



... where the Web was born

CERN and the LHC data challenge

Wolfgang von Rueden, IT Department Head

IBM Study Tour, March 2004, Geneva



What is CERN?

- CERN is the **world's largest particle physics centre** funded by 20 European member states
- Particle physics is about:
 - **elementary particles** of which all matter in the universe is made
 - **fundamental forces** which hold matter together
- Particle physics requires:
 - **special tools** to create and study new particles

CERN is:


- **2500 staff scientists** (physicists, engineers, ...)
- **Some 6500 visiting scientists** (half of the world's particle physicists)

They come from 500 universities representing 80 nationalities.





... is located in Geneva, Switzerland





CERN's special tools are:

- **ACCELERATORS**

Huge machines able to speed up particles to very high energies before colliding them with other particles

- **DETECTORS**

Massive instruments which register the particles produced when the accelerated particles collide

- **COMPUTERS**

Large clusters in a world-wide grid infrastructure





What is LHC?

- LHC will collide beams of protons at an energy of **14 TeV**
- Using the latest super-conducting technologies, it will operate at about **-270°C**, just above absolute zero of temperature.
- With its **27 km circumference**, the accelerator will be the largest superconducting installation in the world.

LHC is due to switch on in 2007


Four experiments, with detectors as 'big as cathedrals':

ALICE
ATLAS
CMS
LHCb



ATLAS Detector


R&D	1988-1995
Proposal	1994
Approval	1996
Installation	2003-2006
Operation	2007-2020



150 million measurement channels
1 million control points
1 million lines of code

> 500 MCHF, CERN: 15%

LHC Data Challenge





LHC data (simplified)

Per experiment:

- 40 million collisions per second
- After filtering, 100 collisions of interest per second
- A Megabyte of digitised information for each collision = recording rate of 0.1 Gigabytes/sec
- 1 billion collisions recorded = 1 Petabyte/year

Total: ~10.000.000.000.000.000

= 1% of

1 Megabyte (1MB)
A digital photo

1 Gigabyte (1GB)
= 1000MB
A DVD movie

1 Terabyte (1TB)
= 1000GB
World annual book production

1 Petabyte (1PB)
= 1000TB
10% of the annual production by LHC experiments

1 Exabyte (1EB)
= 1000 PB
World annual information production


CMS



LHCb



ATLAS



ALICE






LHC data

LHC data correspond to about
20 million CDs each year


Where will the
experiments store all of
these data?





The LHC Data Challenge


- A **particle collision** = an event
- Provides trivial **parallelism**, hence usage of simple farms
- Physicist's goal is to count, trace and characterize all the particles produced and **fully reconstruct the process**.
- Among all tracks, the presence of “**special shapes**” is the sign for the occurrence of interesting interactions.





The LHC Data Challenge

Starting from this event...




Selectivity: 1 in 10^{13}

Like looking for 1 person in a thousand world populations!

Or for a needle in 20 million haystacks!

You are looking for this “signature”






LHC data processing

LHC data analysis requires a computing power equivalent to ~ 70,000 of today's fastest PC processors


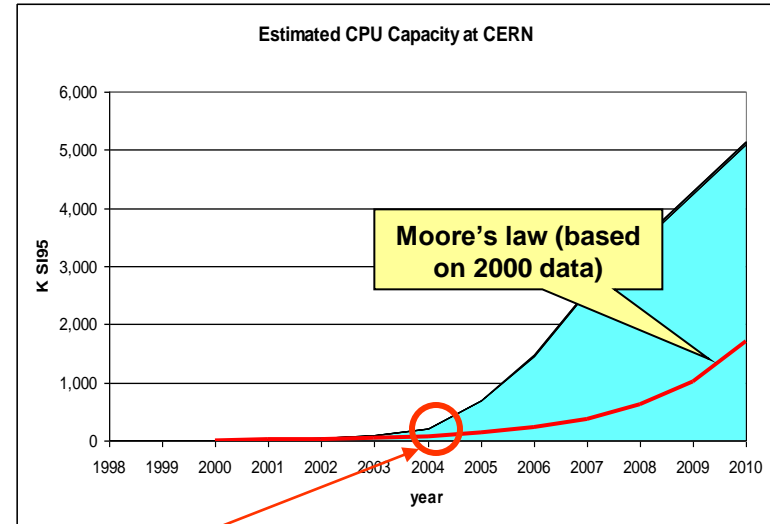
Where will the experiments find such a computing power?






Expected LHC computing needs

Networking:
10 – 40 Gb/s to all big centres



today





Computing at CERN today


- High-throughput computing based on reliable “commodity” technology
- More than 1500 dual processor PCs
- More than 3 Petabyte of data on disk (10%) and tapes (90%)

Nowhere near enough!





Computing at CERN today



CPU servers

The new computer room is being populated...



Disk servers



Tape silos and servers



Computing at CERN today



...while the existing computer centre is being cleared for renovation...



...and an upgrade of the power supply from 0.5MW to 2.5MW is underway.

...and the services continue to run...




Computing for LHC

Europe:
~270 institutes
~4500 users

Elsewhere:
~200 institutes
~1600 users

- **Problem:** even with computer centre upgrade, 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!**





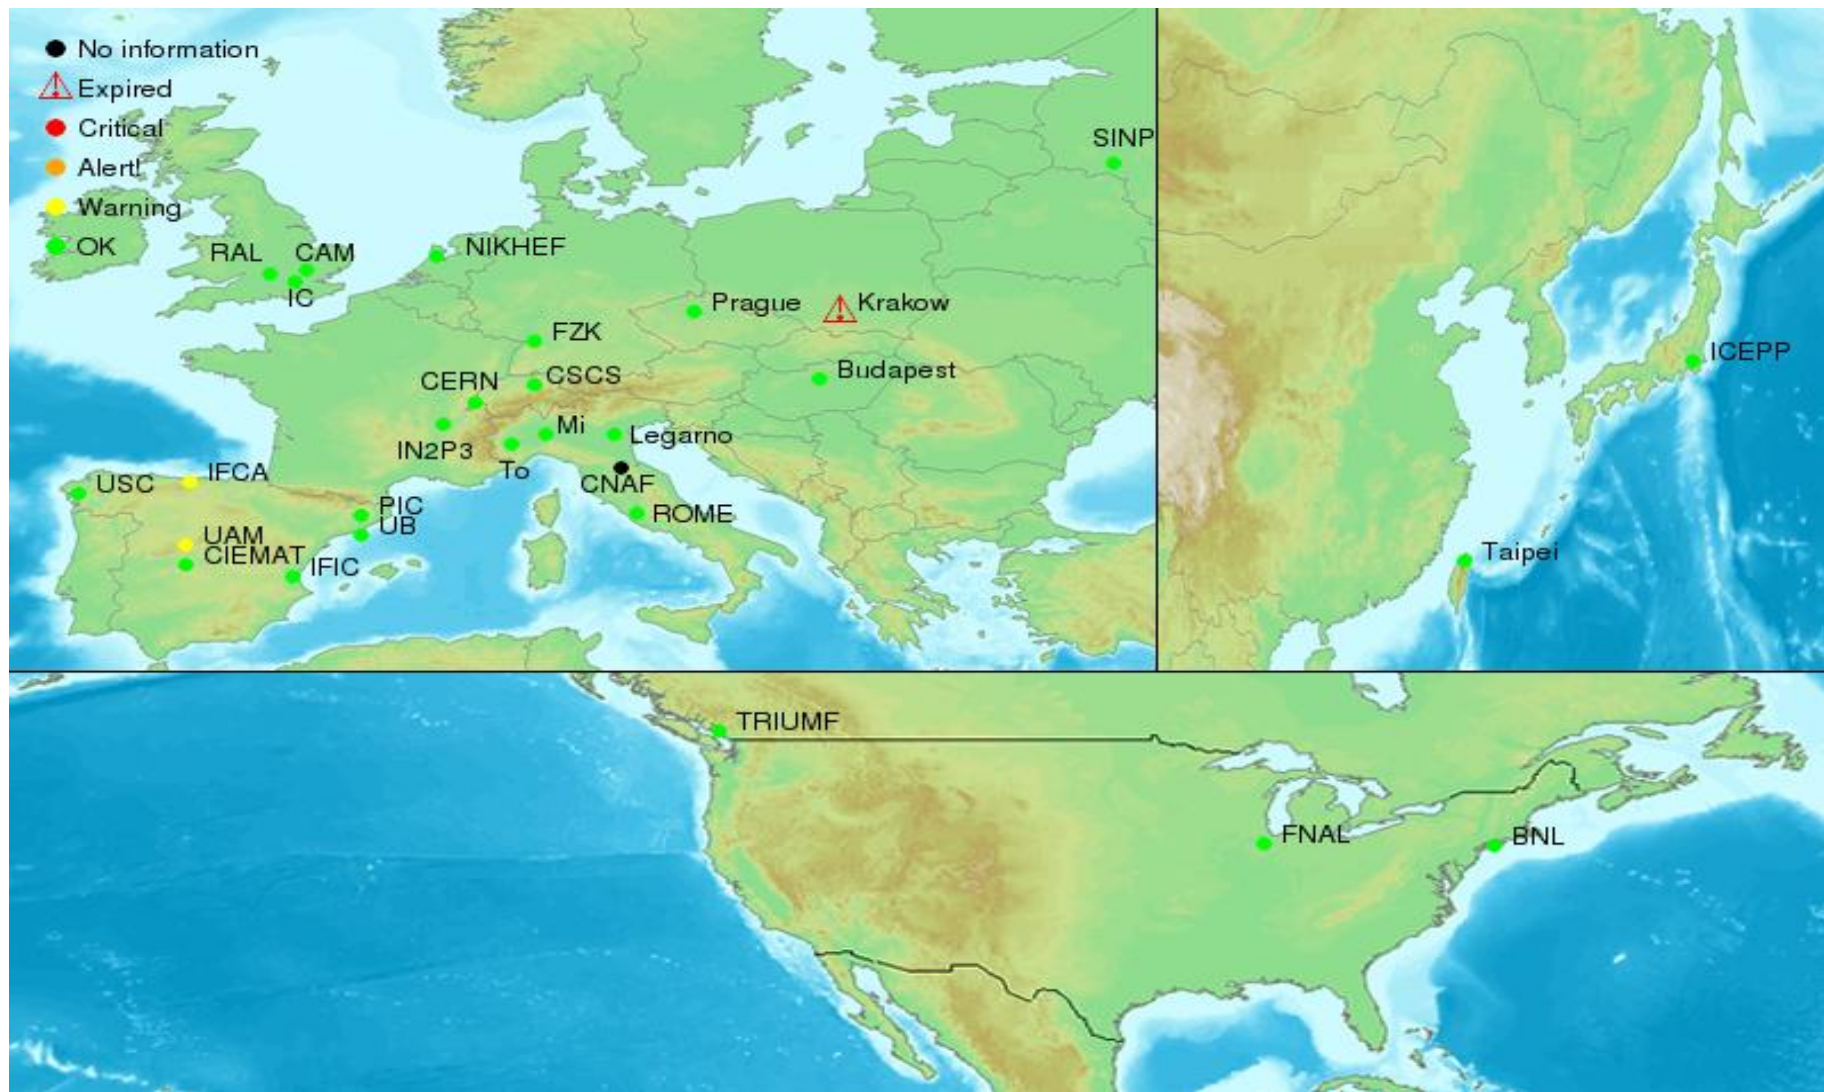
Grid @ CERN

- CERN projects:
LHC Computing Grid (LCG)
- EC funded projects led by CERN:
European DataGrid (EDG), terminates 3/04
European DataTAG (EDT), terminates 3/04
Enabling Grids for E-Science in Europe (EGEE), starts 4/04
- Industry funded projects:
CERN openlab for DataGrid applications





LHC Computing Grid (LCG)



CERN

openlab for DataGrid applications



sponsored by

IBM[®]

intel[®]





i n v e n t

ENTERASYS

NETWORKS[™]


ORACLE[®]






Mission of CERN openlab

- **Testbed** for cutting edge hardware and Grid software
- **Industry consortium** for Grid-related technologies of common interest
- **Training ground** for a new generation of engineers to learn about Grid






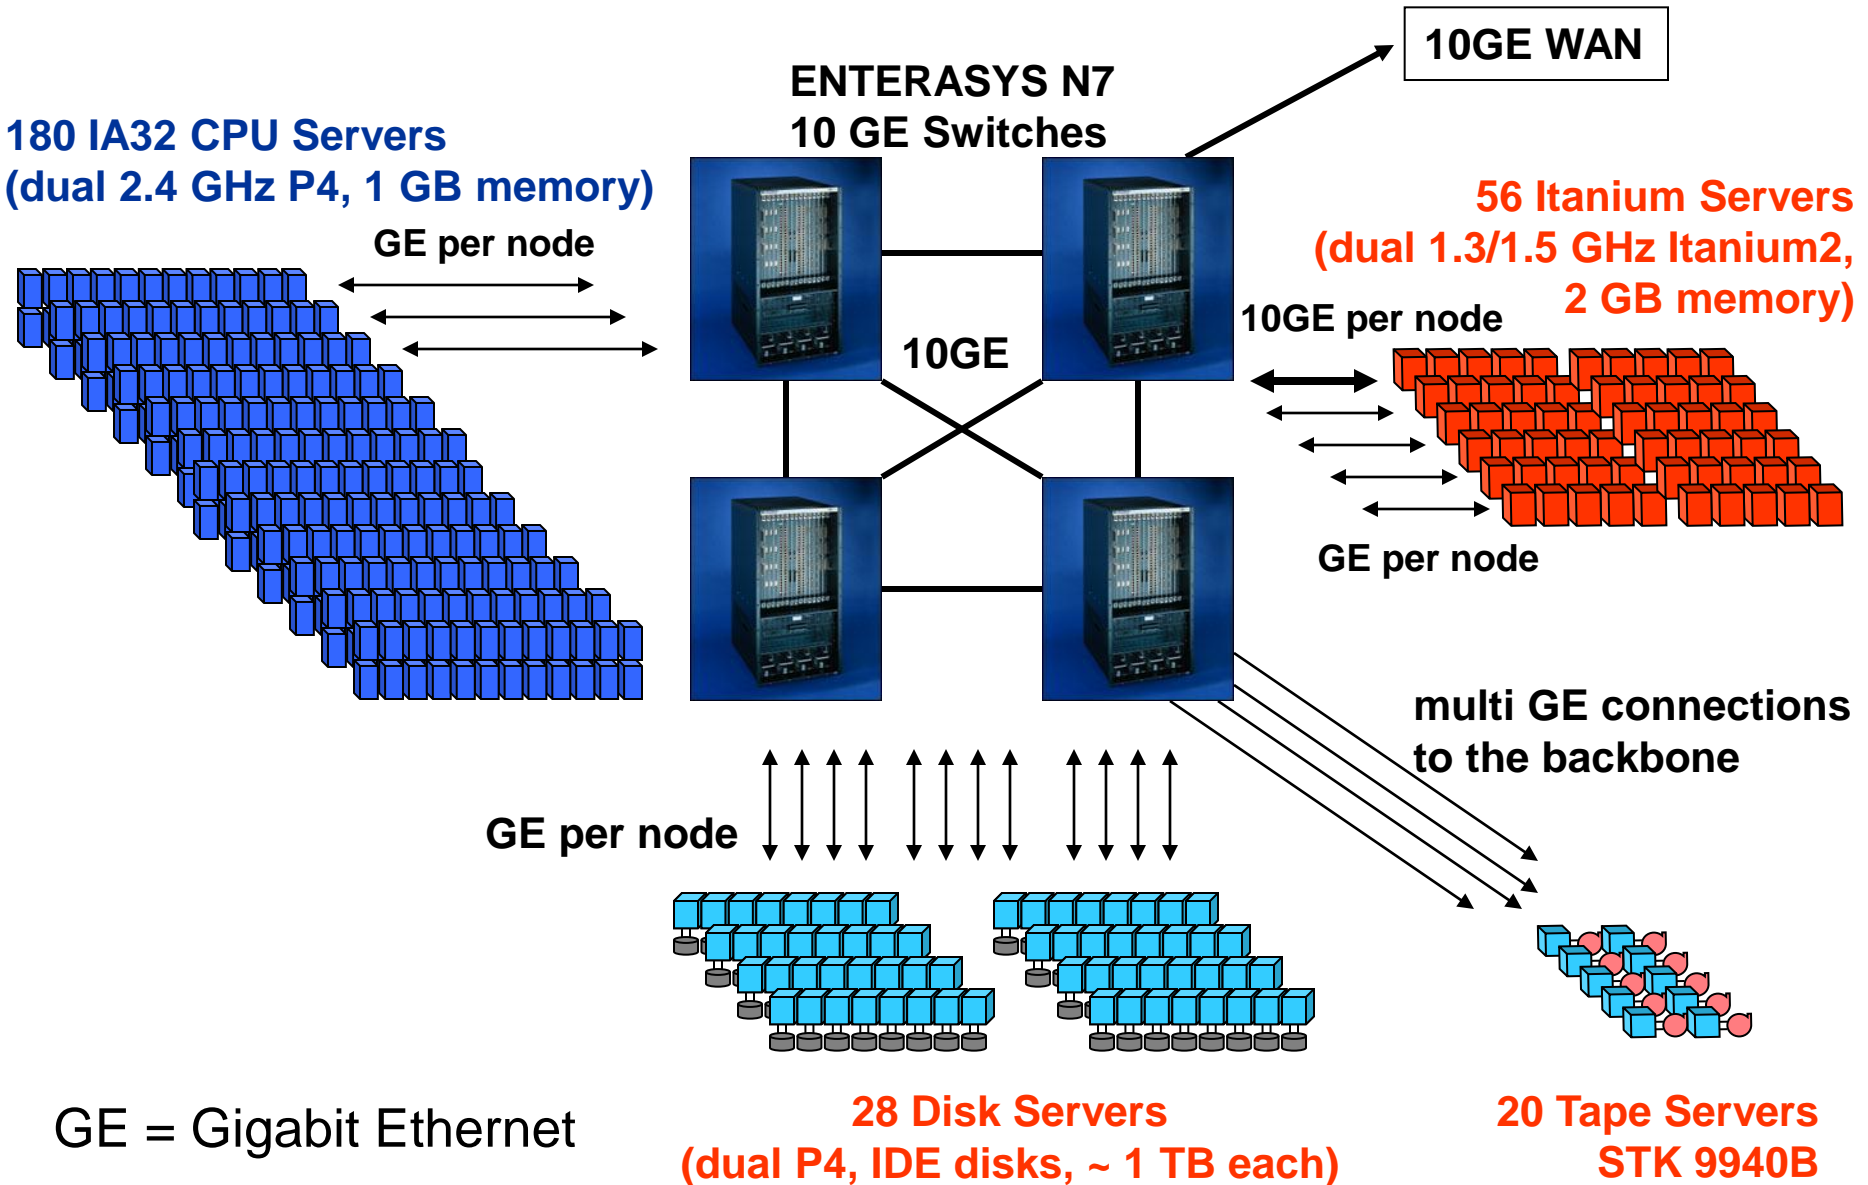
The CERN opencluster


Objectives

- Build an ultrahigh performance computer cluster
- Link it to the DataGrid and test its performance
- Evaluate potential of new technologies for LCG
- Port Grid software and physics applications to new environment



High Throughput Prototype (openlab + LCG prototype)





The CERN opencluster

Sponsorship so far

- 120 **Intel** Itanium™ processors (64-bit technology) and network cards
- **HP** servers for 60-node cluster
- Three 10-Gbit/s switches from **Enterasys** Networks, recently replaced by four large N7-type switches
- 28TB SAN File System System from **IBM**
- **Oracle** joined recently, Oracle 10g now under evaluation
- Dedicated R&D staff from industrial partners
- Funding for four CERN fellows (postdocs) (**HP** & **Oracle**)







The CERN opencluster

IBM's contribution

- IBM joined the CERN openlab in April
- Sponsors SAN File System with
 - 28TB of iSCSI storage
 - 6 node SAN server cluster
- A key feature of the CERN installation is the use of **iSCSI** with Linux clients

IBM delegation at openlab annual sponsors meeting, June 2003





The CERN opencluster


SAN File System status

- **Equipment installed over summer 2003**
- **80 Linux clients connected**
- **Reliability and performance tests currently being run**

Medium-term goals

- **Test SAN File System in CERN environment (very heterogeneous)**
- **Link to opencluster and DataGrid with Distributed SAN File System**
- **Integrate existing disk-servers during 2004**





Preliminary results & goals

- **iSCSI** is a brand-new extension to SAN File System. Initial tests were promising results, but performance was not sufficient


Remember:

This is still an advance development project, not a final product !

ST over long-distance networks

- Find cost-effective solutions matching the needs of LHC





Conclusions



- **CERN's LHC experiments have enormous computing needs**
- **The LCG/EGEE projects are the major driving force for Grid technology developments in the scientific arena**
- **CERN is leading some of the most ambitious data-intensive Grid projects worldwide**
- **CERN is collaborating with leading IT partners via the CERN openlab**
- **IBM is providing key storage technology to the partnership**