Introduction

Grid @ CERN

- **LCG: LHC Computing Grid – the deployment project**
  - **Will run the 24/7 Grid service**
- **EGEE: Enabling Grids for E-Science in Europe**
  - **Started in April 2004 with 70 partners and 32M€ EU funding**
  - **Will provide the next generation middleware for LCG**
- **CERN openlab for DataGrid applications**
  - **Started in 2003 - Funded by Industry and CERN**
  - **Main project: opencluster (including 100 Itanium nodes)**
  - **R&D aimed at deployment in LCG**

Computing for LHC

Europe:
~270 institutes
~4500 users

Elsewhere:
~200 institutes
~1600 users

- **Problem:** even with an upgraded computer centre, CERN can only provide a fraction of the necessary resources
- **Solution:** computing centres, which were isolated in the past, will now be connected, **uniting the computing resources of particle physicists in the world using GRID technologies!**

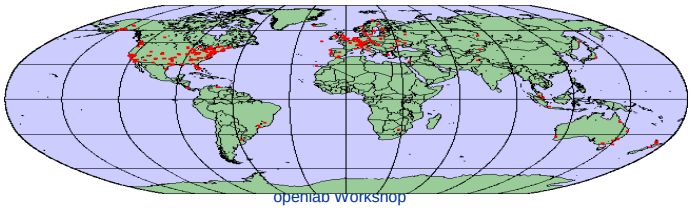

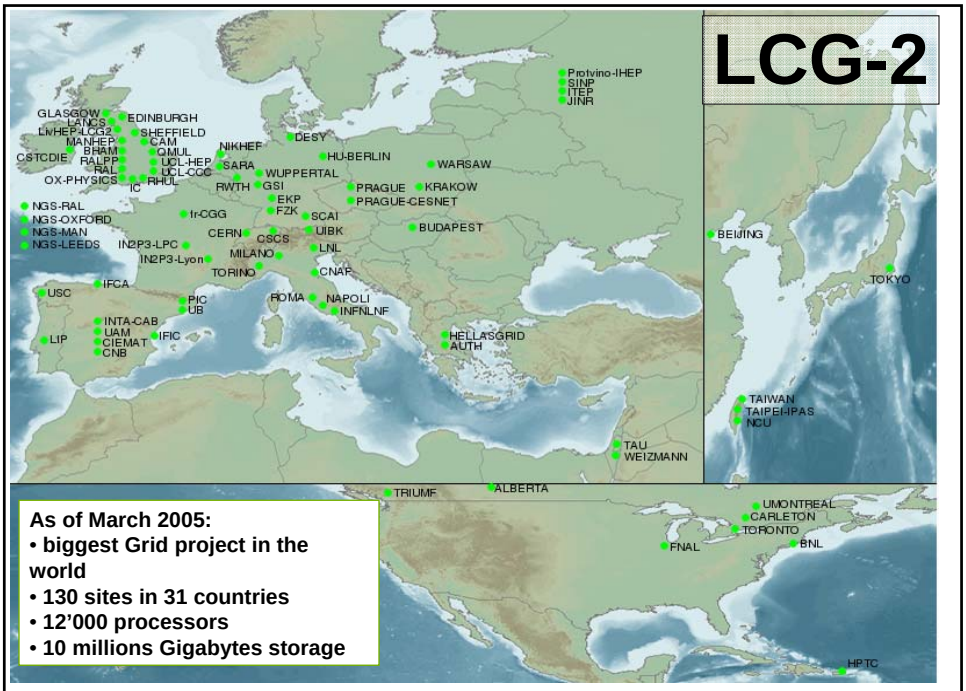

Openlab: Tight integration with the LCG testbed

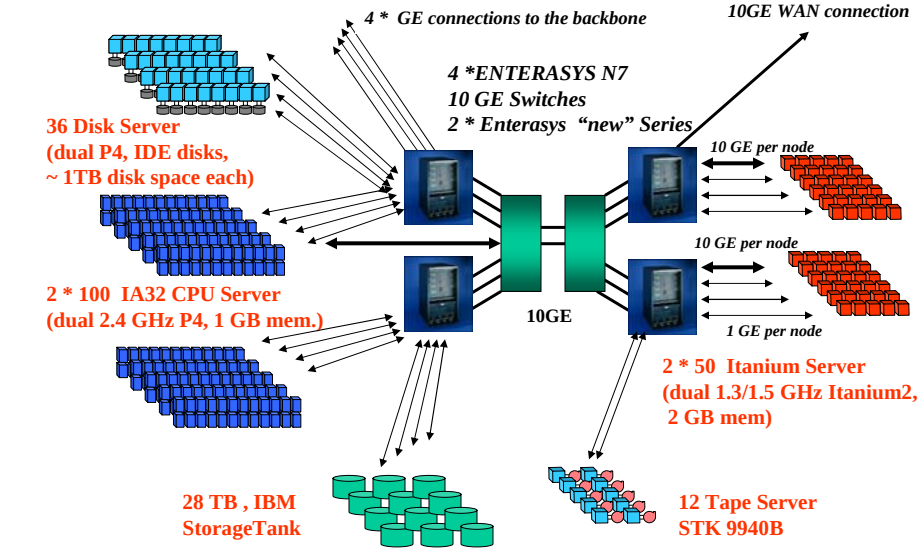

Service Challenge 3

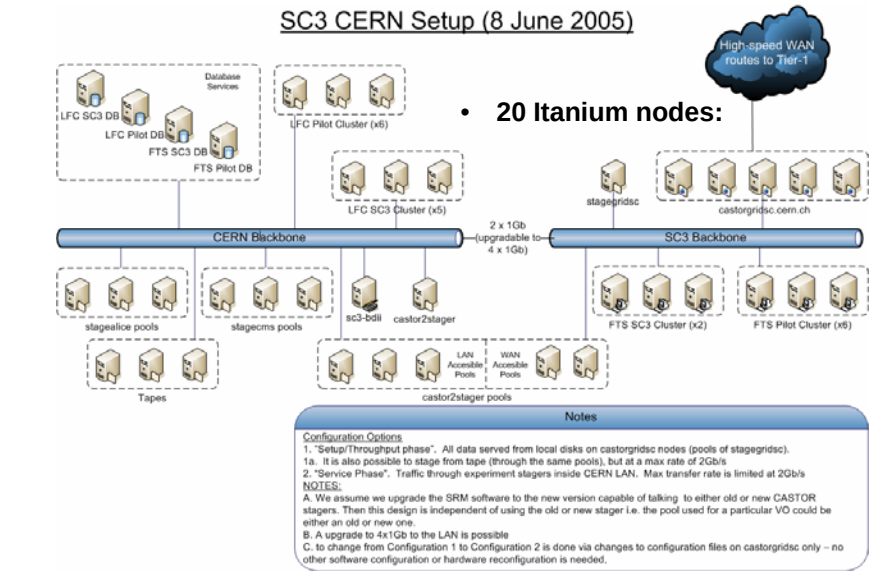

- **20 Itanium nodes:**

64-bit porting project

Itanium / Itanium Processor Family (IPF) / IA-64

The 64-bit issue

- **What exactly is meant?**
- **Simple:**
  - Linux on 32-bit hardware uses “ILP32”
    - Int = Long = Pointer (32 bit)

  - Linux on 64-bit hardware uses I32LP64
    - Int stays 32-bit
    - Pointer = Long (64 bit)

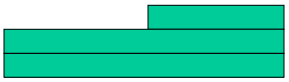

- **As a result:**
  - For instance:
    - Any attempt to cast a pointer to Int (and back) ➔ Fatal error !!!

LCG components

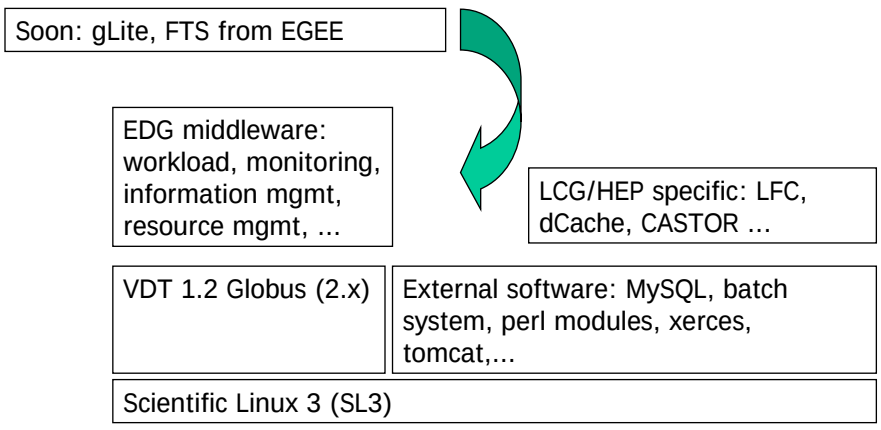

LCG architecture

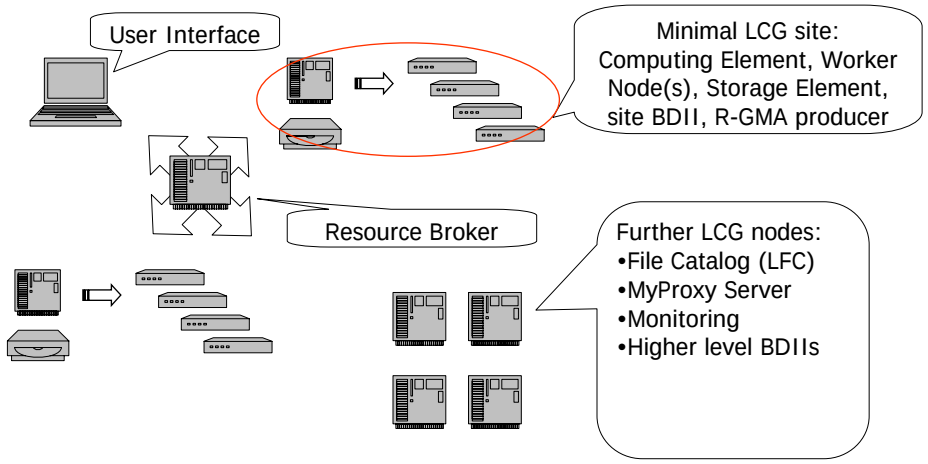

Timeline

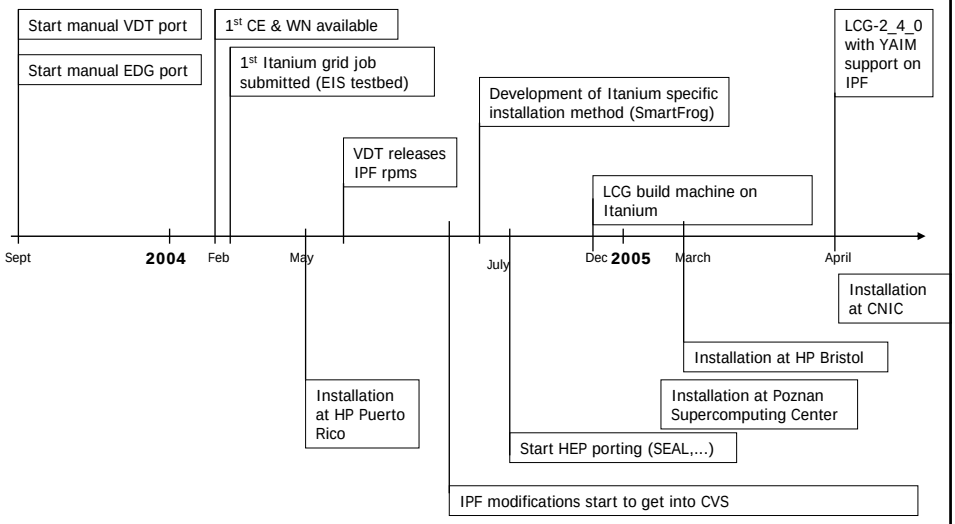

Original LCG build model

Check out from source CVS:

- **EDG software**
- **LCG specific code**

Everything else is "external".

A "build machine" automatically does the build after the checkout (uses GNU autotools)

Initial status (2003)

- **LCG build machine supported only IA-32 with specific version of Red Hat**
- **No binaries available for Itanium/IA-64**
- **Hardly any documentation**
- **Installation of LCG only via LCFGng (fully automatic, IA-32 only)**
  - manual installation was considered to be "extremely difficult" (EDG manual)

Initial strategy

- **Started to port everything on our own**
- **One doctoral student & one fellow**
  - Stephen Eccles (now: Lancaster University)
  - Andreas Unterkircher (CERN)
- **After 6 months we were able to install a minimal (CE,WN,SE) Itanium LCG site and successfully submit jobs.**

Initial obstacles (1)

- **VDT has its own (not documented) build procedure. We had to do “reverse engineering“.**
- **It was often difficult to find the original sources of rpms. EDG used sometimes “special“ versions of well known libraries (e.g. Boost).**

Initial obstacles (2)

- **EDG build procedure was hard-coded for IA-32 on RH 7.3.**
- **As our changes did not get back into the CVS it was difficult for us to keep track with the latest releases**
- **The code had, indeed, several 64-bit issues but the complicated build procedures (EDG as well as VDT) caused us much more trouble.**

Lessons learnt (1)

- **Initial effort was necessary to get noticed by the community.**
  - E.g. when VDT saw that we are serious they started to provide Itanium rpms on their own.
- **Vital: Always get changes back into the CVS on a regular basis.**

Lessons learnt (2)

- **Support for different compilers, OS's and architectures should be considered in the build procedure from the beginning and used for testing on a regular basis.**
- **From a first proof of concept to a fully supported official release can take a long time:**
  - In our case: ~1 year

Lessons learnt (3)

- **Porting LCG to Itanium was a "chicken and egg problem":**
  - LCG was not considering porting as there is no HEP software for Itanium
  - Physicists did not port to Itanium as there were no such resources in LCG.
- **Thus we also started porting of major HEP software (SEAL, POOL, etc.)**
- **Note that ALICE has all its software 64-bit clean!**
  - Mainly an issue of "initial mindset"

Porting to EM64T/AMD64 (1)

- **Should be much easier as IA64 (64bit) code changes are also valid for these platforms.**
  - Exactly the same “I32LP64“ model
- **First one has to ensure that the basic packages (VDT, external software) are available.**
- **Getting modifications back to CVS immediately will be important.**

Porting to EM64T/AMD64 (2)

- **Build procedures not recognizing the architecture could be again the source of much trouble – this must be addressed immediately.**
- **Hopefully EGEE/gLite will prove to be better in this respect than EDG**
  - We (and others) are providing platforms for testing
- **Finally worth mentioning:**
  - Some ports of EDG to other platforms (e.g. PowerPC) are available on the Grid-Ireland homepage.

Overview of 64-bit porting

- **Phase 1 Completed:**
  - ROOT (Data analysis framework)
    - http://root.cern.ch/
  - Geant4 (Physics simulation framework)
    - http://cern.ch/geant4
  - CLHEP (C++ Class Library)
    - http://proj-clhep.web.cern.ch/proj-clhep/
  - CASTOR (CERN Hierarchical Storage Manager)
    - http://cern.ch/castor
  - LCG-2 Grid middleware
    - Originated from EDG (European Data Grid)
      - http://lcg.web.cern.ch/LCG/Sites/releases.html
    - Itanium version:
      - http://openlab-mu-internal.web.cern.ch/openlab-mu-internal/Projects/LCGonIA64/LCGonIA64.asp

64-bit porting (cont'd)

- **Next aim:**
  - Allow the simulation stack of one of the LHC experiments (LHCb) to work on Itanium
    - Set of external packages (Boost, etc.): OK
    - Base set of CERN packages (Geant4, ROOT, CLHEP): OK
    - HEP/LCG packages (SEAL, POOL, PI): In progress
    - Specific packages from the experiment (Gaudi, Gauss, Ganga): In progress
  - Once this experiment's stack is complete, ATLAS and CMS frameworks should also be within range
  - By the way, Intel, Munich is apparently also working on the ATLAS software

Virtualization project

Virtualization

- **Our history**
  - Xen benchmarked with CERN simulation workload on IA-32
    - Work done by summer student 2004
  - Project work on IO workloads under Xen
    - Two students
  - Project work on Xen on Itanium
    - One of the two students (Master thesis in this semester)
      - Collaboration with HP Labs
  - Additionally:
    - One openlab fellow is continuing the work on IA32 w/Linux Fedora version
      - Aim at IO intensive workloads (ROOT analysis, etc.)
- **Rationale: next generation processors (such as IPF Montecito) will have hardware support for virtualization**
- **Question: Will virtualization be one of the underpinnings of future Grid security?**

Conclusions

- **The 64-bit port to Itanium has laid the foundation for:**
- **The inclusion of Itanium systems in LCG-2**
- **A new architectural dimension in the Grid**
  - Heterogeneity
- **A foundation for porting other 64-bit/Linux systems**
- **A multi-platform strategy for Grid middleware development**

BACKUP