

LCG: the LHC Computing Grid project


Overview

May 2007



Les Robertson
LCG Project Leader





The LHC Accelerator



The accelerator generates 40 million particle collisions (events) every second at the centre of each of the four experiments' detectors






LHC DATA

This is reduced by online computers that filter out a few hundred “good” events per sec.




Which are recorded on disk and magnetic tape at 100-1,000 MegaBytes/sec

→ ~15 PetaBytes per year for all four experiments





Data Handling and Computation for Physics Analysis



detector

event filter
(selection & reconstruction)

reconstruction

raw data

event reprocessing

event summary data

analysis

batch physics analysis

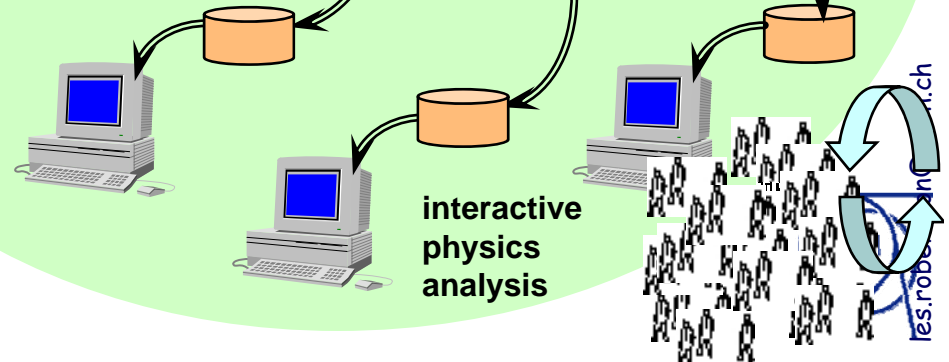
processed data

event simulation

simulation

analysis objects
(extracted by physics topic)

interactive physics analysis





WLCG Collaboration

- **The Collaboration**

- 4 LHC experiments
- ~140 computing centres
- 12 large centres
(Tier-0, Tier-1)
- 38 *federations* of smaller
"Tier-2" centres
- ~35 countries

- **Resources**

- Contributed by the countries participating in the experiments
- Commitment made each October for the coming year
- 5-year forward look

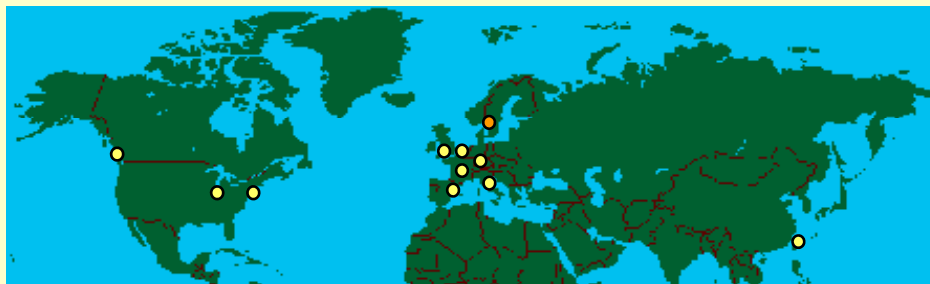




LCG Service Hierarchy

Tier-0 - the accelerator centre

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



Canada – Triumf (Vancouver)
France – IN2P3 (Lyon)
Germany – Forschungszentrum Karlsruhe
Italy – CNAF (Bologna)
Netherlands – NIKHEF/SARA (Amsterdam)
Nordic countries – distributed Tier-1

Spain – PIC (Barcelona)
Taiwan – Academia Sinica (Taipei)
UK – CLRC (Oxford)
US – FermiLab (Illinois)
– Brookhaven (NY)

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 - ~130 centres in ~35 countries

- **End-user (physicist, research group) analysis** –
where the discoveries are made
- Simulation





Distribution of Computing Services




Summary of Computing Requirements

All experiments - 2008
From LCG TDR - June 2008

New data will grow at about 15 PetaBytes per year – with two copies

about 100,000 CPU cores

	CERN	All Tier-1s	All Tier-2s	Total
CPU (MSPECint2000s)	25	56	61	142
Disk (PetaBytes)	7	31	19	57
Tape (PetaBytes)	18	35		53



Significant fraction of the resources distributed over more than 120 computing centres



Solution: the Grid

- Use the Grid to unite computing resources of particle physics institutes around the world

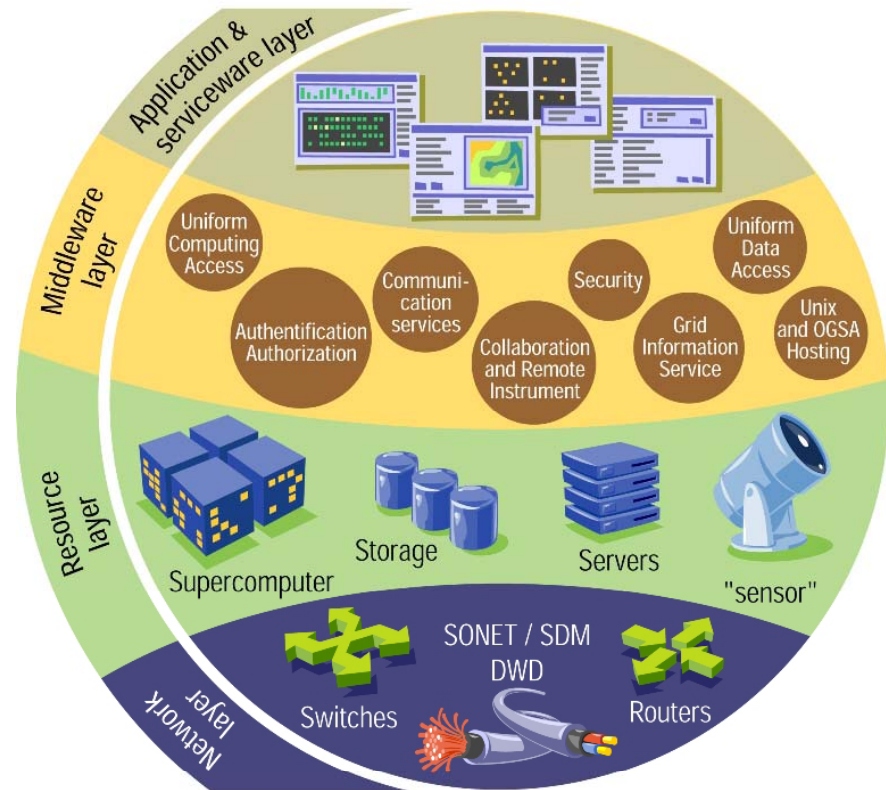
The **World Wide Web** provides seamless access to information that is stored in many millions of different geographical locations

The **Grid** is an infrastructure that provides seamless access to computing power and data storage capacity distributed over the globe



How does the Grid work?

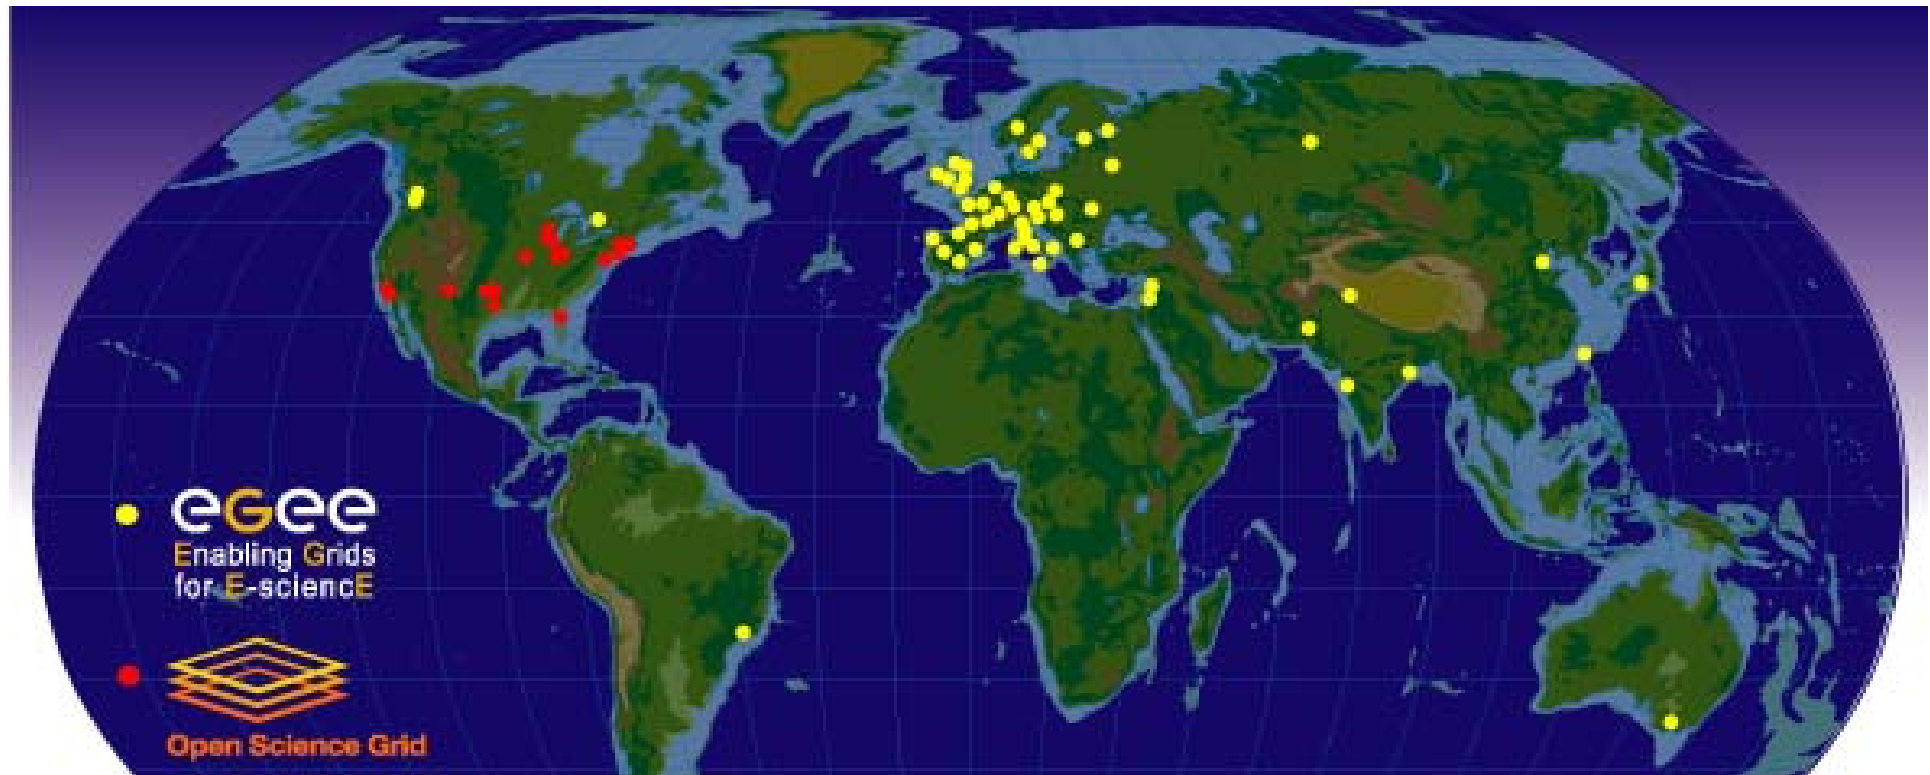
- It relies on special software, called **middleware**.
- Middleware automatically finds the **data** the scientist needs, and the **computing power** to analyse it.
- Middleware balances the load on different resources. It also handles **security, accounting, monitoring** and much more.





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.





LHC Computing → Multi-science Grid

- 1999 - MONARC project
 - First LHC computing architecture - hierarchical distributed model
- 2000 - growing interest in grid technology
 - HEP community main driver in launching the DataGrid project
- 2001-2004 - EU DataGrid project
 - middleware & testbed for an operational grid
- 2002-2005 - LHC Computing Grid - LCG
 - deploying the results of DataGrid to provide a production facility for LHC experiments
- 2004-2006 - EU EGEE project phase 1
 - starts from the LCG grid
 - shared production infrastructure
 - expanding to other communities and sciences



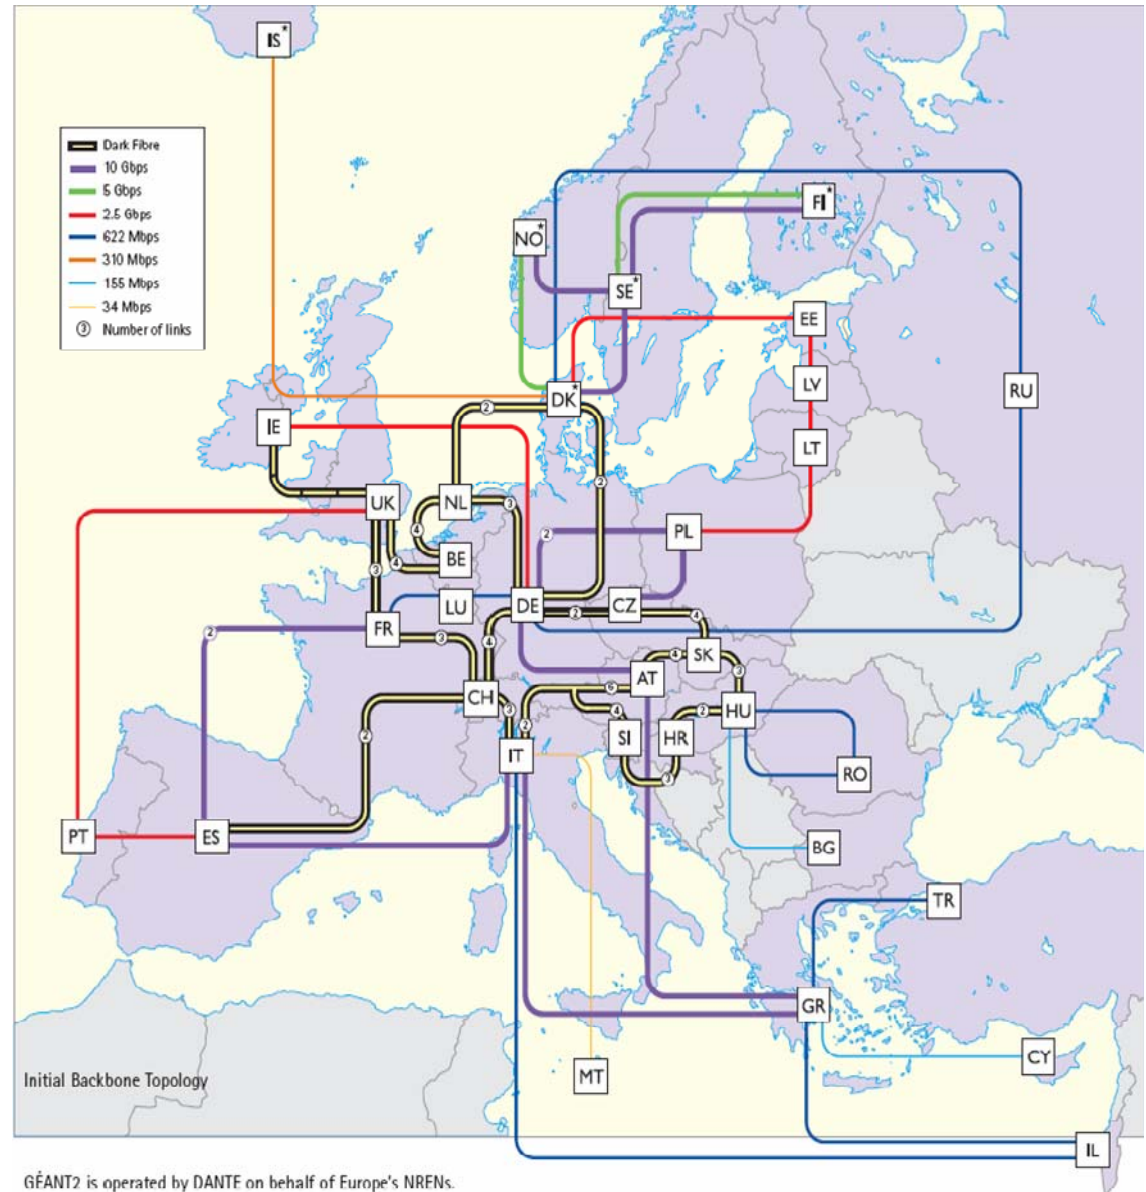
- **More than 20 applications from 7 domains**
 - High Energy Physics (**Pilot domain**)
 - 4 LHC experiments
 - Other HEP (DESY, Fermilab, etc.)
 - Biomedicine (**Pilot domain**)
 - Bioinformatics
 - Medical imaging
 - Earth Sciences
 - Earth Observation
 - Solid Earth Physics
 - Hydrology
 - Climate
 - Computational Chemistry
 - Fusion
 - Astronomy
 - Cosmic microwave background
 - Gamma ray astronomy
 - Geophysics
 - Industrial applications



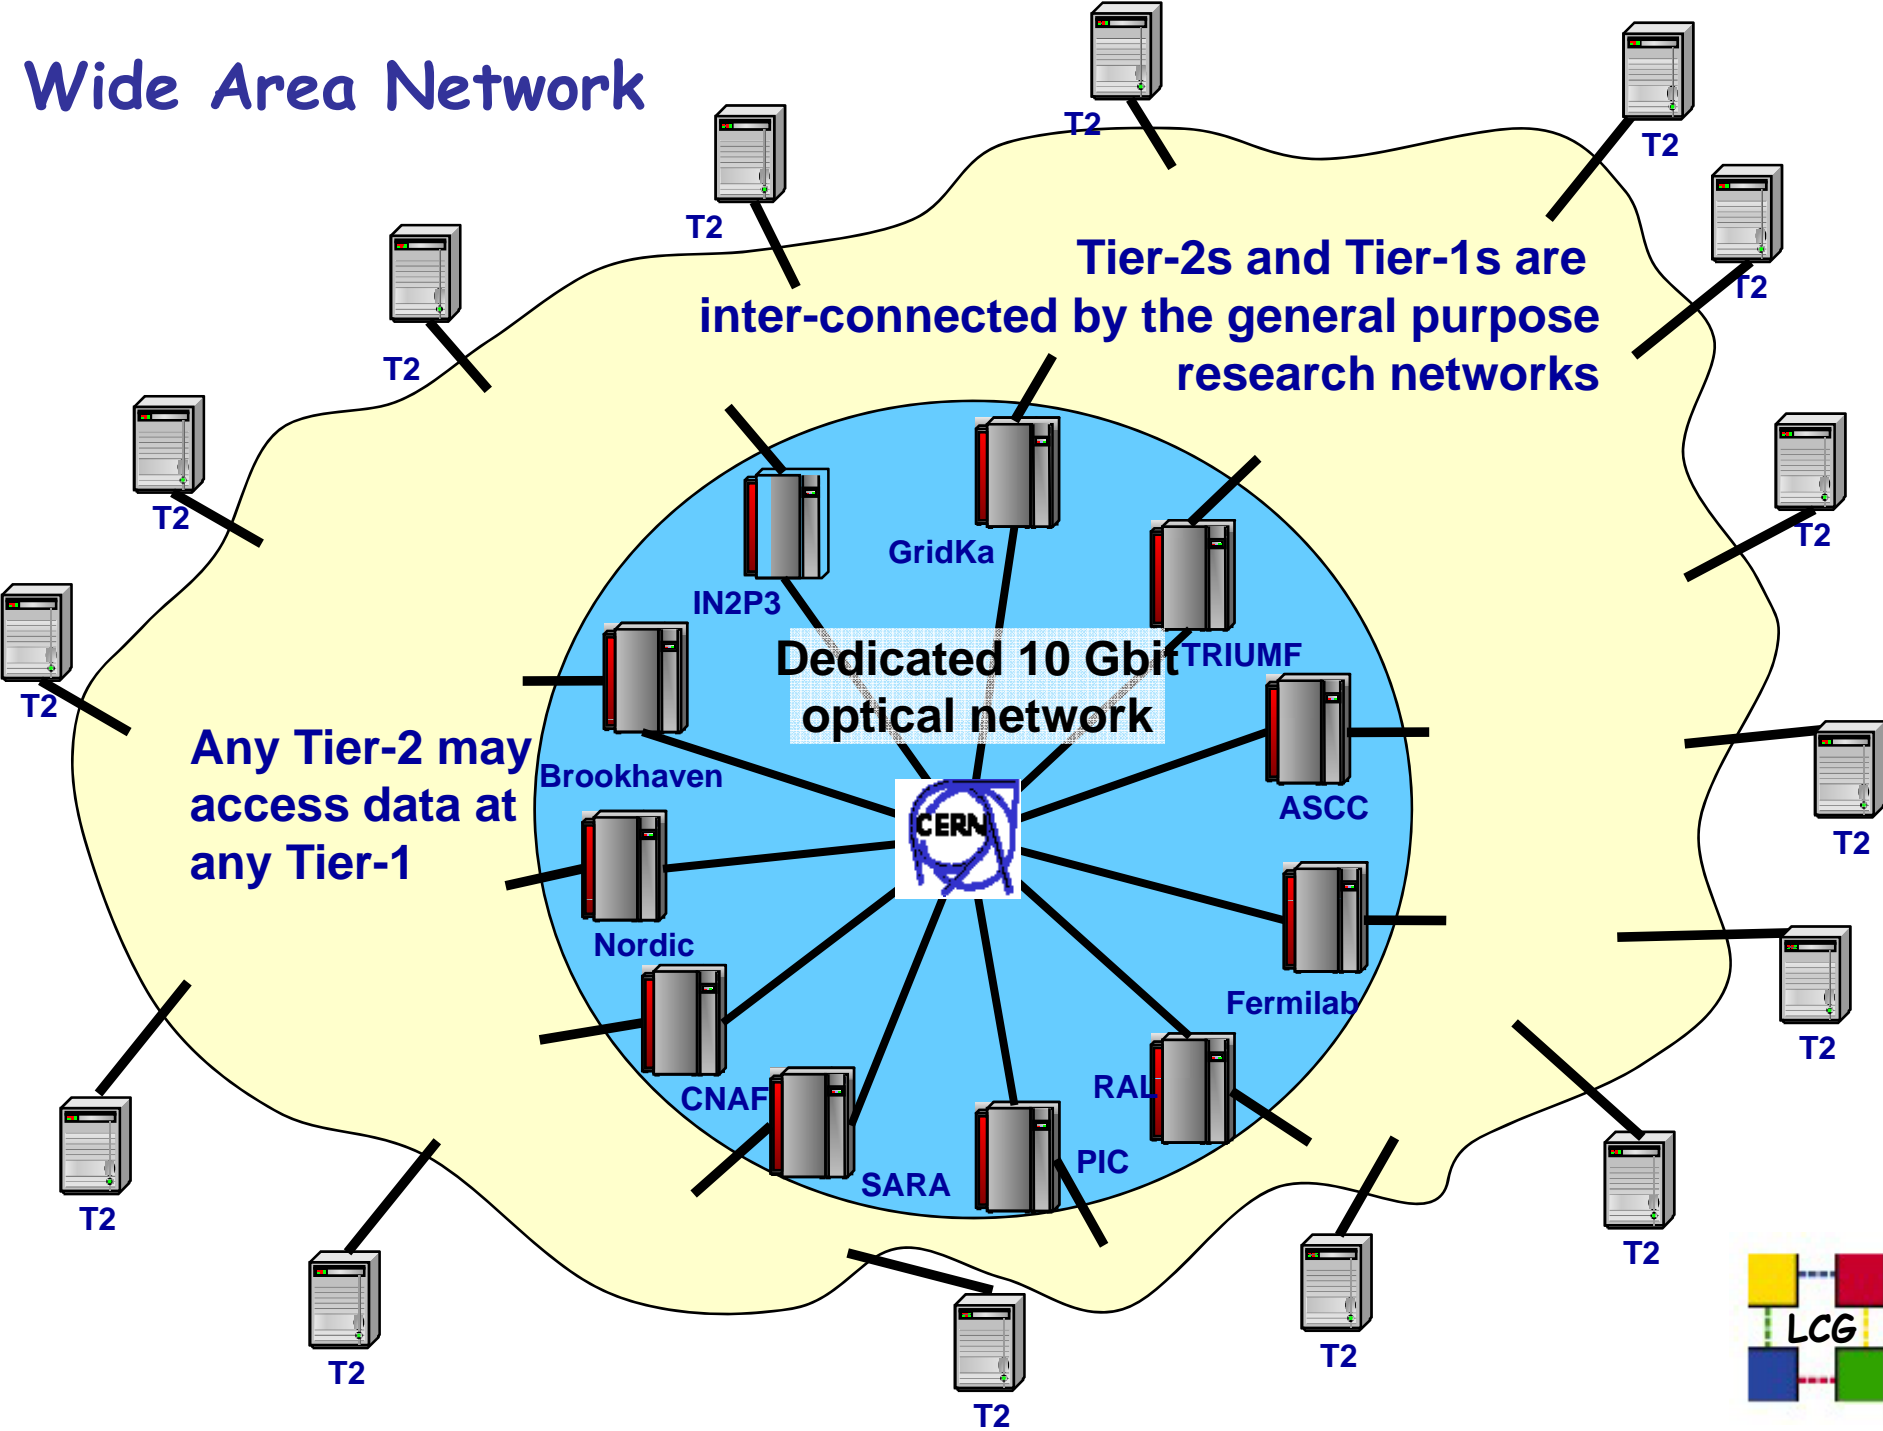


The new European Network Backbone

- LCG working group with Tier-1s and national/regional research network organisations
- New GÉANT 2 - research network backbone
 - Strong correlation with major European LHC centres (Swiss PoP at CERN)
 - Core links are fibre
- Two 622 Mbps circuits to Israel




Wide Area Network






CPU Usage - LHC Experiments March 2007




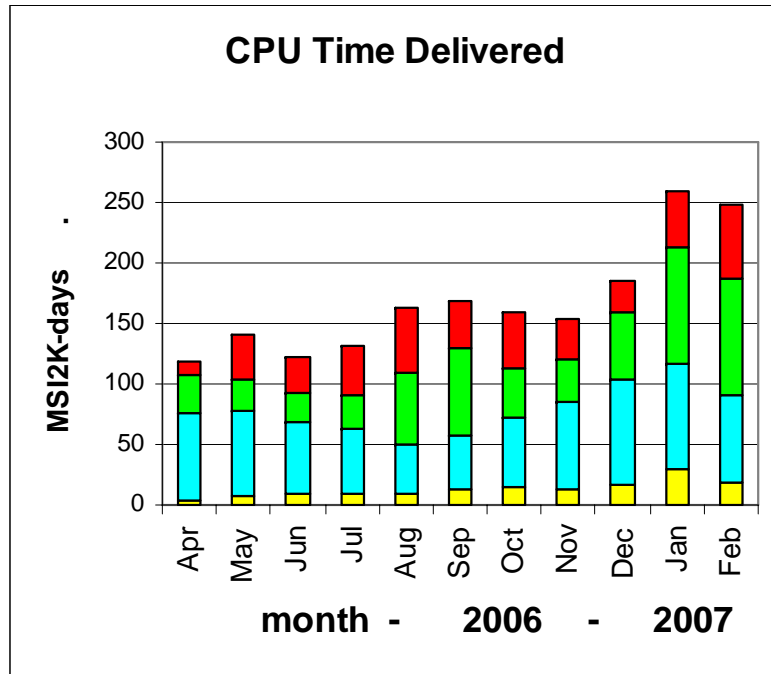
CERN **20%**
11 Tier-1s **60%**
140 Tier-2s **20%**

Tier-2 Sites Contributing in
March 2007





Tier-1s and CERN





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





Ramp-up to First Physics Run



Evolution of installed capacity from April 06 to June 07
Target capacity from MoU pledges for 2007 (due July07)
and 2008 (due April 08)





Traffic load generator: **LoadTest**



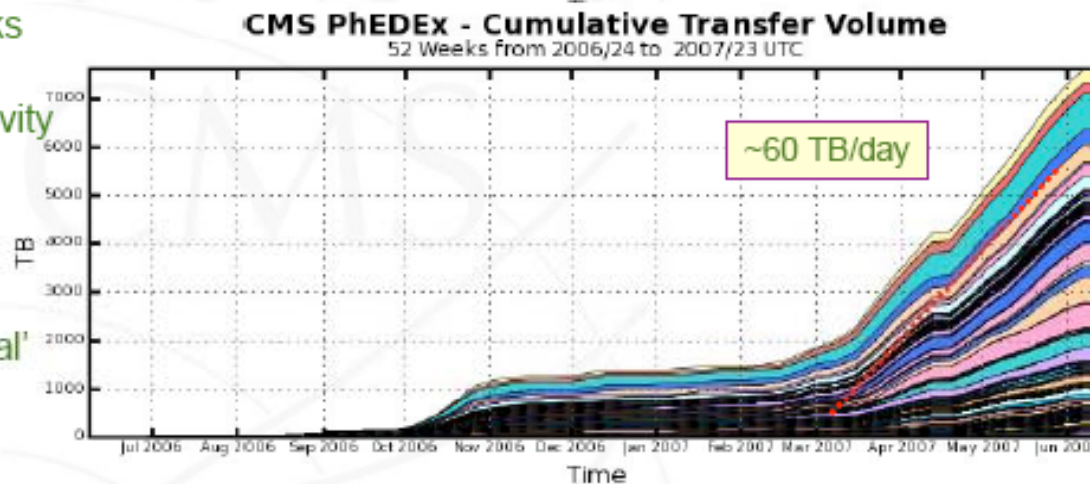

➤ A flexible infrastructure to generate data transfer traffic among CMS Tiers

❑ “fake” but “real”

- ❖ No real physics files
- ❖ Fully PhEDEx-compliant

❑ The planning:

- ❖ 3 full cycles (5-weeks each) before Jun07
- ❖ Basically a 24/7 activity since mid-Feb07
- ❖ T0→T1(tape),
T1↔T1,
T1↔T2 'regional',
T1↔T2 'non-regional'



➤ Consistent data transfer volume, and very useful for Tiers


❑ Being extended into CSA07 preparation activities






Site Reliability - CERN+Tier-1s

- Set of ~30 basic grid tests - CE, SE, BDII
- All tests must succeed or site is considered "down"
- CMS has developed experiment specific tests using the SAM framework to have a VO-specific view
- Over past 6 months:
 - average (11 sites) 83%
 - 8 best site average 90%
- Target for best 8 sites
 - May 06-May 07 88%
 - June 07-Nov 07 93%
 - From Dec 07 95%

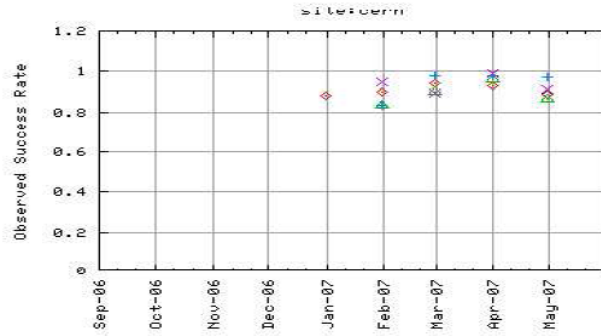




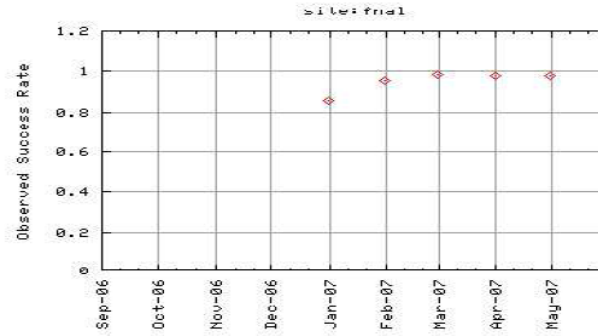
Job success rates

End-user analysis: Crab (CMS), Ganga (ATLAS)
Pilot Jobs (LHCb), Job Agents (ALICE)
Only jobs submitted via EGEE Resource Brokers

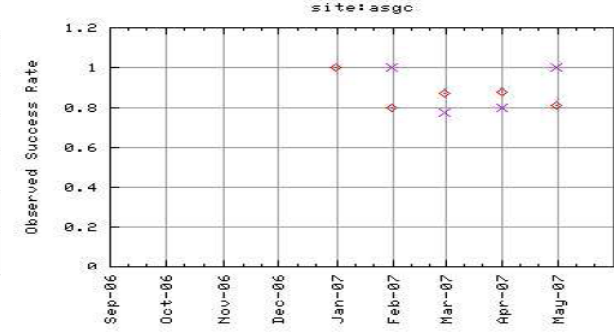
CERN




FNAL




ASGC



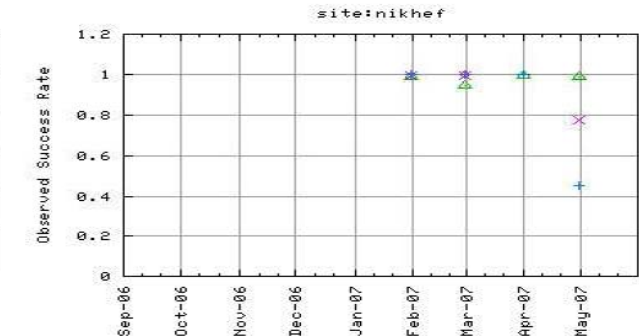
RAL



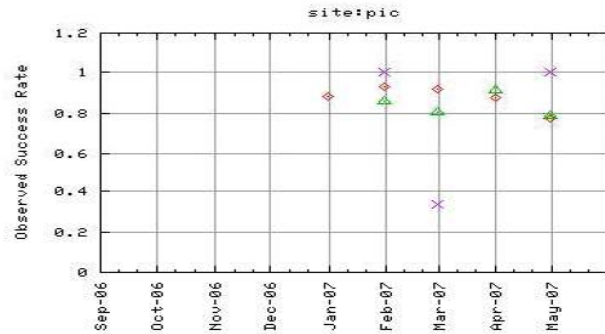
FZK




NIKHEF




PIC



CNAF



IN2P3

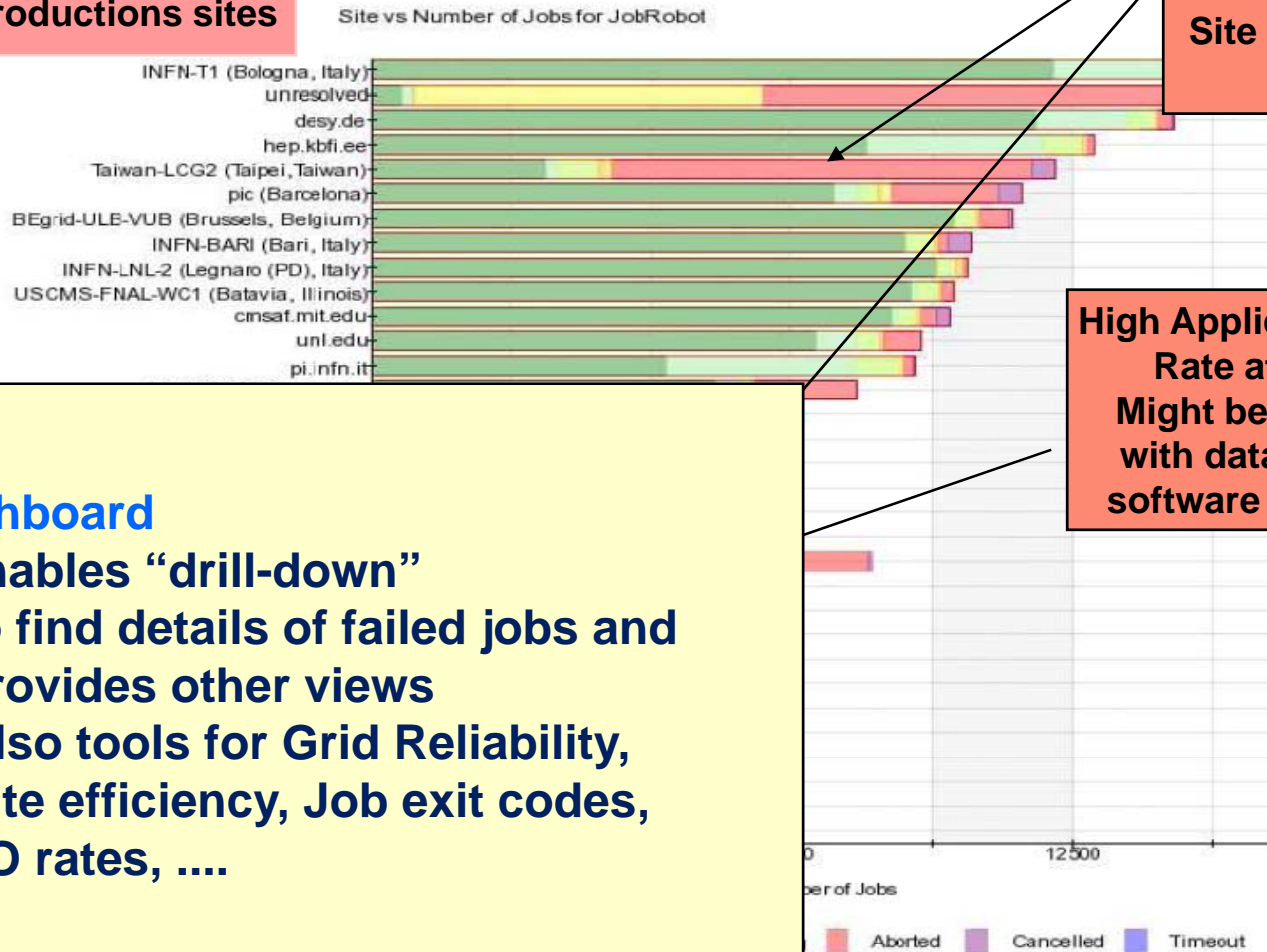




Experiment "Dashboard" - example CMS

CMS JobRobot
Sends test analysis jobs
to all CMS productions sites

High Grid failure rate
At the site
Site misconfiguration





The Worldwide LHC Computing Grid

- The LHC physics data analysis service distributed across the world
 - CERN, 11 large *Tier-1* centres,
~ 140 active *Tier-2* centres
- Status at May 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 1-2 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

