

LCG - The Worldwide LHC Computing Grid

openlab

Board of Sponsors

WLCG Status update


CERN

26 April 2007

Les Robertson

LCG Project Leader

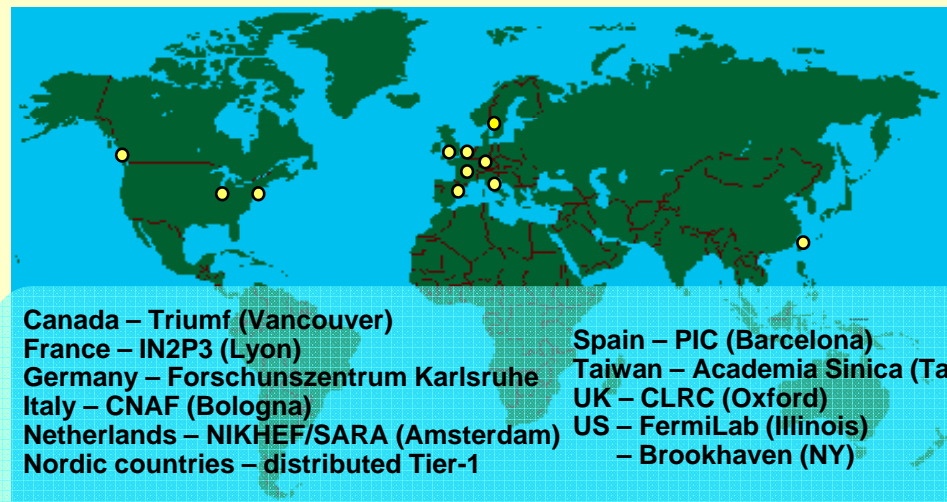




LCG Service Hierarchy

Tier-0 – the accelerator centre

- Data acquisition & initial processing
- Long-term data curation
- Distribution of data → Tier-1 centres




Tier-1 – “online” to the data acquisition process → high availability

- Managed Mass Storage –
→ grid-enabled data service
- Data-heavy analysis
- National, regional support


Tier-2 – ~100 centres in ~40 countries

- Simulation
- End-user analysis – batch and interactive






CERN Tier-0 Progress




Peak transfer of incoming data to tape at over 2GB/s




LCG depends on two major science grid infrastructures

EGEE - Enabling Grids for E-Science


OSG - US Open Science Grid



A map of the worldwide LCG infrastructure operated by EGEE and OSG.

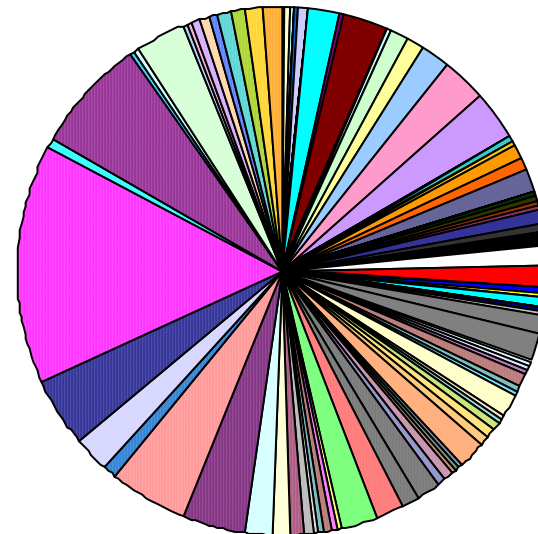


CPU Usage - LHC Experiments March 2007



CERN **20%**
11 Tier-1s **60%**
140 Tier-2s **20%**


Tier-2 Sites Contributing in
March 2007







Grid Activity

- Steady increase in usage of the EGEE and OSG grids
- Example shows LHC experiment jobs run on the EGEE grid
- 3 x increase in past twelve months
- Need a further 5 x increase by mid-2008
- Similar growth on OSG grid





Grid Activity





Tier-1s and CERN


- CPU usage increased by factor of 2 over past year
- Disk usage by a factor of 4.9



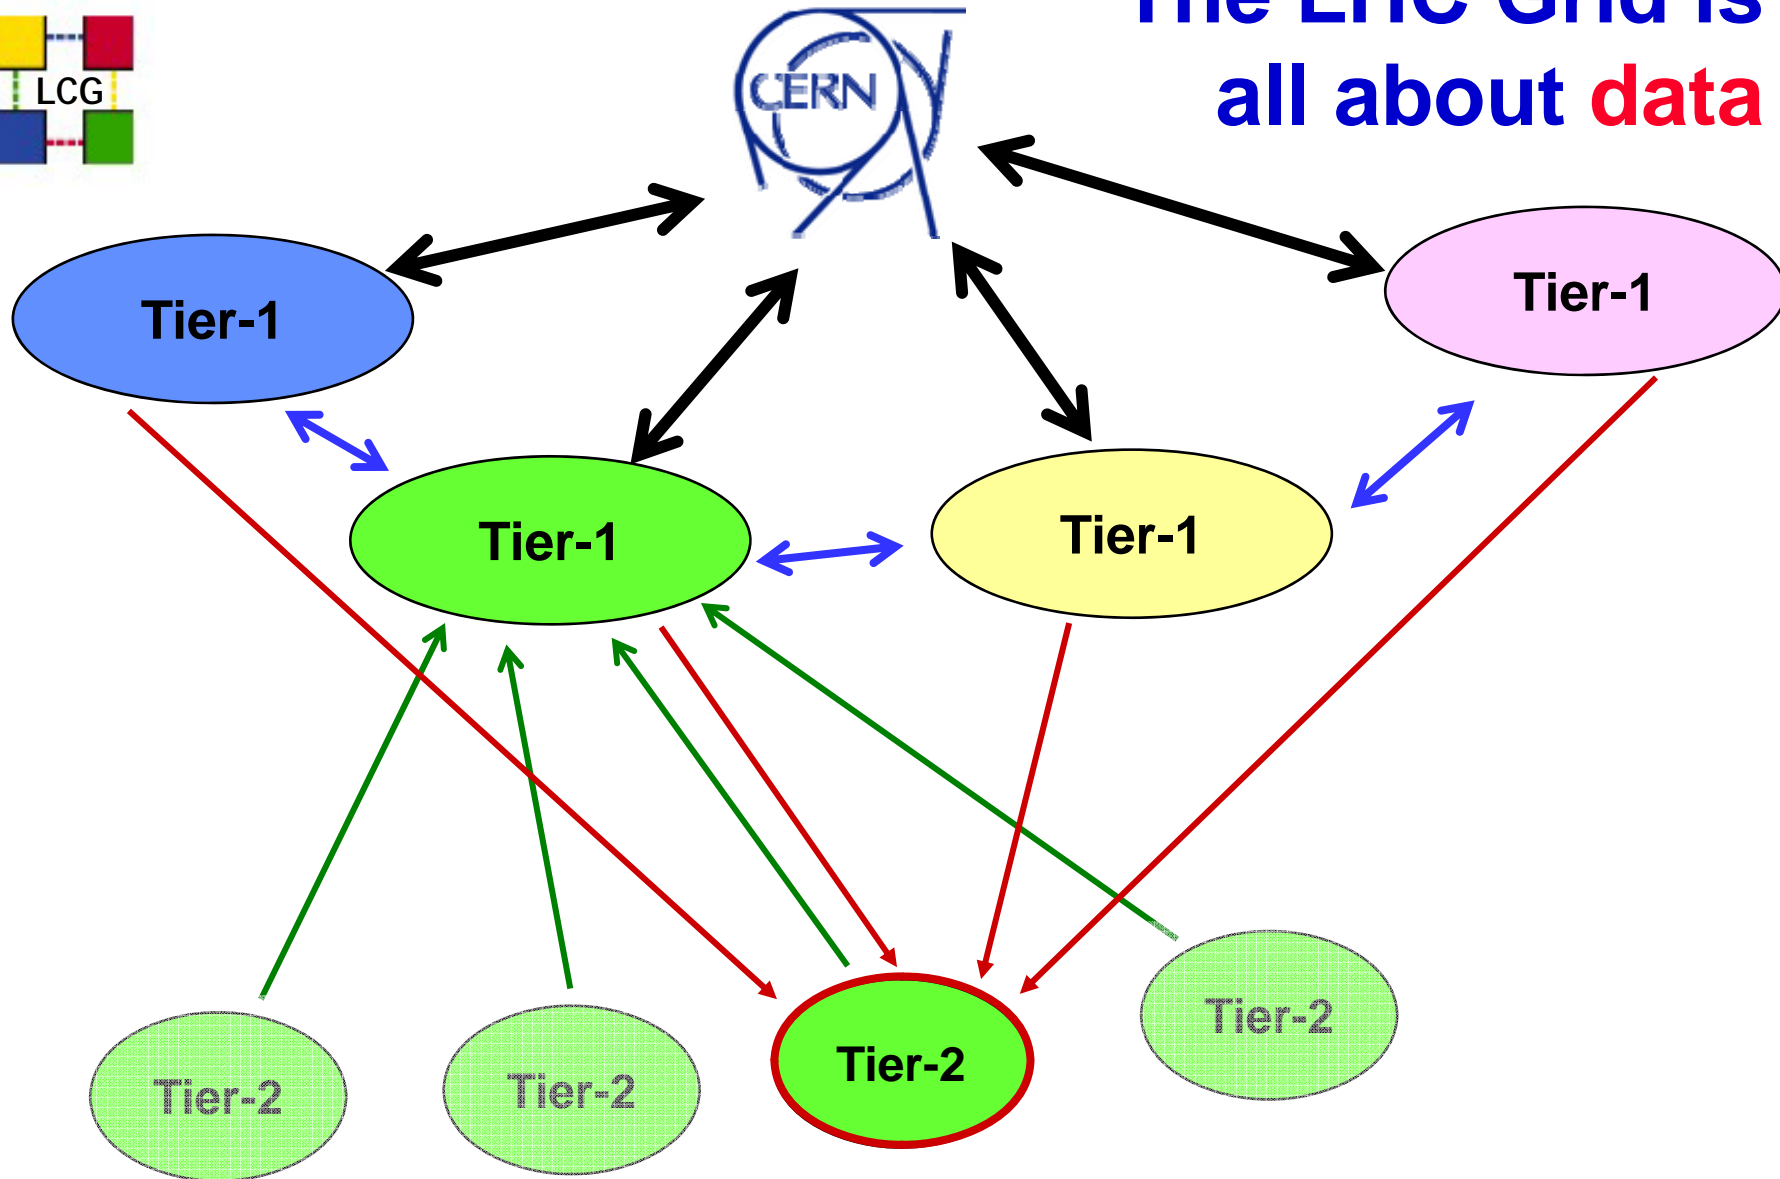


Growth to 2008





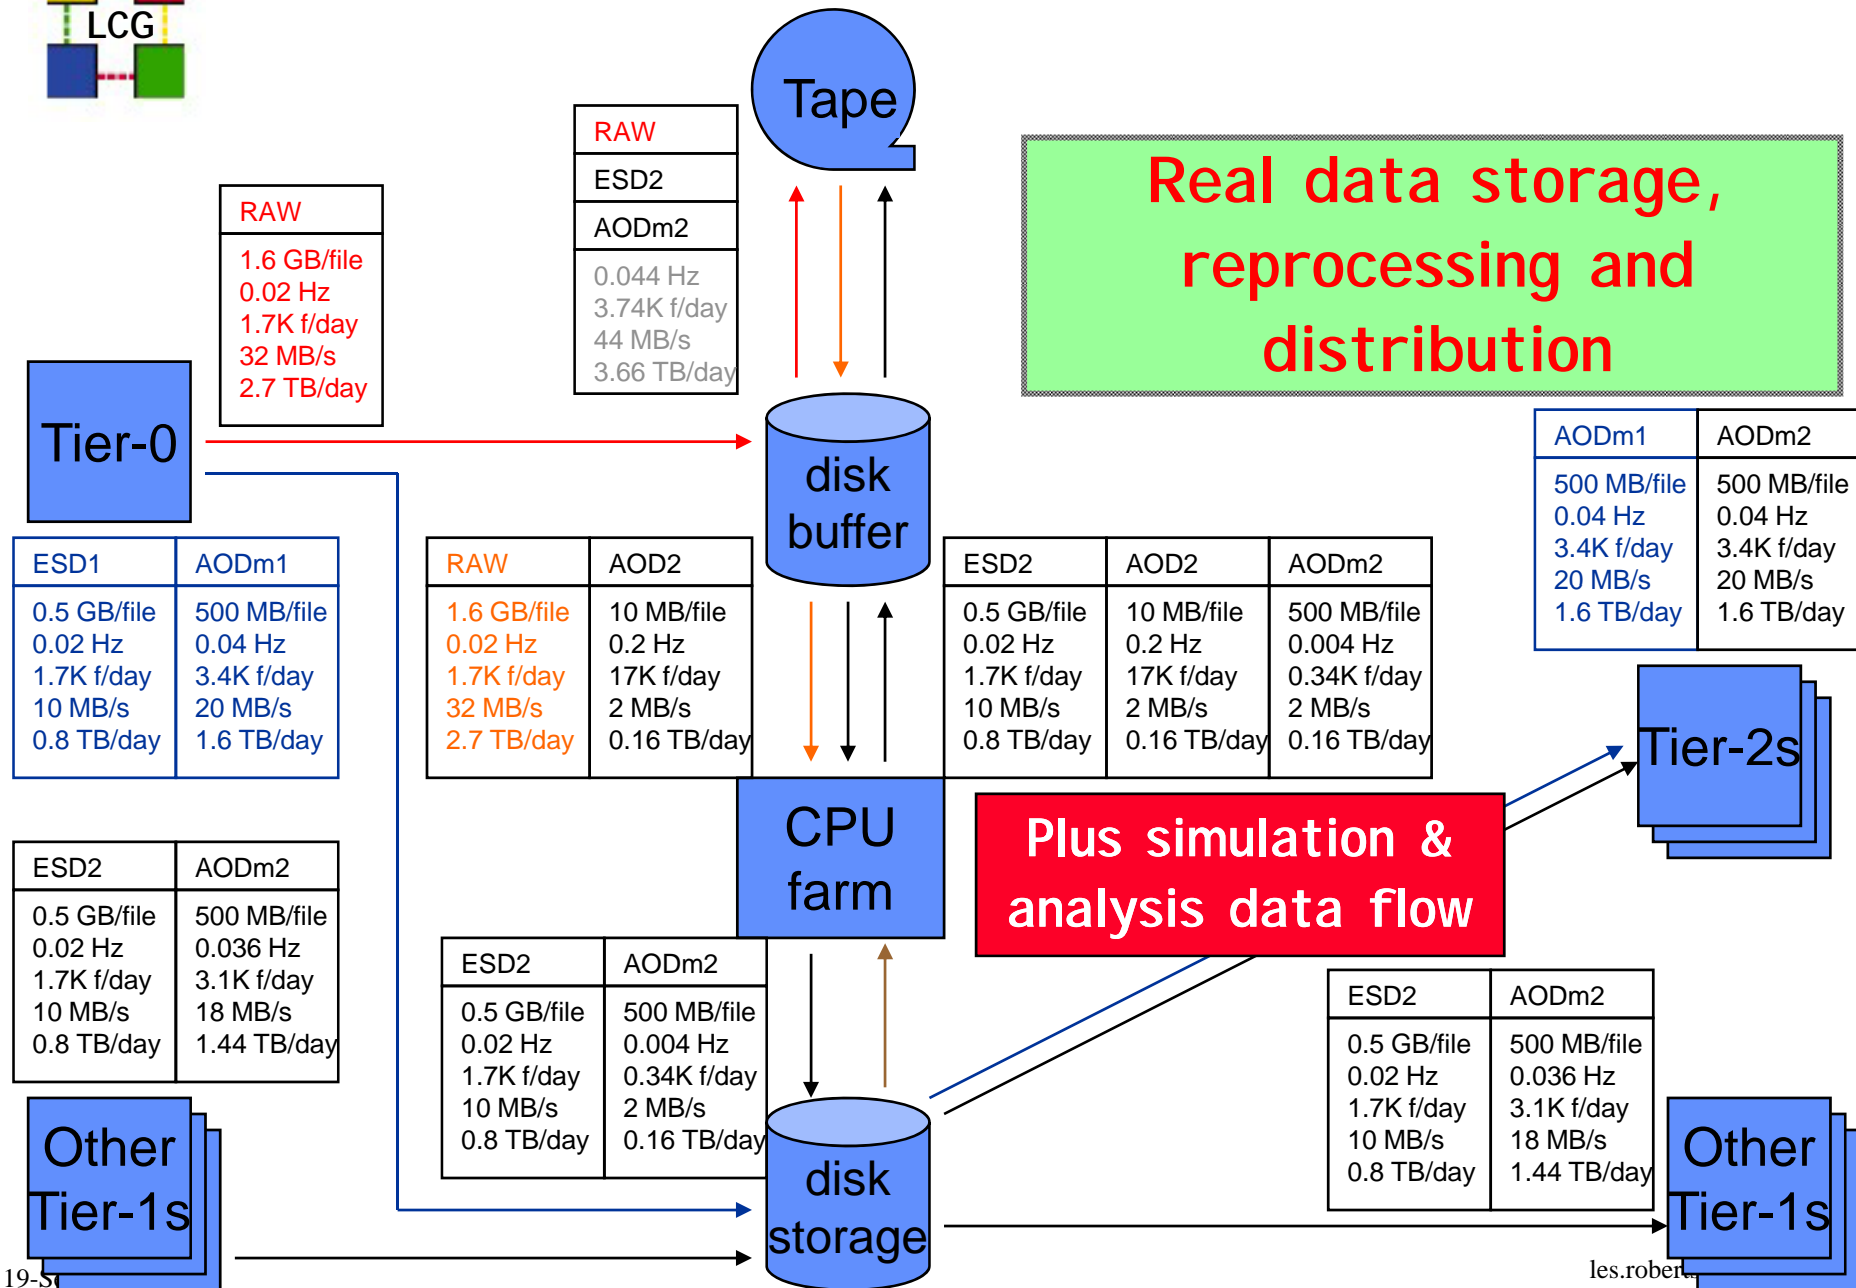
The LHC Grid is
all about **data**



Experiment computing models define specific
data flows between CERN, Tier-1s and Tier-2s



ATLAS "average" Tier-1 Data Flow (2008)






Data Distribution Tests 2006 - CERN → Tier-1s

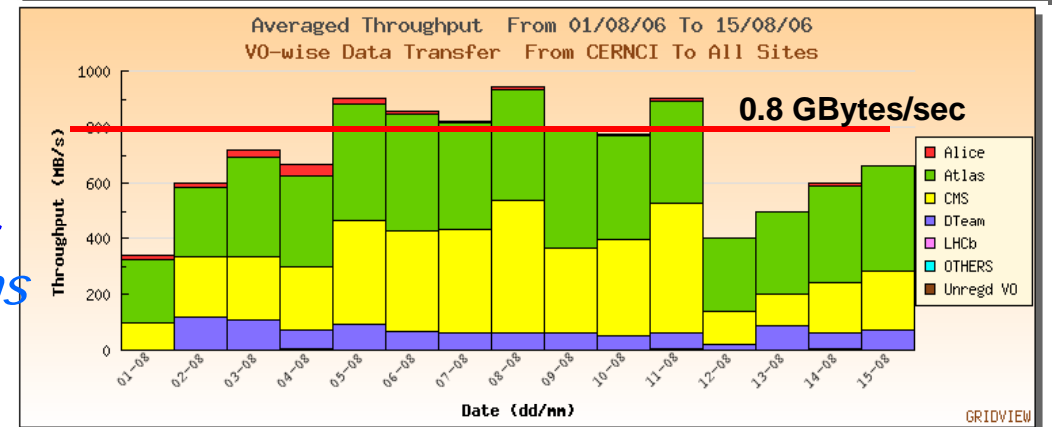
April 06 test period

- Sustained operation at the “nominal” rate when LHC is running - 1.3 GB/s
- Peak daily rate 1.6 GB/s




August

- experiment-driven transfers (ATLAS and CMS) sustained 60% of the SC4 target under *much more realistic conditions*



- CMS transferred a steady 1 PByte/month between Tier-1s & Tier-2s during a 90 day period
- ATLAS distributed 1.25 PBytes from CERN during a 6-week period

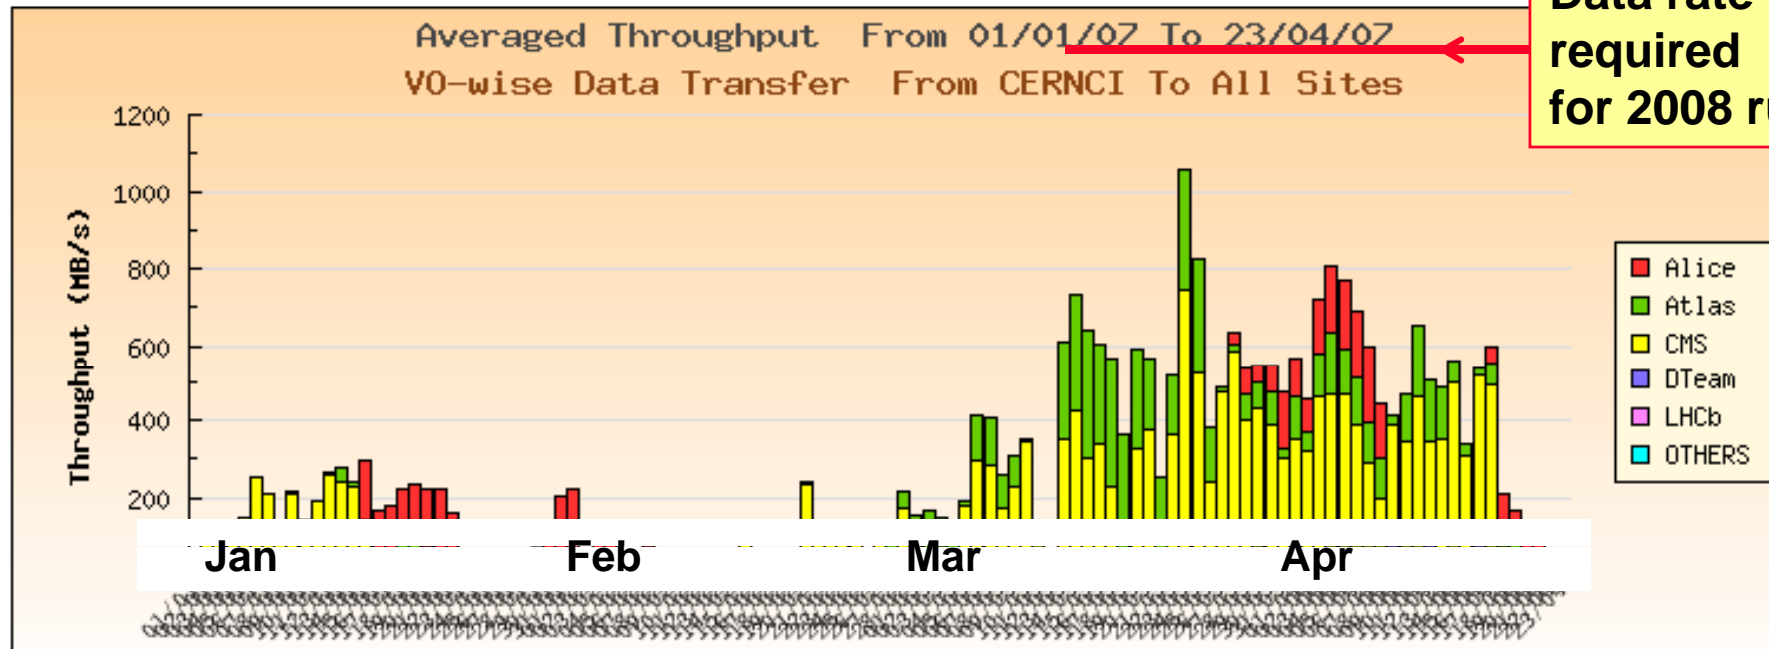


2007 - CERN → Tier-1 Data Distribution



Daily Report

(VO-wise Data Transfer From CERNCI To All Sites)






Data Transfers

Comparison with CSA06 – weekly




Jens Rehn
April 2007
Computing workshop



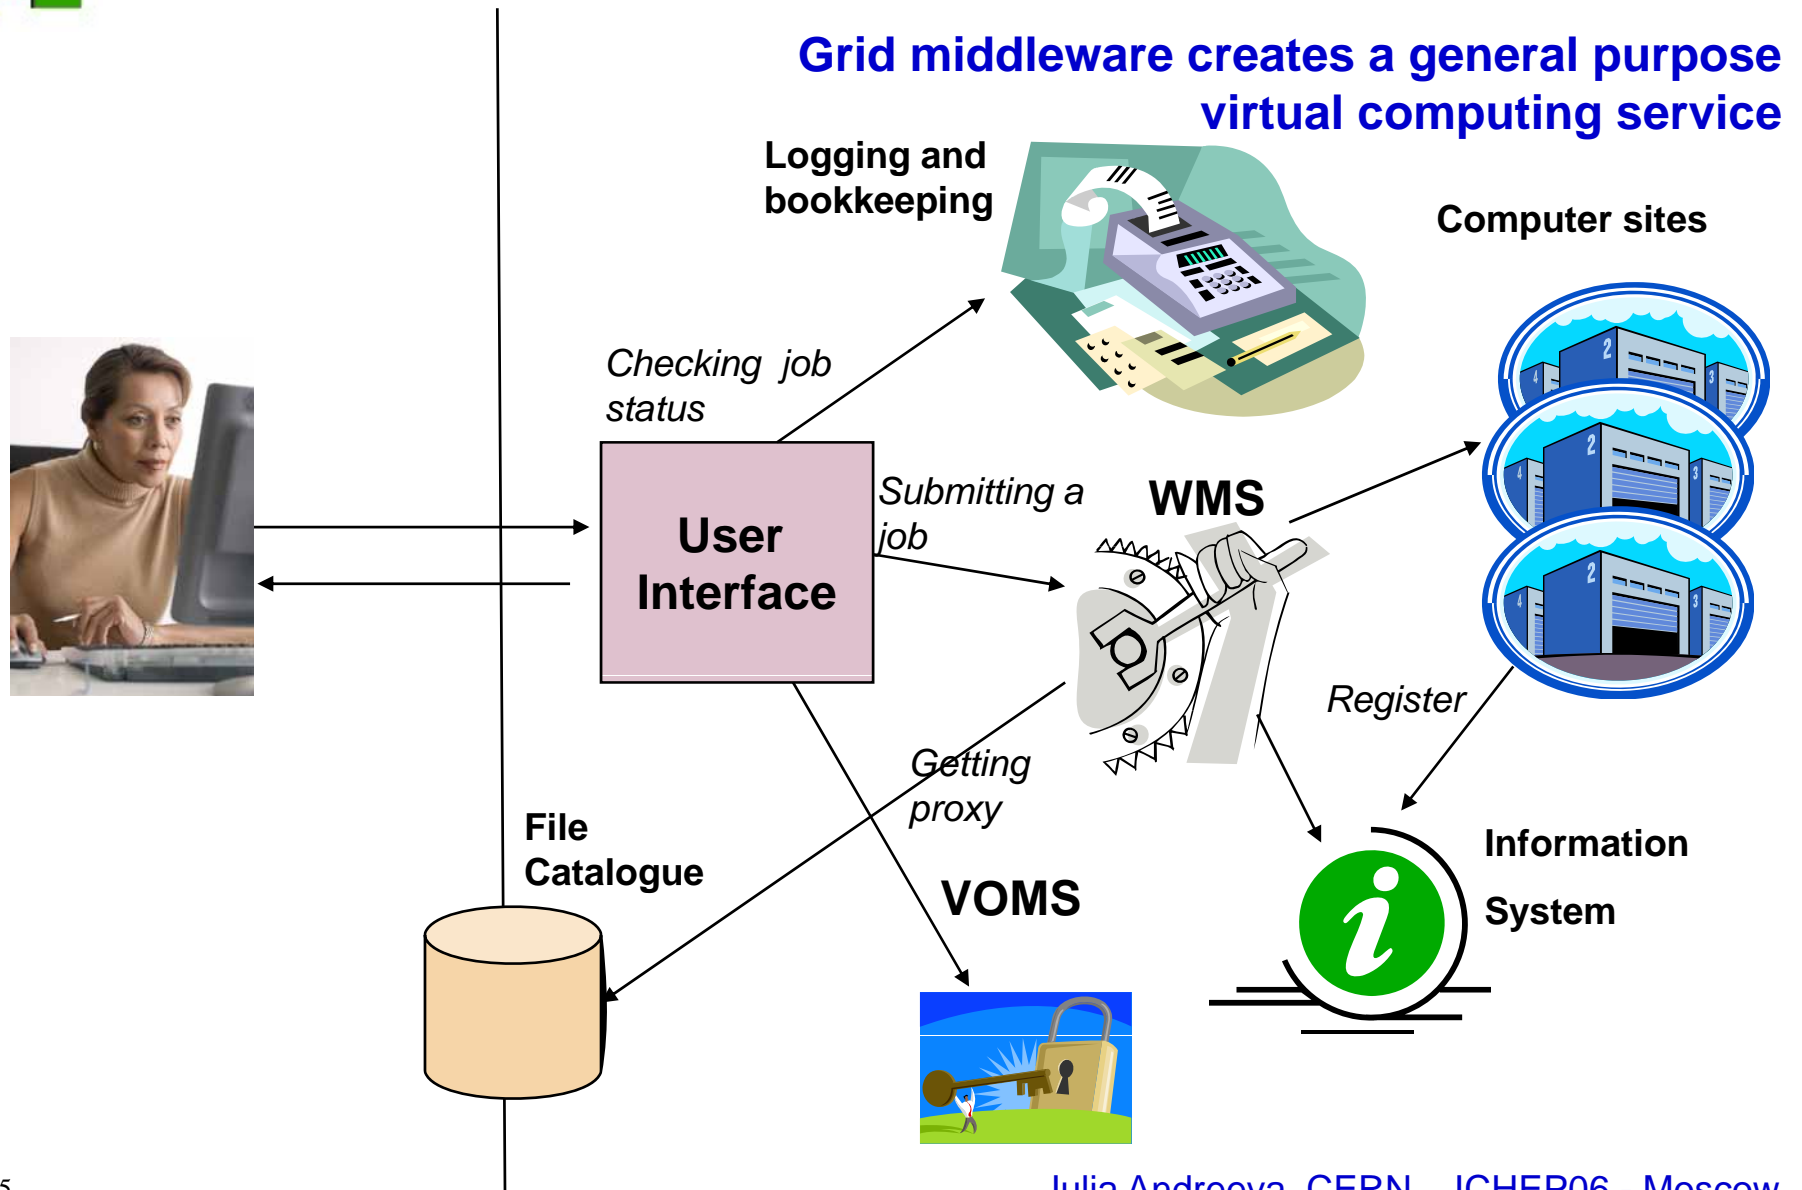
- | | | | | |
|------------------|----------------------|---------------------|-------------------|---------------------|
| T1_ASGC_Buffer | T1_CERN_Buffer | T1_CNAF_Buffer | T1_FNAL_Buffer | T1_FZK_Buffer |
| T1_IN2P3_Buffer | T1_PIC_Buffer | T1_PIC_Disk | T1_RAL_Buffer | T1_RAL_Stage |
| T2_Bari_Buffer | T2_Beijing_Buffer | T2_Belgium_IHE | T2_Belgium_UCL | T2_Budapest_Buffer |
| T2_CSCS_Buffer | T2_Caltech_Buffer | T2_DESY_Buffer | T2_Estonia_Buffer | T2_Florida_Buffer |
| T2_GRIF_DAPNIA | T2_GRIF_LAL | T2_GRIF_LLZ | T2_GRIF_LPNHE | T2_HEPGRID_UERJ |
| T2_IHEP_Disk | T2_IHEP_Buffer | T2_JINR_Buffer | T2_KNU_Buffer | T2_Legnaro_Buffer |
| T2_London_Brunel | T2_London_IC_HEP | T2_London_RHUL | T2_MIT_Buffer | T2_Nebraska_Buffer |
| T2_Pisa_Buffer | T2_Purdue_Buffer | T2_RWTH_Buffer | T2_Rome_Buffer | T2_SINGAPORE_Buffer |
| T2_SPRACE_Buffer | T2_SouthGrid_Bristol | T2_SouthGrid_RALPPO | T2_Spain_CIEMAT | |


Maximum: 898.73 MB/s, Minimum: 0.02 MB/s, Average: 1.07.49 MB/s, Current: 529.55 MB/s

all sites ↔ all sites



Data Analysis on the Grid

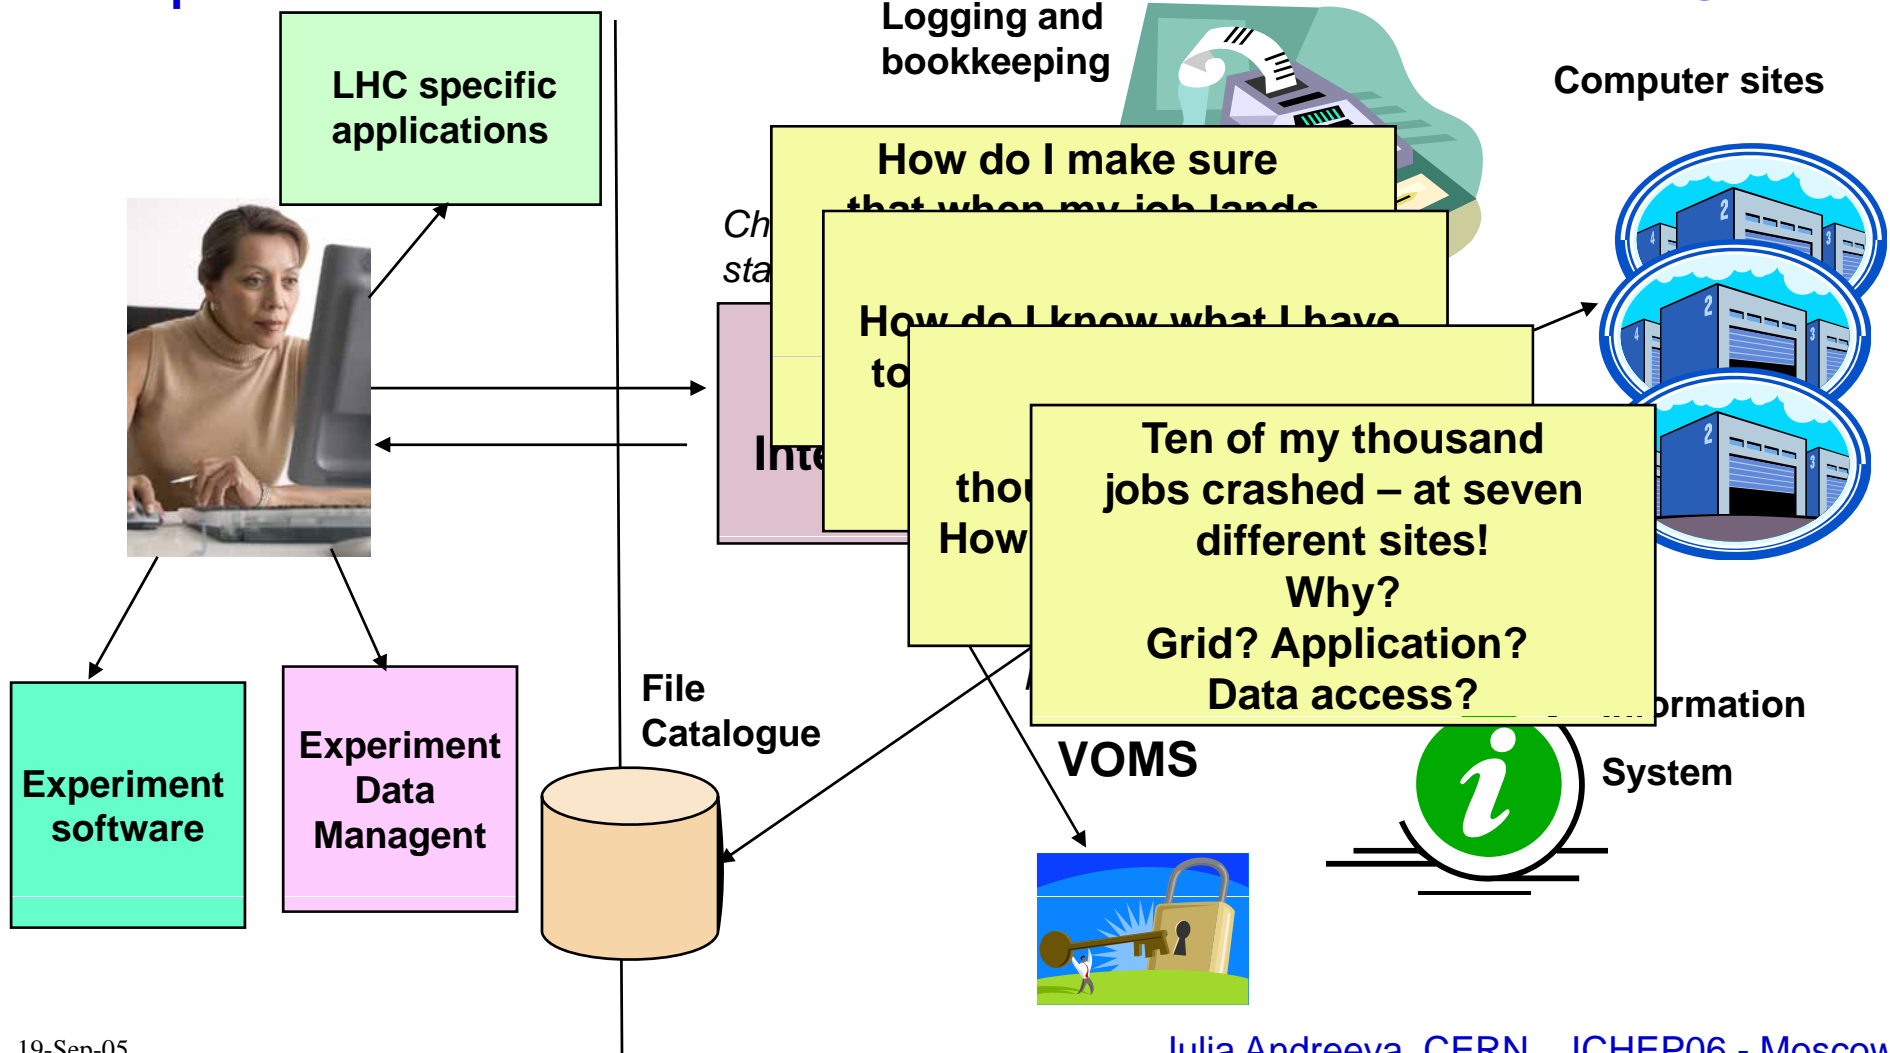





Data Analysis on the Grid

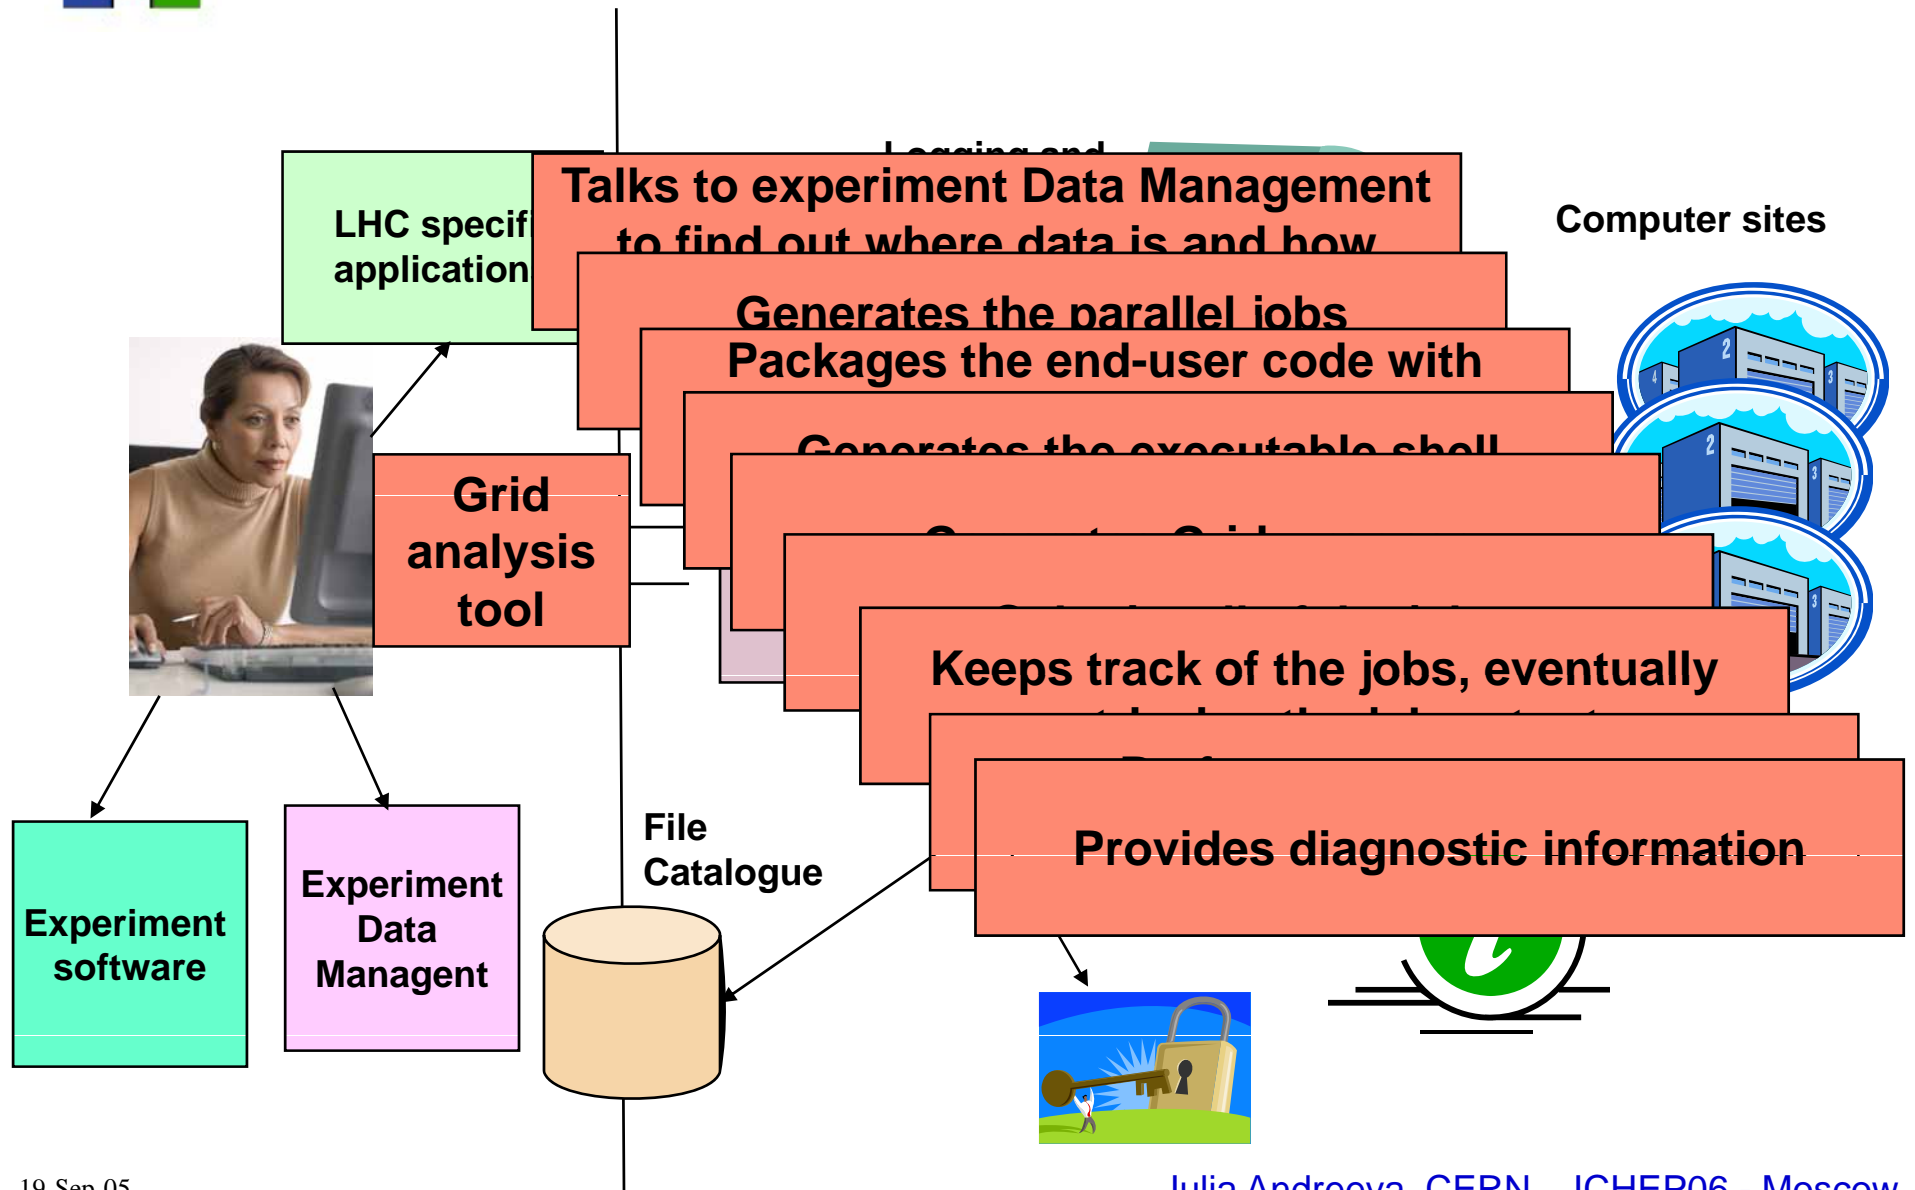
But there is also a very complex VO specific environment


Grid middleware creates a general purpose virtual computing service





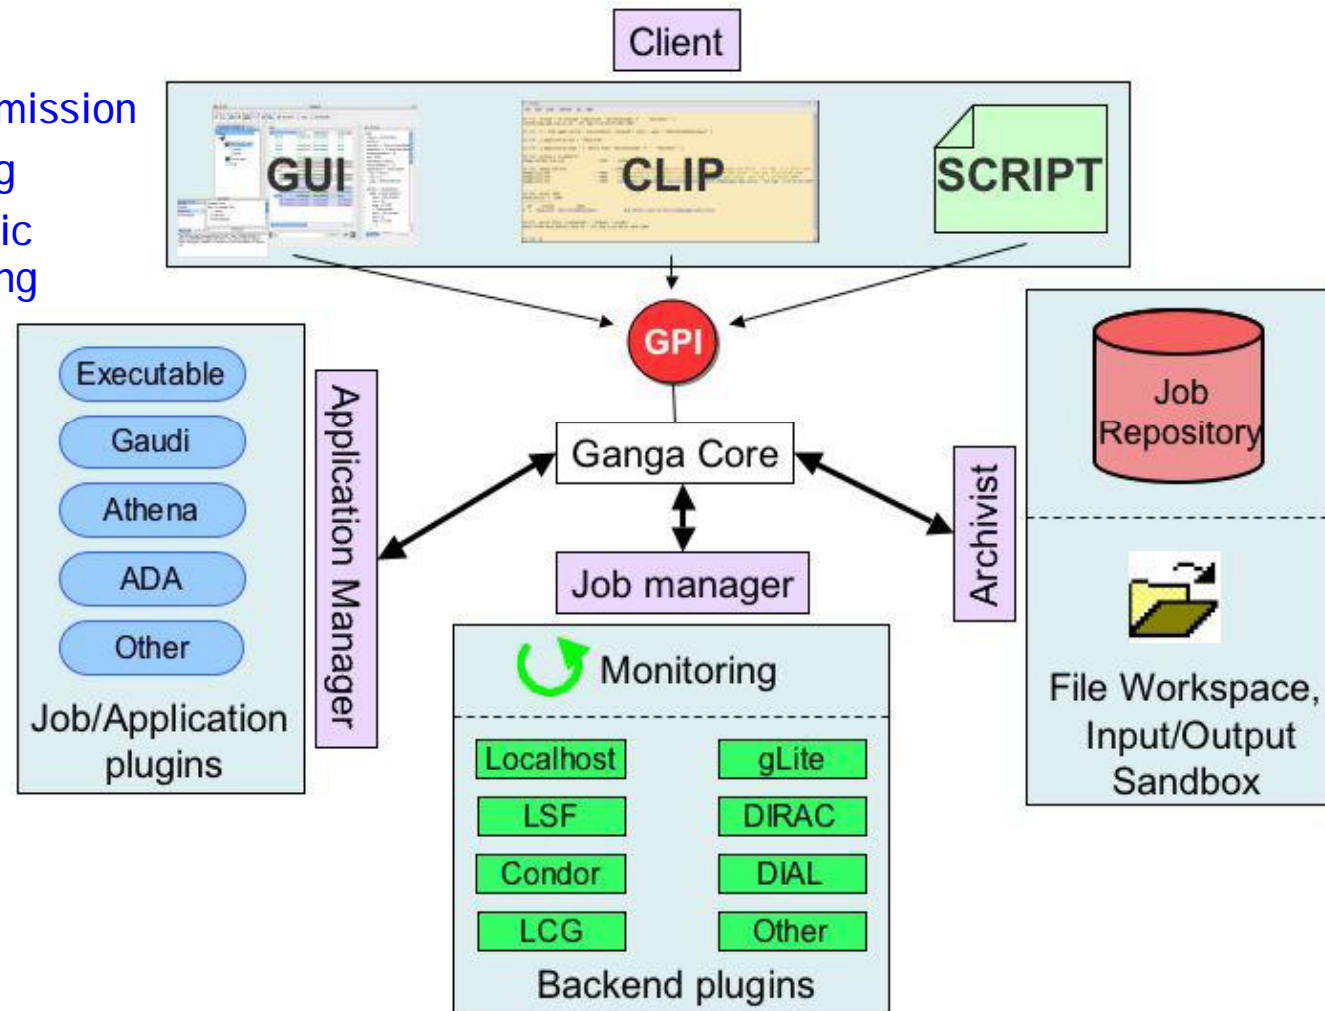

Each VO needs tools to hide the complexity of the distributed application environment





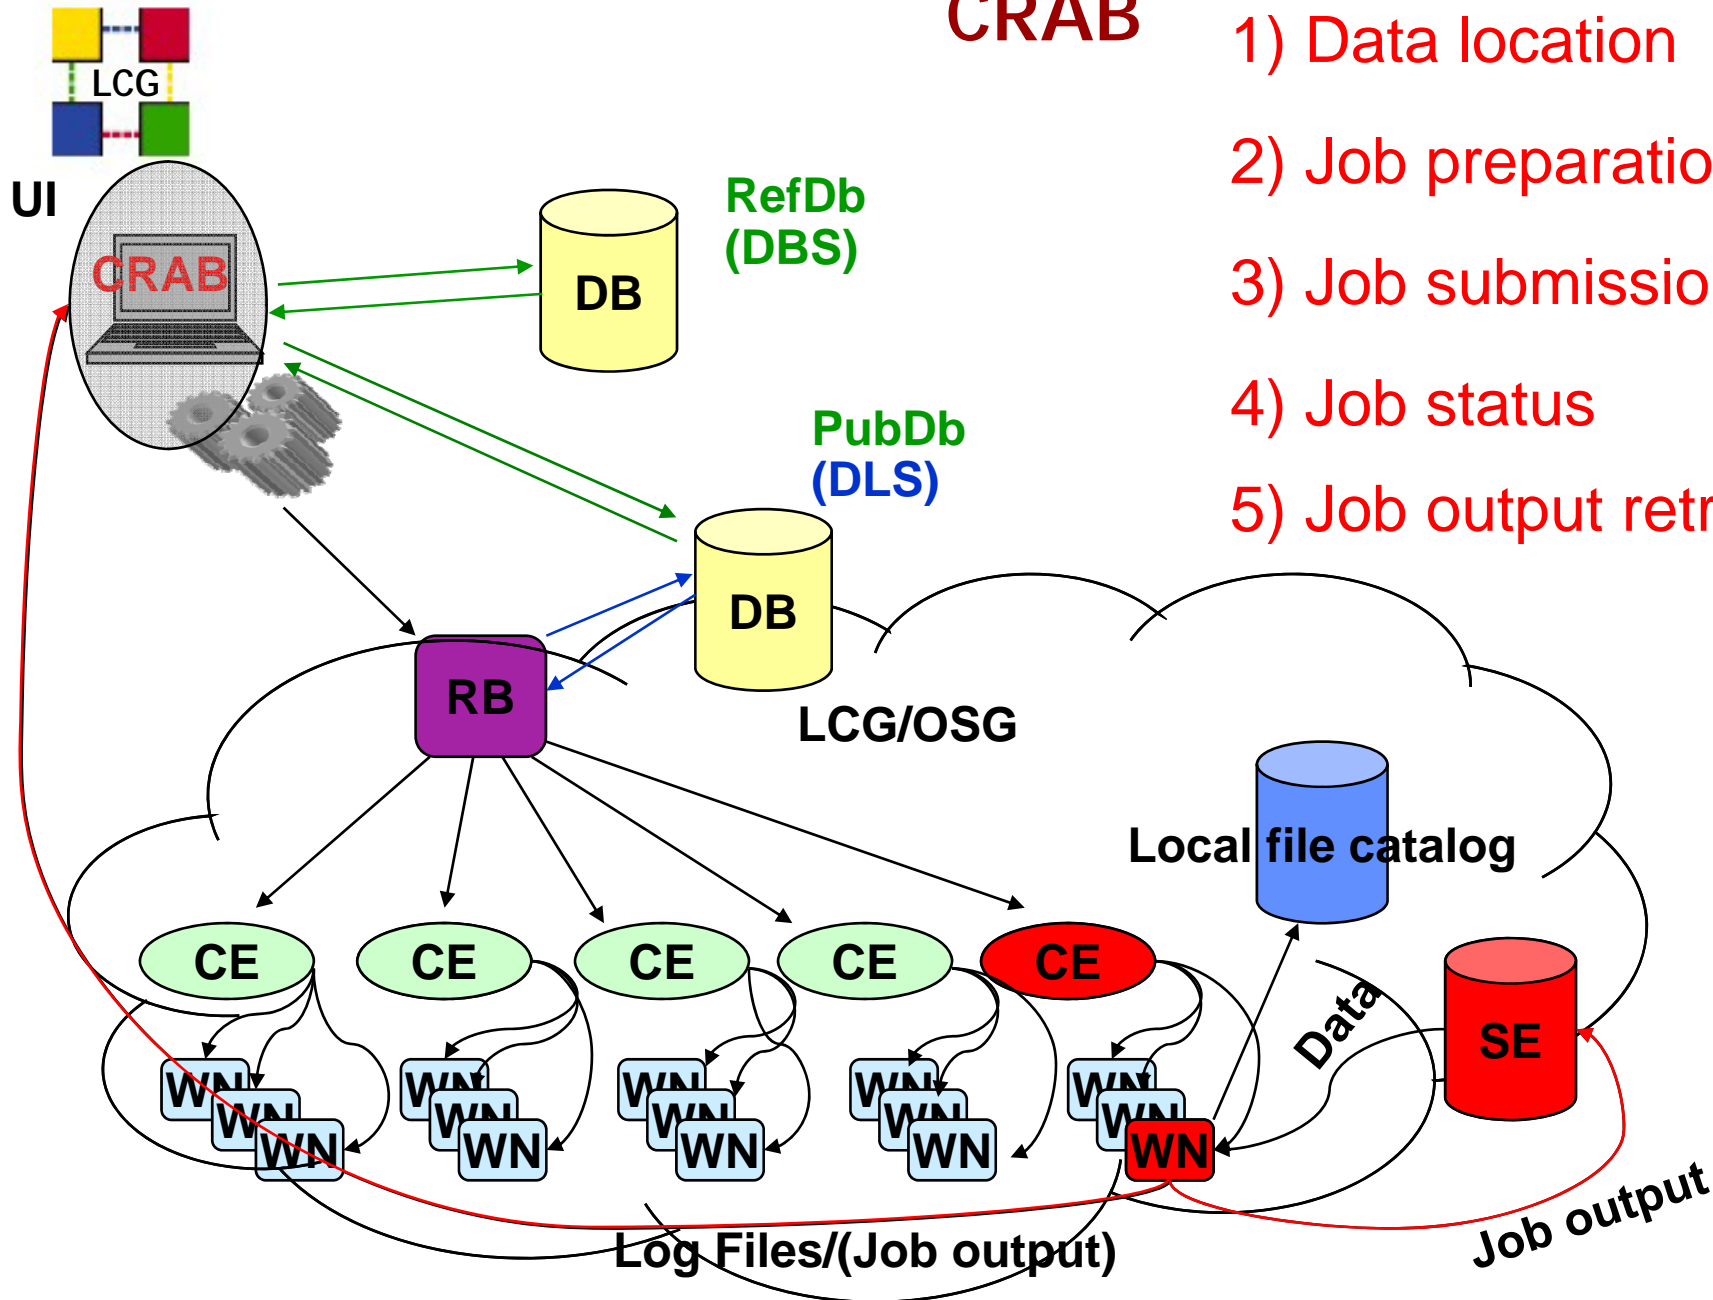
Ganga


- Job Submission
- Archiving
- Automatic monitoring



CRAB

- 1) Data location
- 2) Job preparation
- 3) Job submission
- 4) Job status
- 5) Job output retrieval

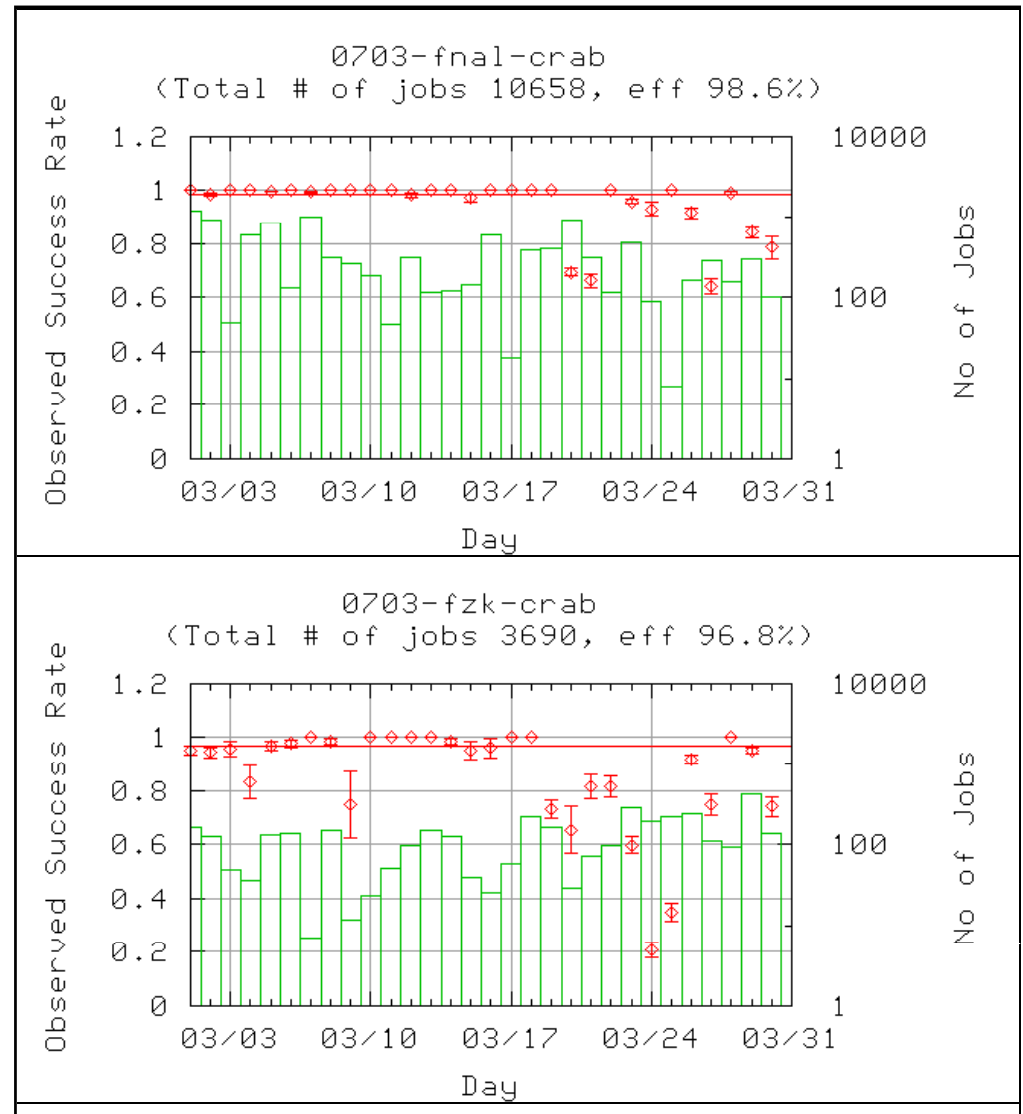





Analysis Job Success Rates

**CMS CRAB jobs
submitted through EGEE
Resource Brokers - March 2007**

**User view – after automatic job
re-submission by the Resource
Brokers**





The Worldwide LHC Computing Grid

- The LHC physics data analysis service distributed across the world
 - CERN, 11 large *Tier-1* centres, over 100 active *Tier-2* centres
- Status at April 2007
 - Established the 10 Gigabit/sec optical network that interlinks CERN and the Tier-1 centres
 - Demonstrated data distribution from CERN to the Tier-1 centres at 1.3 GByte/sec – the rate that will be needed in 2008
 - ATLAS and CMS can each transfer 1 PetaByte of data per month between their computing centres
 - Running ~2 million jobs each month across the grid
 - The distributed grid operation, set up during 2005, has reached maturity, with responsibility shared across 7 sites in Europe, the US and Asia
 - End-user analysis tools enabling “real physicists” to profit from this worldwide data-intensive computing environment

