


LCG - The Worldwide LHC Computing Grid

Resource Review Board

Project Status

CERN
24 October 2006
Les Robertson
LCG Project Leader





Applications Support

- All products now in production use by experiments
- Merger of the ROOT and SEAL packages largely complete - experiments have started to migrate to the new version (ROOT 5)
- Production release of the re-engineered common relational database interface package (CORAL)
- FLUKA released with new licence model
- The HepMC package installed in the LCG external area and maintained by FNAL effort

Persistency Framework

POOL

Conditions Database

Core libraries and services - ROOT

analysis framework

components for experiment frameworks

maths library

dictionary, ..

Simulation

Simulation framework

GEANT4


Fluka

Physics validation


Garfield

MC generator services

Software Process & Infrastructure



Worldwide LHC Computing Grid





Grid Operations - Key Points (i)

- Baseline services from the Technical Design Report are in operation
- gLite 3 - New middleware distribution for the EGEE grid
 - Basis for LHC startup
 - Introduced for Service Challenge 4 - SC4 - started in June
 - New Workload Management System (grid-wide job scheduler) - now entering production
- Standard data distribution services - File Transfer Service (FTS)
- Distributed Database service entering operation
 - Still to take place at 4 of the Tier-1s
- A few key services still to be introduced
 - Implementations of a *new version* of the Storage Resource Manager (SRM) -provides a standard interface to storage at all WLCG sites, supporting the storage classes needed for LHC startup
 - Job scheduling taking account of the user's group and role within an experiment



Grid Operations – Key Points (ii)

- WLCG standard metrics for usage, reliability and performance being introduced
 - accounting introduced for Tier-1s and CERN (cpu and storage)
 - site reliability and availability measurement system introduced - reporting for Tier-1s & CERN from May
 - data transfer performance and reliability
 - job failure analysis
- Grid operations
 - All major LCG sites active
 - Daily monitoring and operations now mature -
 - in the EGEE grid coordination taken in turn by 5 sites
 - Evolution of EGEE regional operations support structure




Grid Usage

Between EGEE and OSG -

- ~60K jobs/day during data challenges this summer
- ~15K *simultaneous* jobs during prolonged periods


Jobs/day – June-05 → Aug-06



- 3 X increase in past twelve months

but still some way to go -
only ~20% of the 2008 target

Jobs/day– Aug-Sep 06



CMS

ATLAS
lesrobertson@cern.ch





Grid Usage & Reliability

CPU Usage - CERN+Tier-1s EGEE + OSG

- 160 K processor-days/month
- 66% in Tier-1s

Site Reliability – seen from the Grid
CERN+Tier-1s
 averaging 83% of the 2006 target
 best 8 sites average 91% of target





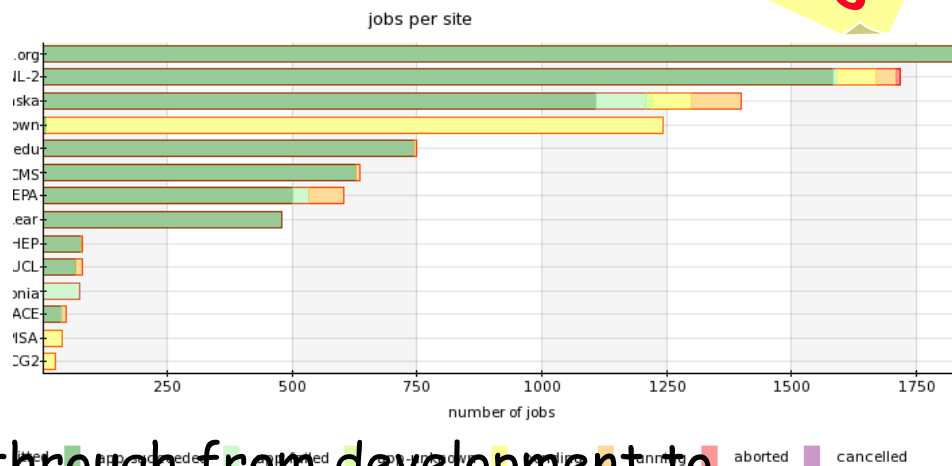
Job Reliability Monitoring

Work in progress

- System to process and analyse job logs implemented in a *dashboard* for some of the major activities in ATLAS and CMS

- Errors identified, frequency reported to developers, sites

- Expect to see results feeding through from development to products in a fairly short time



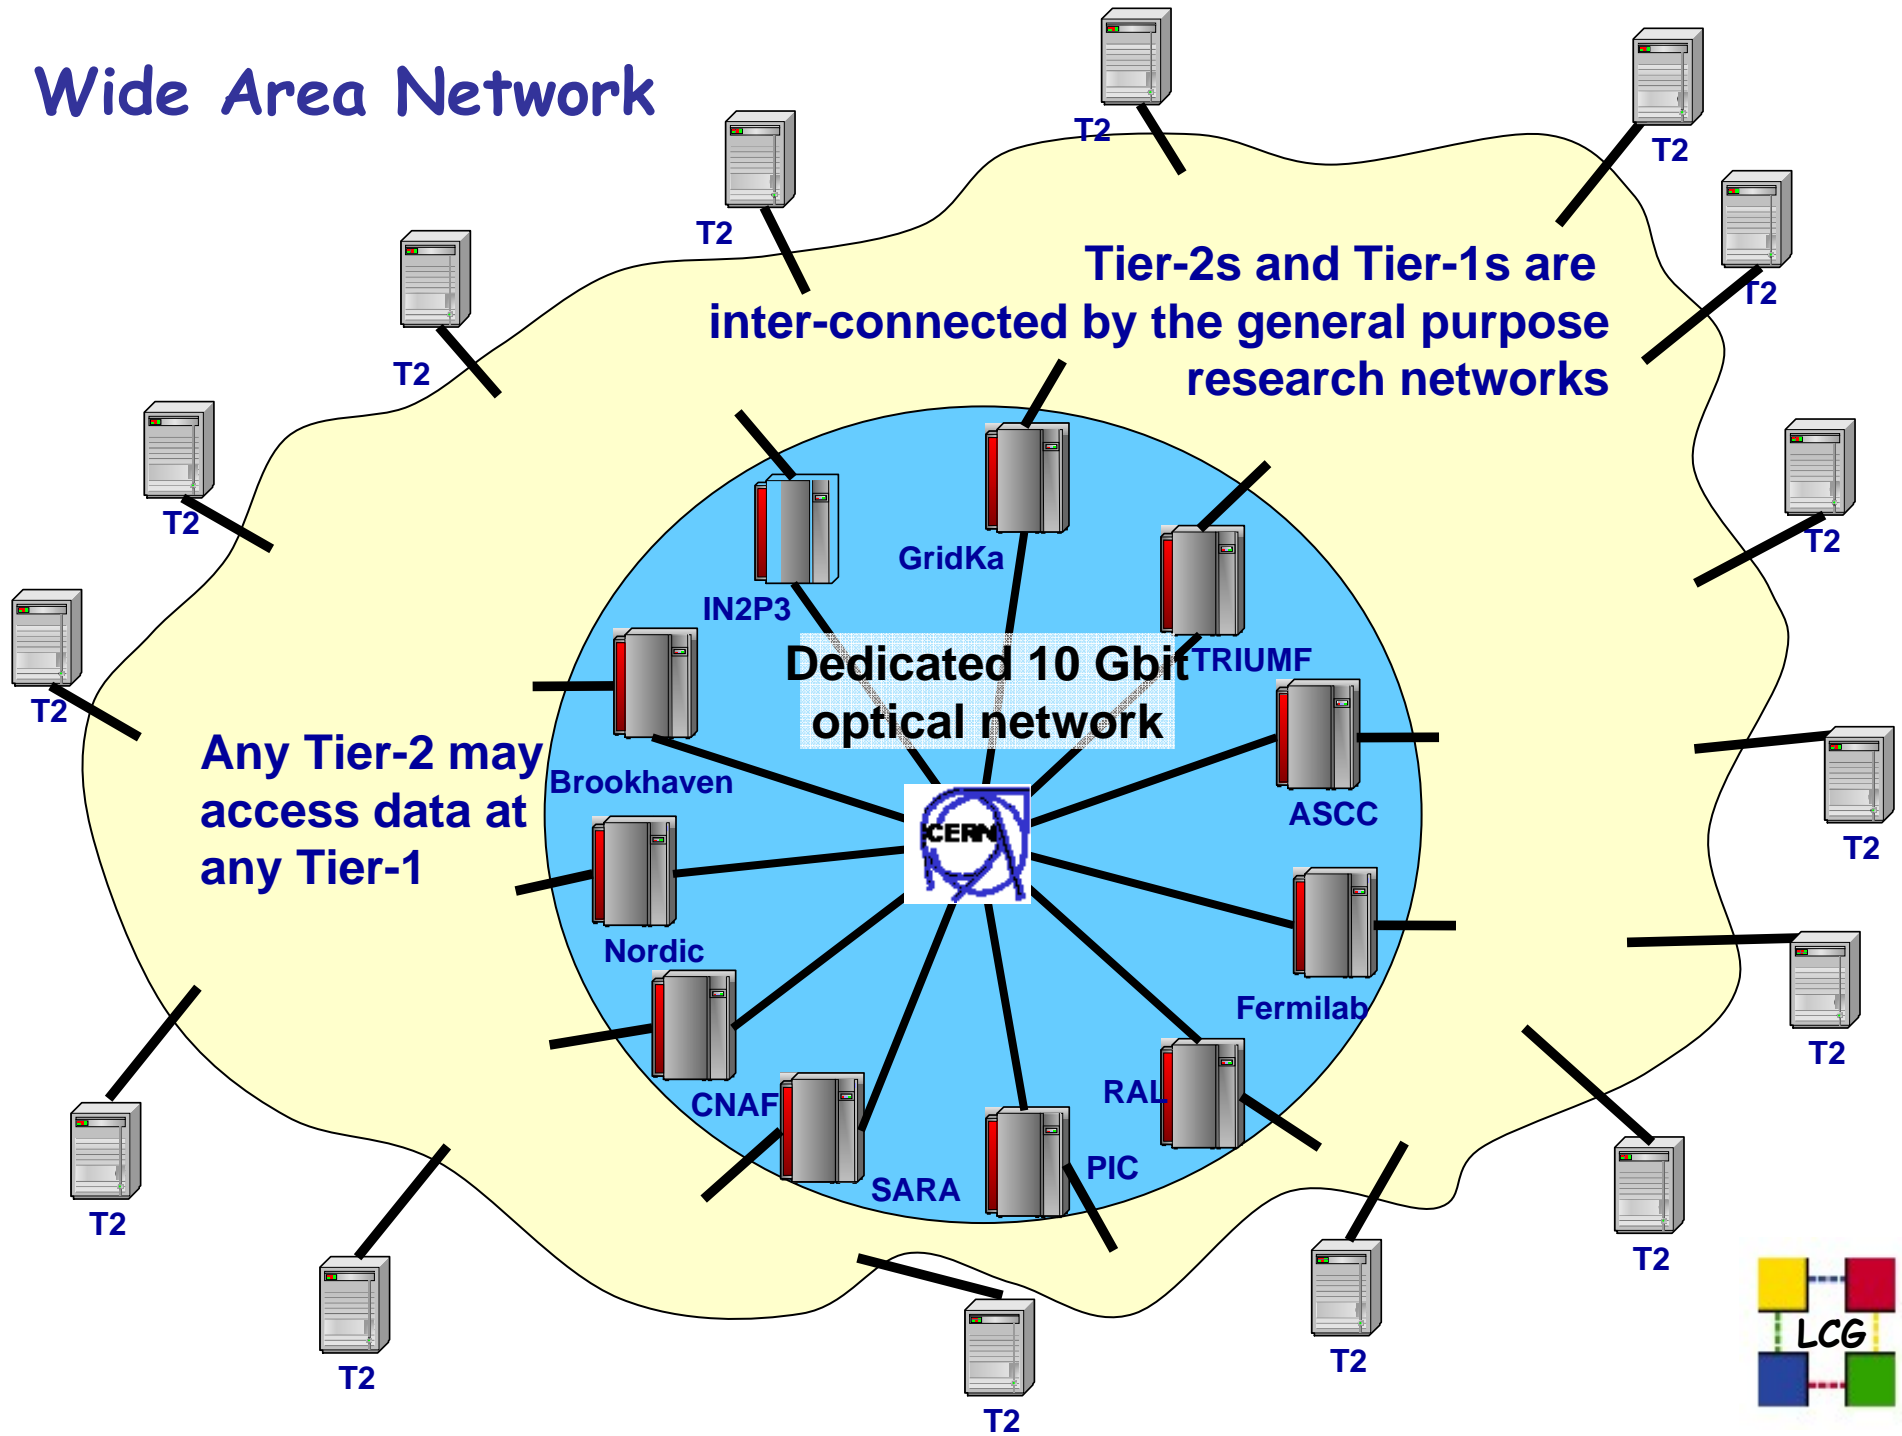
pic	116	163	41.58%
ce04.pic.es:2119/jobmanager-lcgpbs-cms	116	163	41.58%
Jobids	Number of jobs	Successful?	
See all the jobids...	103	No	
See all the jobids...	22	No	
See all the jobids...	11	No	
See all the jobids...	11	No	
See all the jobids...	Submitted####Waiting####Ready####Scheduled#Job	0	
See all the jobids...	successfully submitted to Globus####Done#Got a job	0	
See all the jobids...	held event reason: Unspecified gridmanager	0	
See all the jobids...	error####Done#Job got an error while in the CondorG	0	
See all the jobids...	queue.	es	
See all the jobids...	15	Yes	
See all the jobids...	1	Yes	
hep.kbfi.ee	101	106	48.79%
desy.de	128	93	57.92%
USCMS-FNAL-WC1	668	313	68.09%
ITEP	141	61	69.80%

- Daily report on most important site problems

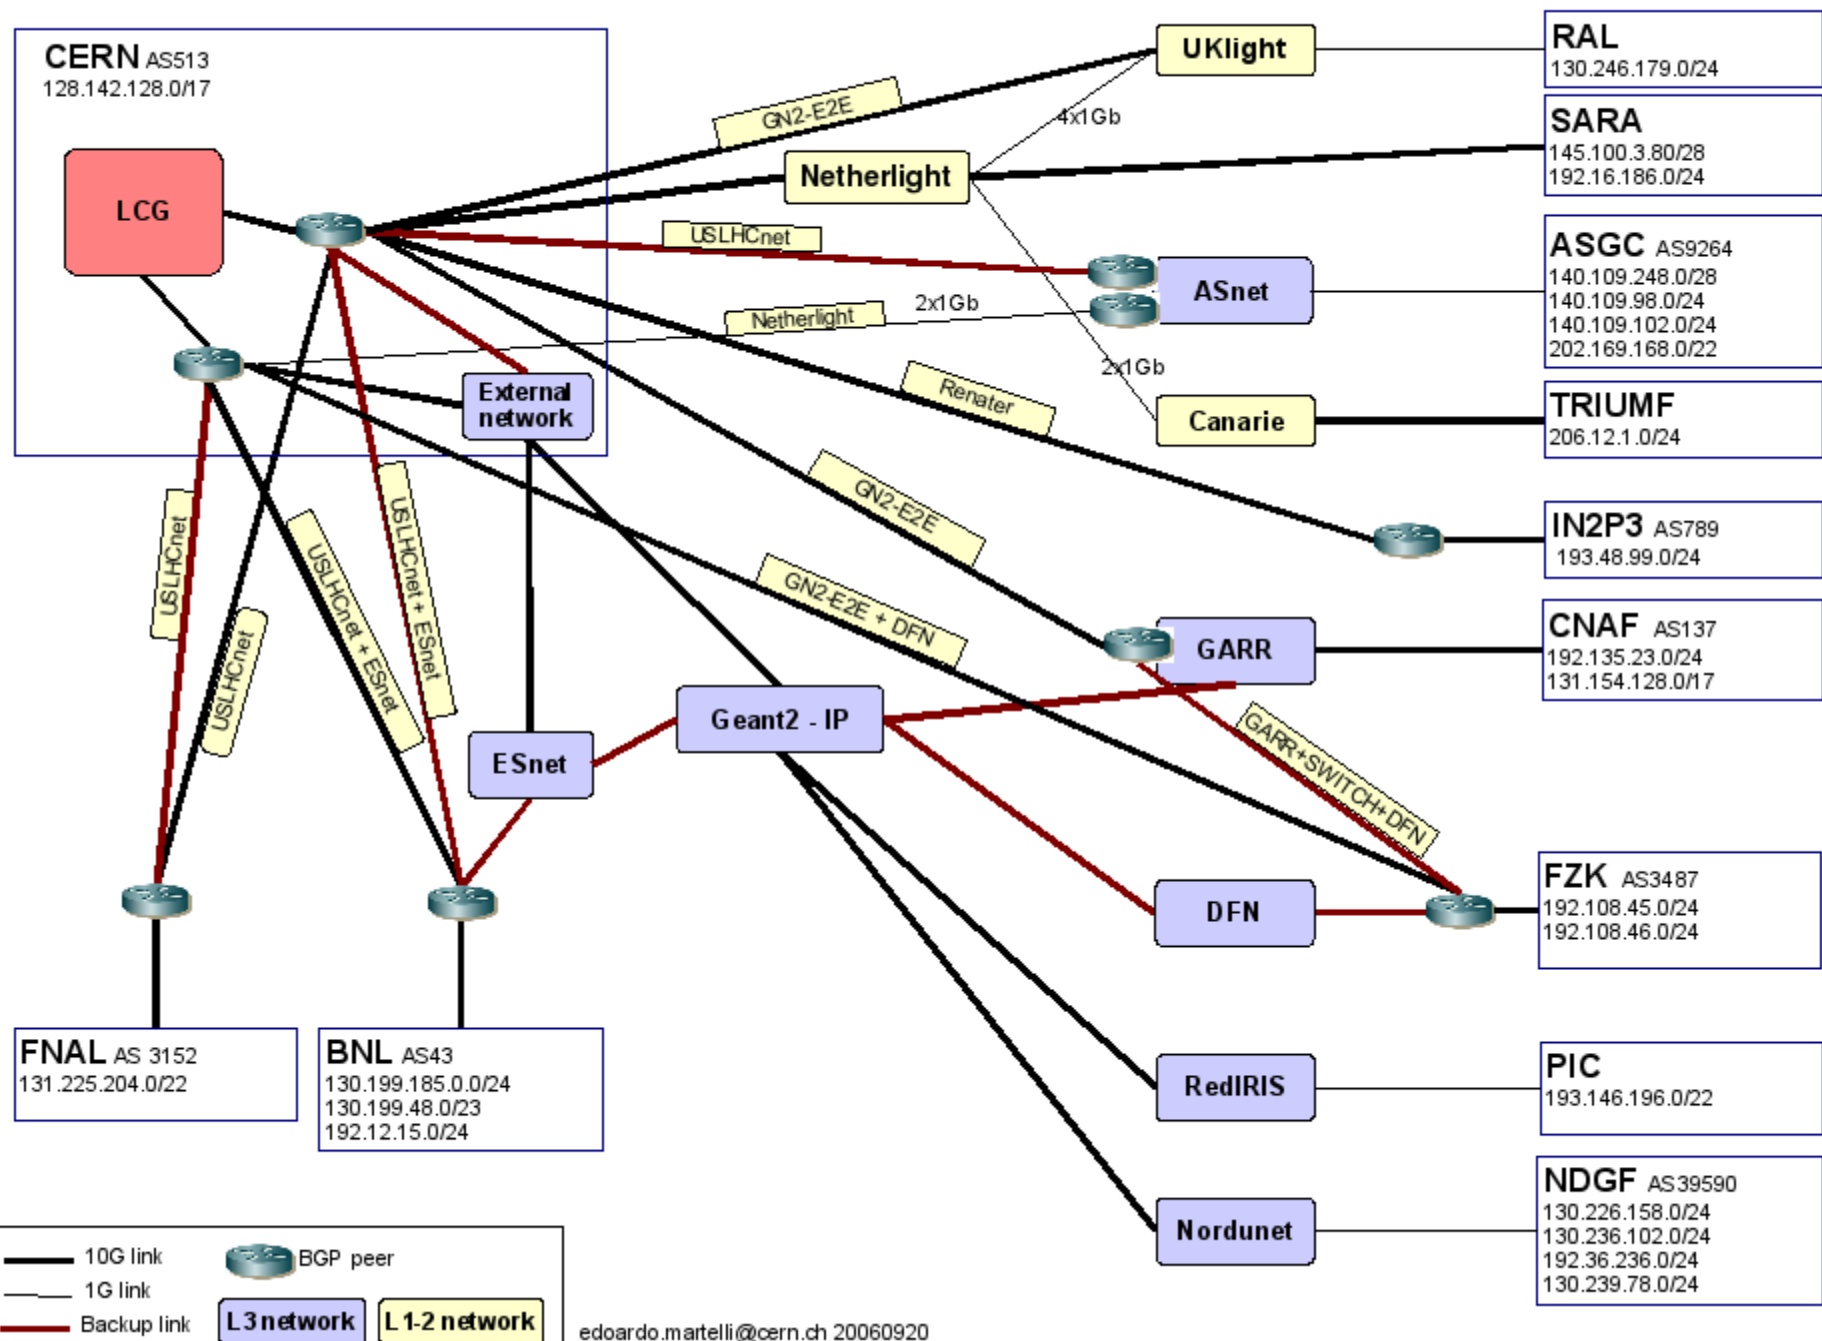
- allows the operation team to drill down from site, to computing elements to worker nodes

- The intention is to provide reports showing longer term trends by site and experiment

Wide Area Network



LHCOPN – current status





Defining the Inter-Site Relationships


- A small group is generating a table defining the inter-site relationships
 - input from the Computing Models
 - quantifying
 - the Tier-1 storage and network services required to support each Tier-2
 - the inter-Tier-1 network services
 - the requirements by storage class at each Tier-1
- The resulting data will enable each site to
 - identify the other sites with which it has to communicate
 - verify the network & storage capabilities
 - set up the required services and verify their performance
- These are minimum requirements for the first year of data taking - network requirements will certainly grow as distributed data analysis develops



Data Distribution


April 06 test period

- CERN → T1s -
- "nominal" rate when LHC is operating - 1.6 GB/s -
- though only for one day
- But sustained operation at 80% of the target



August


- experiment-driven transfers (ATLAS and CMS) sustained 50% 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






CERN Fabric ATLAS T0 test



ATLAS T0 tests with Castor2

- Simulated DAQ – reconstruction – recording
- 220 CPU nodes + 24 disk servers
12 tape drives
- up to 3000 batch jobs
- reached more than the ATLAS nominal speed - 320 MB/s








Full T0 tests – Castor 2

emulated DAQ + Reconstruction +
Data export + Tape recording

- 120 CPU nodes + 48 disk servers
+ 28 tape drives
- 2.1 GByte/s input data rate
+ 2.2 GByte/s output data rate
- reached 4.3 GByte/s disk
throughput with 250 streams
- pp running – expected disk
throughput is 4.5 Gbytes/s






Dependence on e-Science Projects

LCG Operation depends on EGEE and OSG for grid infrastructure, and EGEE, Globus, Condor for grid middleware

- EGEE & OSG are still heavily HEP - but both are evolving as multi-science infrastructures
- EGEE
 - Current project funded to end-March 2008 - but preparations starting to prepare a proposal for a third phase
 - Early discussions on a possible long term science grid infrastructure (cf. the GEANT research network backbone)
- OSG
 - Funded by DoE and NSF for next 5 years at ~\$6M per year - facility development and operations
 - Supports the Virtual Data Toolkit (VDT) - which packages Globus, Condor and other middleware components





Commissioning Schedule

Continued testing of computing models, basic services

Testing **DAQ→Tier-0 (??)** & integrating into DAQ→Tier-0→Tier-1 data flow

Building up end-user analysis support

Exercising the computing systems, ramping up job rates, data management performance,

Experiments

2006

2007

2008

SC4 – becomes initial service when reliability and performance goals met

Introduce residual services
Full FTS services; 3D;
SRM v2.2; VOMS roles

Initial service commissioning – increase reliability, performance, capacity to target levels, experience in monitoring, 24 X 7 operation,

01 jul07 - service commissioned
- full 2007 capacity, performance

first physics

Sites & Services



Summary

- Grids are now operational and heavily used
 - ~200 sites between EGEE and OSG
 - Grid operations centres mature
 - > 60K LCG jobs per day,
long periods with >15K simultaneous jobs
- Core Tier-0 and Tier-1 services progressing well -
 - Data recording and Tier-0 testing at CERN on target
 - Data distribution from CERN to Tier-1s within 80% of 2008 requirements
- Many Tier-2s actively involved in the test programmes during 2006
- Focus is now on commissioning the service, supporting the experiments' system tests and integrating experiment and service operations
 - A steep road ahead to ramp up the capacity over the next year
 - Substantial challenges to achieve the targets for reliability and performance