LCG - The Worldwide LHC Computing Grid

Computing Resource Review Board

Project Status

CERN
24 April 2007

Les Robertson
LCG Project Leader

Applications Support

Preparations for the first LHC run

- New versions of main components released at end of year
- Nightly build/test of all components

Simulation

- Continued progress with physics validation of GEANT4
- New Generator Services system delivered as agreed with MC generator developers

Persistency Framework

- Conditions Database in use by experiments, second version about to be released

ROOT

- Much new functionality
- PROOF (parallel ROOT facility) in use by ALICE at CERN

Persistency Framework
- POOL
- Conditions Database
- Common Relational Access Layer

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

les.robertson@cern.ch

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

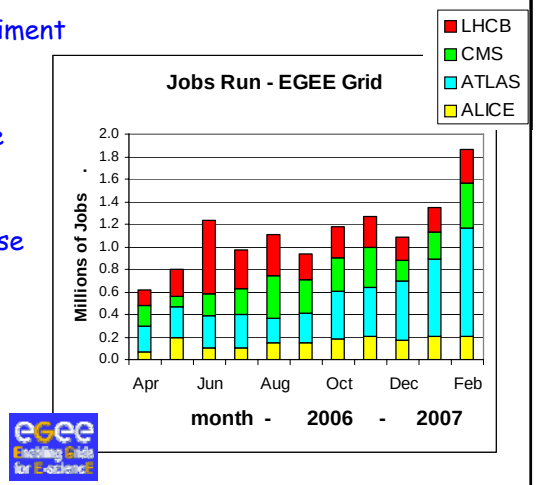

les.robertson@cern.ch

Grid Activity

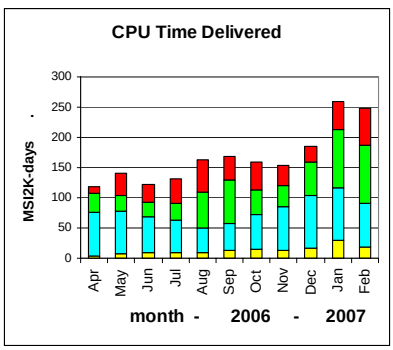

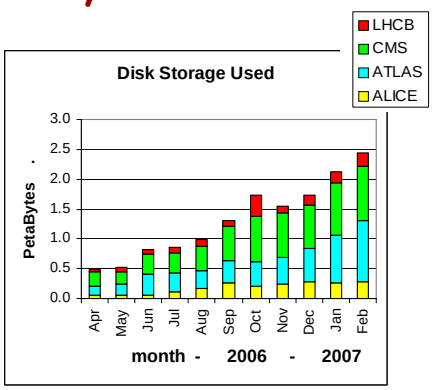

Tier-1s and CERN

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

les.robertson@cern.ch

Growth to 2008

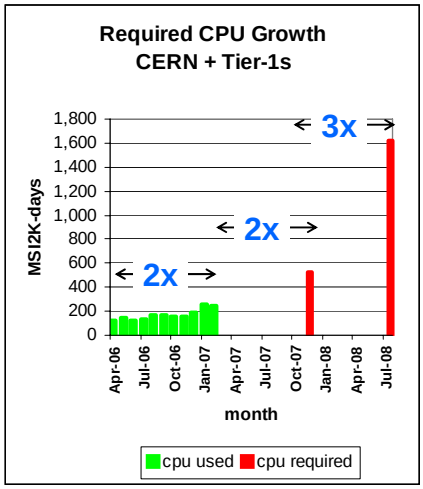

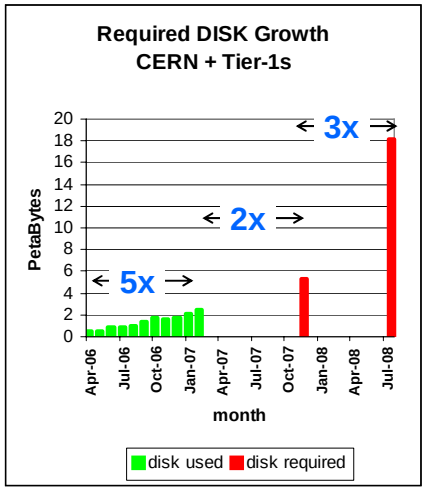

les.robertson@cern.ch

2007 – CERN →Tier-1 Data Distribution

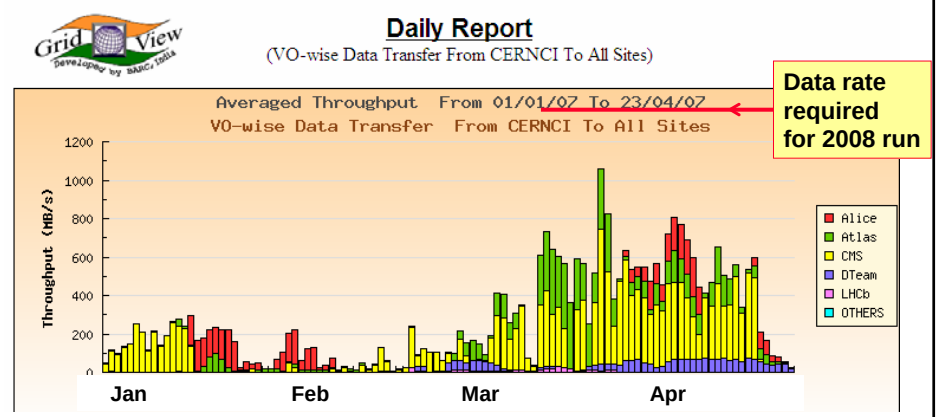

Data rate required for 2008 run

les.robertson@cern.ch

Data Transfers
Comparison with CSA06 — weekly

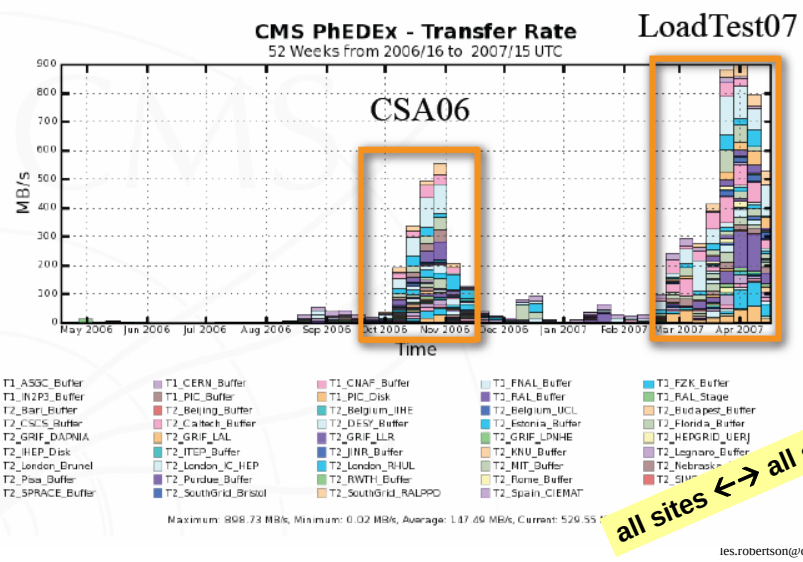

all sites ↔ all sites

Site Reliability – seen from the Grid

CERN+Tier-1s

Average of 8 best sites now meets target for 1H06 – but only just

Average of all sites fluctuating around 90% of the target

Plan is increase the target for average of 8 best sites to
-- June - 93%
-- December – 95%

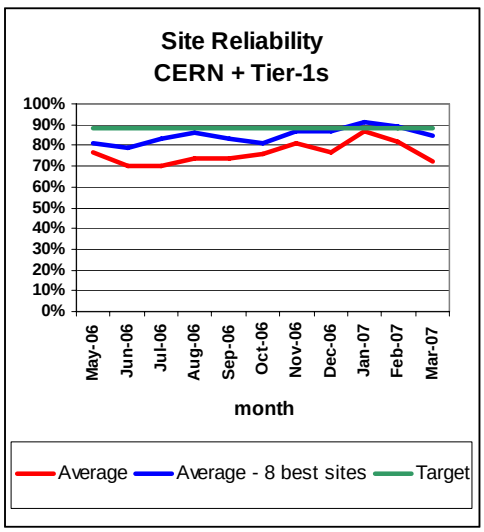

Job Success Rates

Developing a standard monthly report by site for specific job sets

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

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

Will begin to publish job success rates for CERN and Tier-1 sites

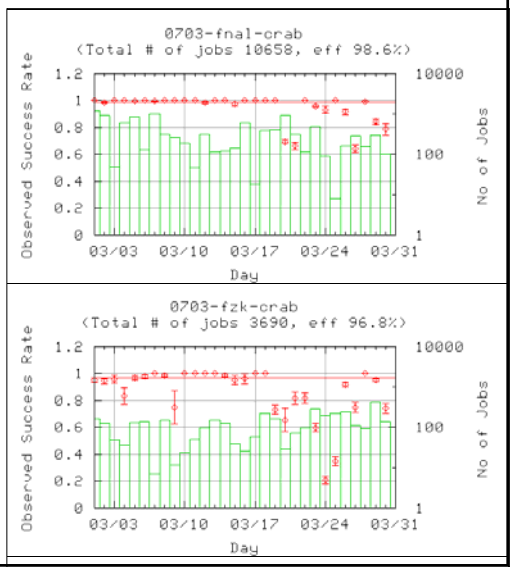

Data & Storage Services

- One of the top concerns
  - Complex integration of experiment data management systems, inter-site data transfer services, site-specific mass storage configurations and services, distributed databases and catalogs
  - Mass storage services not yet thoroughly tested at all sites, especially sustained tape traffic, Tier-1 <-> Tier-2 data movement
  - Mass storage problems hard to debug –
    -- errors passed up through several layers of system and experiment software

→ Testing, debugging, learning how to run these services reliably will take time and (scarce) operational effort
-- for both sites and experiments

les.robertson@cern.ch

CASTOR (CERN's mass storage system)

- Concern over stability and performance
  - Last year Castor achieved all of the Tier-0 and Data Distribution milestones – emphasising throughput
  - BUT – a significant software weakness was uncovered when handling large numbers of requests, mixed workloads (CASTOR is able to process 500K requests per day)
  - A prototype fix showed promising results in September – - but production ready code took longer to develop than expected
  - ATLAS testing in 2007 immediately generated overload on the ATLAS stager – making this issue critical
- Task Force set up to focus on this and other stability and performance issues in Castor
  - Working closely with ATLAS, bringing in additional expertise

  BUT --
  - May delay other Castor developments, including SRM 2.2 deployment

les.robertson@cern.ch

SRM v2.2

- Three implementations of the SRM v1.1 standard are used in WLCG –

  dCache (DESY, FNAL), Castor (CERN), DPM (CERN),
- Originally scheduled for availability end of 2006, full deployment April 2007
- Various delays – agreeing specifications, interpreting the standard, implementation difficulties
- Testing started only Nov-Dec 2006
- Current Status:
  - All implementations have passed basic functionality, use case and stability tests – now in stress testing
  - DPM released
  - dCache being tested at FNAL
    support integrated in new version (1.8) – pilot deployment at DESY, BNL
  - CASTOR being tested at CERN
    Potential scheduling conflicts with critical enhancements to Castor (referred to earlier)
- Experiment beta testing not before May/June
- Full deployment not before September
- Note – interoperation of v1.1and v2.2 has always been planned

les.robertson@cern.ch

Status of other new, upgraded services

- **3D – Synchronised Distributed Oracle Database Service**
  - All but 3 sites now certified or awaiting certification by experiments.
- **FTS – File Transfer Service**
  - Version 2 in test at CERN, deployment to Tier-1s planned for May
  - Tier-1s are responsible for operation of the Tier-1/Tier-2 traffic – experiments organizing testing of all connections during first half of year
- **VOMS-based Job Scheduling Priorities**
  - Being deployed at Tier-1s
- **gLite Workload Management System**
  - In final stages of test – agreed performance and stability requirements achieved – entering certification
- **gLite Compute Element**
  - Problems encountered during tests – discussions on changing strategy
- **Port of EGEE Middleware to Scientific Linux 4**
  - Worker node components ready for distribution
  - Problems with other some of the other components

les.robertson@cern.ch

High Level Milestones since last C-RRB

Level 1 Milestones due October 2006 to March 2007

ID	Date	Milestone
DBS-1	30.09.06	**Full LCG database service in place** Completed by the end of the year for Frontier/Squid (CMS Tier-1s) and for Oracle/Streams at 6 of the 10 Tier-1s concerned (ASCC, BNL, CNAF, GridKA, IN2P3 and RAL), all of which have been tested at the replication rates estimated to be required by ATLAS and LHCb. A new milestone (WLCG-07-09) has been defined for the final step in getting the Oracle/Streams into production including the remaining Oracle Tier-1s (NIKHEF, NDGF, PIC, TRIUMF).
IS-1	30.09.06	**Initial LHC Service in Operation** Began at the end of SC-4. The service will be used for extended testing of the computing systems of the four experiments, for simulation and for processing of cosmic-ray data.
WLCG-07-01	28.02.07	**24 X 7 support defined at CERN and Tier-1s** Definitions of the level of support provided outside of normal working hours, including processes for monitoring and reporting problems, according to the urgency. ASGC | IN2P3 | CERN | FZK | INFN | NDGF PIC | RAL | NIKHEF | TRIUMF | BNL | FNAL
WLCG-07-08	31.03.07	**Accounting Data published into the APEL repository** CERN and Tier-1 sites automatically publish their CPU accounting data to the central repository at the Grid Operations Centre (GOC) provided by Rutherford Lab. This data will be used as the sole source for the accounting reports from May 2007. ASGC | IN2P3 | CERN | FZK | INFN | NDGF PIC | RAL | NIKHEF | TRIUMF | BNL | FNAL
WLCG-07-09	31-03-07	**3D Oracle/Streams Service in Production** Oracle services in production and certified by the experiments. ASGC | IN2P3 | CERN | FZK | INFN | NDGF PIC | RAL | ~~NIKHEF~~ | ~~TRIUMF~~ | BNL NIKHEF has been certified by ATLAS, and is awaiting certification by LHCb. TRIUMF is awaiting certification by ATLAS.

achieved

not achieved < 1 month late

not achieved > 1 month late

les.robertson@cern.ch

Summary

- Applications support in good shape
- WLCG service in operation, with continuously increasing workload
  -- and a steep ramp-up ahead to the capacity needed for the 2008 run
- Experiment testing progressing – involving more sites, and becoming increasingly realistic
  → uncovering new problems to be understood and addressed
- Data and storage, along with general site reliability
  → are major concerns
- Sites & experiments working well together to tackle the problems

les.robertson@cern.ch