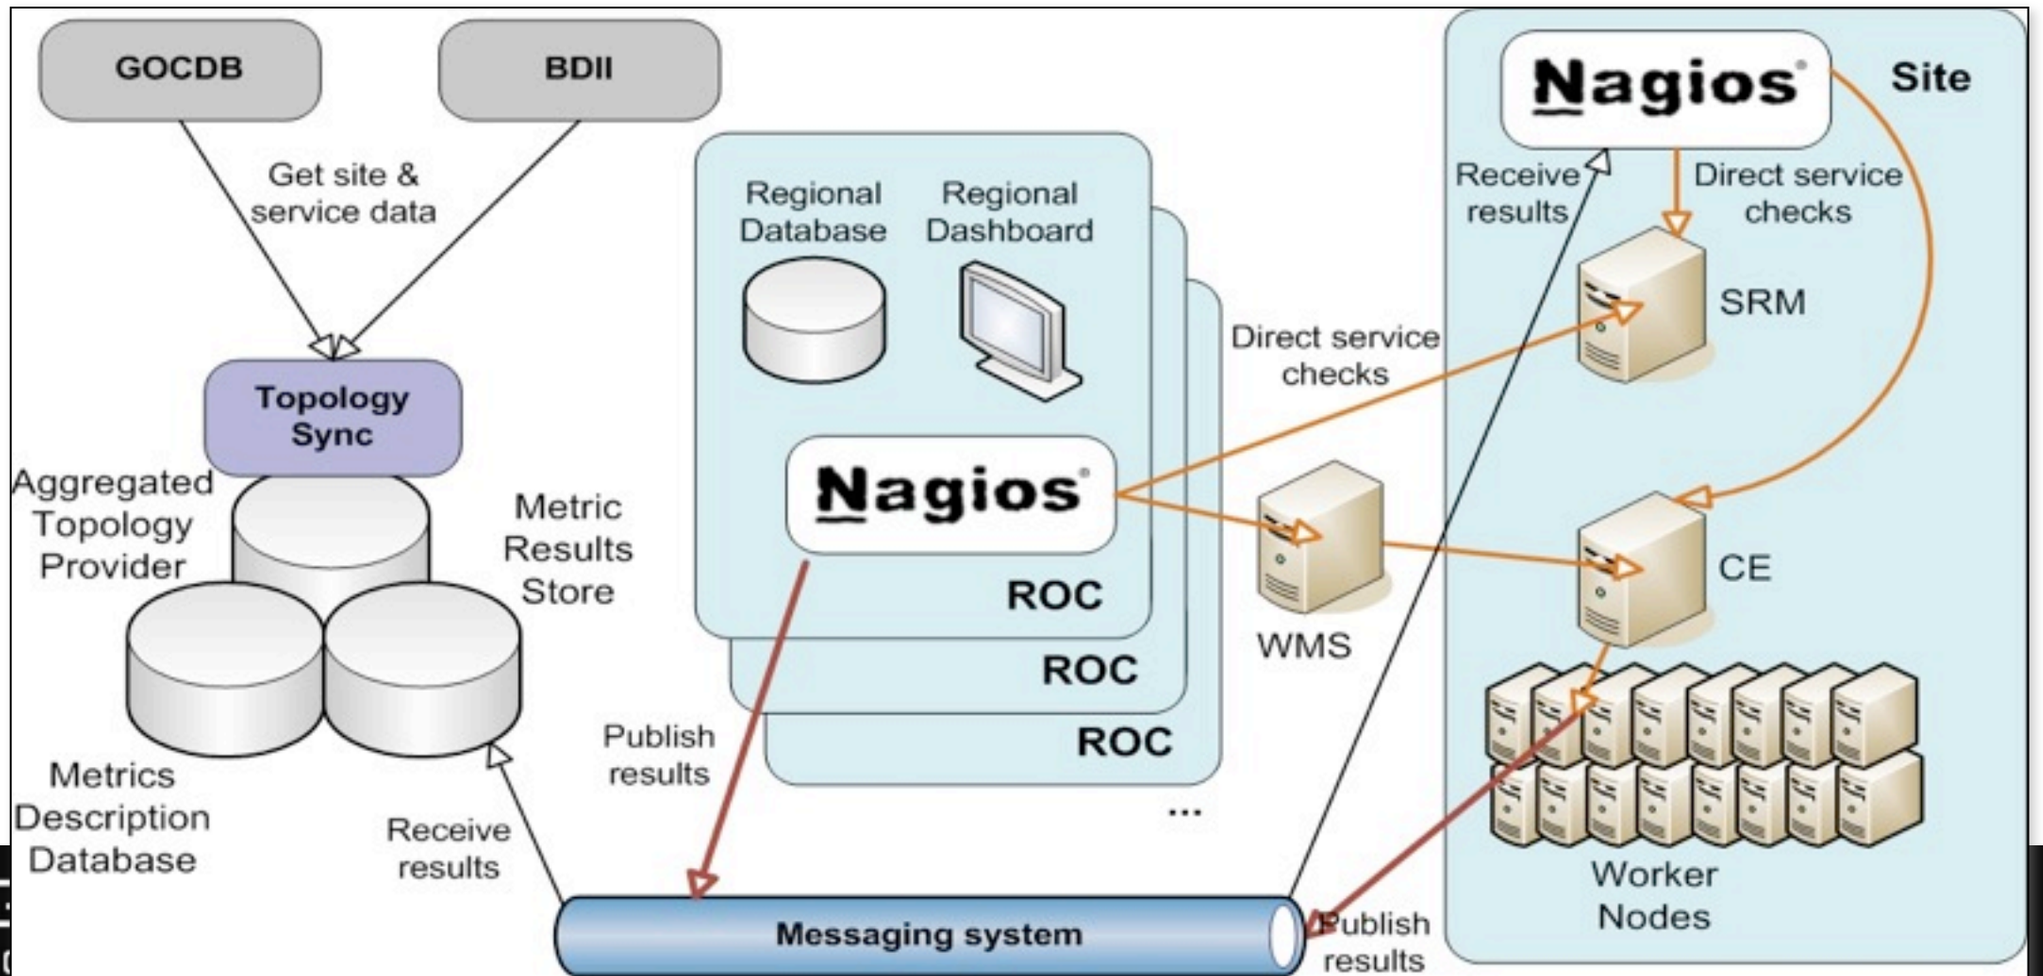

Security & Policy

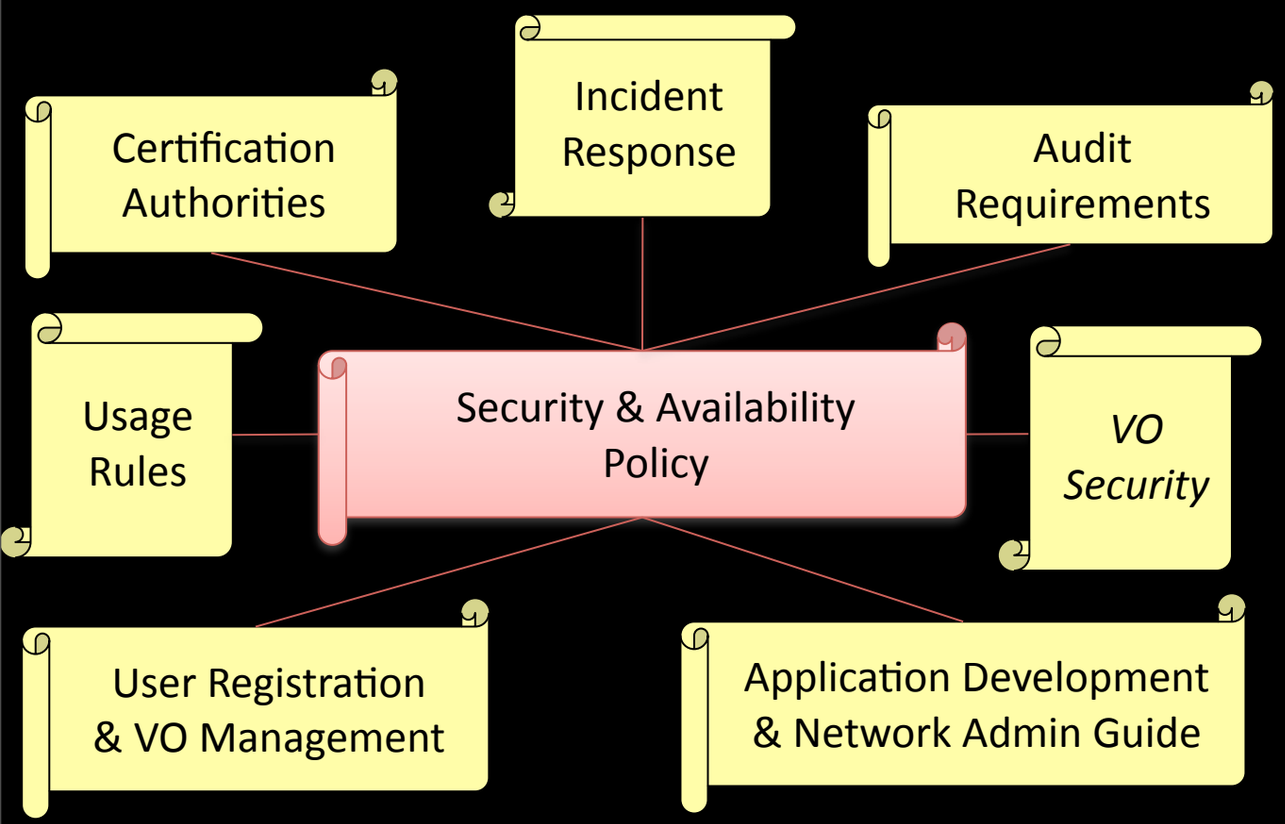

Collaborative policy development

- Many policy aspects are collaborative works; e.g.:

- Joint Security Policy Group
- Certification Authorities
  - EUGridPMA → IGTF, etc.
- Grid Acceptable Use Policy (AUP)
  - common, general and simple AUP
  - for all VO members using many Grid infrastructures
    - EGEE, OSG, SEE-GRID, DEISA, national Grids…
- Incident Handling and Response
  - defines basic communications paths
  - defines requirements (MUSTs) for IR
  - not to replace or interfere with local response plans

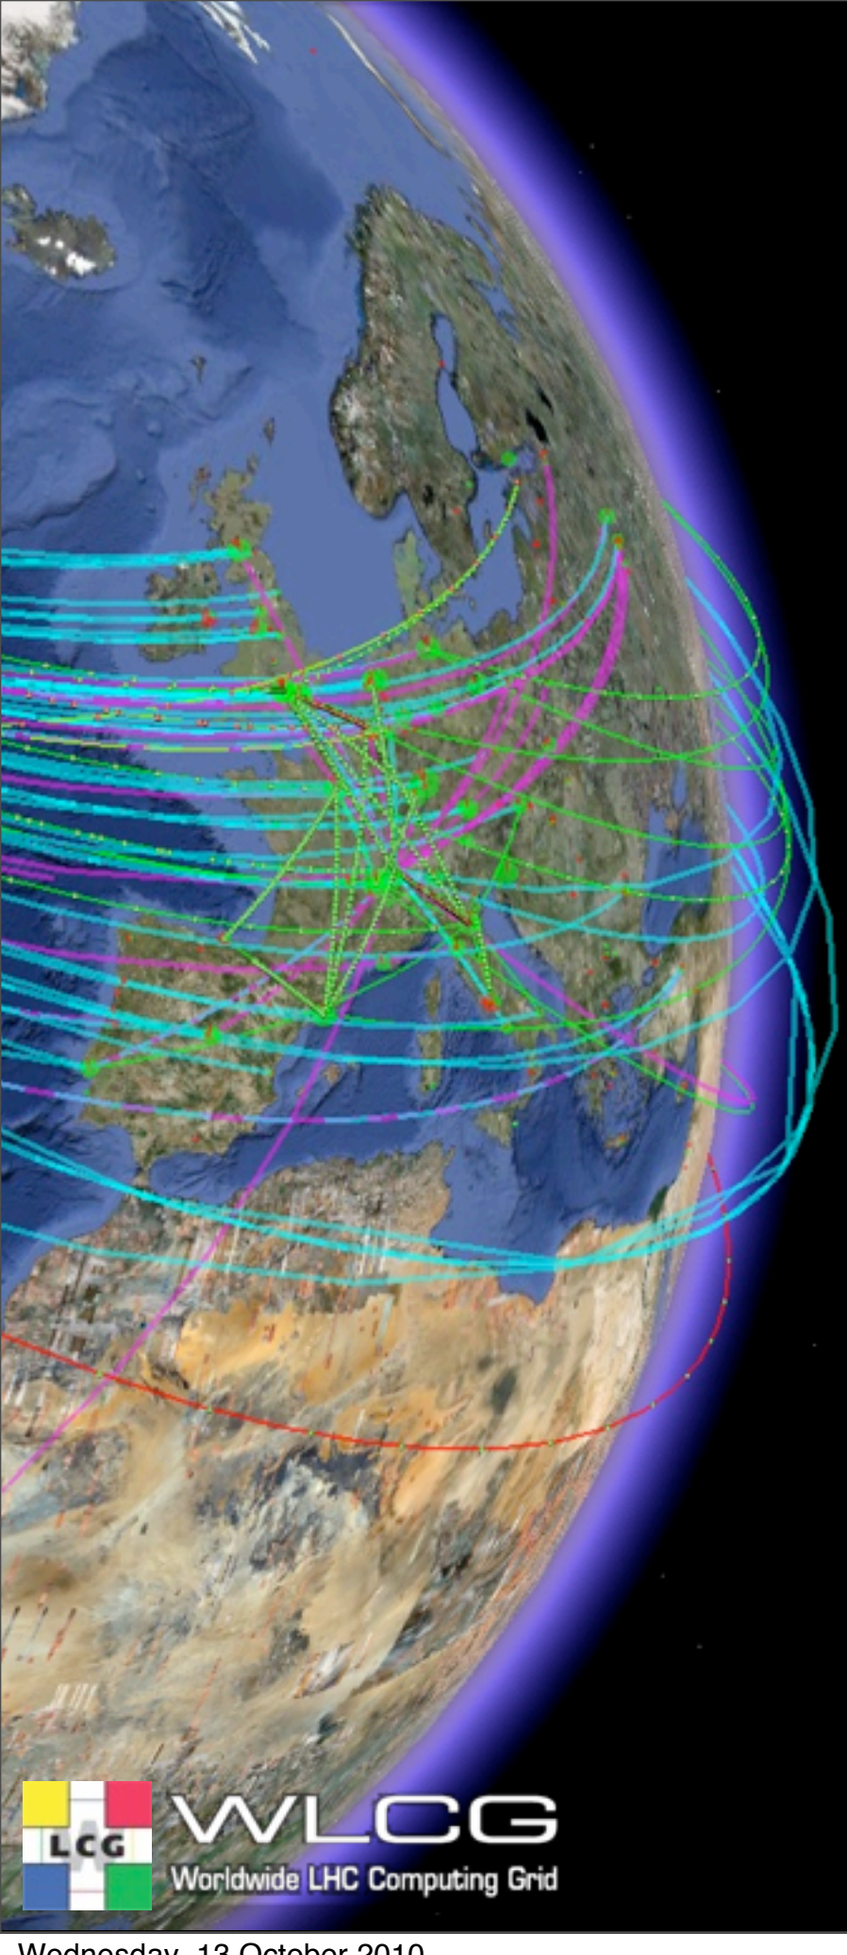

Grid Usage

WLCG
Worldwide LHC Computing Grid

EGEE Achievements - Applications

- >270 VOs from several scientific domains
  - Astronomy & Astrophysics
  - Civil Protection
  - Computational Chemistry
  - Comp. Fluid Dynamics
  - Computer Science/Tools
  - Condensed Matter Physics
  - Earth Sciences
  - Fusion
  - High Energy Physics
  - Life Sciences
- Further applications under evaluation

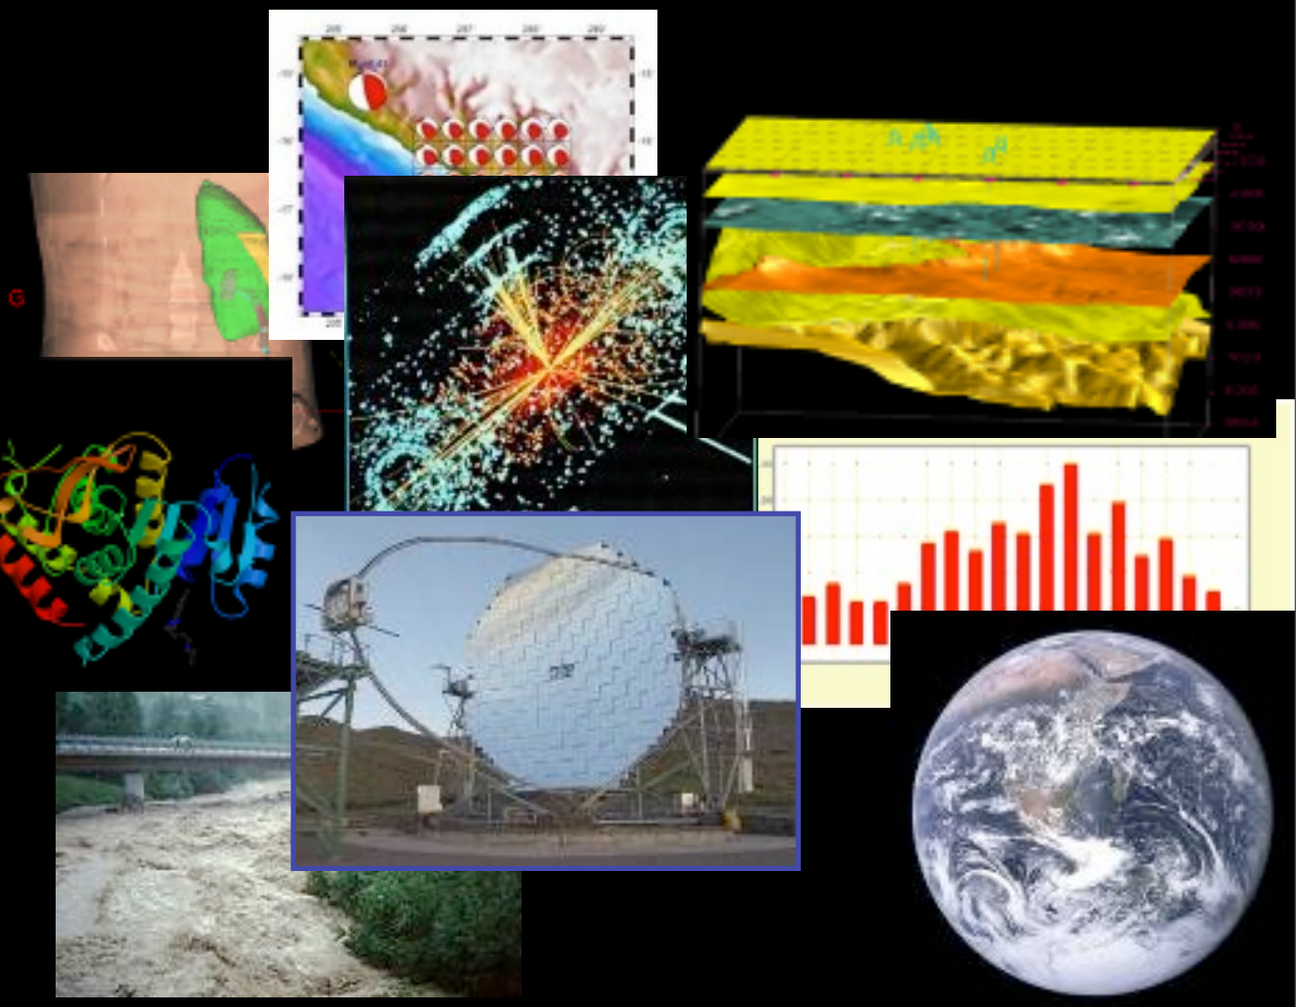

Applications have moved from testing to routine and daily usage

EGEE Registered Collaborating Projects

web page

Infrastructures

geographical or thematic coverage

Applications

improved services for academia, industry and the public

Support Actions

key complementary functions

Global Multi Science Infrastructure mission critical for many communities

Enabling Grids for E-scien

Open Science Grid

NDGF
Nordic DataGrid Facility

Arche
Astro
Astro
Civil P
Comp
Earth
Finan
Fusio
Geopl
High I
Life S
Multir
Mater
…

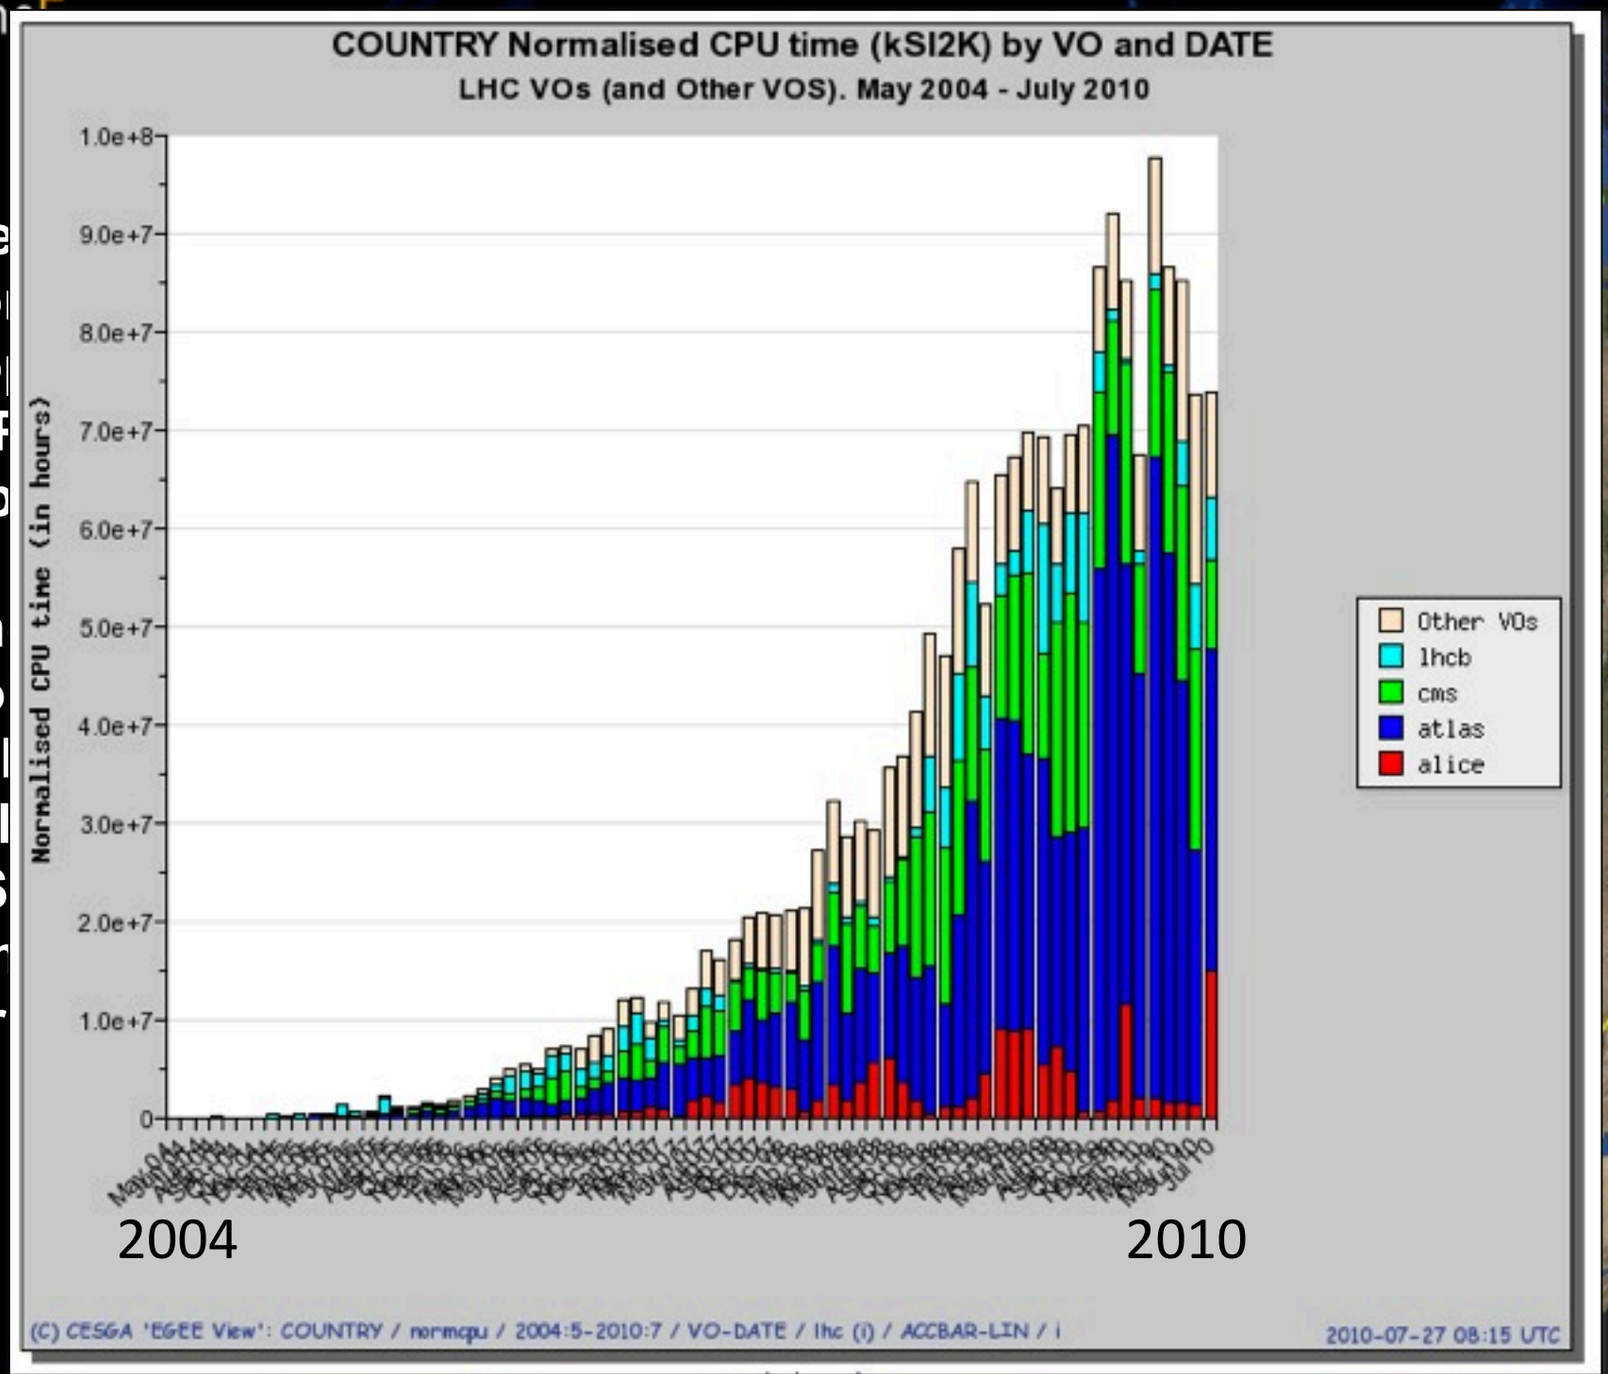

>340sites
48 countries
>170,000 CPUs
>25 PetaBytes disk
>10,000 users
>170 communities
>500,000 jobs/day

GridPP
UK Computing for Particle Physics

21:13:50 UTC

CERN, IT Department

Grid Activity – delivered CPU

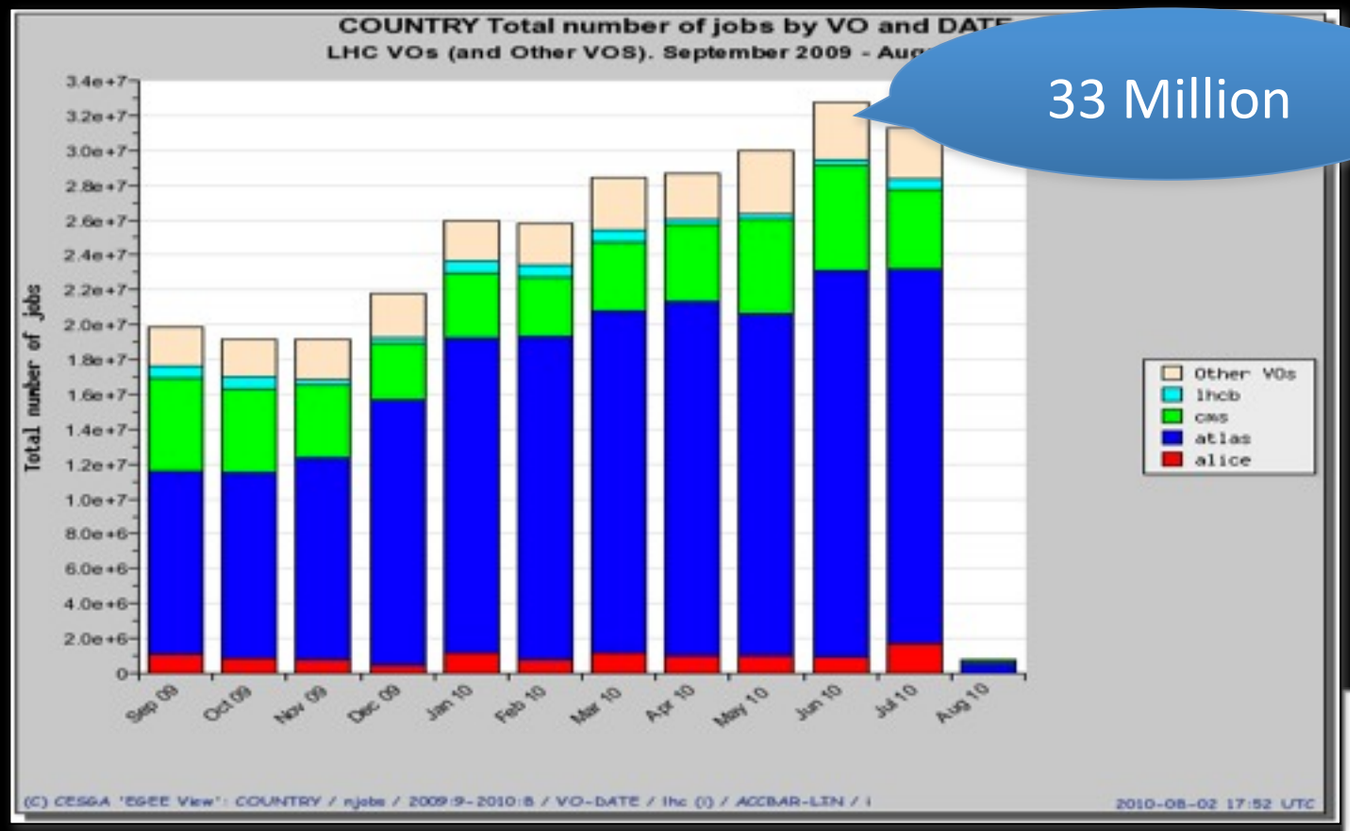

33 Million

- Distribution of work across Tier0/ Tier1/Tier 2 really illustrates the importance of the grid system
  - Tier 2 contribution is ~ 50%;
  - >85% is external to CERN

Tier 2 sites

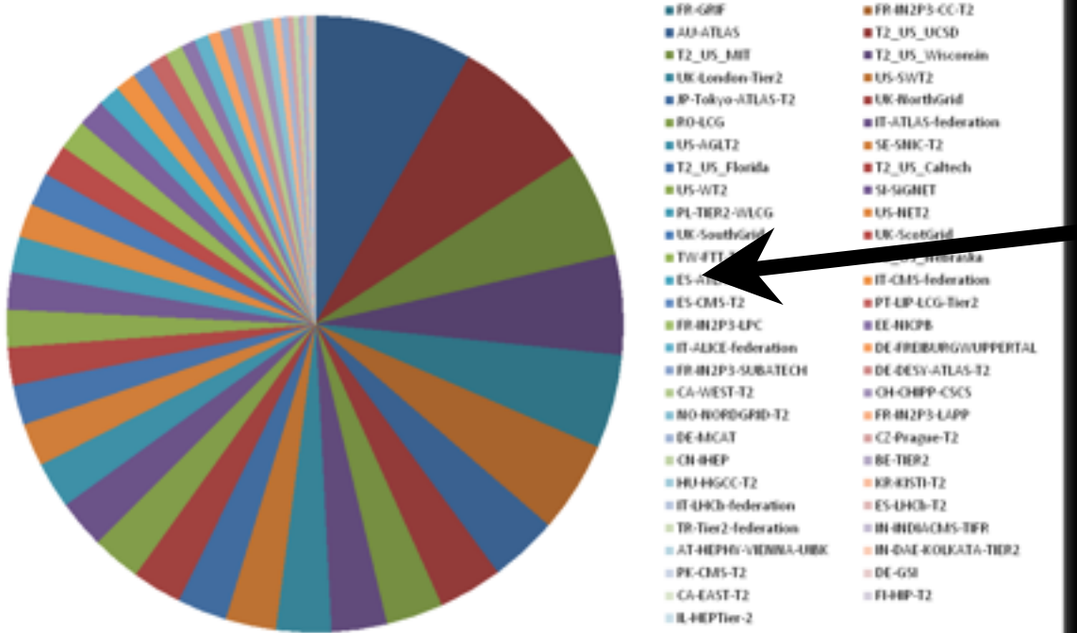

Tier 0 + Tier 1 sites

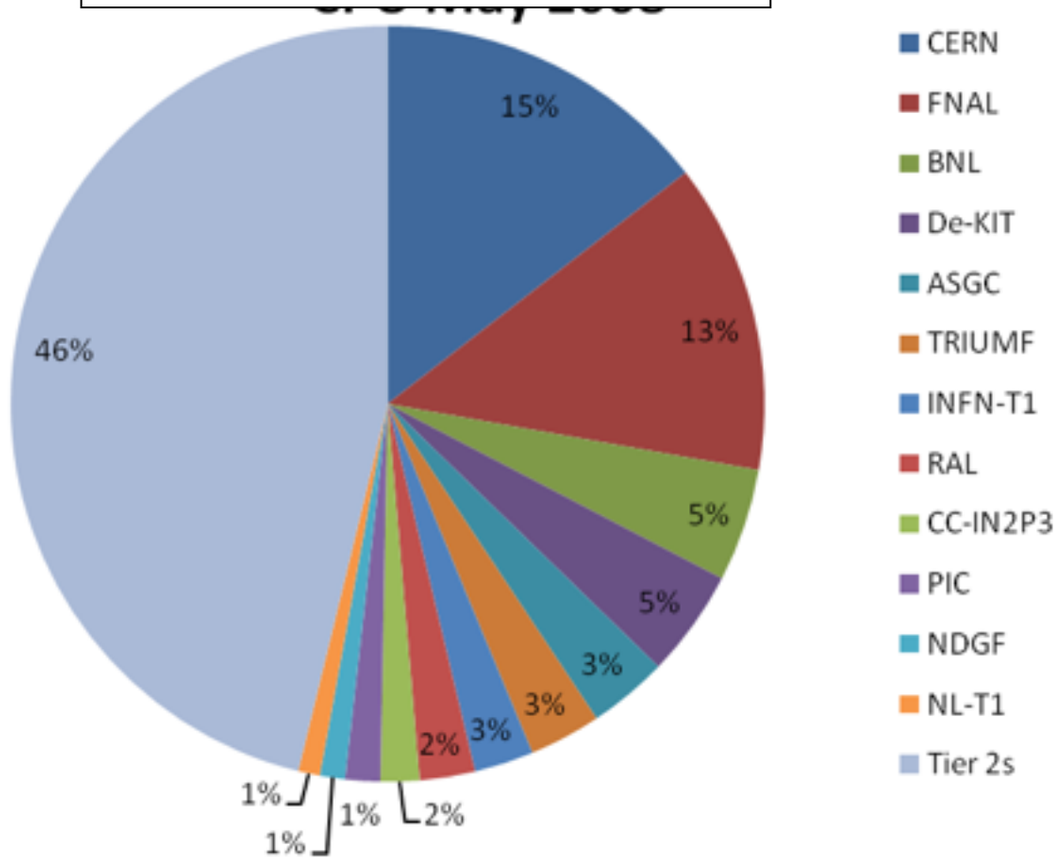

Data Transfers

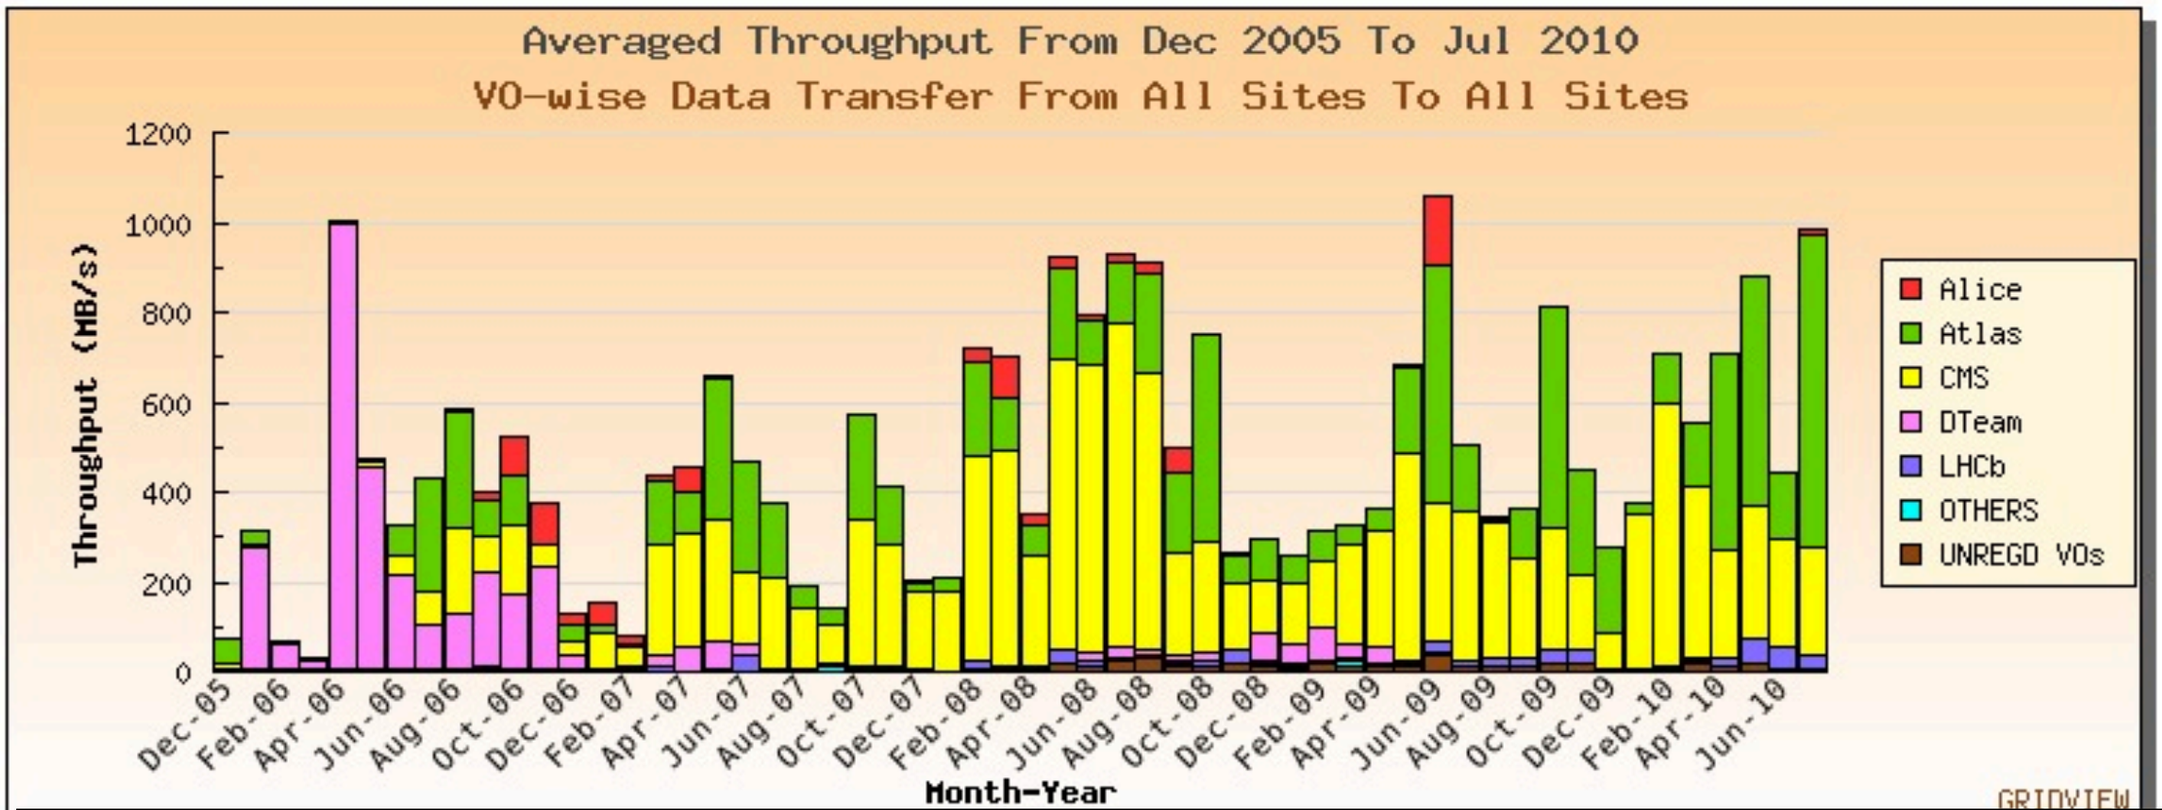

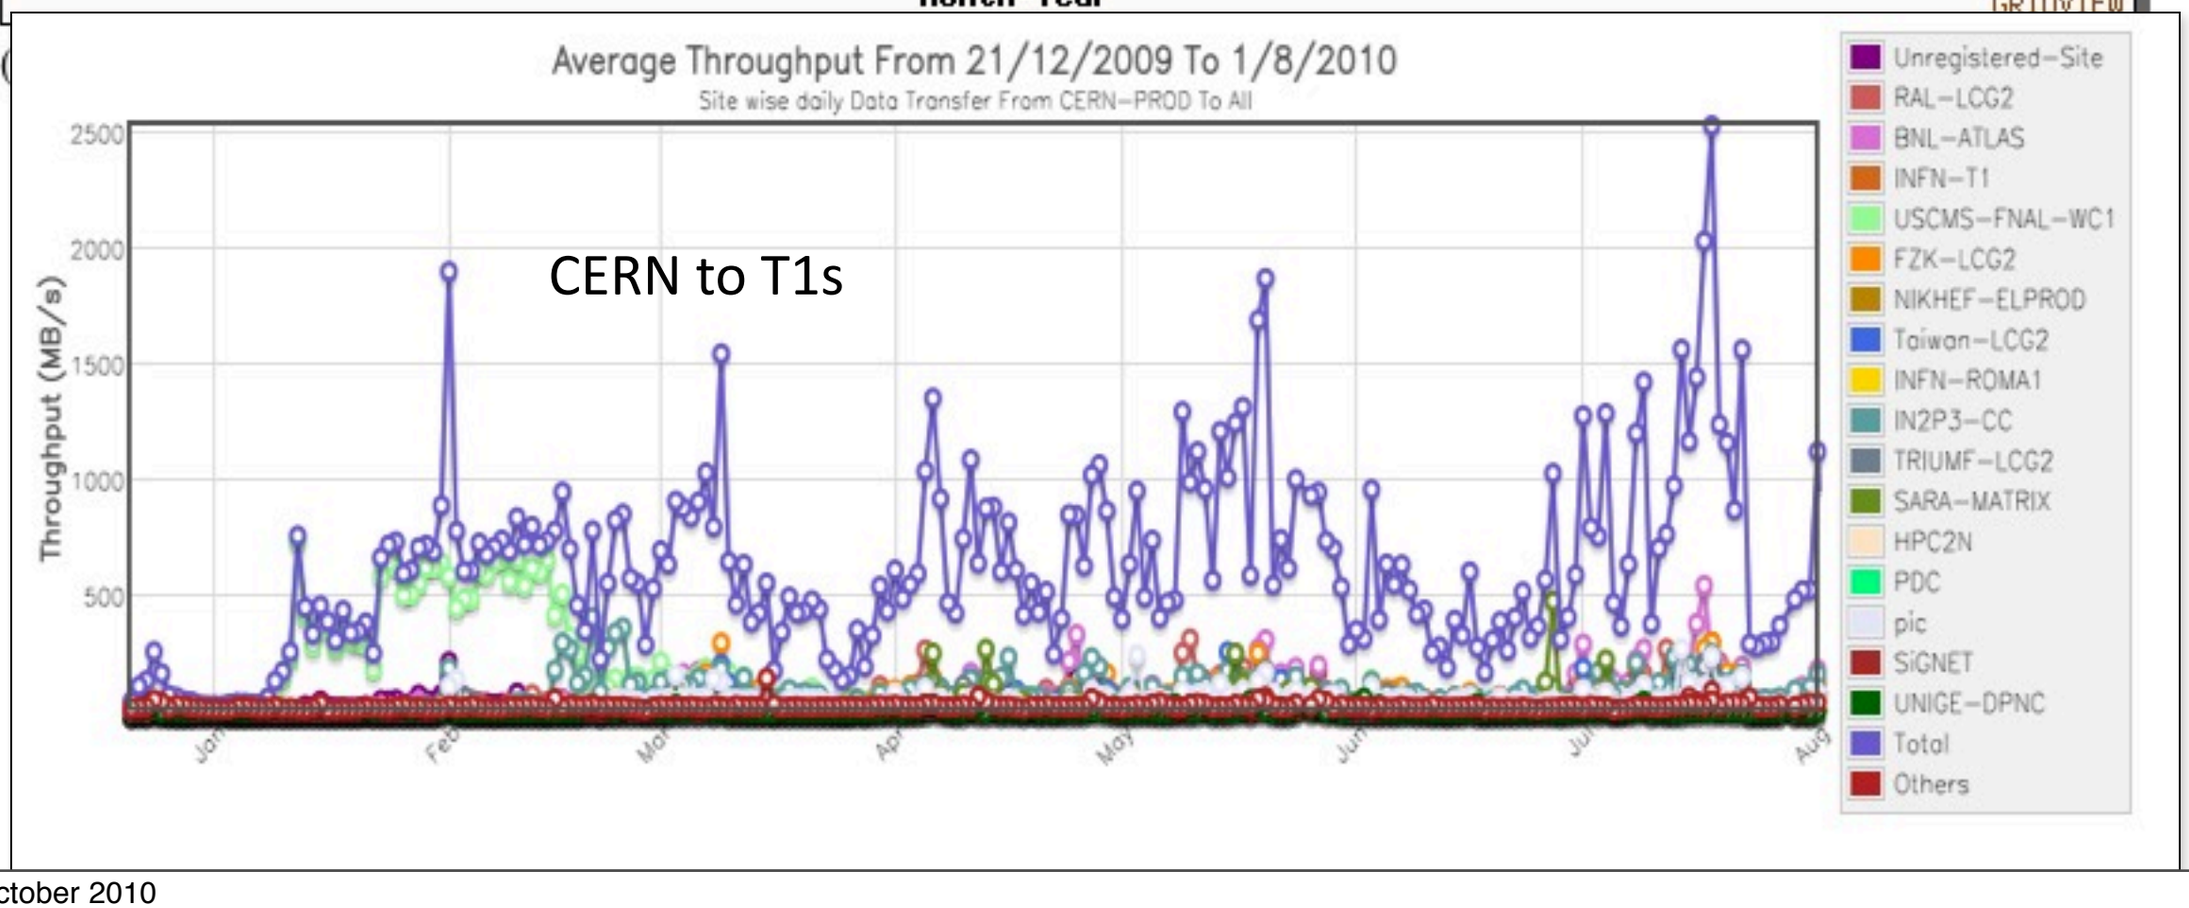

What is Grid Middleware?

- For today:
- The software that allows to build a system that full fills the above requirements

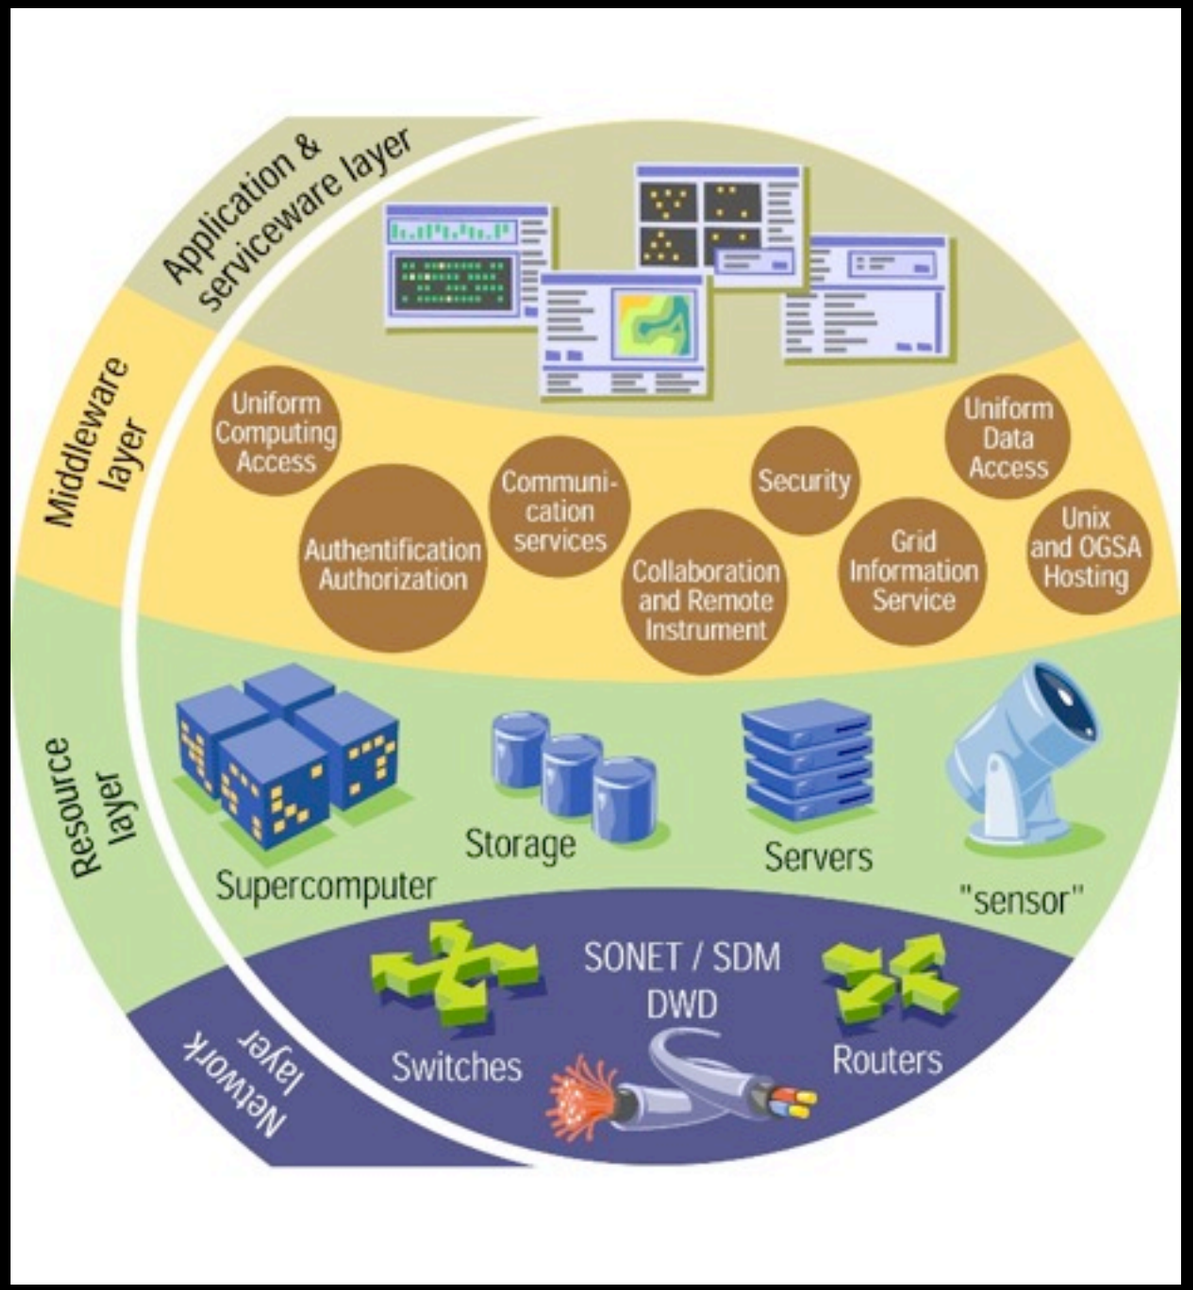

Software Approach

- Identify an AAA system that all can agree on
  - Authentication, Authorization, Auditing
  - –That doesn't require local user registration
  - –That delegates "details" to the users (Virtual Organizations)
- Define and implement abstraction layers for resources
  - –Computing, Storage, etc.
- Define and implement a way to announce your resources
- Build high level services to optimize the usage
- Interface your applications to the system

gLite as an example

gLite Middleware Distribution

- Combines components from different providers
  - Condor and Globus 2 (via VDT)
  - LCG
  - EDG/EGEE
  - Others
- After prototyping phases in 2004 and 2005 convergence with LCG-2 distribution reached in May 2006
  - gLite 3.0
- Focus on providing a deployable MW distribution for EGEE production service

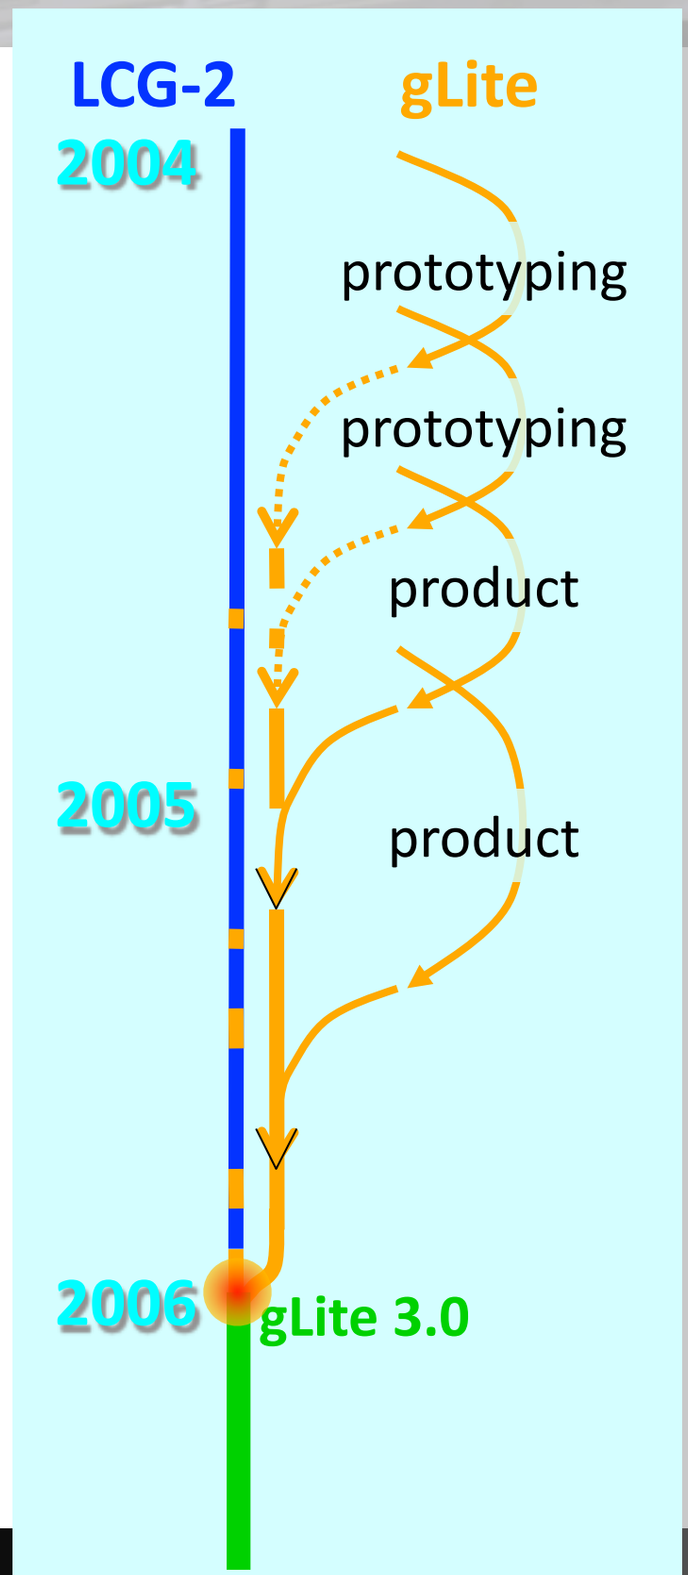

gLite
Lightweight Middleware for Grid Computing

gLite middleware

Data Services

- Storage Element
- File and Replica Catalog
- Metadata Catalog

Job Management Services

- Computing Element
- Worker Node
- Workload Management
- Job Provenance

Security Services

- Authorization
- Authentication

Information & Monitoring Services

- Information System
- Job Monitoring
- Accounting

Access Services

- User Interface
- API

Computing Access

- Computing Elements (CE)
  - gateways to farms

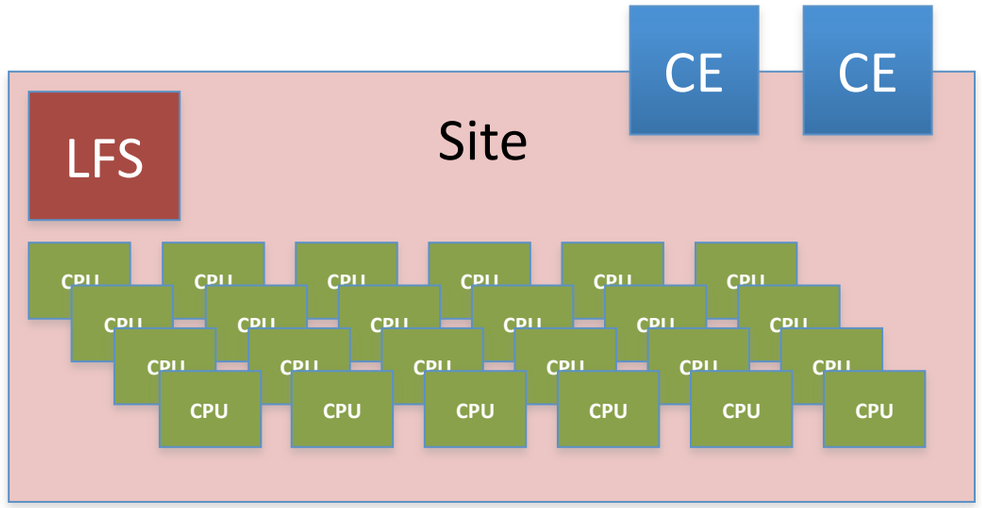

- EGEE:
  - LCG-CE (450 instances)
    - Minor work on stabilization/scalability (50u/4KJ) , bug fixes
    - LEGACY
  - CREAM-CE (120 instances (up from 26))
    - Significant investment on production readiness and scalability
    - Handles direct submission (pilot job friendly)
    - SL4/SL5
    - BES standard compliant, parameter passing from grid <-> batch

Workload Management

- EGEE WMS/LB
  - Matches resources and requests
    - Including data location
  - Handles failures (resubmission)
  - Manages complex workflows
  - Tracks job status

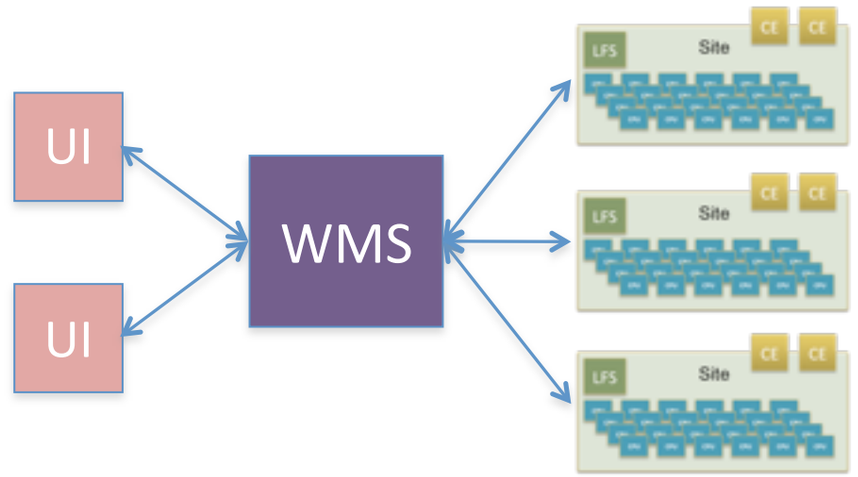

- EGEE WMS/LB (124 Instances)
  - Fully supports LCG-CE and CREAM-CE
    - Early versions had some WMS<->CREAM incompatibilities

MultiUserPilotJobs

- Idea: Matching resources and jobs by Vos
- Pilot is a placeholder for the real job
- Identity is changed on demand on the WN

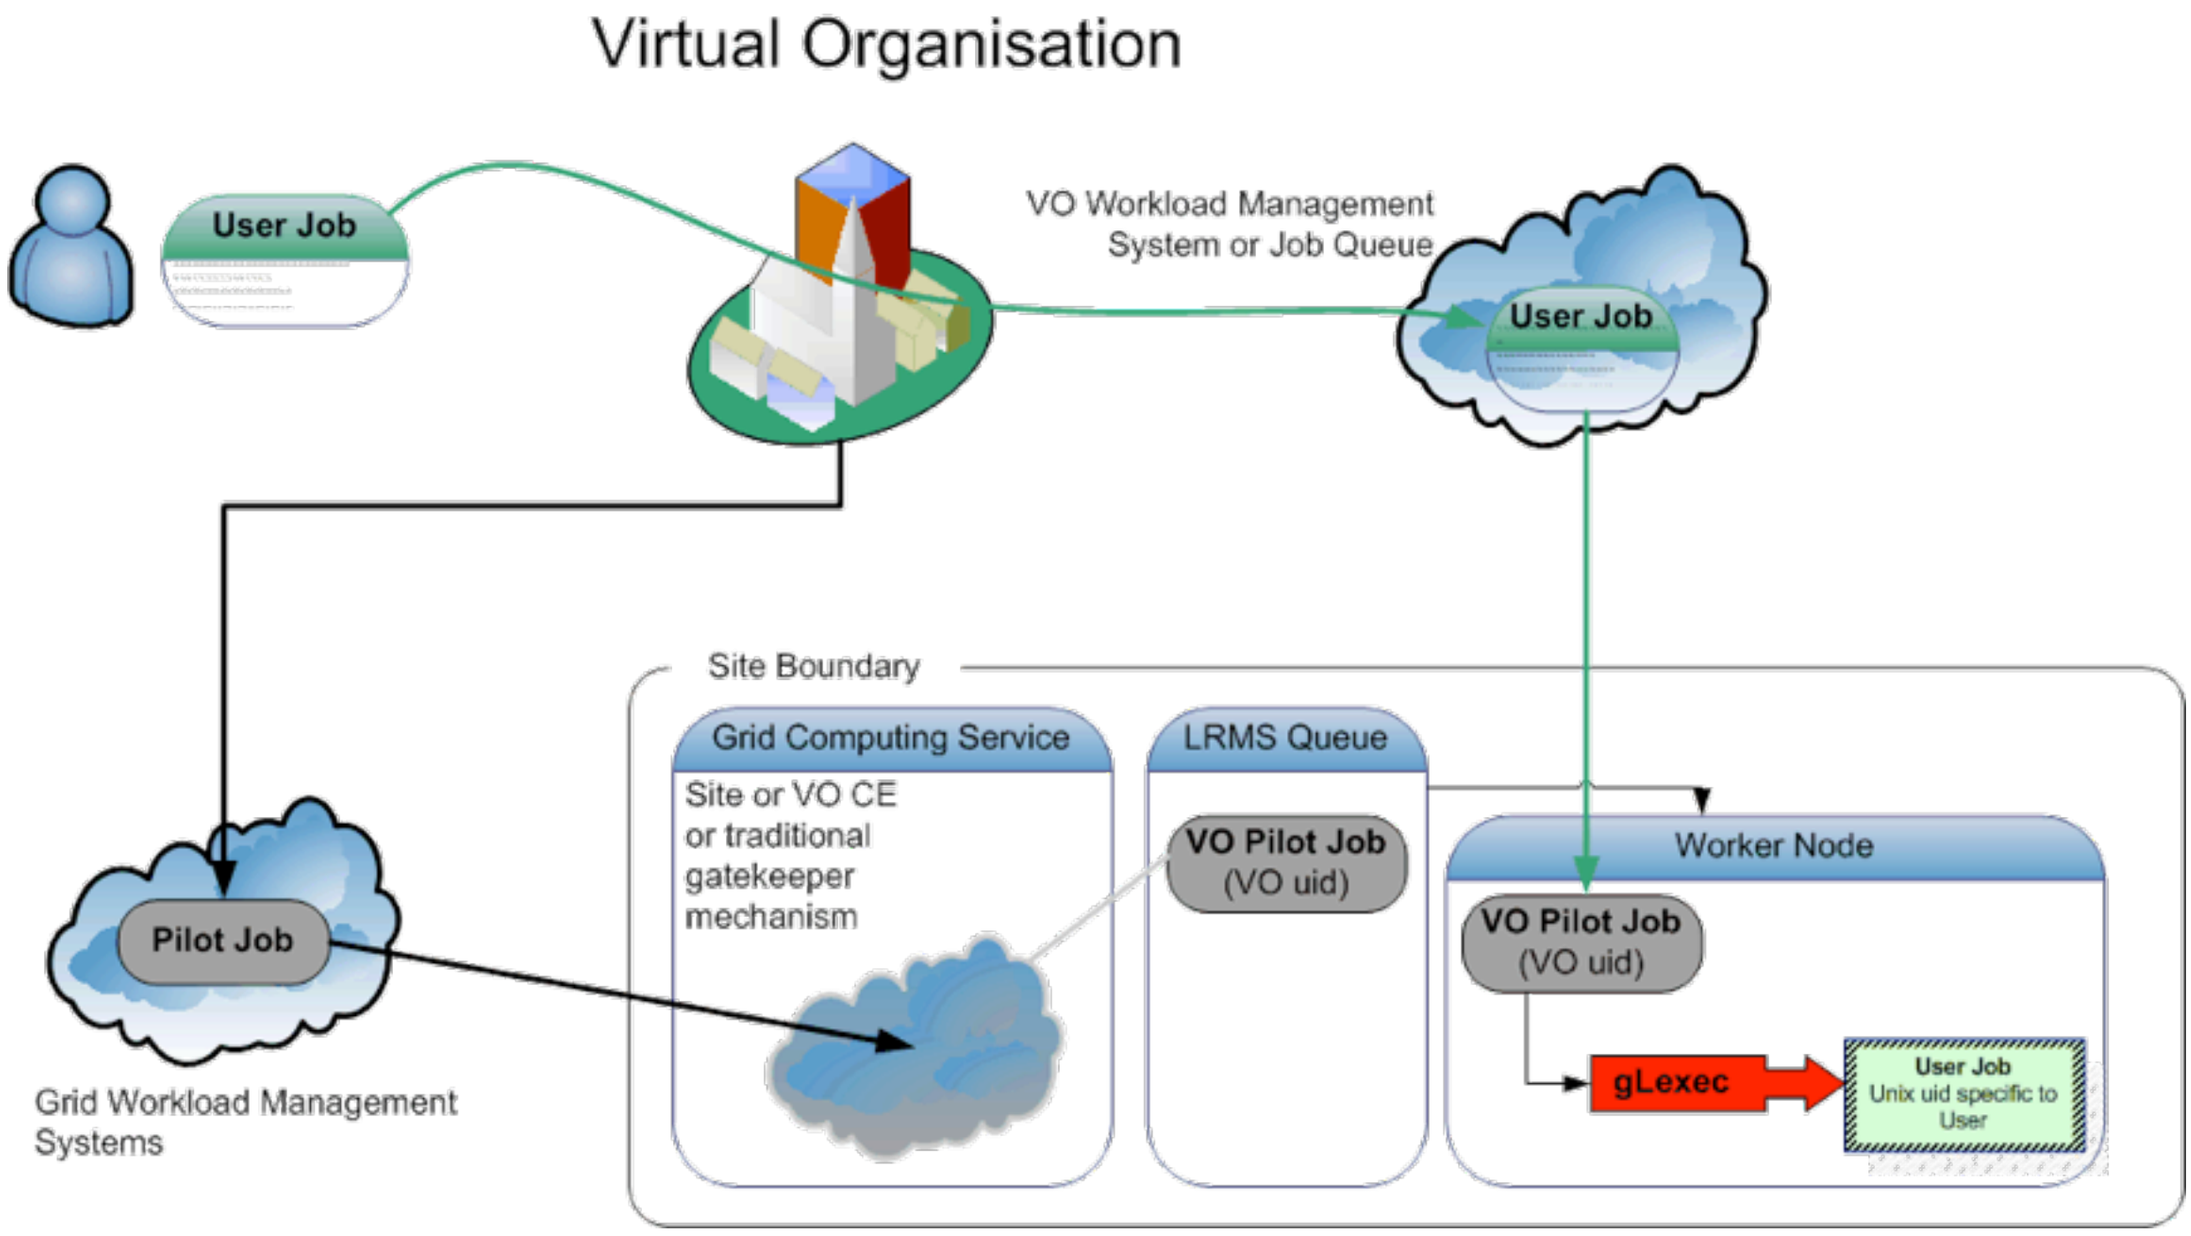

Data Management

- Storage Elements (SEs)
  - External interfaces based on SRM 2.2 and gridFTP
  - Local interfaces: POSIX, dcap, secure rfio, rfio, xrootd
  - DPM (241)
  - dCache (82)
  - STORM (40)
  - BestMan (26)
  - CASTOR (19)
  - "ClassicSE" (27) → legacy since 2 years….

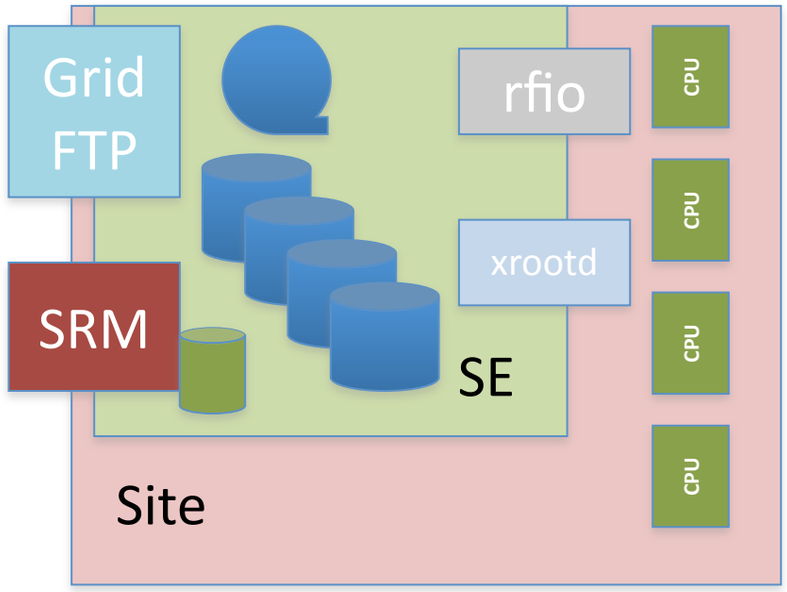

- Catalogue: LFC (local and global)
- File Transfer Service (FTS)
- Data management clients gfal/LCG-Utils

Information System

- BDII
- Light weight Database
- LDAP protocol
- GLUE 1.3 (2) Schema
  - Describes resources and their state
  - Approx 100MB
- Several hundred instances

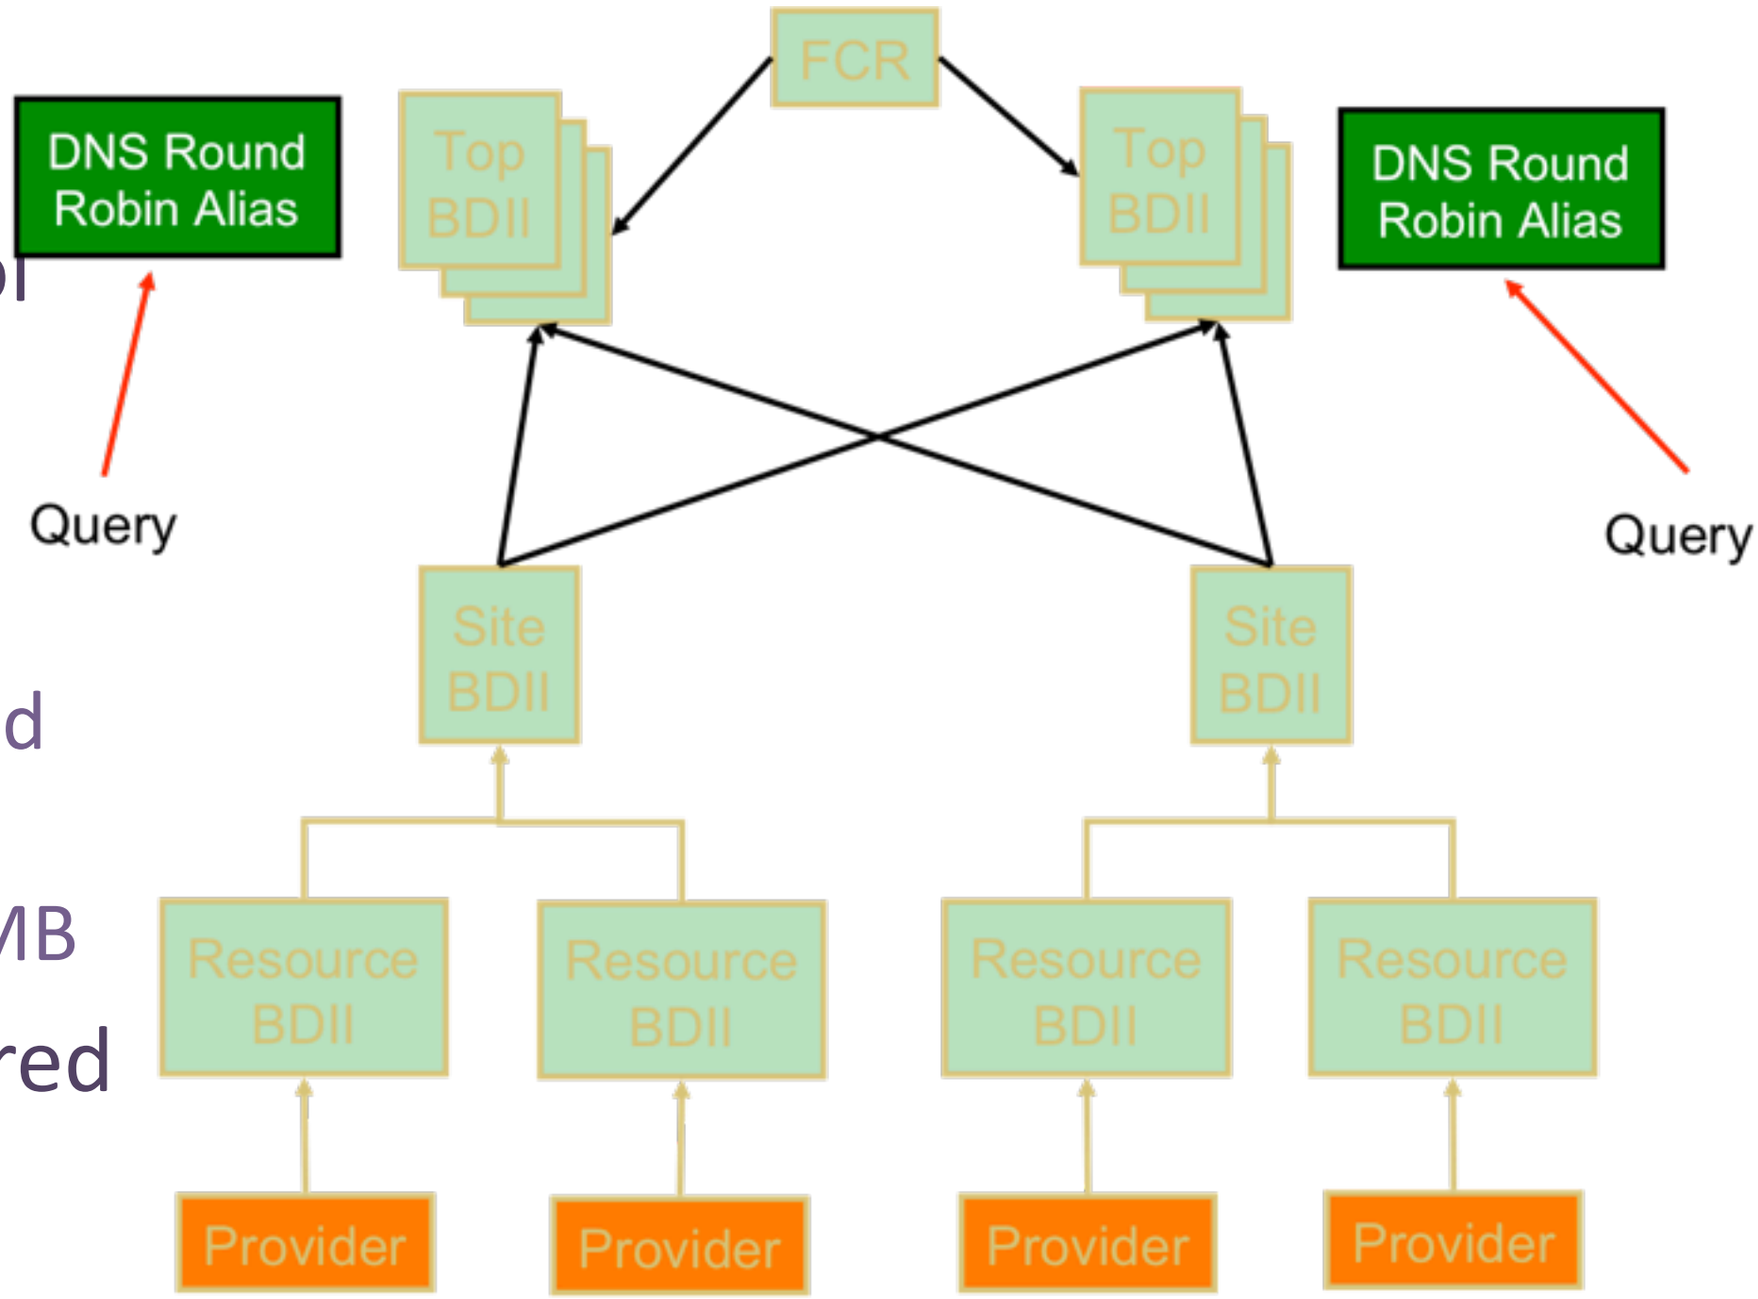

Authentication

- Authentication is based on X.509 PKI infrastructure (Public Key)
  - Certificate Authorities (CA) issue (long lived) certificates identifying individuals (much like a passport)
    - Commonly used in web browsers to authenticate to sites
  - Trust between CAs and sites is established (offline)
  - In order to reduce vulnerability, on the Grid user identification is done by using (short lived) proxies of their certificates
- Short-Lived Credential Services (SLCS)
  - issue short lived certificates or proxies to its local users
    - e.g. from Kerberos or from Shibboleth credentials
- Proxies can
  - Be delegated to a service such that it can act on the user's behalf
  - Be stored in an external proxy store (MyProxy)
  - Be renewed (in case they are about to expire)
  - Include additional attributes -→ Authorization

Public Key Based Security

- How to exchange secret keys?
  - –340 Sites (global)
    - With hundreds of nodes each?
  - –200 User Communities (non local)
  - –10000 Users (global)
- And keep them secret!!!

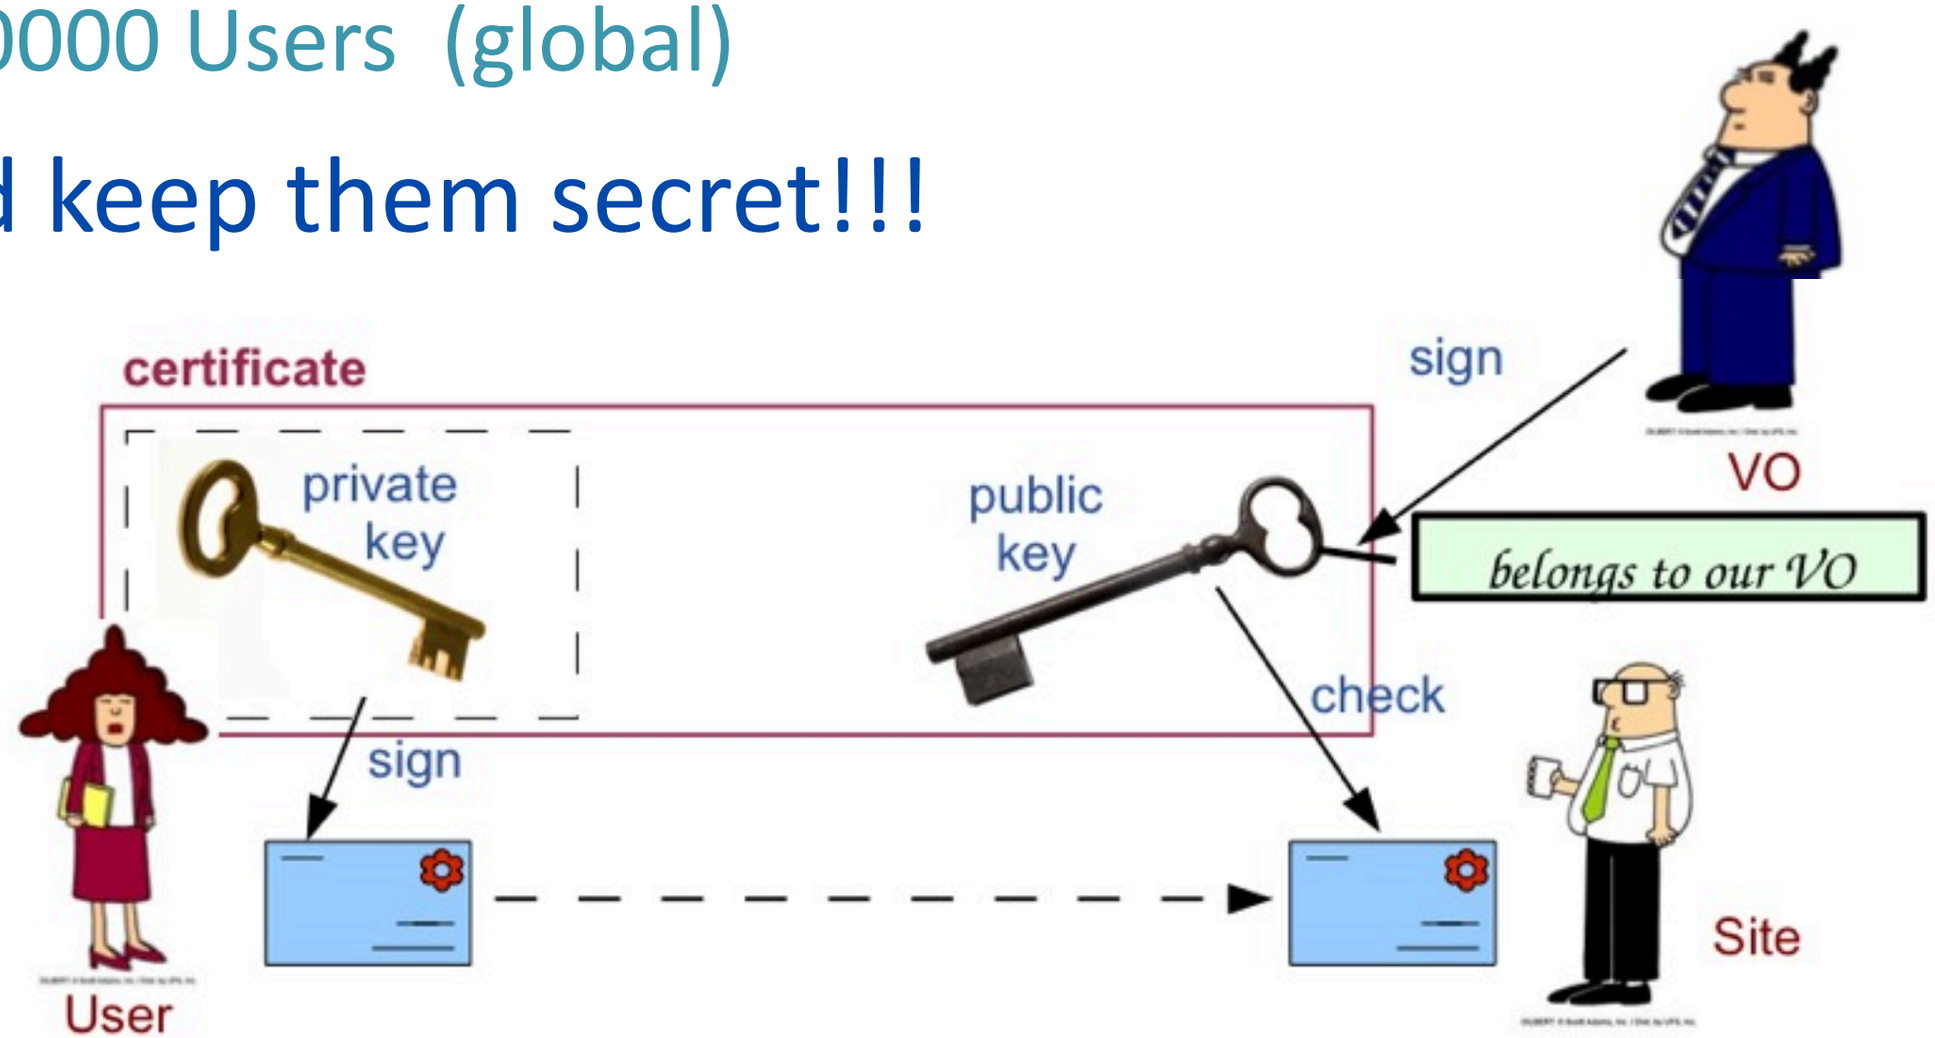

Security - overview

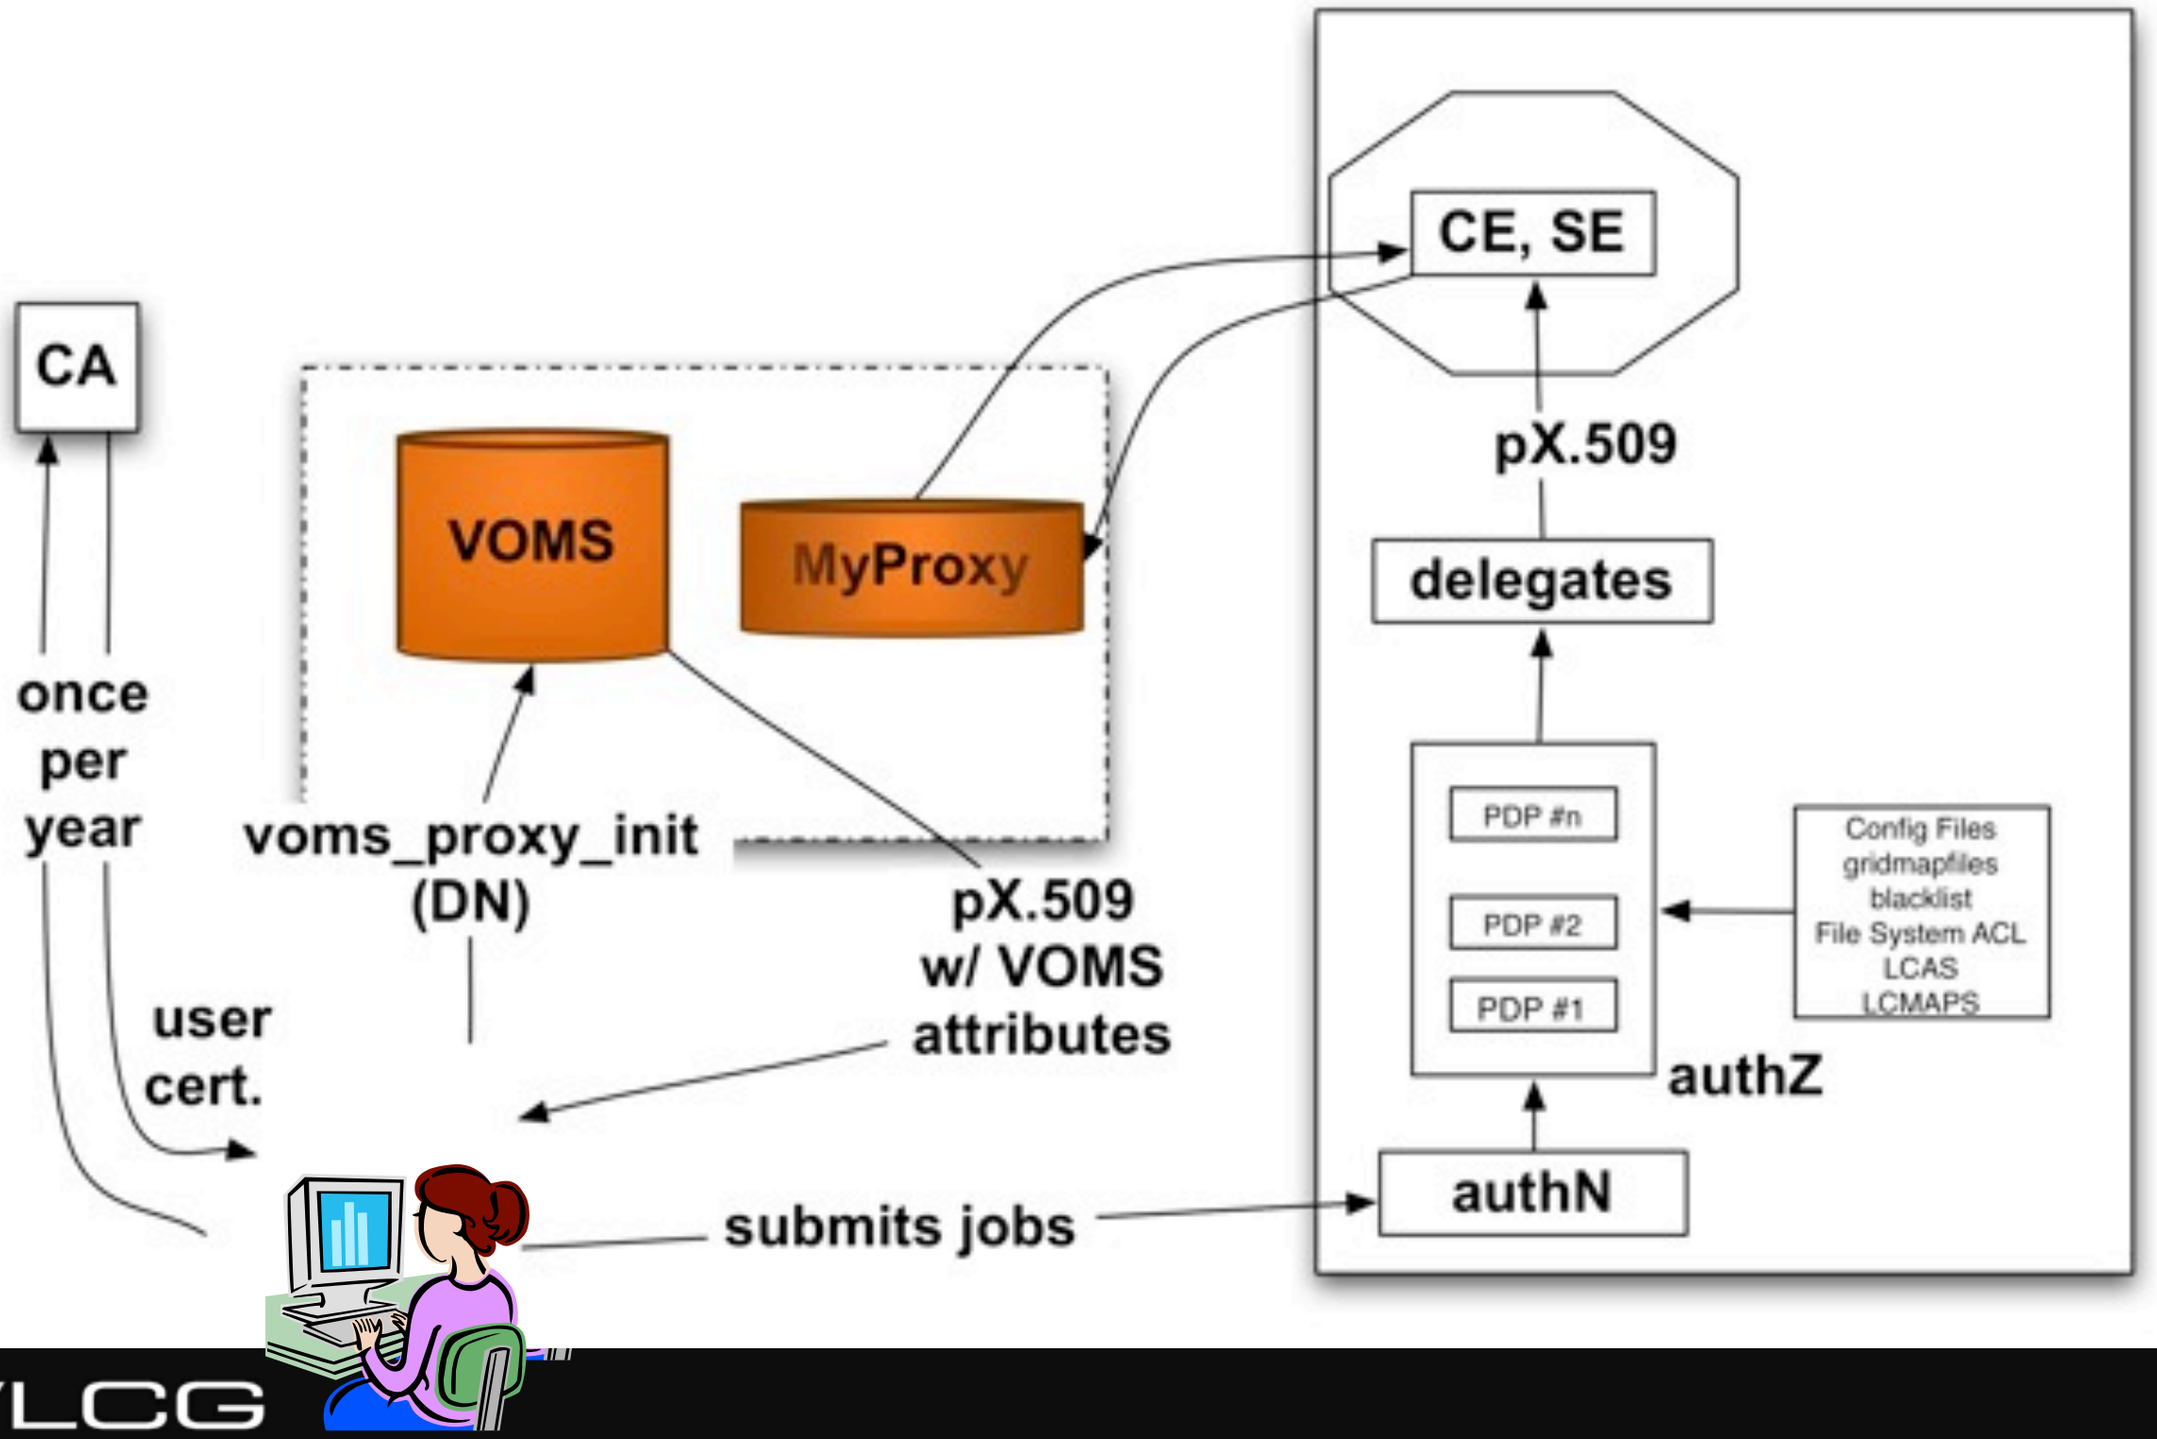

Authorization

- VOMS is now a de-facto standard
  - Attribute Certificates provide users with additional capabilities defined by the VO.
  - Basis for the authorization process
- Authorization: currently via mapping to a local user on the resource
  - glexec changes the local identity (based on suexec from Apache)
- Designing an authorization service with a common interface agreed with multiple partners
  - Uniform implementation of authorization in gLite services
  - Easier interoperability with other infrastructures
  - Prototype being prepared now

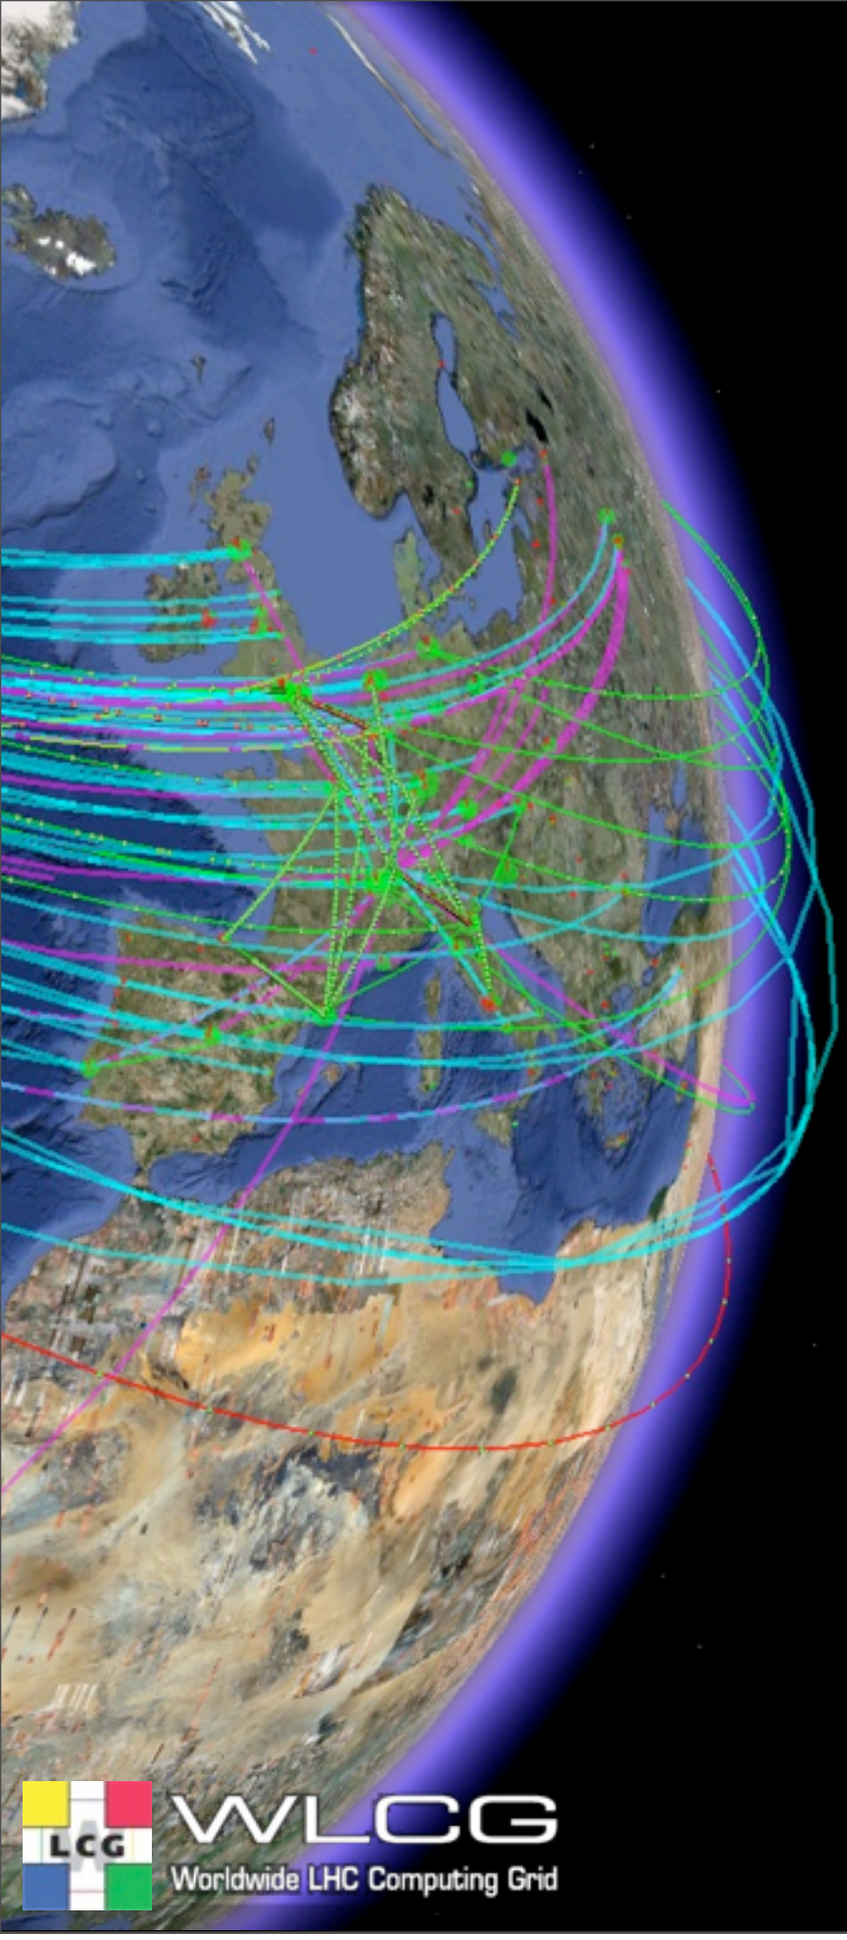

Impossible to discuss all components

- Illustrate complexity
- Some Data Management Details

Workload Management (compact)

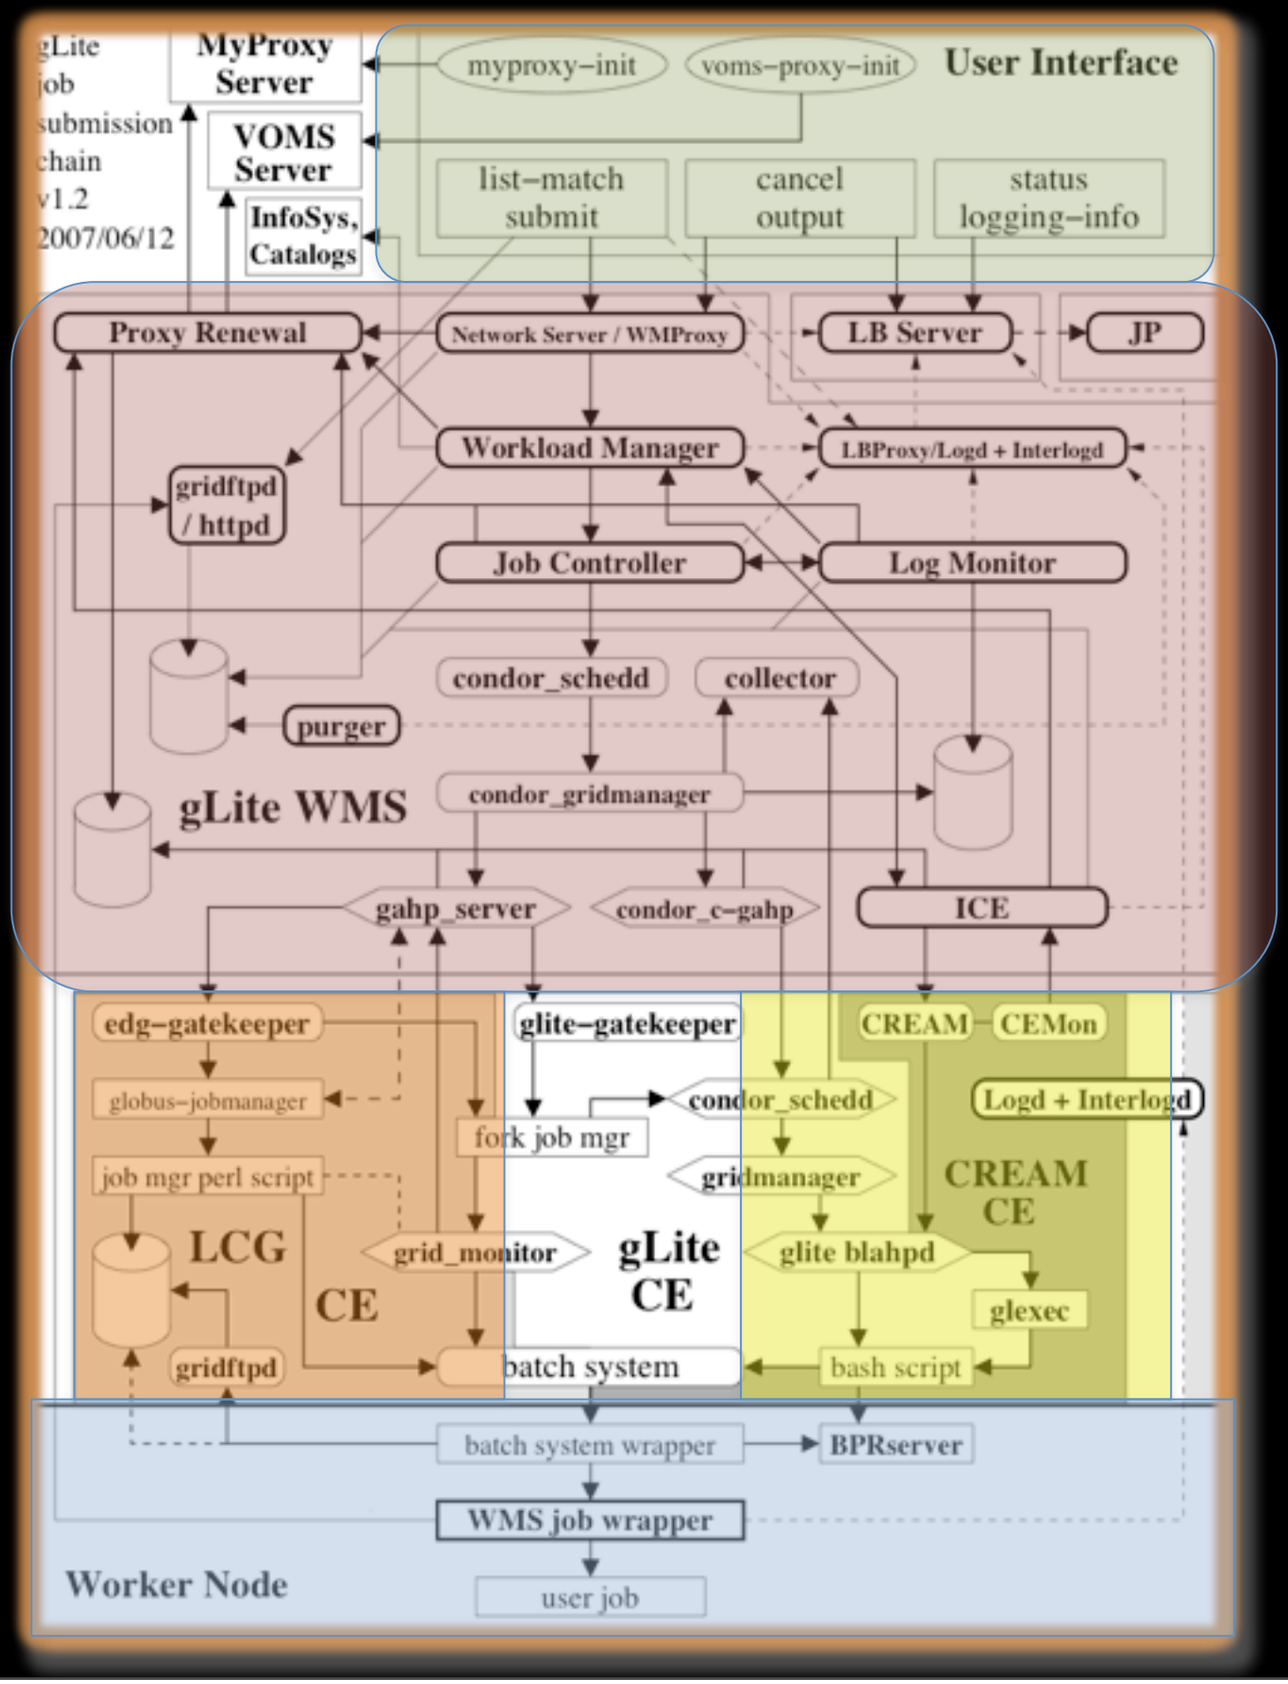

Desktops

A few
~50 nodes

1-20 per site

1-24000
per site

Job Description Language

```
[
  Executable = "my_exe";
  StdOutput = "out";
  StdError = "err";
  Arguments = "a b c";
  InputSandbox = {"/home/giaco/my_exe"};
  OutputSandbox = {"out", "err"};
  Requirements = Member(
    other.GlueHostApplicationSoftwareRunTimeEnvironment,
    "ALICE3.07.01"
  );
  Rank = -other.GlueCEStateEstimatedResponseTime;
  RetryCount = 3
]
```

Data Management

VO Frameworks

lcg_utils FTS

GFAL

Vendor Specific APIs

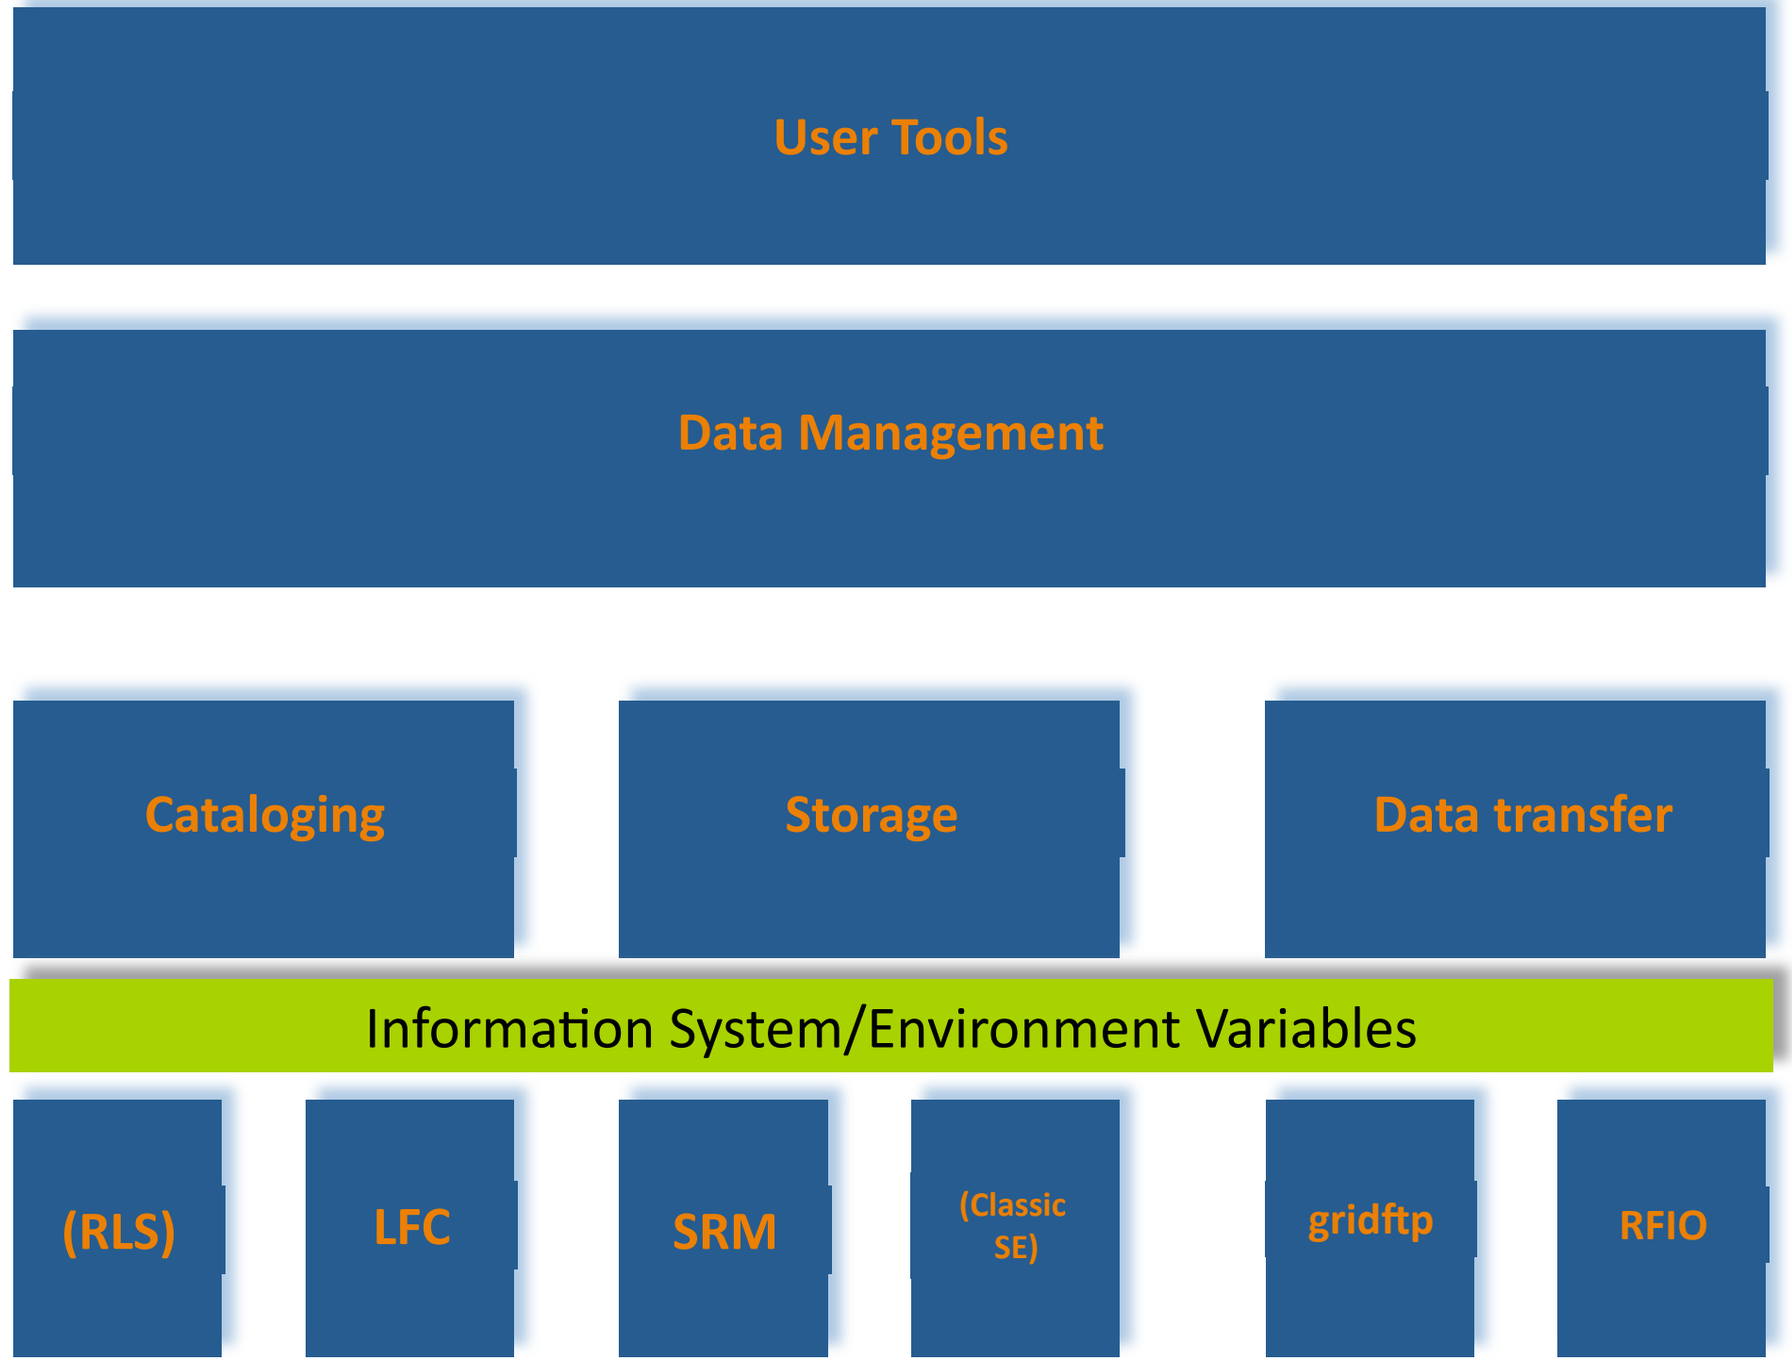

General Storage Element

- Storage Resource Manager (SRM)
  - hides the storage system implementation (disk or active tape)
  - handles authorization
  - translates SURLs (Storage URL) to TURLs (Transfer URLs)
  - disk-based: DPM, dCache,+; tape-based: Castor, dCache
  - Mostly asynchronous
- *File I/O*: *posix-like* access from local nodes or the grid

  ➔GFAL (Grid File Access Layer)

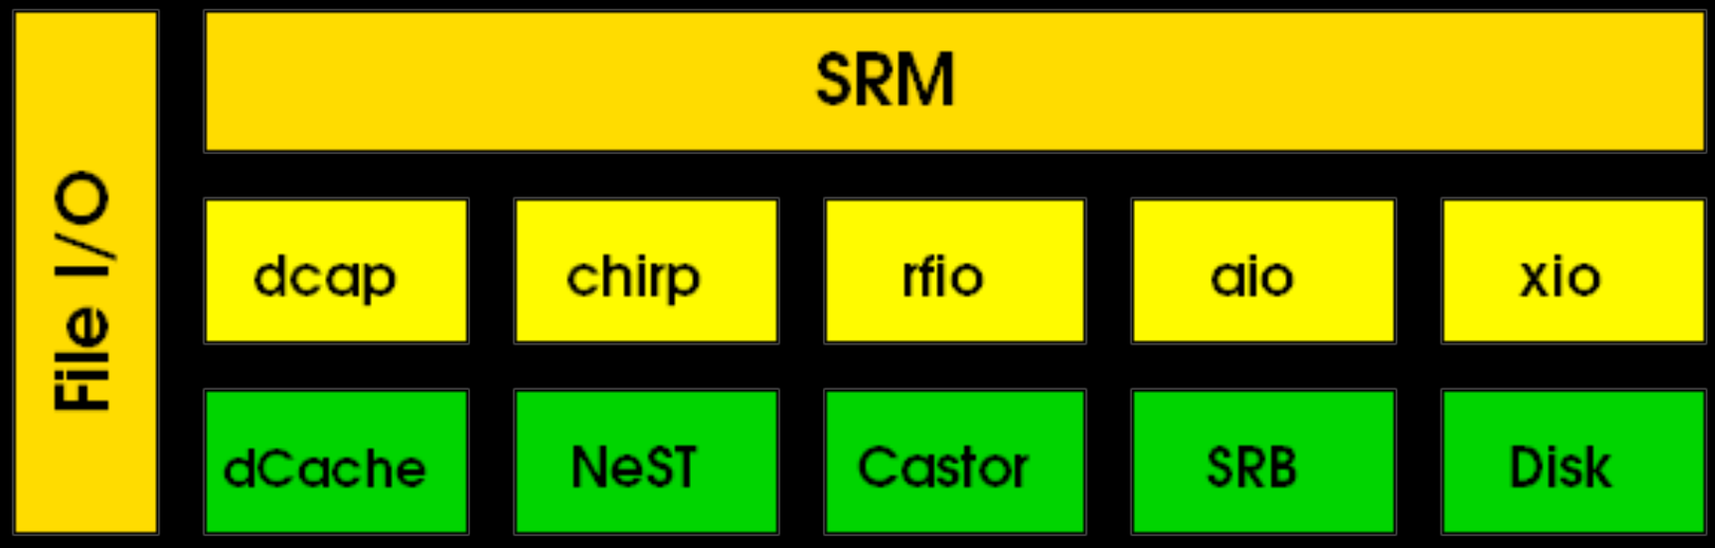

Approach to SRM

- An abstraction layer for storage and data access is necessary
  - Guiding principle:
  - Non-interference with local policies
- Providing all necessary user functionality and control
  - Data Management
  - Data Access
  - Storage management
  - Control:
    - Pinning files
    - Retention Policy
    - Space management and reservation
  - Data Transfers
- Grid enabled and based on current technology
  - Interface technology (gSOAP)
  - Security Model (gsi security)
  - To integrate with the grid infrastructure

Implementation(s)

- WSDL to generate client stubs
- gSOAP
- GSI for A(A)

- External interface is the “easy” part
- Integration with the existing storage system is the challenge
- Different approaches concerning the accepted level of entanglement between SRM and storage
  - New storage systems have less problems
    - Designed for SRM

SRM basic and use cases tests

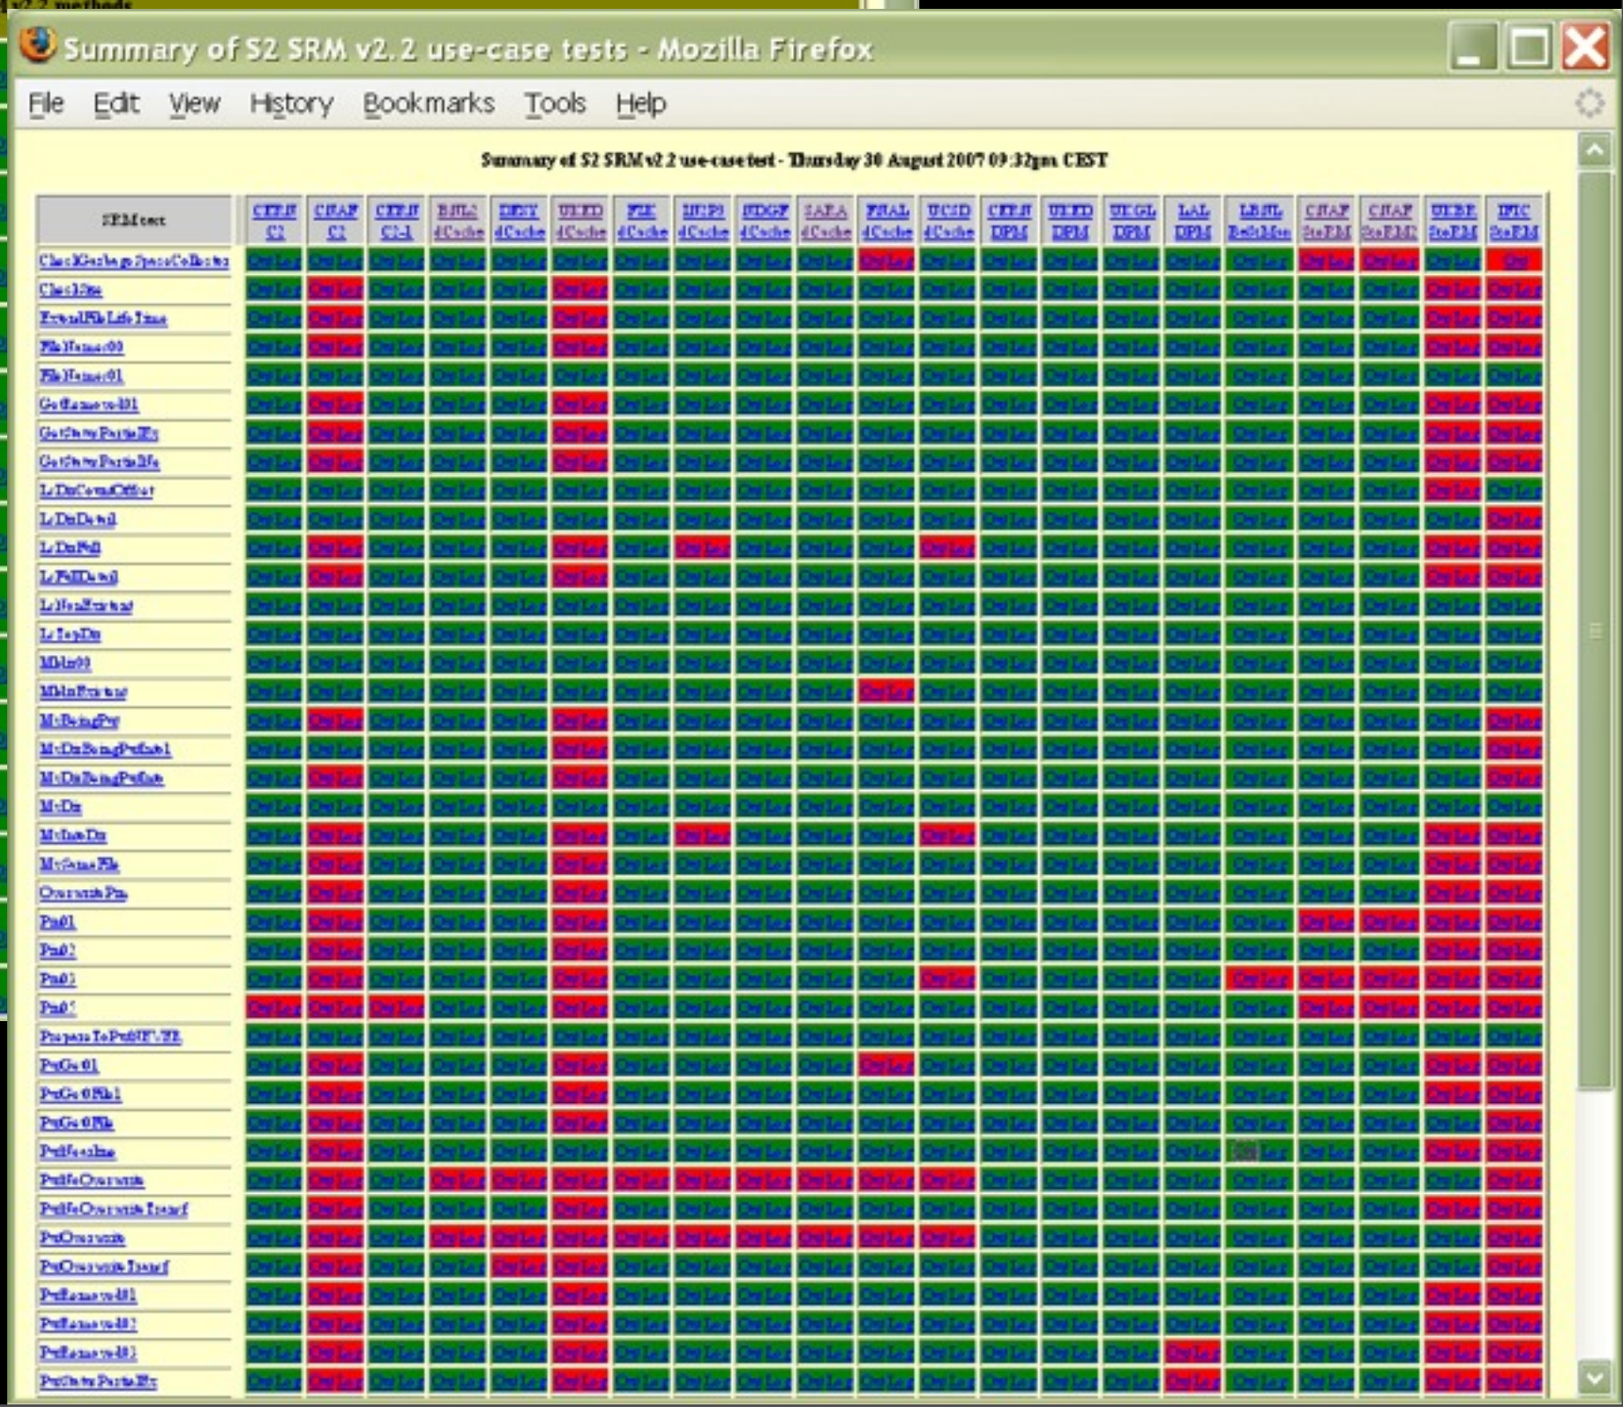

LCG “File” Catalog

- The LFC stores mappings between
  - Users’ file names
  - File locations on the Grid

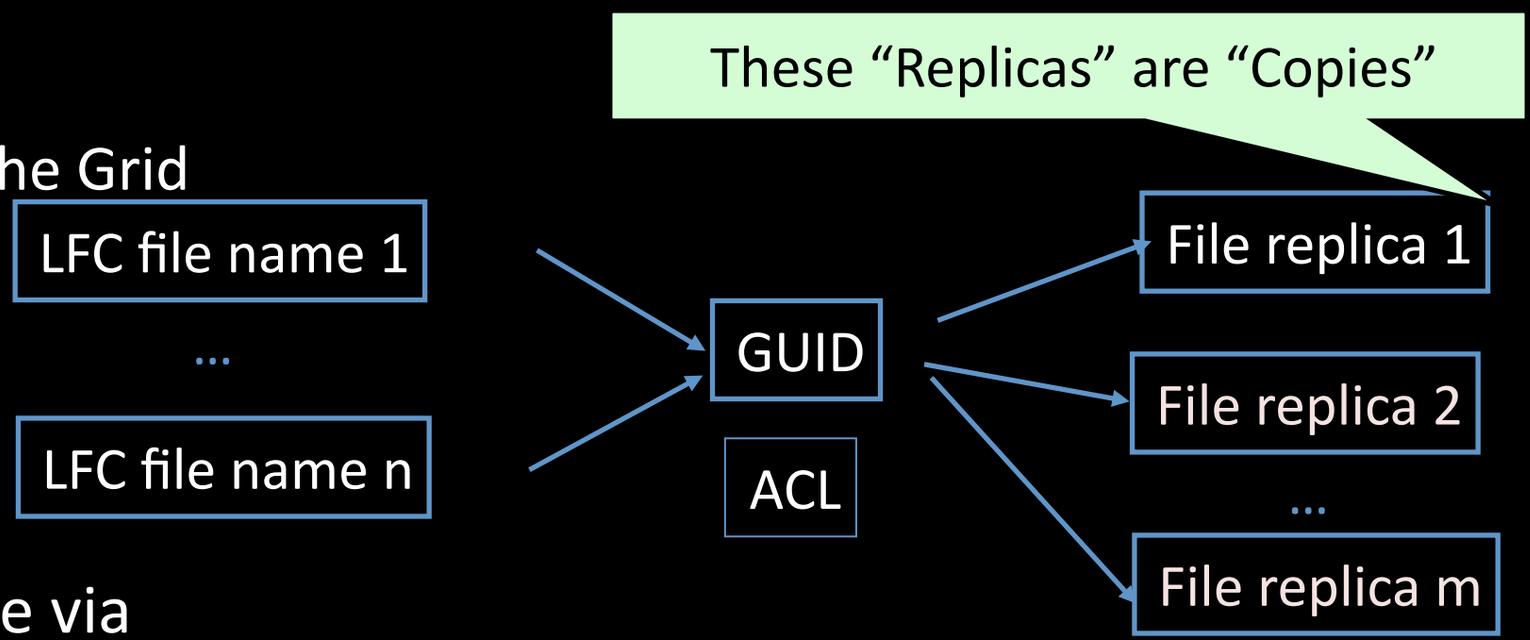

- The LFC is accessible via
  - CLI, C API, Python interface, Perl interface
    - Supports sessions and bulk operations
  - Data Location Interface (DLI)
    - Web Service used for match making:
      - given a GUID, returns physical file location
- ORACLE backend for high performance applications
  - Read-only replication support

All files are “Write Once Read Many”

LFC features

LFC features

Hierarchical Namespace

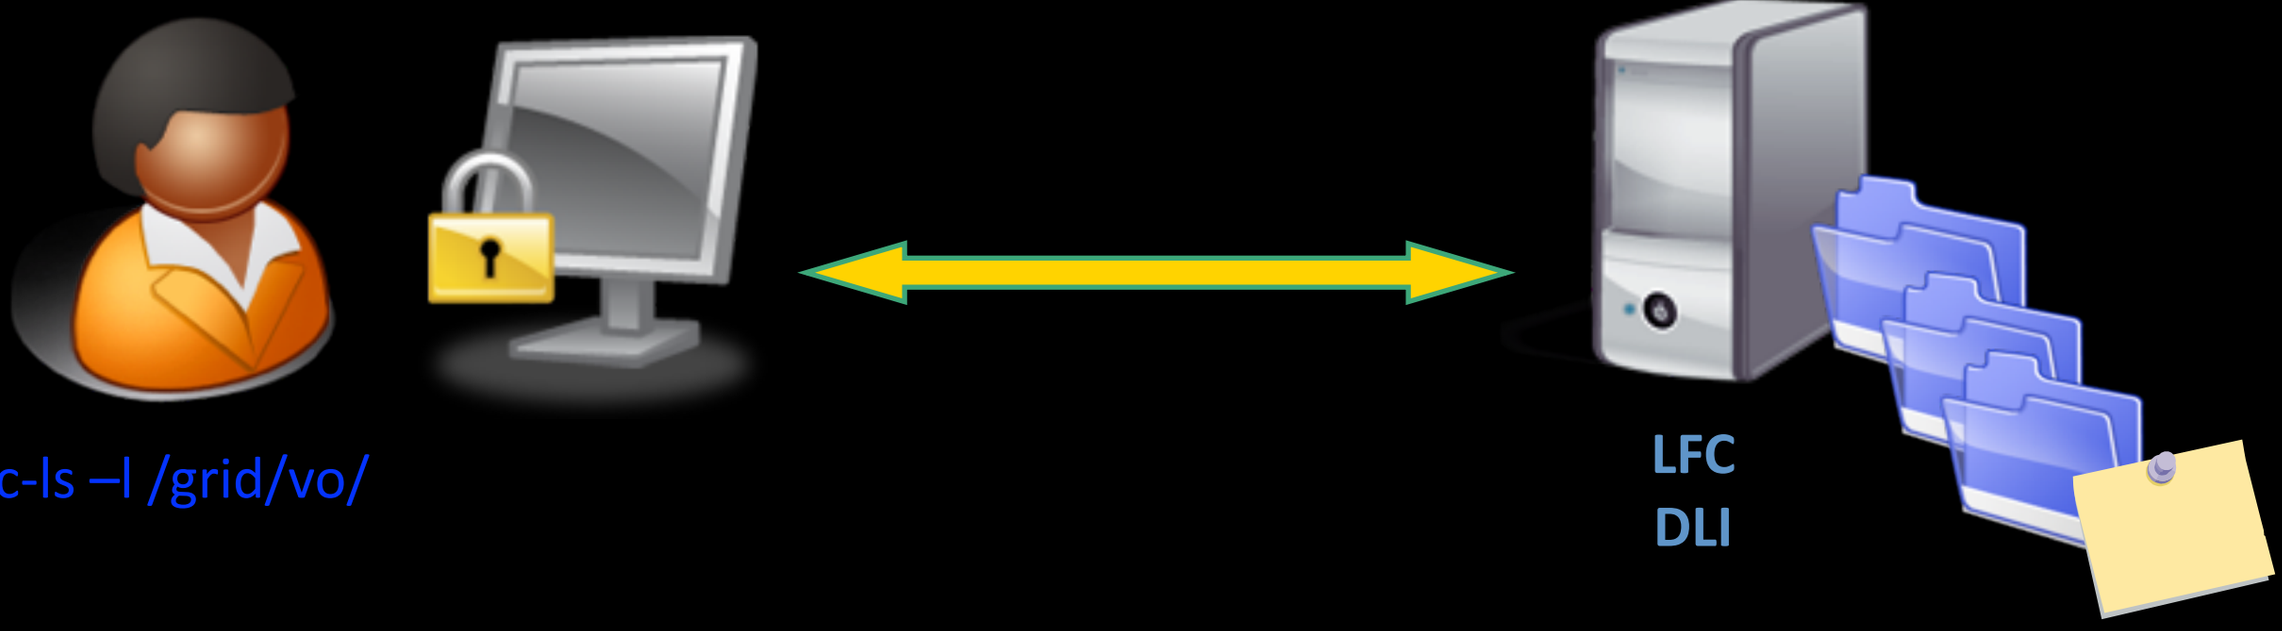

LFC features

Hierarchical Namespace
GSI security
Permissions and ownership
ACLs (based on VOMS)

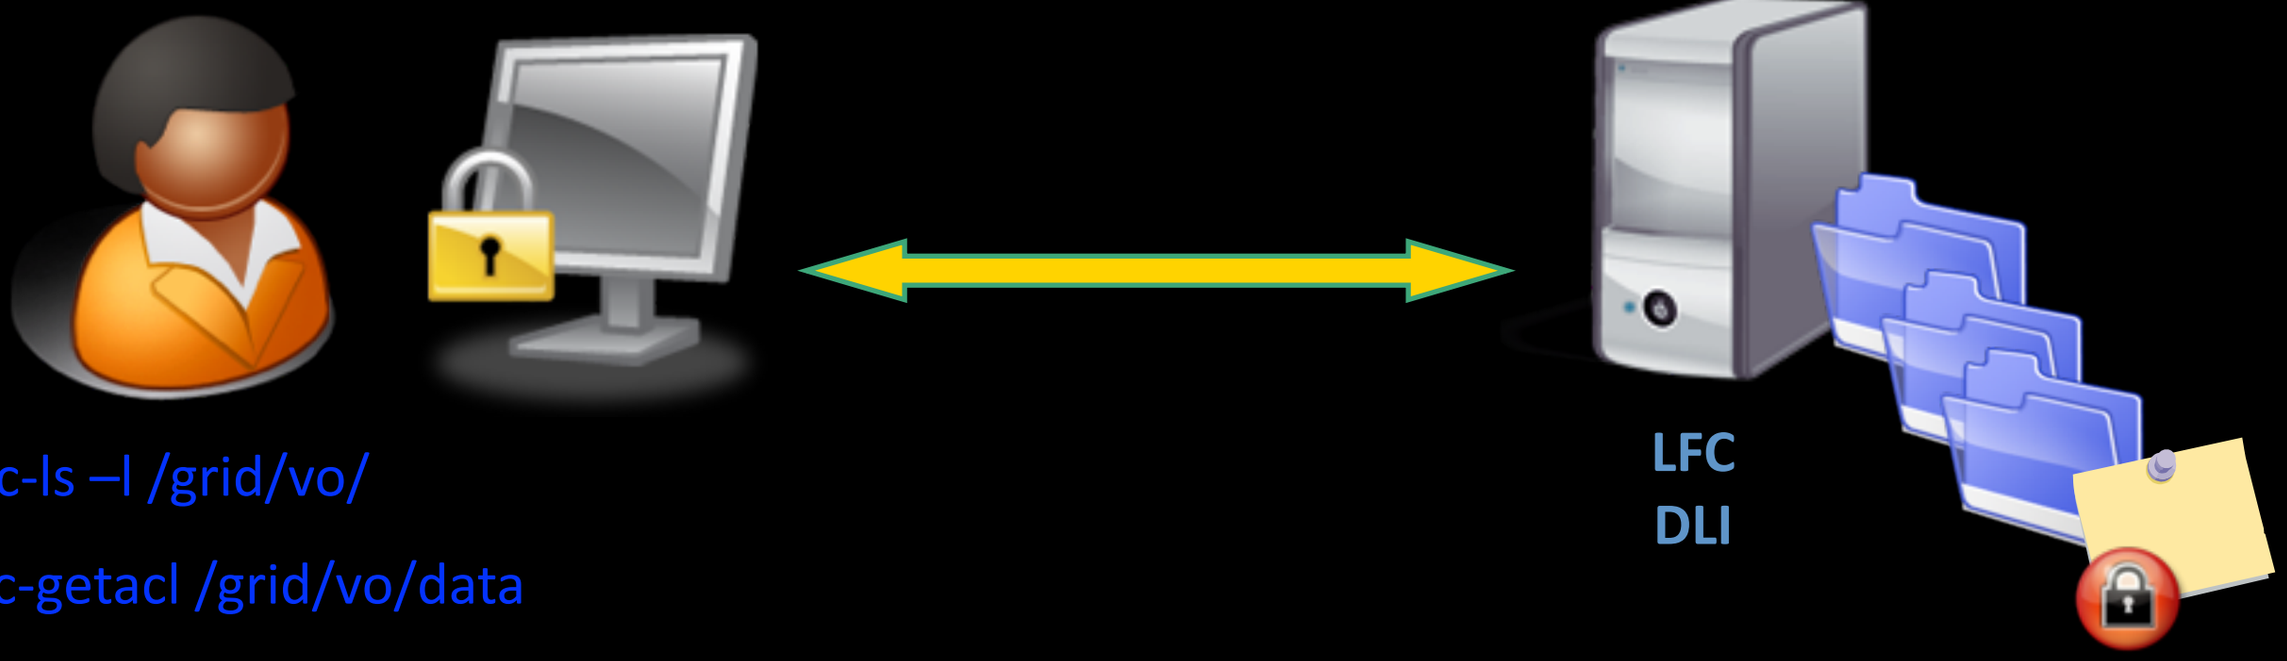

lfc-ls –l /grid/vo/

lfc-getacl /grid/vo/data

LFC features

Hierarchical Namespace
GSI security
Permissions and ownership
ACLs (based on VOMS)
Virtual ids
- Each user is mapped to (uid, gid)

VOMS support
- To each VOMS group/role corresponds a virtual gid

Bulk operations




LFC
DLI

lfc-ls –l /grid/vo/

lfc-getacl /grid/vo/data

DPM

- Disk Pool Manager
  - –Manages storage on disk servers
  - –SRM support
    - 1.1
    - 2.1 (for backward compatibility)
    - 2.2 (released in DPM version 1.6.3)
  - –GSI security
  - –ACLs
  - –VOMS support
  - –Secondary groups support (see LFC)

DPM strengths

- Easy to use
  - Hierarchical namespace
    - $ dpns-ls /dpm/cern.ch/home/vo/data
  - Many protocols supported (including HTTPS)
- Easy to administrate
  - Easy to install and configure
  - Low maintenance effort
  - Easy to add/drain/remove disk servers
- Target: small to “medium” sites (1.6 PB)
  - Single disks --> several disk servers

DPM: user's point of view

CLI, C API,
SRM-enabled client, etc

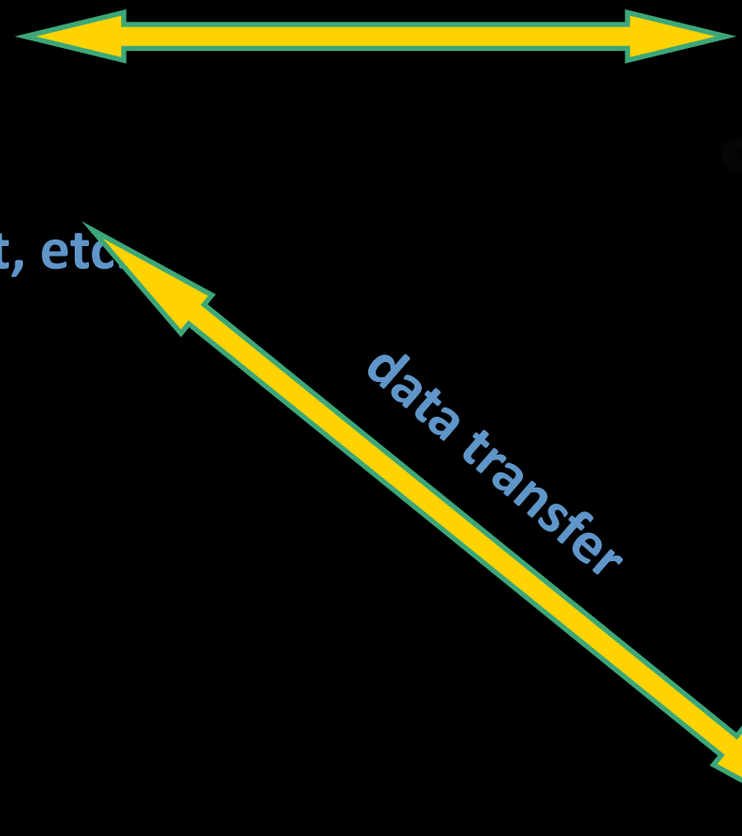

data transfer

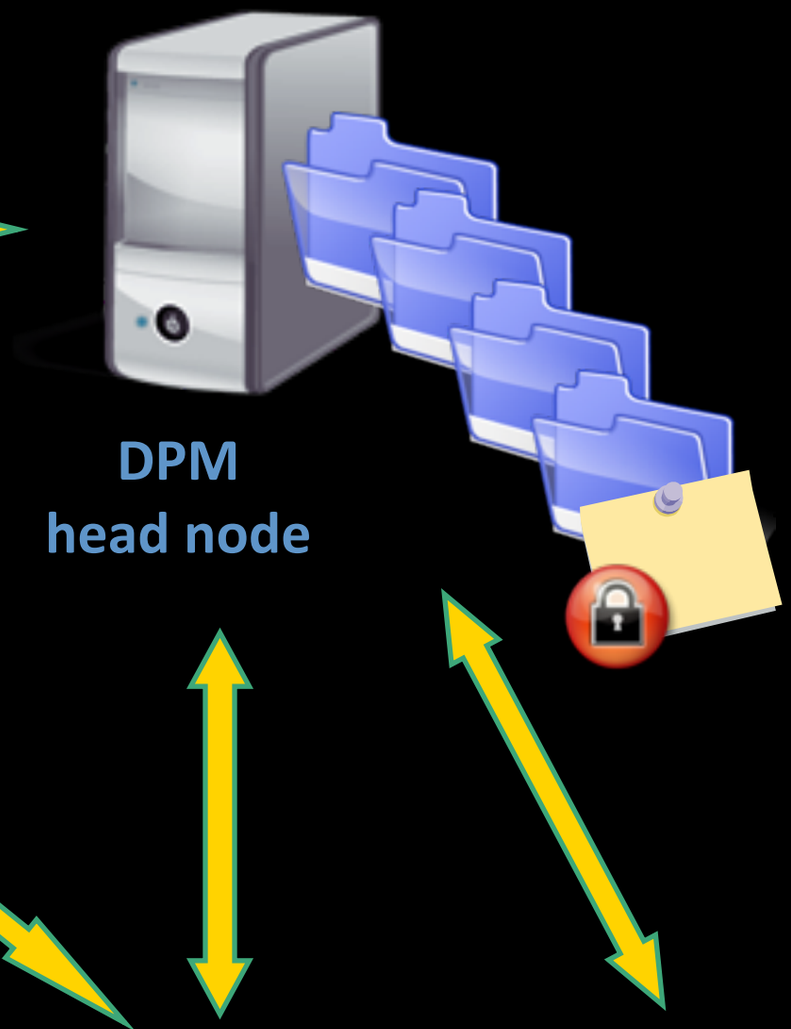

DPM
head node

DPM Name Server

- Namespace
- Authorization
- Physical files location

Disk Servers

- Physical files

Direct data transfer from/to disk server (no bottleneck)

External transfers via gridFTP

DPM
disk servers

File Transfer Service

- FTS: Reliable, scalable and customizable file transfer
  - Multi-VO service, used to balance usage of site resources according to the SLAs agreed between a site and the VOs it supports
  - WS interface, support for different user and administrative roles (VOMS)
  - Manages transfers through channels
    - mono-directional network pipes between two sites
  - File transfers handled as jobs
    - Prioritization
    - Retries in case of failures
  - Automatic discovery of services
- Designed to scale up to the transfer needs of very data intensive applications
  - Demonstrated about 1 GB/s sustained
  - Over 9 petabytes transferred in the last 6 months (> 10 million files)

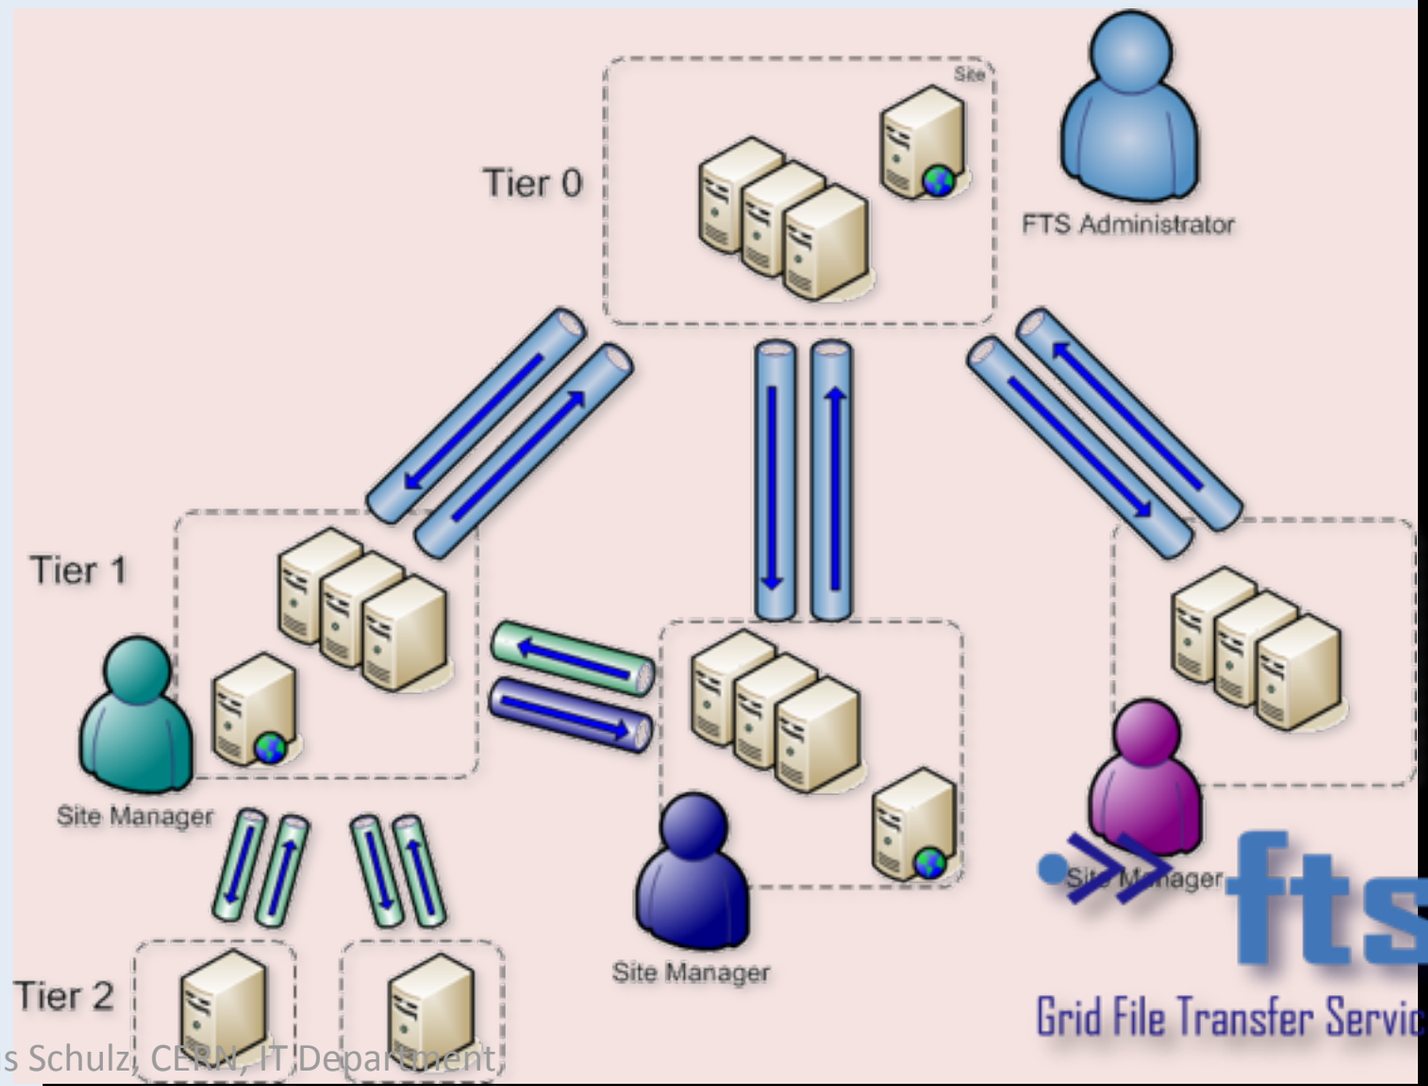

FTS: key points

- Reliability
  - It handles the retries in case of storage / network failures
    - VO customizable retry logic
  - Service designed for high-availability deployment
- Security
  - All data is transferred securely using delegated credentials with SRM / gridFTP
  - Service audits all user / admin operations
- Service and performance
  - Service stability: it is designed to efficiently use the available storage and network resources without overloading them
  - Service recovery: integration of monitoring to detect service-level degradation

FTS: key points

- Reliability
  - It handles the retries in case of storage / network failures
    - VO customizable retry logic
  - Service designed for high-availability deployment
- Security
  - All data is transferred securely using delegated credentials with SRM / gridFTP
  - Service audits all user / admin operations
- Service and performance
  - Service stability: it is designed to efficiently use the available storage and network resources without overloading them
  - Service recovery: integration of monitoring to detect service-level degradation

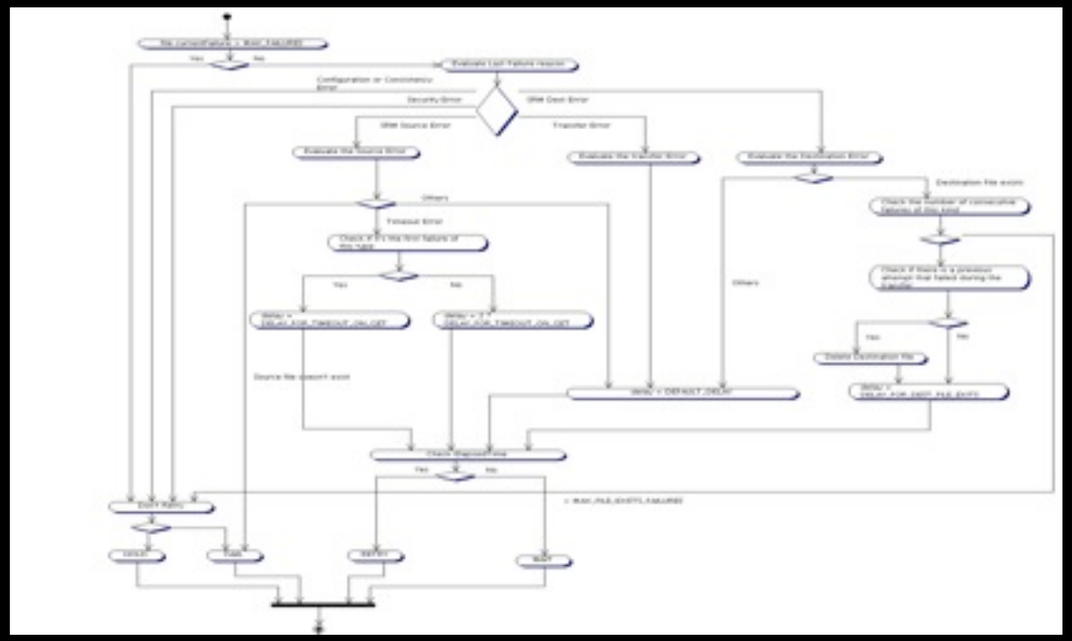

FTS: key points

- Reliability
  - It handles the retries in case of storage / network failures
    - VO customizable retry logic
  - Service designed for high-availability deployment
- Security
  - All data is transferred securely using delegated credentials with SRM / gridFTP
  - Service audits all user / admin operations
- Service and performance
  - Service stability: it is designed to efficiently use the available storage and network resources without overloading them
  - Service recovery: integration of monitoring to detect service-level degradation

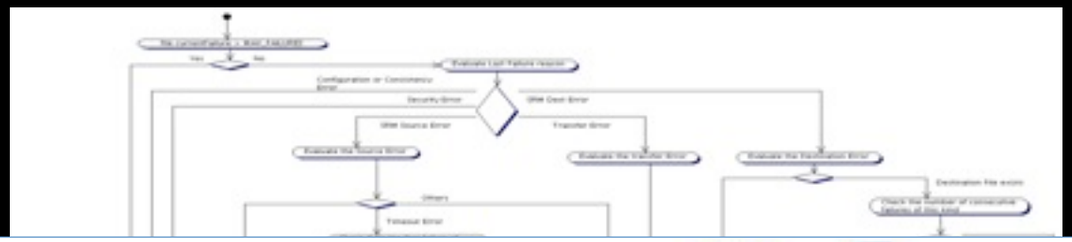

FTS Report

Disclaimer
This page contains a report generated from information stored in the FTS Database and is intended for reporting purposes only. Since the format will probably change in the future, it's therefore recommended not to use parsing robots on it.

Statistics concerning all the transfers performed yesterday
Between 2006-10-12 08:00:00 +02:00 and 2006-10-13 08:00:00 +02:00

CERN-.* Filter Show VO details

Channel Name	VO Name	Total	% Failures	# Succ.	# Fail.	1st Failure Reason	% 1st Failure Reason	2nd Failure Reason	% 2nd Failure Reason	Avg. Size (GB)	Avg. Duration (sec)	Avg. Tx Rate (MB/sec)	Eff. Tx Bytes (GB)	Tx Bytes (GB)
CERN-PIC	[All]	12262	73.97	3192	9070	Dest SRM	56.22	Other	37.53	0.53	263.03	1.62	1700.41	1700.41
	atlas	8932	99.92	7	8925	Dest SRM	57.13	Other	38.08	0	220	0	0	0
	cms	208	0	208	0					2.7	767.55	3.64	561.26	561.26
	dteam	974	0.51	969	5	Other	80	Source SRM	20	0.95	356.31	2.88	923.83	923.83
	lhcb	2145	6.53	2005	140	Source SRM	99.29	Other	0.71	0.11	165.85	0.81	215.32	215.32
	ops	3	0	3	0					0	202.67	0	0	0
CERN-RAL	[All]	8699	59.26	3544	5155	Other	84.91	Source SRM	14.88	0.85	478.22	2.59	3026.81	3027.57
	alice	1155	82.6	201	954	Other	99.58	Dest SRM	0.31	1.86	1805.05	1.11	372.95	372.95
	atlas	4512	88.52	518	3994	Other	84.85	Source SRM	15.15	1.79	1428.94	1.49	926.26	926.57
	cms	227	3.08	220	7	Dest SRM	85.71	Source SRM	14.29	2.53	348.65	10.08	555.61	555.61
	dteam	1077	3.99	1034	43	Other	86.05	Source SRM	9.3	0.95	276.64	4.01	980.47	980.91
	lhcb	1725	9.1	1568	157	Source SRM	99.36	Other	0.64	0.12	146.03	1.16	191.52	191.52
	ops	3	0	3	0					0	27	0.01	0	0
CERN-SARA	[All]	8792	42.55	5051	3741	Dest SRM	83.77	Source SRM	12.22	1.34	108.02	15.4	6777.95	6784.92
	alice	3134	15.12	2660	474	Source SRM	57.17	Dest SRM	41.14	1.66	109.53	18.43	4426.44	4430.29
	atlas	2018	53.32	942	1076	Dest SRM	72.4	Source SRM	16.54	1.15	144.44	9.42	1085.07	1087.6
	dteam	3488	61.32	1349	2139	Dest SRM	98.74	Other	0.98	0.93	81.91	14.66	1260.74	1261.32
	lhcb	148	35.14	96	52	Dest SRM	92.31	Other	3.85	0.06	76.1	0.93	5.7	5.7
	ops	4	0	4	0					0	97.25	0.02	0	0
CERN-INFN	[All]	11492	42.31	6630	4862	Dest SRM	43.85	Other	37.7	1.13	395.77	3.21	7514.29	7614.84
CERN-CERN	[All]	1536	39.71	926	610	Source SRM	58.36	Dest SRM	15.9	0.07	287.71	0.38	67.89	69.08
CERN-ASCC	[All]	6851	23.54	5238	1613	Source SRM	50.84	Other	28.89	1.14	1098.6	1.08	5955.81	6080.58
CERN-GRIDKA	[All]	12755	21.38	10028	2727	Source SRM	64.36	Other	32.53	0.87	371.97	3.19	8762.02	8767.53
CERN-TRIUMF	[All]	2244	20.63	1781	463	Other	61.77	Source SRM	31.1	1.04	395.15	3.63	1847.25	1917.13
CERN-BNL	[All]	13975	19.42	11261	2714	Source SRM	69.97	Other	24.24	0.44	190.38	3.41	4951.59	4960.34
CERN-IN2P3	[All]	11697	13.76	10087	1610	Source SRM	48.57	Other	47.45	1.22	296.21	5.33	12329.63	12329.63
CERN-FNAL	[All]	917	4.58	875	42	Transfer	97.62	Other	2.38	0	379.88	0	0	0

Click on the Channel Name to show the VO details

FTS server architecture

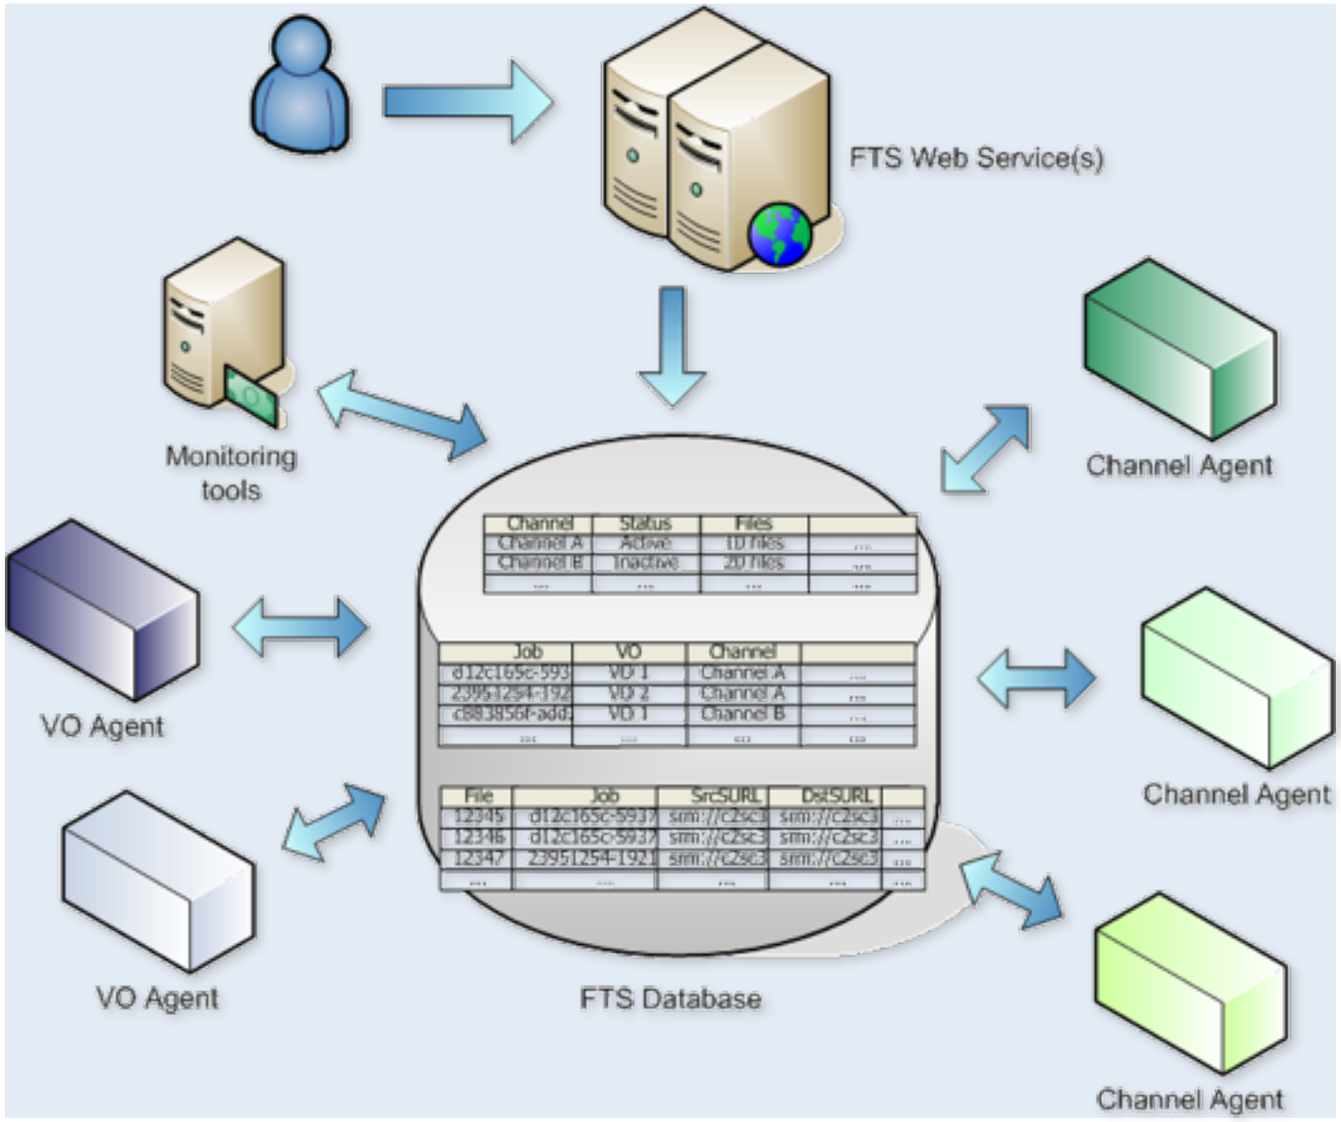

Experiments interact via a web-service

VO agents do VO-specific operations (1 per VO)

Channel agents do channel specific operation (e.g. the transfers)

Monitoring and statistics can be collected via the DB

- All components are decoupled from each other
  - Each interacts only with the (Oracle) database

gLite code base

Total Physical Source Lines of Code (SLOC)

SLOC = 1622714

Total SLOC grouped by language (dominant language first)

Language	Total SLOC
ansic	578598 (35%)
cpp	491801 (30%)
java	251382 (15%)
sh	191798 (11%)
python	54510 (3%)
perl	39258 (2%)
yacc	7445 (0%)
jsp	4444 (0%)
lex	2274 (0%)
csh	701 (0%)
awk	307 (0%)
fortran	124 (0%)
sed	68 (0%)
asm	4 (0%)

- **Distributed under an open source license.**
- **Main platform is Scientifc Linux (recompiled RH EL).**
- **Many 3rd party dependencies**
  - tomcat, log4*,gSOAP , ldap etc.

- **~ 20 FTEs, 80 people, 12 institutes (mostly academic)**
- **Geographically distributed, independent**
  - Coding conventions, Documentation, Naming Conventions
  - Testing and quality, dependency management

Stability of the software

- All components still see frequent changes
- Many developments started 2002
  - –Why do we still need changes?
    - Scale of the system increased rapidly
    - Number of user and use cases increased
  - –Deeper code coverage
  - –New functional requirements
- Less tolerance to failures
  - –Implementation of fail over
- Emerging standards
  - –Project started when no standards where available
  - –Incremental introduction

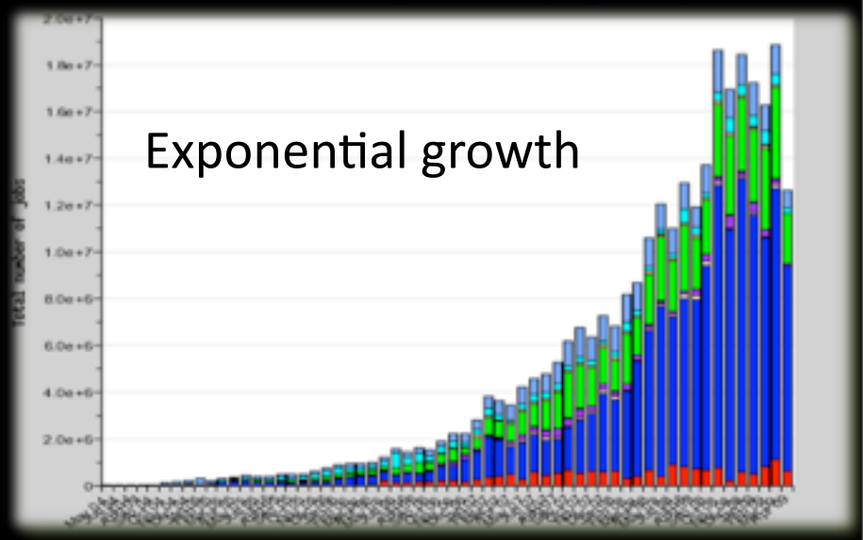

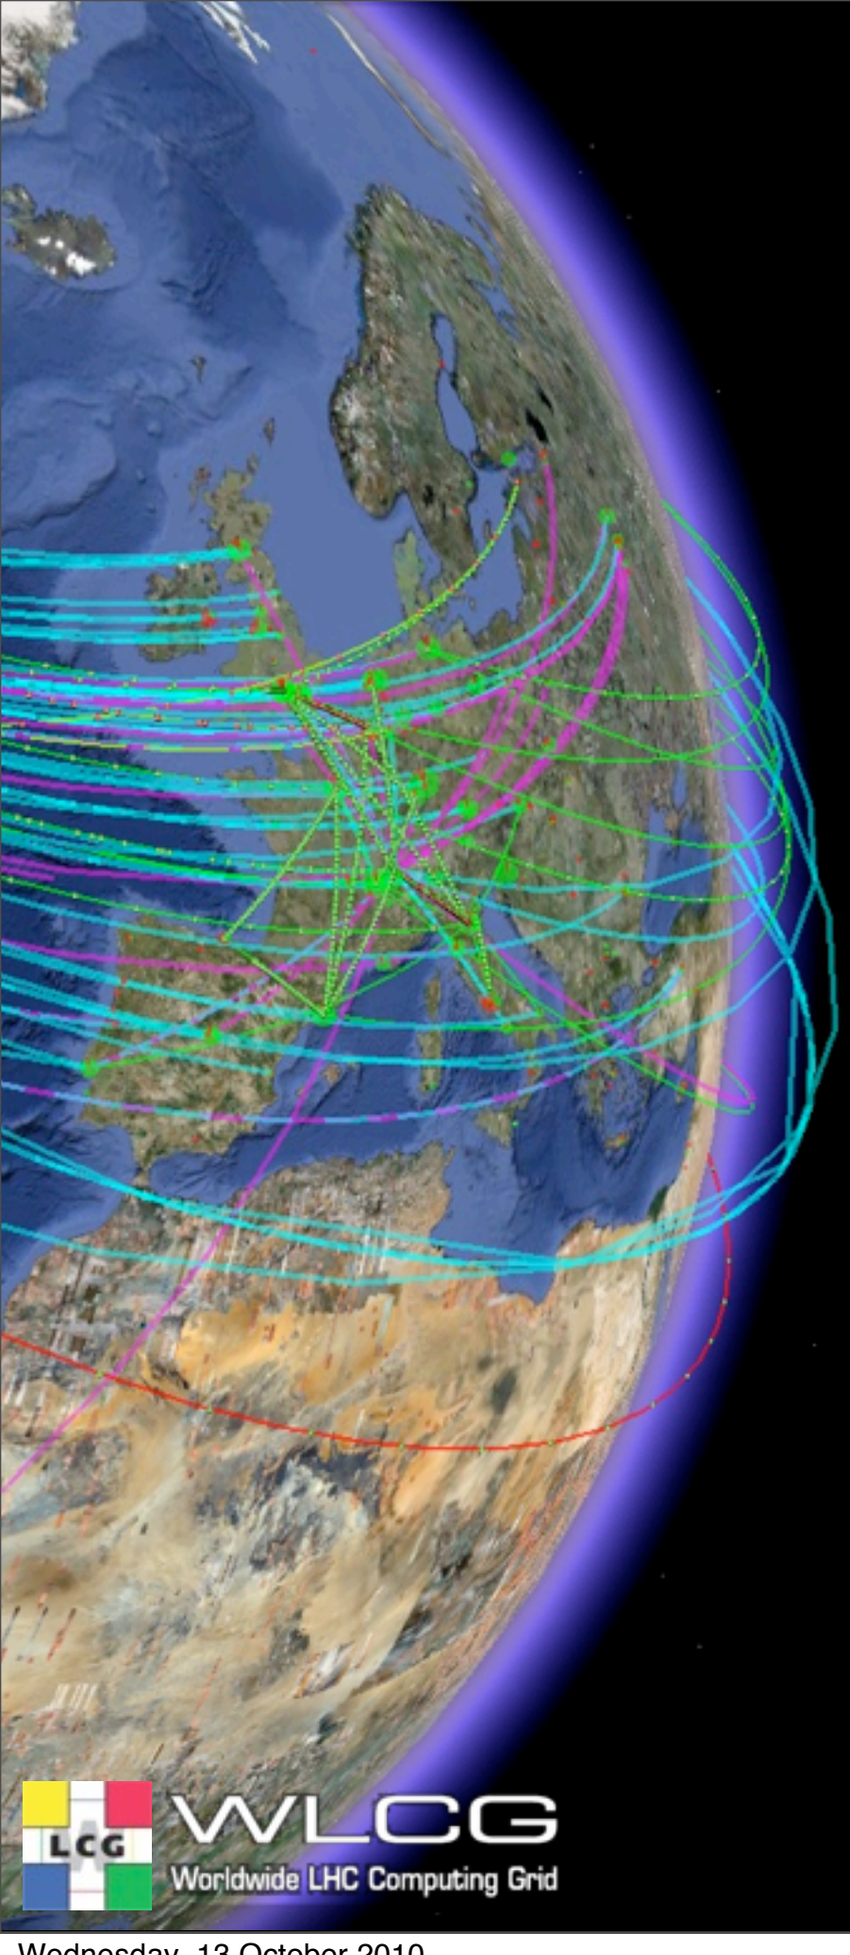

How do we manage the code?

- Builds and test are mostly managed by ETICS
- Configuration management by YAIM
  - –modular bash shell script
    - >37 000 lines, >30 modules
- Complex certification and release process

YAIM

Error rates / Usage

- Bug rate almost flat
- Exponential increase in usage
- Example: gLite

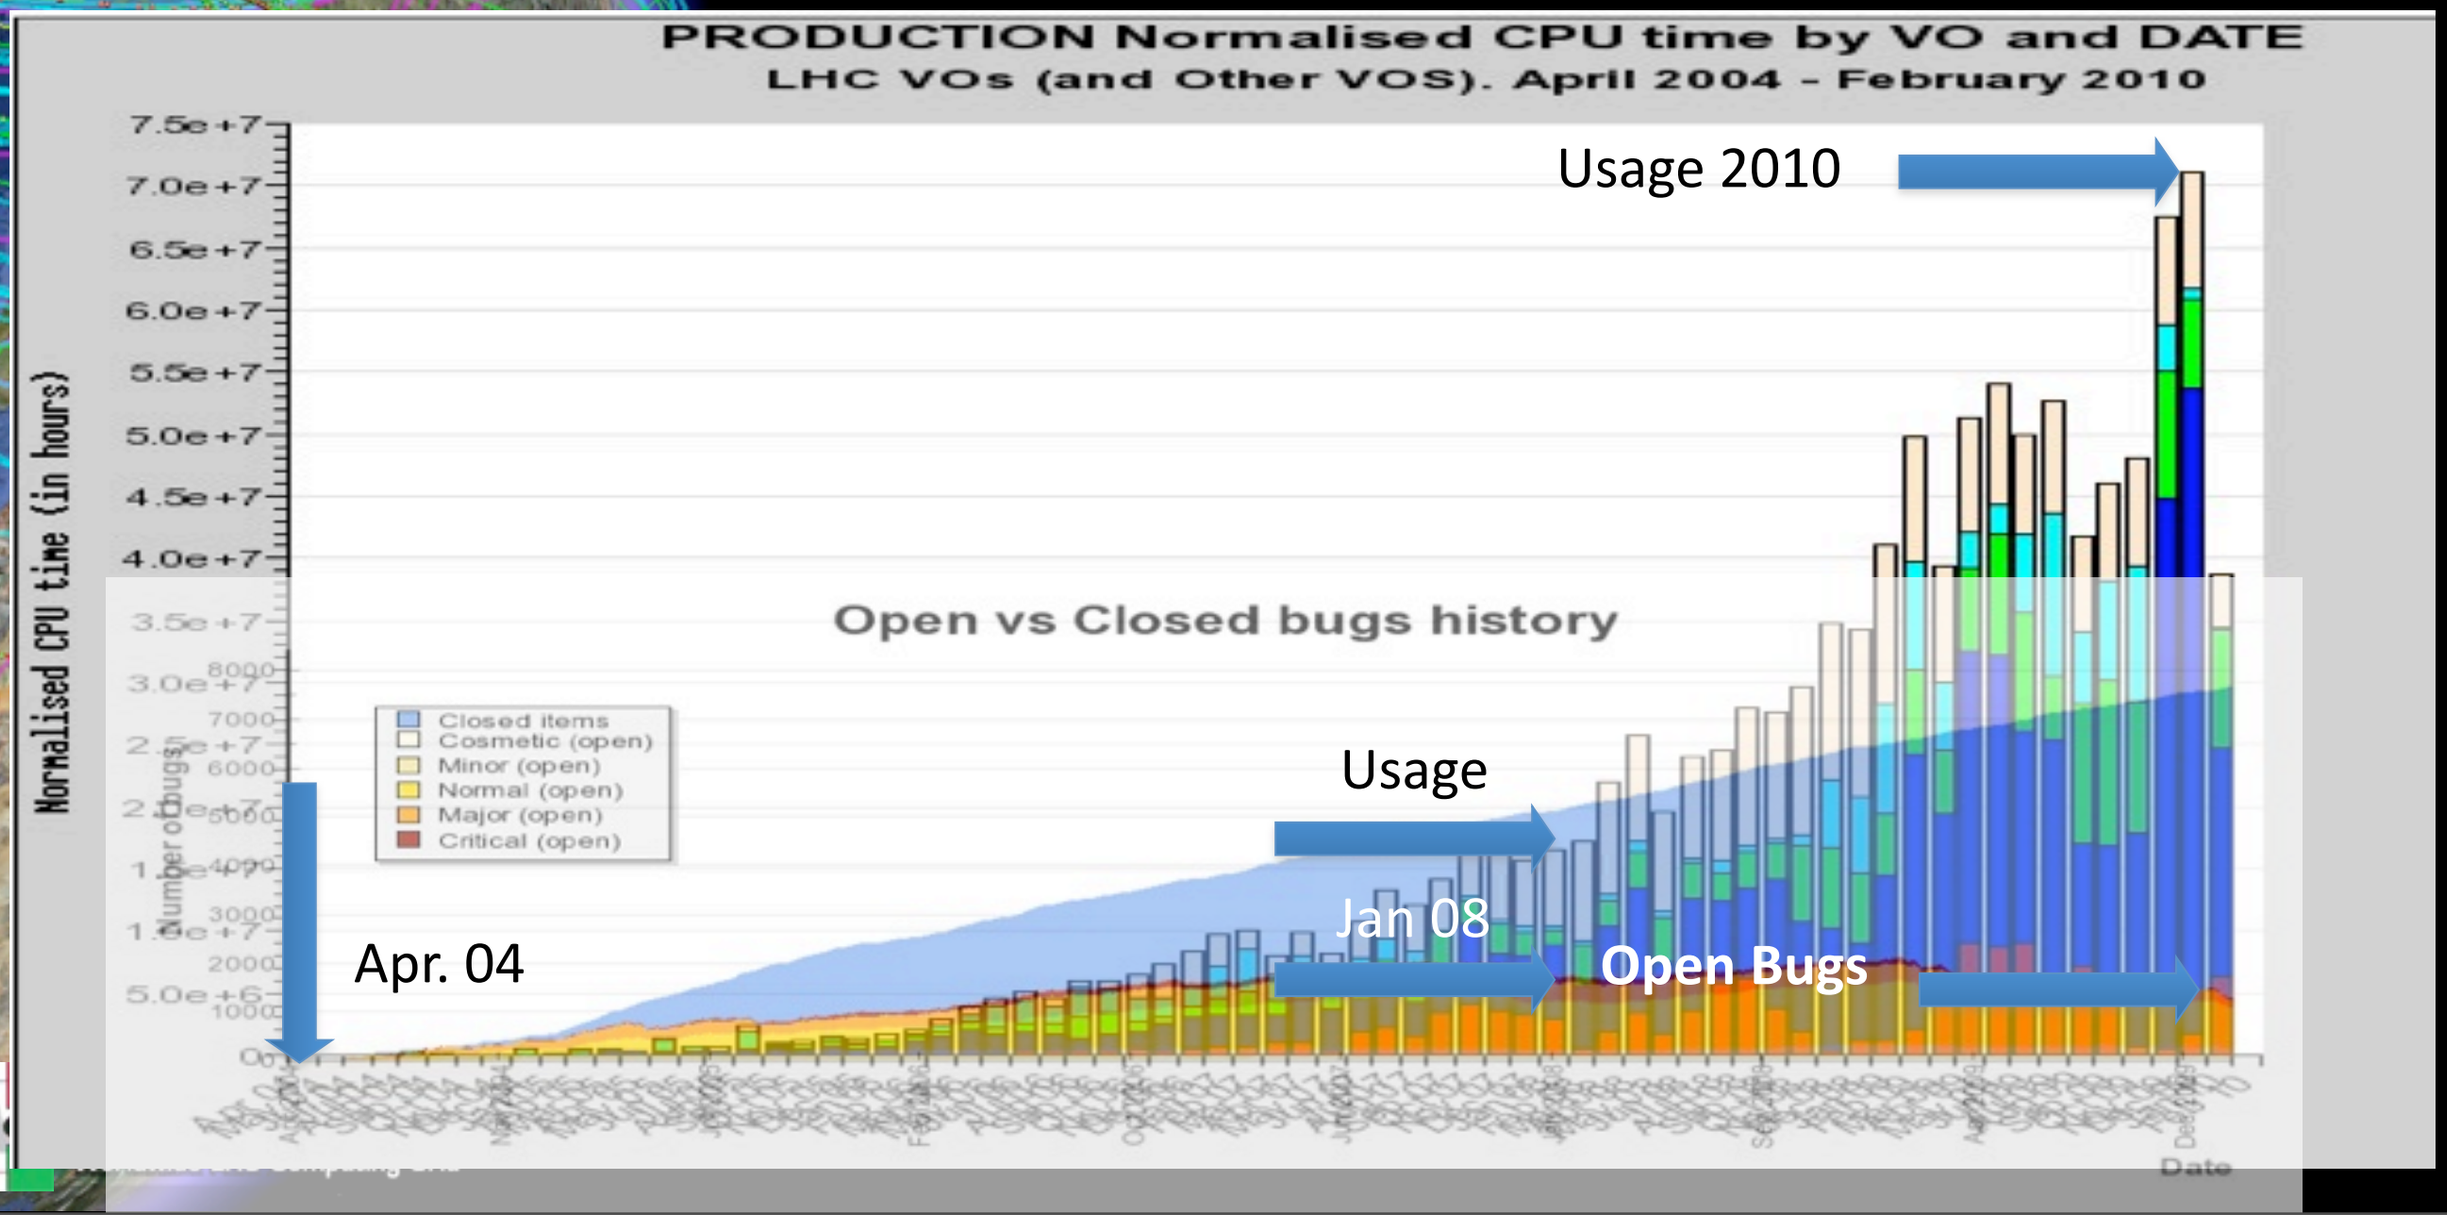

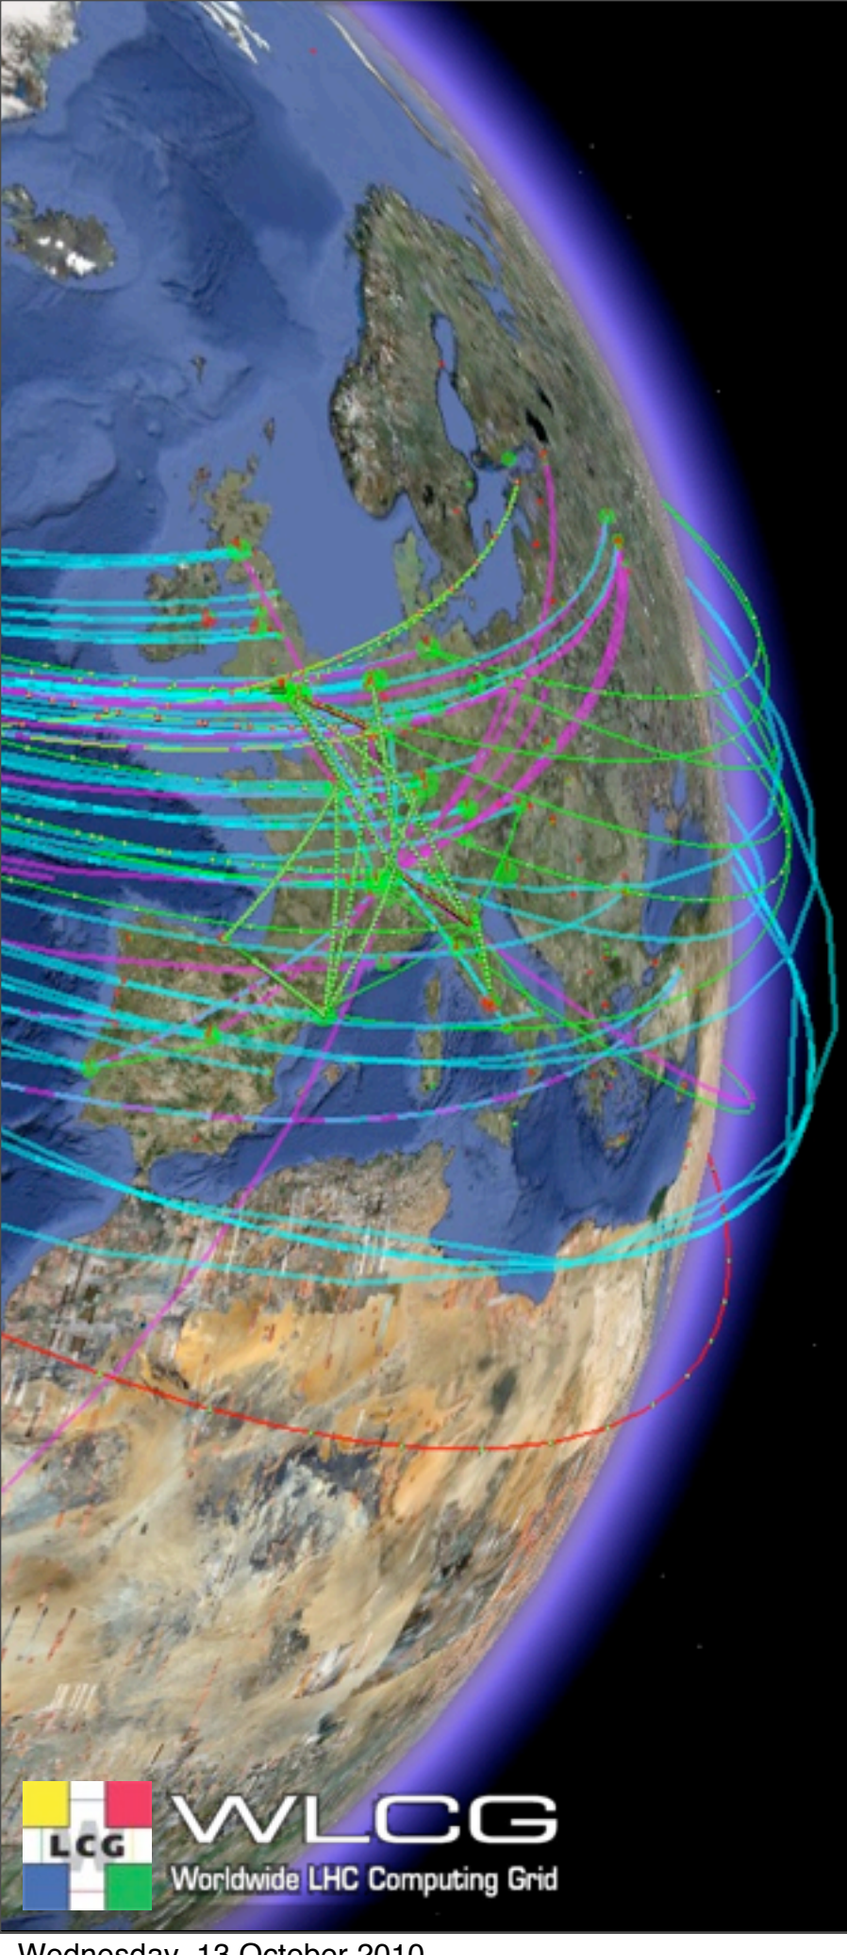

How does the future look like?

- Focus on standardization and interoperation
  - Driving the process
    - OGF etc.
- Focus on stability
- Simplifying the system
- Integrating virtualization
- Integrating Clouds
- Moving to EMI

WLCG Worldwide LHC Computing Grid

Partners (26)


Lund University
Uppsala Universitet
University of Oslo
Manchester 1824
The University of Manchester
Helsinki Institute of Physics
Science & Technology Facilities Council
JÜLICH
Forschungszentrum
icm
NIKHEF
CERN
Technische Universität Dresden
KiSTi
www.kisti.re.kr
CESNET
CESGA
SWITCH
INFN
CINECA
grnet
CSIC
ITY
IS
NO
SE
FI
RU
EE
LV
LT
GB
IE
BE
DE
PL
UA
SK
AT
HU
HR
BA
RS
ME
MK
AL
GR
IT
PT
ES
GE
AM
AZ
South Korea
Taiwan

Partners (26)


Lund University
Uppsala Universitet
University of Oslo
Manchester 1824
The University of Manchester
Helsinki Institute of Physics
Google
Science & Technology Facilities Council
JÜLICH
Forschungszentrum
icm
NIKHEF
CERN
Technische Universität Dresden
KiSTi
www.kisti.re.kr
CESNET
CESGA
SWITCH
INFN
CINECA
grnet
CSIC
IS
NO
SE
FI
RU
EE
LV
LT
GB
IE
BE
DE
PL
UA
SK
HU
AT
CZ
PT
ES
IT
HR
BA
RS
ME
MK
AL
GR
GE
AM
AZ
MT
South Korea
Taiwan