




CERN and LCG report



10 September 2004

Jürgen Knobloch
IT Department, CERN

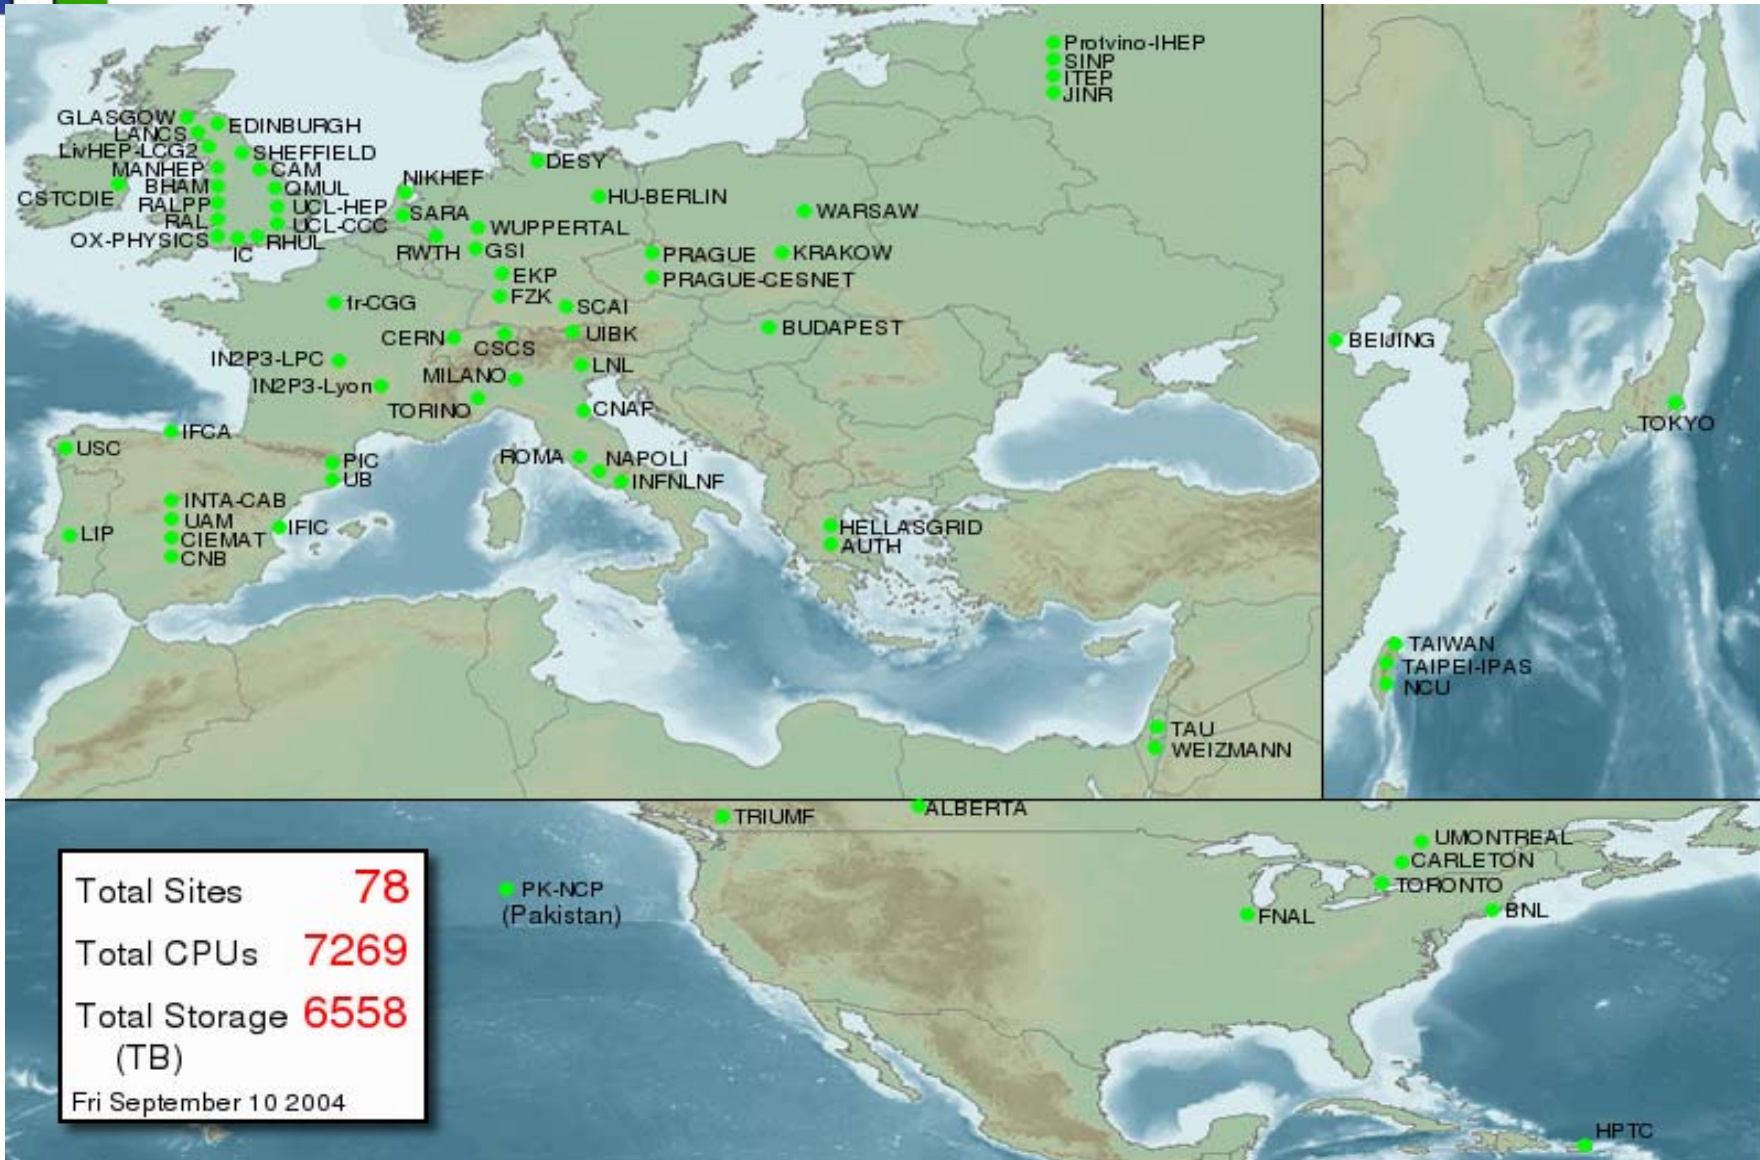
This file is available at:

http://cern.ch/lcg/presentations/LCG_JK_HTASC_2004_09.ppt





LCG-2 sites - status 10 September 2004





LCG

- Experiment data challenges phase-1 terminated - a lot has been learned - improvements have been and are being implemented
- Data challenges are being complemented by service challenges - trying high throughput data transfer - sustained operation between Tier-0 and selected Tier-1s
- ARDA & Scientific Linux - see other talks
- Phase-2 planning is progressing
- Issues on missing resources are being addressed
- Grid deployment and gLite development proceed in parallel
- The following slides are summaries from the last GDB meeting: of ALICE, ATLAS, CMS, LHCb and Ian Bird



ALICE - Considerations

- ✦ LCG is providing a (more and more) coherent infrastructure with a large quantity of resources
- ✦ Relations with LCG are good
 - ✦ With the inevitable strains of time-constrained exercise
 - ✦ Special thanks to P.Mendez for playing the difficult and ungrateful role of interface between the unruly ALICE mob and LCG support and site admins
 - ✦ Thanks to LCG for providing additional storage resources for a local SE at CERN exclusively for ALICE (3.5 TB).
- ✦ The middleware is improving within the limits of its current design and architecture
 - ✦ GAG is preparing a detailed feedback based on the experience of the Data Challenges
 - ✦ ALICE observations are similar to those of the other experiments



ALICE - Conclusions


- ★ The ALICE DC04 started out with (almost unrealistically) ambitious objectives
- ★ We are coming very close to reach this objectives and LCG has played an important role in this
- ★ We are ready and willing to move to gLite as soon as possible and contribute to its evolution with our feedback




CPU usage & Jobs




ATLAS DC2 - CPU usage




ATLAS DC2 - Grid3 - September 7



ATLAS DC2 - LCG - September 7



NorduGrid



ATLAS - Summary (1)



- ❑ Major efforts on the past few months
 - Redesign of the ATLAS Event Data Model and Detector Description
 - Integration of the LCG components (G4; POOL; ...)
 - Introduction of the **Production System**
 - Interfaced with 3 Grid flavors (and "legacy" systems)
- ❑ Delays in all activities have affected the schedule of DC2
 - Note that Combined Test Beam is ATLAS 1st priority
 - And DC2 schedule was revisited
 - To wait for the readiness of the software and of the Production system

ATLAS - Summary (2)



□ DC2

- About 80% of the Geant4 simulation foreseen for Phase I has been completed using only Grid and using the 3 flavors coherently; Pile-up just starting
- The 3 Grids have been proven to be usable for a real production and this is a major achievement

□ BUT

- Phase I progressing slower than expected and it's clear that all the involved elements (Grid middleware; Production System; deployment and monitoring tools over the sites) need improvements
- It's one of the goals of the Data Challenges to identify these problems as early as possible.

□ Phase II (Tier0 exercise) is scheduled for October



Continuous Operation


- ❖ We have been in full time production for a year now
 - ◆ Some of this on GRID resources and some on dedicated resources
- ❖ We are aggressively working on:
 - ◆ Enabling complex tasks like Digitization to run on GRID resources
 - ◆ Smoothing the production operations (continuous delivery rather than large blocks at a time;
 - ◆ End-to-end delivery;
 - ◆ Data movement;
 - ◆ Publication;
 - ◆ Enabling user jobs
- ❖ New CCS Tasks of Data Management, Workload management and Production Operations and cross-project APROM to coordinate and to do the work.
- ❖ Continuous user access to data at GRID sites is a big issue.
 - ◆ Analysis issues being implemented and tested and made to work



CMS Plan / Status

DST's : software and production

- ❖ **DST (v1)** – CCS aspects worked in DC04 but usefulness to P



- ❖ **DST (v2)** - OK for 10 million $W \rightarrow e\nu$ calibration sample
- ❖ **DST (v3)** - PRS have made physics objects “generally good enough”



Physics TDR

- ❖ Physics TDR scheduled for December 2005
 - ◆ Not a “yellow report”, but a detailed study of methods for initial data taking and detector issues such as calibration as well as physics reach studies.
- ❖ Current Simulated samples more or less adequate for low luminosity
 - ◆ About to enter re-reconstruction phase with new DST version
 - ◆ Estimate similar sample sizes for high luminosity
- ❖ Target 10M events/month throughput
 - ◆ Generation, Simulation, Digitization, Reconstruction, Available for analysis
 - ◆ New production operations group in CMS to handle this major workload
- ❖ In light of DC04 experience, DC05 is cancelled as a formal computing exercise
 - ◆ Not possible to serve current physics requirements with data challenge environment
 - ◆ However, specialized component challenges are foreseen



Continuous Operation Scale and Scope

- ❖ Instead of DC05 we need to progressively bring up a full time operation, including data access
 - ◆ Not a test or a challenge.
 - ◆ Physicist access to data required.
- ❖ Generic GRID resources (for Generation/Simulation)
 - ◆ ~750 CPU Continuous
 - CPU means current generation ~2.4+GHz
- ❖ CMS T1 resources (Grid with significant Disk and MSS)
 - ◆ Needed for data intensive Digitization and Reconstruction steps
 - ◆ ~750 CPU Continuous
 - ◆ Now 60TB +20TB/month (Spread across T1 centers)
- ❖ T1/T2 resources (probably not generic)
 - ◆ ~150 CPU Continuous
 - ◆ Now 40TB Analysis disk space to grow by about 2-5TB/month (T1+T2)
- ❖ We intend to run all this within LCG
 - ◆ "LCG" being all those resources available to CMS being steady migration from/between LCG2, GRID3, gLite, ...
- ❖ We need to reach the 10M event/month level soon (Autumn)
- ❖ We need to make the resources available to a growing CMS user base in the same time scale

- LHCb DC'04 Phase 1 is over.
- The Production Target has been achieved:
 - 186 M Events in 424 CPU years.
 - ~ 50% on LCG Resources (75-80% at the last weeks).
- **Right LHCb Strategy:**
 - Submitting “empty” DIRAC Agents to LCG has proven to be very flexible allowing a success rate above LCG alone.
- Big room for improvements, both on DIRAC and LCG:
 - DIRAC needs to improve in the reliability of the Servers:
 - big step already during DC.
 - LCG needs improvement on the single job efficiency:
 - ~40% aborted jobs, ~10% did the work but failed from LCG

- Phase 2
 - Stripping in last steps of preparation
 - Need to Run over **65 TB of Data distributed** over 4 Tier1 Sites (CERN, CNAF, FZK, PIC), with “small” CPU requirements.
- Phase 3
 - **End user analysis** will follow
 - GANGA tools in preparation
- (Phase 1)
 - Keep a continuous rate of production activity with programmed mini DC (i.e., few days once a month).



Summary of issues (I. Bird at last GDB)

- LCG-2 services have been supporting the data challenges
 - Many middleware problems have been found - many addressed
 - Middleware itself is reasonably stable
- Biggest outstanding issues related to providing and maintaining stable operations
- Has to be addressed in large part by management buy-in to providing sufficient and appropriate effort at each site
- Future middleware has to take this into account:
 - Must be more manageable, trivial to configure, etc
 - Management and monitoring must be built into services from the start



Other points on CERN computing

- Security
 - High rate of incidents since July
 - Lionel Cons nominated Deputy Computing Security Officer
 - ACB stopped July 1st
 - Windows: reset password of local admin
 - W-XP: service pack 2 - pilot
 - Recently, LCG and MACs were also hit
- IP telephony pilot project running

