



where the Web was born

CERN and the LHC Computing Challenge

by
Wolfgang von Rueden
Head, IT Department

CERN 50th anniversary with openlab partners
19th October 2004



What is CERN?

- CERN is the world's largest **particle physics** centre
- Particle physics is about:
 - **elementary particles**, the constituents all matter in the Universe is made of
 - **fundamental forces** which hold matter together
- Particles physics requires:
 - **special tools** to create and study new particles

CERN is also:

-2500 staff
(physicists, engineers,
technicians, ...)

- Some **6500 visiting
scientists** (half of the
world's particle
physicists)


They come from
500 universities
representing
80 nationalities.





What is CERN?

- Physicists smash particles into each other to:
 - identify their **components**
 - **create** new particles
 - reveal the nature of the **interactions** between them
 - recreate the environment present at the origin of our Universe (big bang)
- **What for?** To answer fundamental questions like:
*how did the Universe begin? What is the origin of mass?
What is the nature of antimatter?*





What is CERN?


The special tools for particle physics are:

- **ACCELERATORS**, huge machines able to speed up particles to very high energies before colliding them into other particles
- **DETECTORS**, massive instruments which register the particles produced when the accelerated particles collide
- **COMPUTING**, to re-construct the collisions, to extract the physics data and perform the analysis





The CERN Site





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 the 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



Typical Experiment Layout




- Complex system of detectors centred around the beam interaction point



ATLAS, one of the four LHC experiments

ATLAS has 150 million measurement channels



As tall as our main building !





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





The LHC Computing Grid Project

Les Robertson
LCG Project Leader
CERN, IT Department

les.robertson@cern.ch - www.cern.ch/lcg





LHC Computing Grid Project

Aim of the project

To prepare, deploy and operate the **computing environment**

for the **experiments** to analyse the data from
the LHC


detectors

Applications development environment, common tools and frameworks

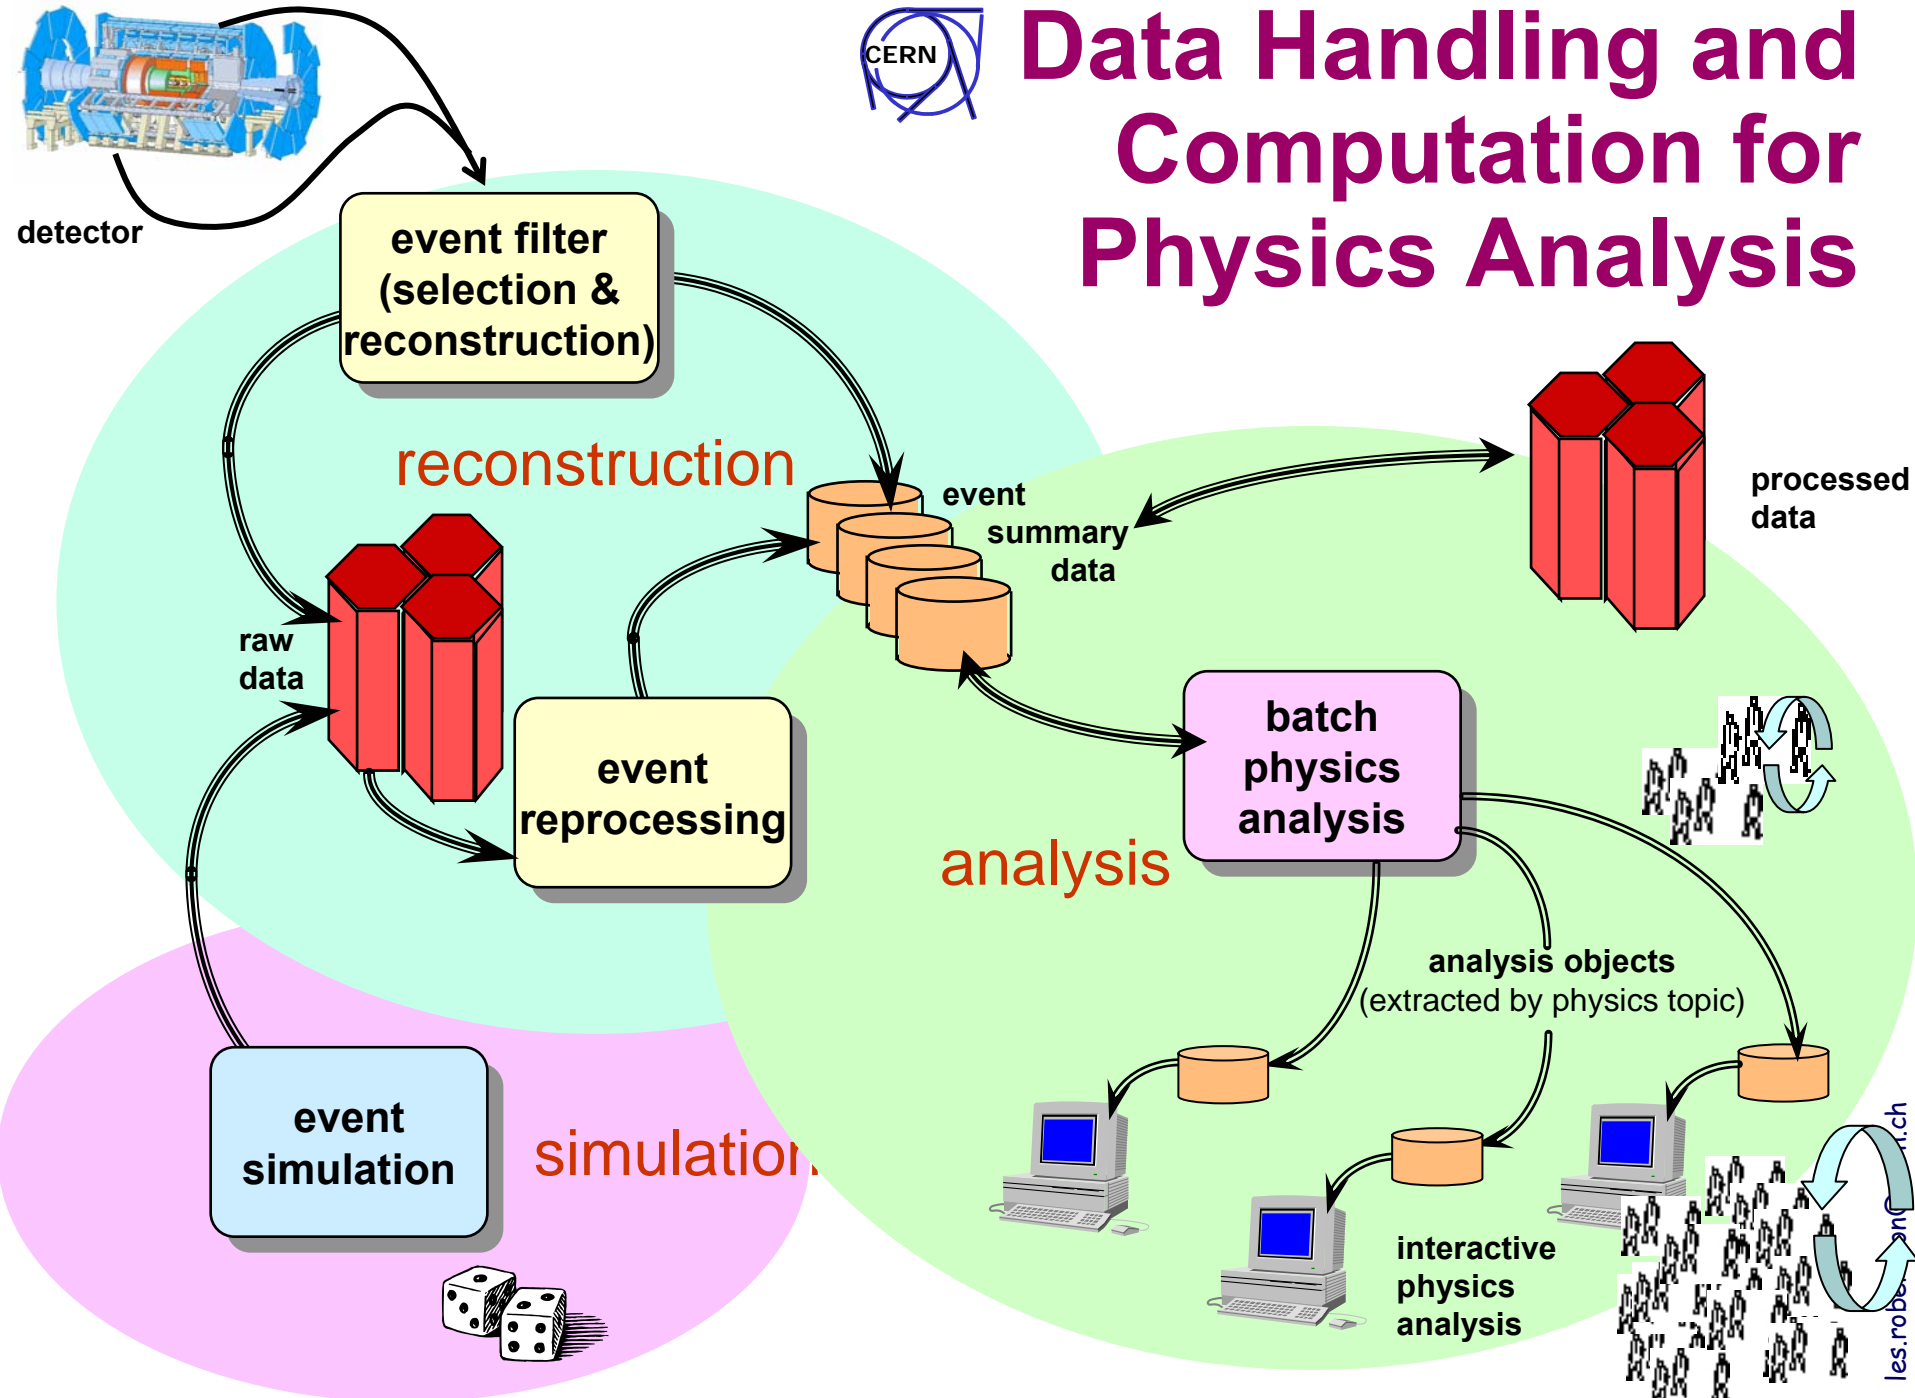
Build and operate the LHC computing service

The Grid is just a tool towards achieving this goal



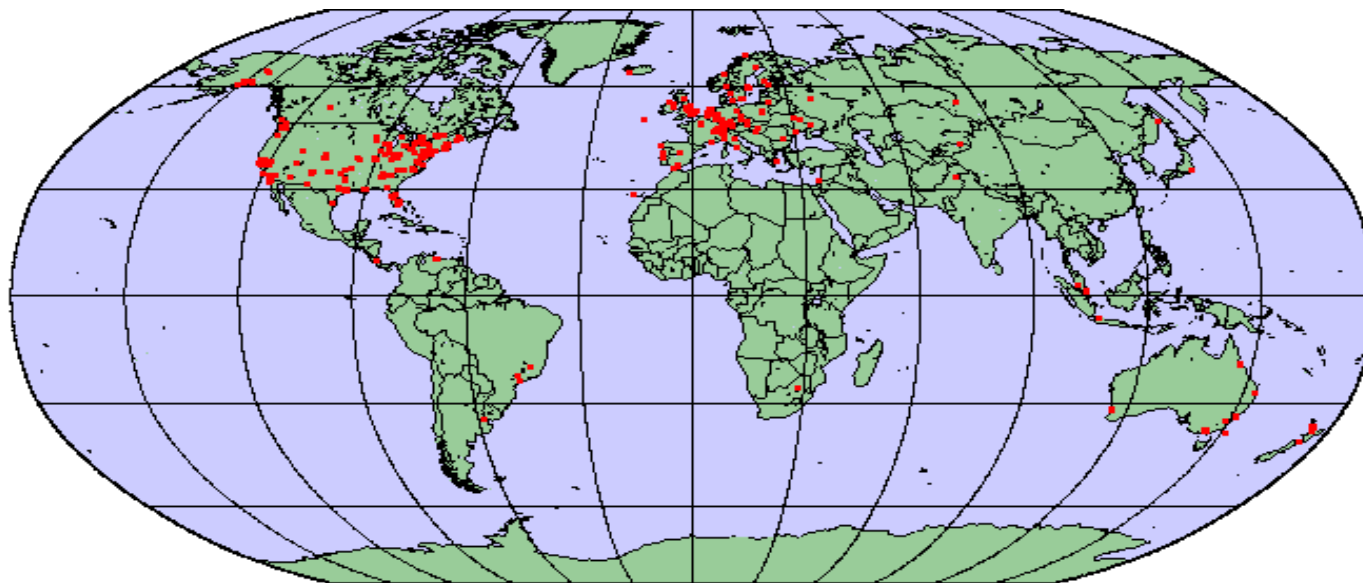


Data Handling and Computation for Physics Analysis





The CERN Community



Europe:	267 institutes	4603 users
Elsewhere:	208 institutes	1632 users



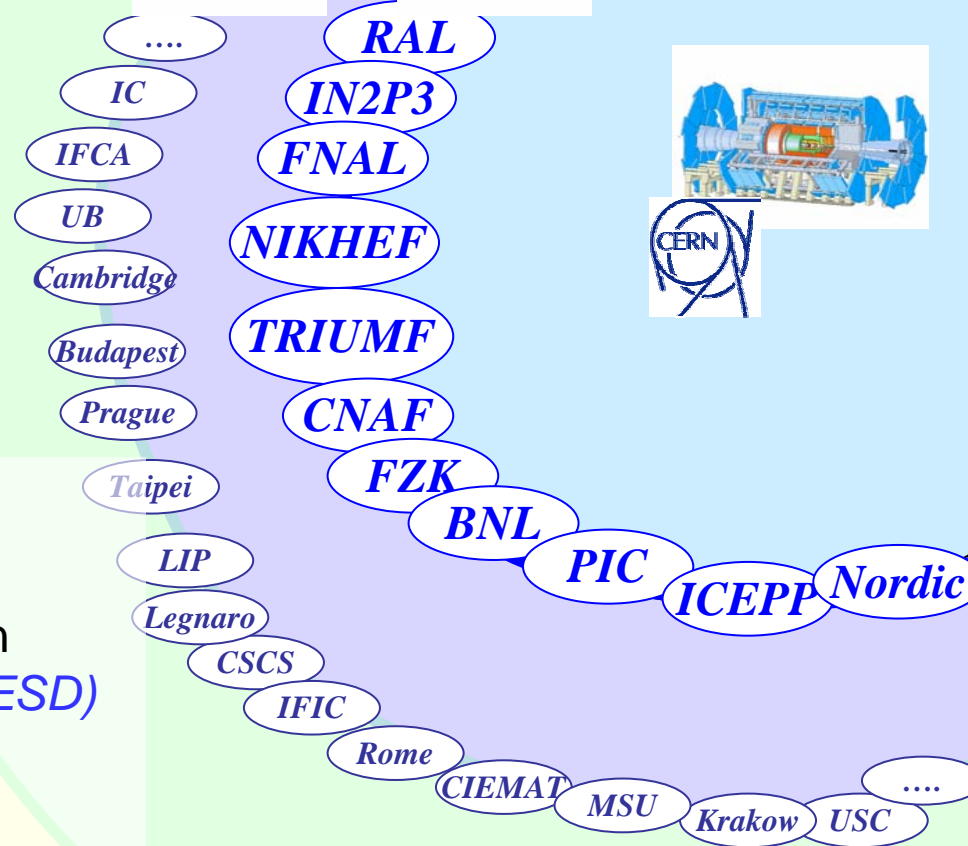


desktops
portables

small
centres

Tier-2

Tier-1



LHC Computing Model (**simplified!!**)

- Tier-0 – the accelerator centre
 - Filter → *raw data* → reconstruction
→ *event summary data (ESD)*
 - Record and distribute the data to Tier-1s
- Tier-1 –
 - **Managed Mass Storage** –
→ **grid-enabled data service**
 - Data-heavy, batch analysis
 - National, regional support
 - “online” to the data acquisition process
high availability, long-term commitment

- Tier-2 –
 - Well-managed, grid-enabled disk storage
 - End-user analysis – batch and interactive
 - Simulation




desktops
portables

small
centres

Tier-2

Tier-1



Current estimates of Computing Resources needed at Major LHC Centres

First full year of data - 2008

Processing
M SI2000**

Disk
PetaBytes

Mass
Storage
PetaBytes

CERN

20

5

20

Major data
handling centres
(Tier 1)

50

22

17

Other large
centres (Tier 2)

40

12

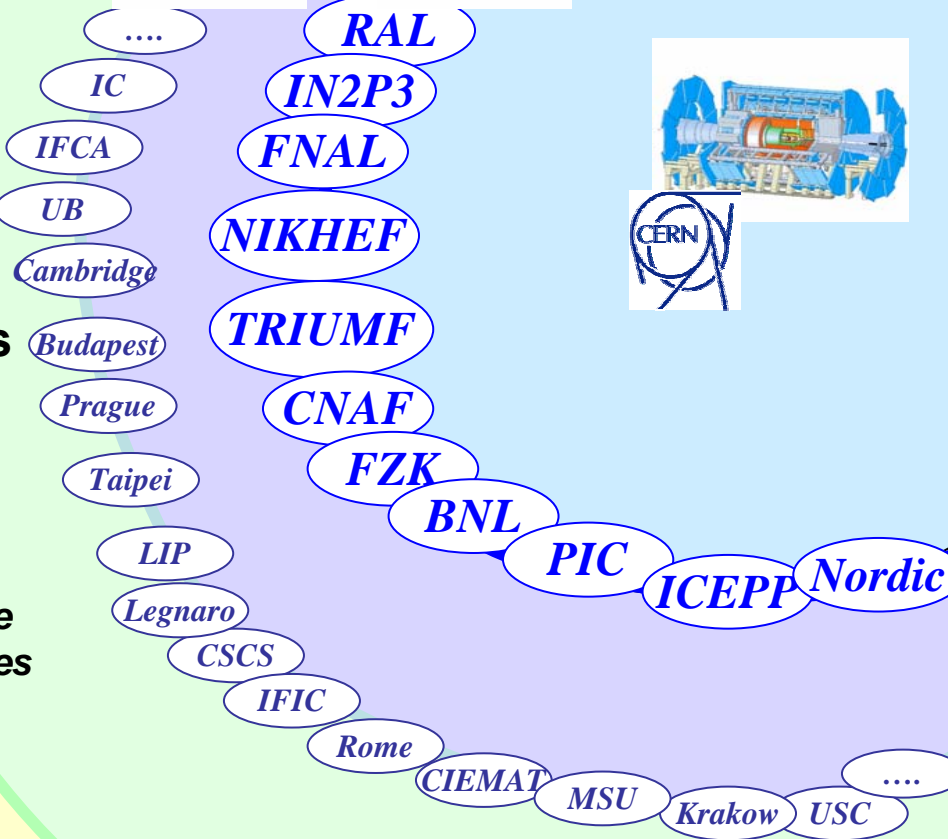
5

Totals

110

39

42



** Current fast processor ~1K SI2000



LHC Computing Grid Project - a Collaboration

Building and operating the LHC Grid – a collaboration between

- The physicists and computing specialists from the LHC experiments
- The projects in Europe and the US that have been developing Grid middleware
- The regional and national computing centres that provide resources for LHC
- The research networks

Researchers

Computer Scientists & Software Engineers

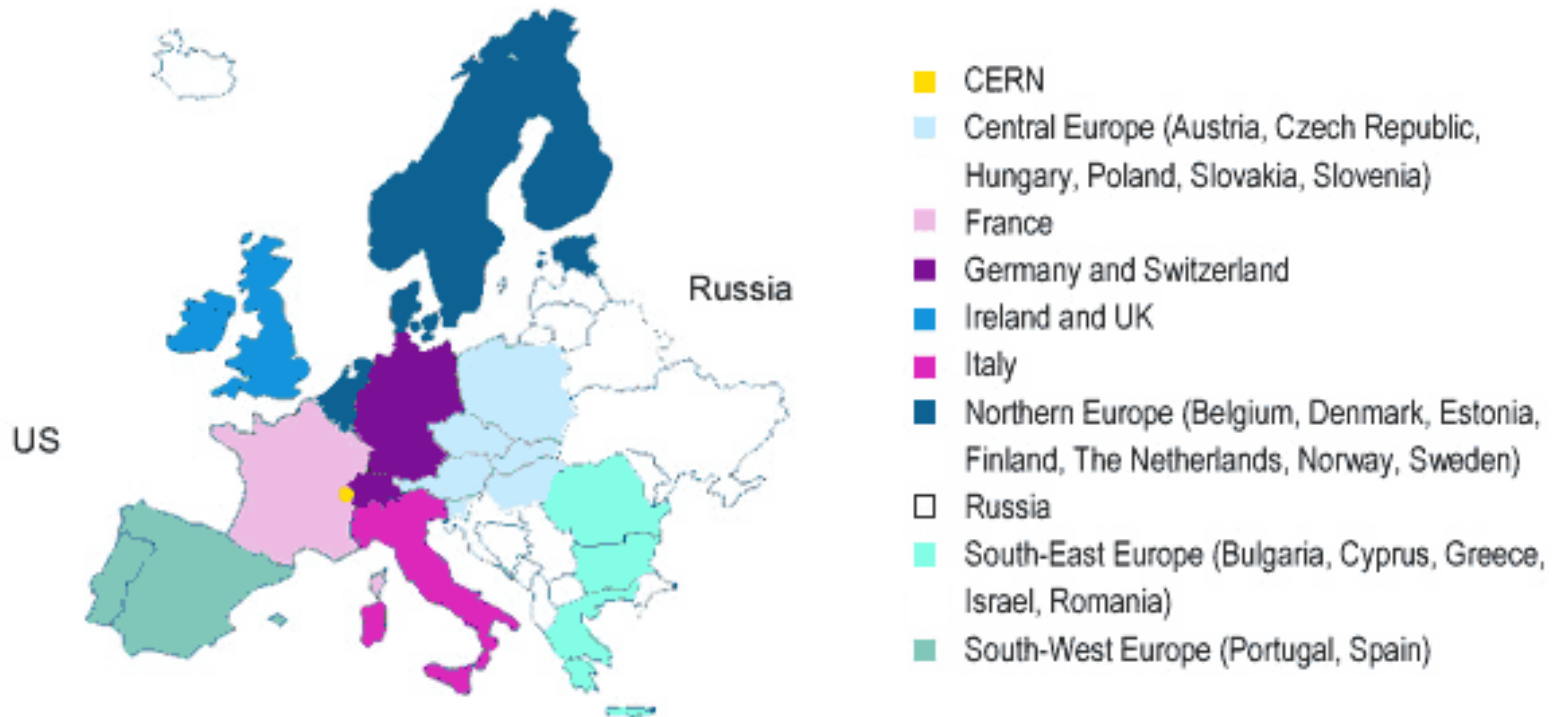
Service Providers



Virtual Data Toolkit



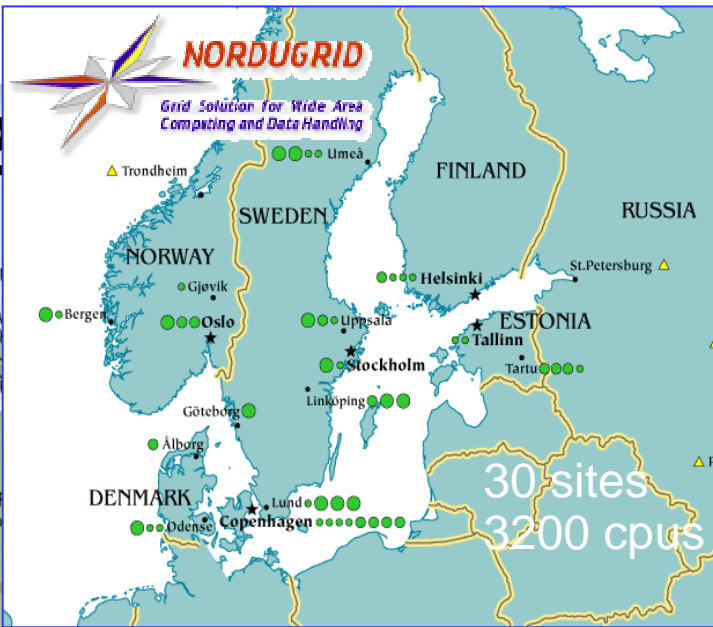
Enabling Grids for
E-science in Europe



70 institutions in 27 countries



LCG-2




- GLASGOW
- LANCS
- LIVHEP
- MANHEP
- BHAM
- CSTCDIE
- RALPP
- RAL
- OX-PHYSICS
- IC
- RH
- EDINB
- SHE
- CA
- Q
- L
- IC
- NGS-RAL
- NGS-OXFORD
- NGS-MAN
- NGS-LEEDS
- IN2P3-LI
- IN2P3
- USC
- IFCA
- INTA-CAB
- LIP
- UAM
- CIEMAT
- IFIC
- CNB
- HELLASGRID
- AUTH

EGEE
Enabling Grids for
E-science in Europe

**25 Universities
4 National Labs
2800 CPUs**




Total Sites	82
Total CPUs	7269
Total Storage (TB)	6558
Wed September 22 2004	





Data Readiness Programme

- 1 Dec04 - Basic data handling verification**
CERN + 3 Tier-1s, 500 MB/sec, physics data sets - sustained for two weeks
- 2 Mar05 - Reliable file transfer service in operation**
mass store (disk) - mass store (disk), CERN+5 sites
500 MB/sec between sites, sustained for one month
- 3 Jul05 - Infrastructure verification**
CERN + 50% of Tier-1s, sustained operation at 300 MB/sec. including tapes
Nov05 - ATLAS and CMS Tier-0/1 model verification at *half scale*
- 4 Apr06 - Infrastructure operational**
ALL Tier-1s, 50% of Tier-2s - full target data rates
Aug06 - All experiments - Tier-0/1/2 model verification at *full scale*
- 5 Nov 06 Infrastructure Ready**
ALL Tier-1s, most Tier-2s - operating at *twice* target data rates
Feb07 - all experiments - full model in operation





Summary

- LHC computing -
 - **Data intensive** - Geographically distributed
 - Independent regional centres

- LHC Grid -
 - Reliable environment for data intensive batch work
 - An early example of a working data-intensive grid
 - Co-existing with multiple grids, other sciences

- Current status
 - Large global grid established - and being used for real work by LHC experiments
 - Middleware - basic functionality, acceptable reliability
 - Beginning now to tackle
 - **Operations management**
 - **Performance**
 - Ambitious schedule to achieve required service level by March 2007
 -

- Long-term expectation -
 - Science grids operated as national/international infrastructure





Enabling Grids for
E-science in Europe

OpenLab sponsors meeting, October 2004

EU EGEE project – status and plans

Bob Jones


EGEE Technical Director

Bob.Jones@cern.ch



In 2 years EGEE will:

- **Establish production quality sustained Grid services**
 - 3000 users from at least 5 disciplines
 - over 8,000 CPU's, 50 sites
 - over 5 Petabytes (10^{15}) storage
- Demonstrate a viable general process to **bring other scientific communities on board**
- **Propose a second phase** in mid 2005 to take over EGEE in early 2006




32 Million Euros EU funding over 2 years starting 1st April 2004

- **48 % service activities** (Grid Operations, Support and Management, Network Resource Provision)
- **24 % middleware re-engineering** (Quality Assurance, Security, Network Services Development)
- **28 % networking** (Management, Dissemination and Outreach, User Training and Education, Application Identification and Support, Policy and International Cooperation)




Emphasis in EGEE is on operating a production grid and supporting the end-users

- BioMedical
 - Bioinformatics (gene/proteome databases distributions)
 - Medical applications (screening, epidemiology, image databases distribution, etc.)
 - Interactive application (human supervision or simulation)
 - Security/privacy constraints
 - Heterogeneous data formats - Frequent data updates - Complex data sets - Long term archiving
- BioMed applications deployed and going live in September
 - **GATE** - Geant4 Application for Tomographic Emission
 - **GPS@** - genomic web portal
 - **CDSS** - Clinical Decision Support System



- Intended to replace LCG-2
- Starts with existing components from AliEN, EDG, VDT etc.
- Aims to address LCG-2 shortcomings and advanced needs from applications
- Prototyping short development cycles for fast user feedback
- Initial web-services based prototypes being tested with representatives from the application groups



Application requirements <http://egee-na4.ct.infn.it/requirements/>

Intellectual Property

- The existing EGEE grid middleware (LCG-2) is distributed under an Open Source License developed by EU DataGrid
 - Derived from modified BSD - no restriction on usage (academic or commercial) beyond acknowledgement
 - Same approach for new middleware (gLite)
- Application software maintains its own licensing scheme
 - Sites must obtain appropriate licenses before installation



Who else can benefit from EGEE?

- EGEE Generic Applications Advisory Panel:
 - 4 applications presented
 - 3 applications (comp. chemistry, earth science, astro-particle) recommended for deployment with allocation of NA4 resources
 - EU GRACE project already tested
- EU projects: MammoGrid, Diligent, SEE-GRID ...
- Expression of interest: Planck/Gaia (astroparticle), SimDat (drug discovery)



<http://agenda.cern.ch/age?a042351>

Next meeting at EGEE conference (November)

User training and induction

- Training material and courses from introductory to advanced level
- Train a wide variety of users both internal to the EGEE consortium and external groups from across Europe
- ~20 courses/presentations already held and many more planned (see roadmap)
- Experience with GENIUS portal and GILDA testbed
- Courses inline with the needs of the projects and applications

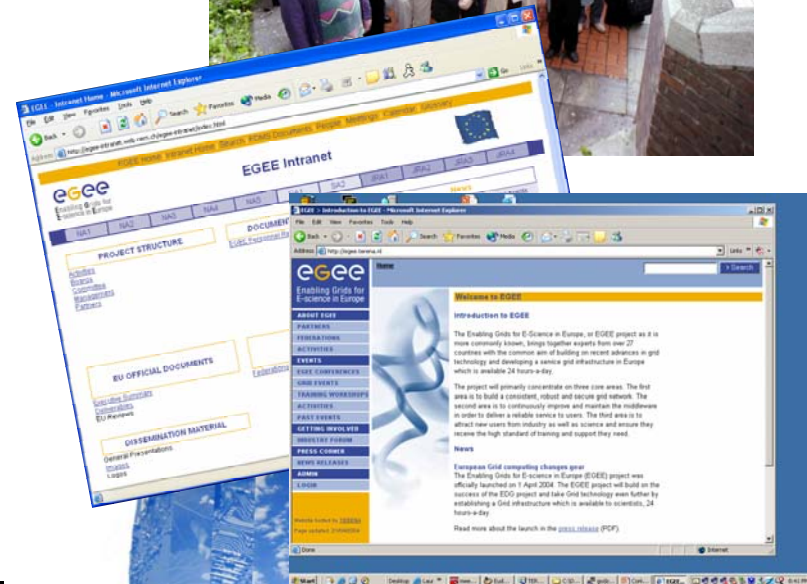


Training: <http://www.egee.nesc.ac.uk/>

Roadmap: <http://www.egee.nesc.ac.uk/schedreg/index.html>


Dissemination

- 1st project conference
 - Over 300 delegates came to the 4 day event during April in Cork Ireland
 - Kick-off meeting bringing together representatives from the 70 partner organisations
- 2nd conference scheduled
 - 22-26 November in The Hague
 - <http://public.eu-egEE.org/conferences/2nd/>
- Websites, Brochures and press releases
 - For project and general public **www.eu-egEE.org**
 - Information packs for the general public, press and industry



eGEE
Enabling Grids for E-science in Europe
<http://public.eu-egEE.org/>

- EGEE Industry Forum
 - raise awareness of the project in industry to encourage industrial participation in the project
 - foster direct contact of the project partners with industry
 - ensure that the project can benefit from practical experience of industrial applications
- For more info:
<http://public.eu-egee.org/industry/>
- Expect Industry to play an important role in follow-on project



EGEE Summary

- **EGEE** is the first attempt to build a worldwide Grid infrastructure for data intensive applications from **many scientific domains**
- A **large-scale production grid service** is already deployed and being used for HEP and BioMed applications with new applications being ported
- Resources & user groups will **rapidly expand** during the project
- A process is in place for **migrating new applications** to the EGEE infrastructure
- A **training programme** has started with events already held
- Prototype “*next generation*” middleware is being tested (**gLite**)
- Plans for a **follow-on project** are being discussed

More information: www.eu-egee.org



CERN



openlab for DataGrid applications

*In partnership with
and sponsored by*

IBM

intel®



ORACLE®

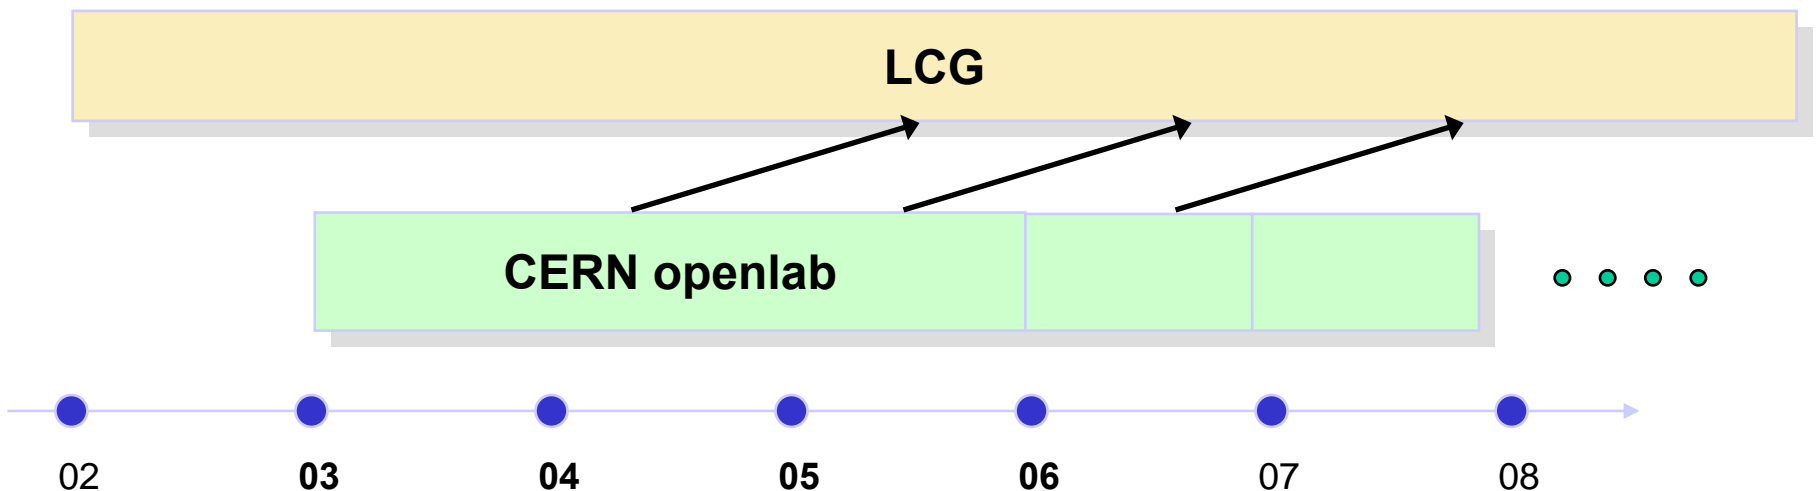


**Wolfgang von Räden
Head, IT Department**



CERN openlab

- IT Department's main R&D focus
- Framework for collaboration with industry
- Evaluation, integration, validation
 - of cutting-edge technologies that can serve LCG
- Initially a 3-year lifetime
 - As of 1.1.2003
 - Later: Annual prolongations
- Slogan: “You make it, we break it”.





openlab participation

- **Five Partners (contributing ≥ 1.5 M€ over 3 years)**
 - Enterasys:
 - 10 GbE core routers
 - HP:
 - Integrity servers (103 * 2-ways, 2 * 4-ways)
 - Two post-doc positions
 - IBM:
 - Storage Tank file system (SAN FS), currently with 28 TB
 - Intel:
 - Large number of 64-bit Itanium processors & 10 Gbps NICs
 - 64-bit Nocona system w/PCI-Express
 - Oracle:
 - 10g Database software w/add-ons
 - Two post-doc positions
- **One contributor (contributing ≥ 170 k€ for 1 year)**
 - Voltaire
 - 96-way Infiniband switch and necessary HCAs

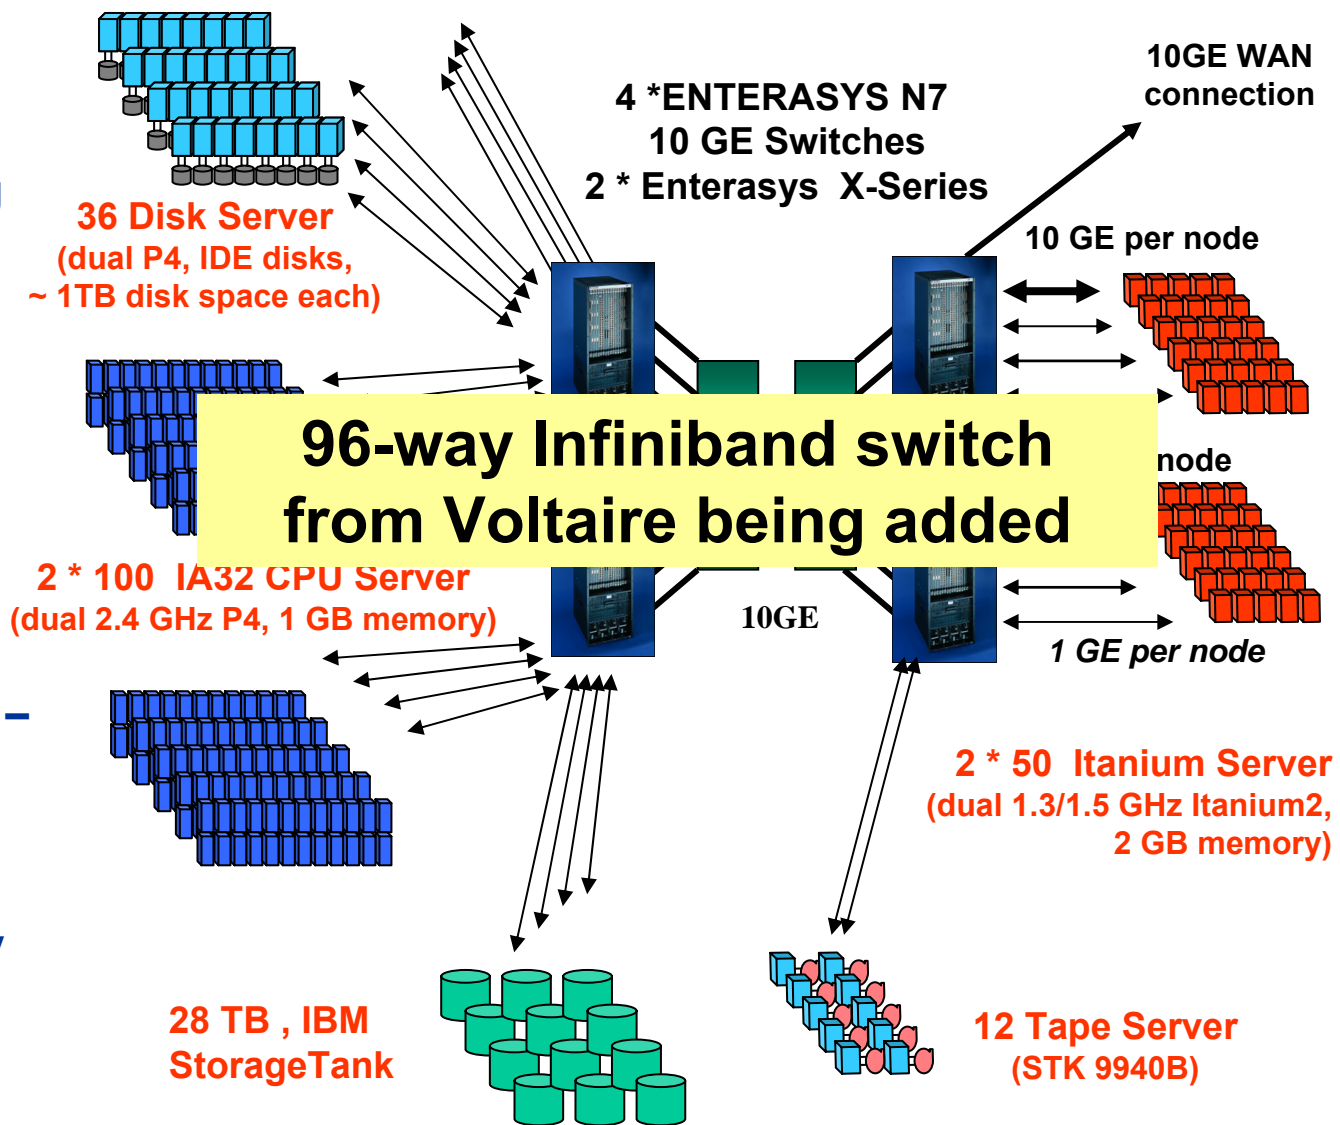


High Throughput Cluster Prototype openlab/LCG

4 * GE connections
to the campus backbone

openlab/LCG

- Experience with likely ingredients in LCG:
 - 64-bit programming
 - next generation I/O (10 Gb Ethernet, Infiniband, etc.)
- High performance cluster used for evaluations, and for data challenges with experiments
- Flexible configuration – components moved in and out of production environment
- Co-funded by industry and CERN





Next project: security

- **Grid projects (LCG/EGEE) address Grid specific security issues**
- **Site security is not included**
- **No Grid security w/o site security**
- **Proposal:**
 - Address site security globally, including all aspects
 - 25 companies invited, 16 expressed interest and came to the first meeting last Thursday
 - Workshop to refine the project on 18/19 November
 - Still open for others to join

Thank you for your attention



Questions?