

Worldwide LHC Computing Grid Project

Getting ready
for LHC start-
up
October 2007



Sverre Jarp
CERN openlab CTO





Briefly about CERN





What is CERN?

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







CERN's tools

The special tools for particle physics are:

- **ACCELERATORS**, huge machines (inside a complex underground structure) - able to accelerate 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 reconstruct the collisions, to extract the physics data and to perform the analysis



CERN in Numbers

- 2500 Staff
- 6500 Users
- 500 Fellows and Associates
- 80 Nationalities
- 500 Universities
- Budget ~1200 MCHF/year
(~730 M€/year)

- 20 Member States:
Austria, Belgium, Bulgaria, the Czech Republic, Denmark, Finland, France, Germany, Greece, Hungary, Italy, Netherlands, Norway, Poland, Portugal, Slovakia, Spain, Sweden, Switzerland and the United Kingdom.
- 8 Observers:
India, Israel, Japan, the Russian Federation, USA, Turkey, the European Commission and UNESCO



What is LHC?

LHC will be switched on in **2007**

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


ALICE

ATLAS

CMS


LHCb

- It is a particle accelerator that will collide beams of protons at an energy of **14 TeV**
- Using the latest super-conducting technologies, it will operate at about **-271°C** , just above the absolute zero of temperature
- With its **27 km circumference**, the accelerator will be the largest superconducting installation in the world.
- Its two proton beams will interact 40 million times per second (3000 bunches of 100 billion protons each)






ATLAS construction






An ATLAS event






PHYSICS COMPUTING






LHC DATA




This is reduced by online computers that filter out a few hundred “good” events per sec.



Which are recorded on disk and magnetic tape

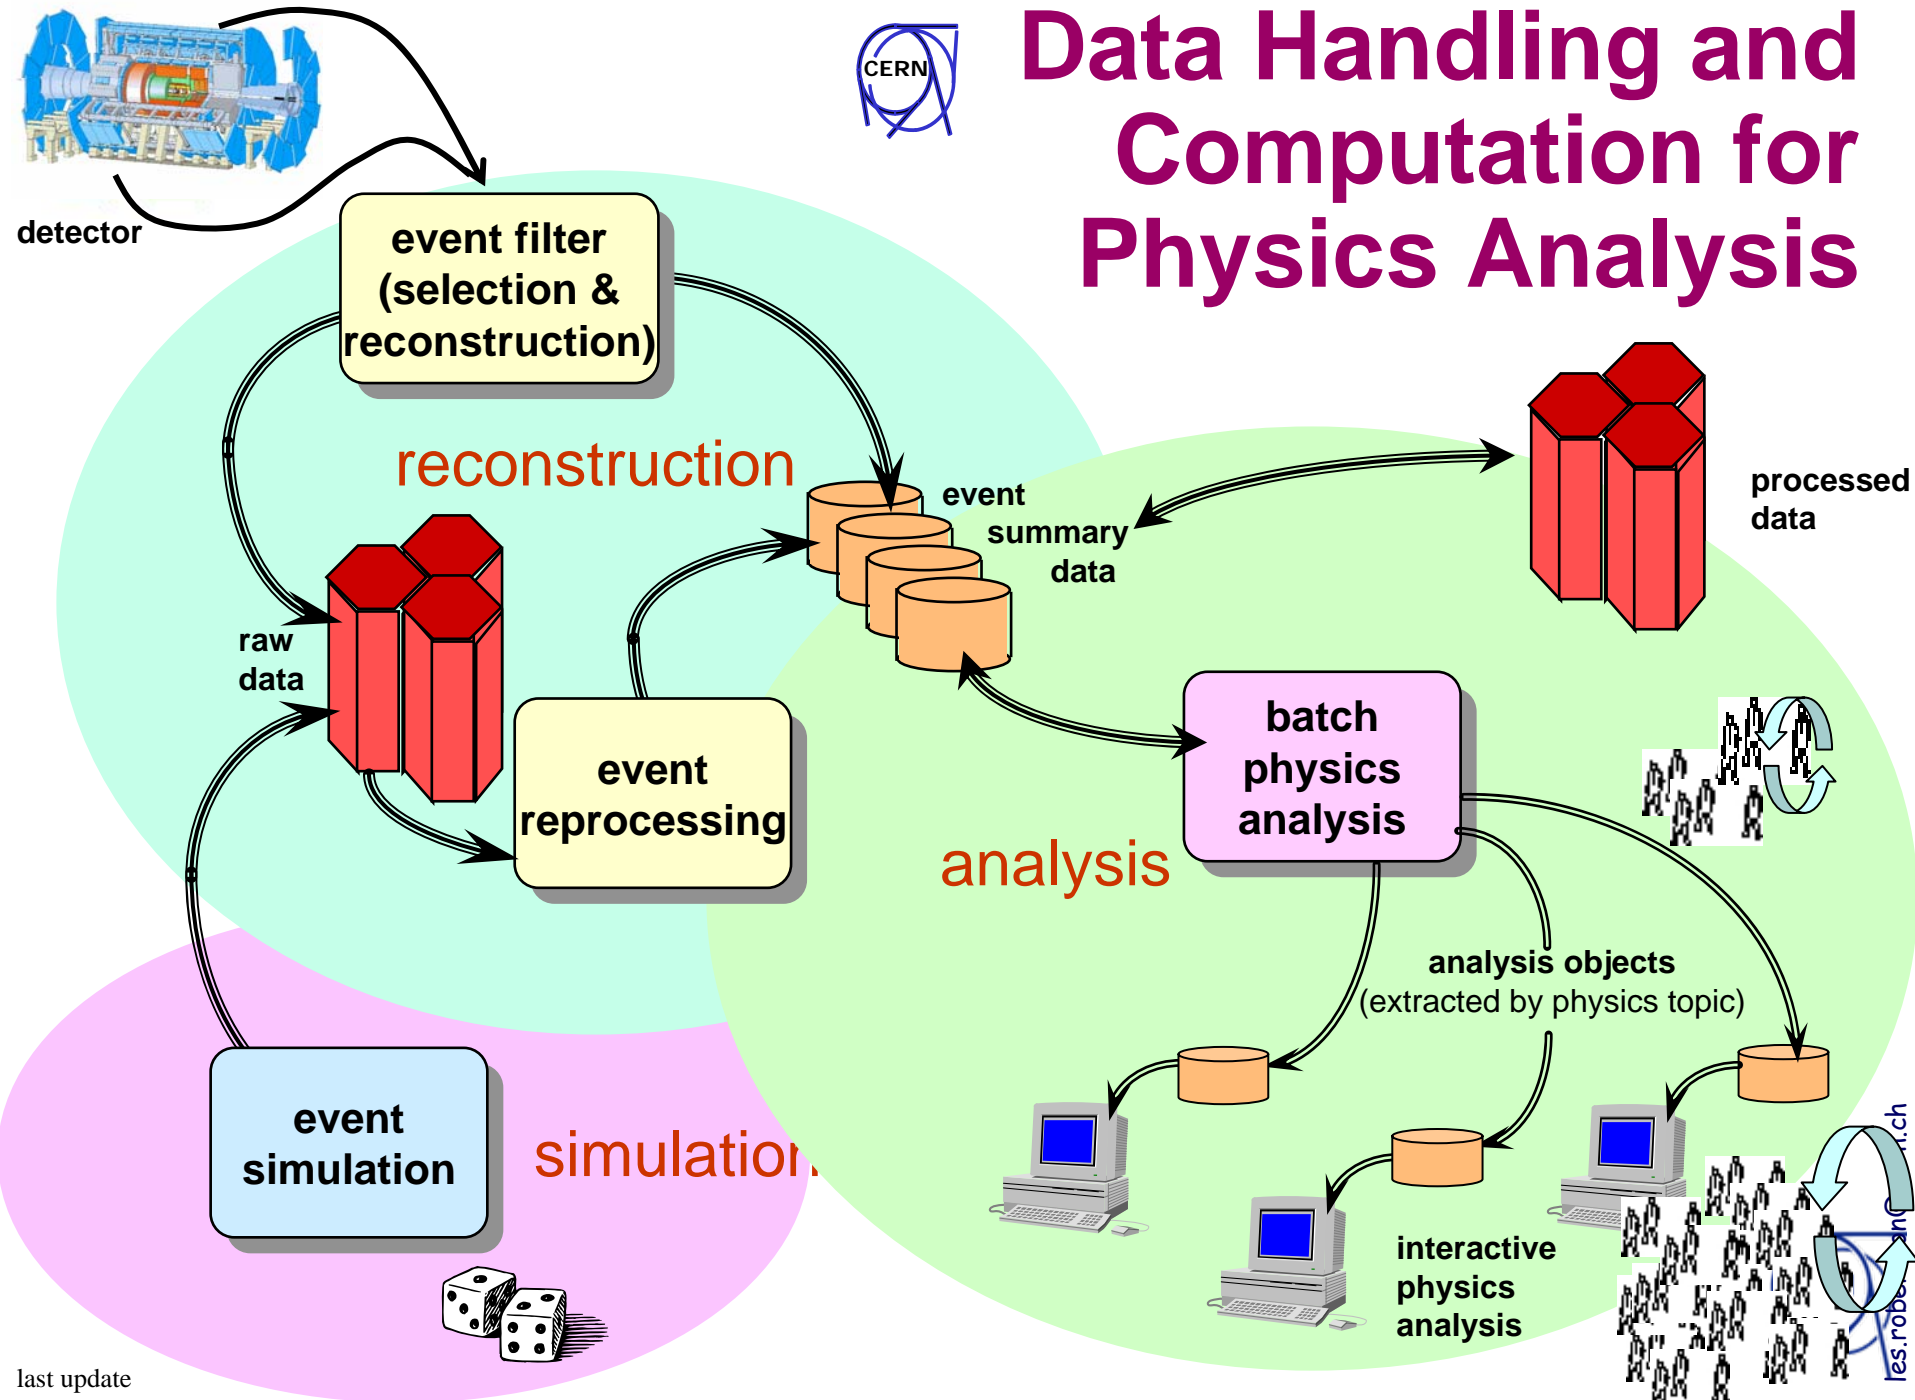
at 100-1,000 MegaBytes/sec

→ ~15 PetaBytes per year
for all four experiments





Data Handling and Computation for Physics Analysis



High Energy Physics Computing Characteristics




- Independent events (collisions of particles)
 - trivial (read: pleasant) parallel processing
- Bulk of the data is read-only
 - versions rather than updates
- Meta-data in databases, but physics data in “flat” files
- Compute power measured in SPECint (rather than SPECfp)
 - But good floating-point is important
- Very large aggregate requirements:
 - computation, data, input/output
- Chaotic workload -
 - research environment - physics extracted by iterative analysis, collaborating groups of physicists
 - Unpredictable → unlimited demand





The Computing Environment

- High-throughput computing (based on reliable “commodity” technology)
 - Around 3000 (dual-socket Xeon) PCs with “Scientific Linux”
 - Now typically also “dual-core”
 - Quad-core included in our next acquisition





The world-wide LHC Computing Grid

- The WLCG Collaboration

- 4 LHC experiments
- ~140 computing centres
- 12 large centres
(Tier-0, Tier-1)
- 38 *federations* of smaller
"Tier-2" centres
- ~35 countries

- Resources

- Contributed by the countries participating in the experiments
- Commitment made each October for the coming year
- 5-year forward look

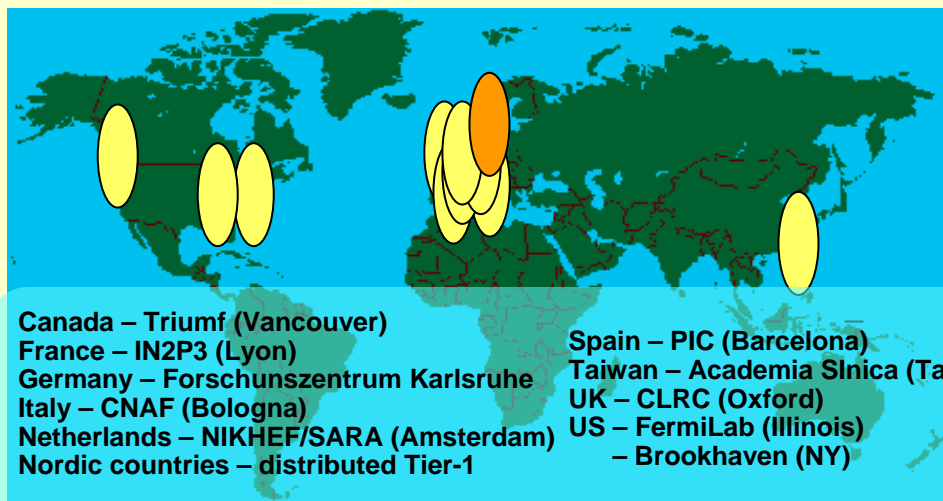




LCG Service Hierarchy

Tier-0 - the accelerator centre

- Data acquisition & initial processing
- Long-term data curation
- Distribution of data → Tier-1 centres



Tier-1 - "online" to the data acquisition process → high availability

- Managed Mass Storage -
→ grid-enabled data service
- Data-heavy analysis
- National, regional support

Tier-2 - ~130 centres in ~35 countries

- **End-user (physicist, research group) analysis** –
where the discoveries are made
- Simulation





Distribution of Computing Services

Summary of Computing Requirements




All experiments - 2008

From LCG TDR - June 2008

New data will grow at about 15 PetaBytes per year – with two copies

about 100,000 cores

	CERN	All Tier-1s	All Tier-2s	Total
CPU (MSPECint2000s)	25	56	61	142
Disk (PetaBytes)	7	31	19	57
Tape (PetaBytes)	18	35		53



Significant fraction of the resources distributed over more than 120 computing centres



Solution: the Grid

- Use the Grid to unite computing resources of particle physics institutes around the world

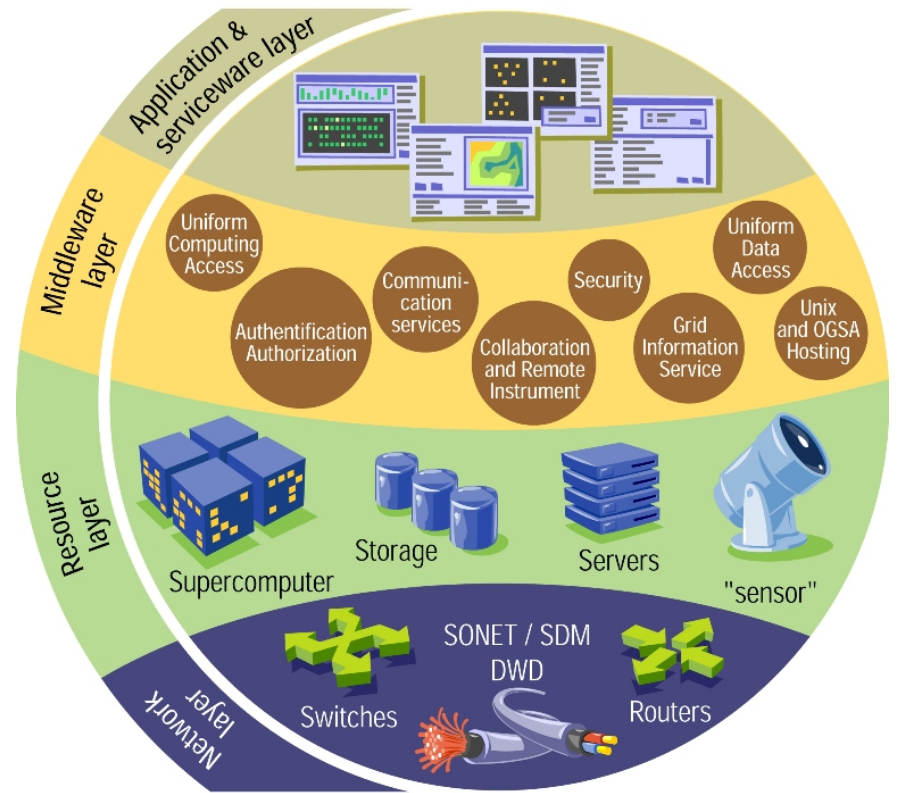
The **World Wide Web** provides seamless access to information that is stored in many millions of different geographical locations

The **Grid** is an infrastructure that provides seamless access to computing power and data storage capacity distributed over the globe



How does the Grid work?

- It relies on special software, called **middleware**.
- Middleware automatically finds the **data** the scientist needs, and the **computing power** to analyse it.
- Middleware balances the load on different resources. It also handles **security**, **accounting**, **monitoring** and much more.





LCG depends on two major science grid infrastructures ...



EGEE - Enabling Grids for E-Science

OSG - US Open Science Grid



A map of the worldwide LCG infrastructure operated by EGEE and OSG.

- **More than 20 applications from 7 domains**
 - High Energy Physics (**Pilot domain**)
 - 4 LHC experiments
 - Other HEP (DESY, Fermilab, etc.)
 - Biomedicine (**Pilot domain**)
 - Bioinformatics
 - Medical imaging
 - Earth Sciences
 - Earth Observation
 - Solid Earth Physics
 - Hydrology
 - Climate
 - Computational Chemistry
 - Fusion
 - Astronomy
 - Cosmic microwave background
 - Gamma ray astronomy
 - Geophysics
 - Industrial applications



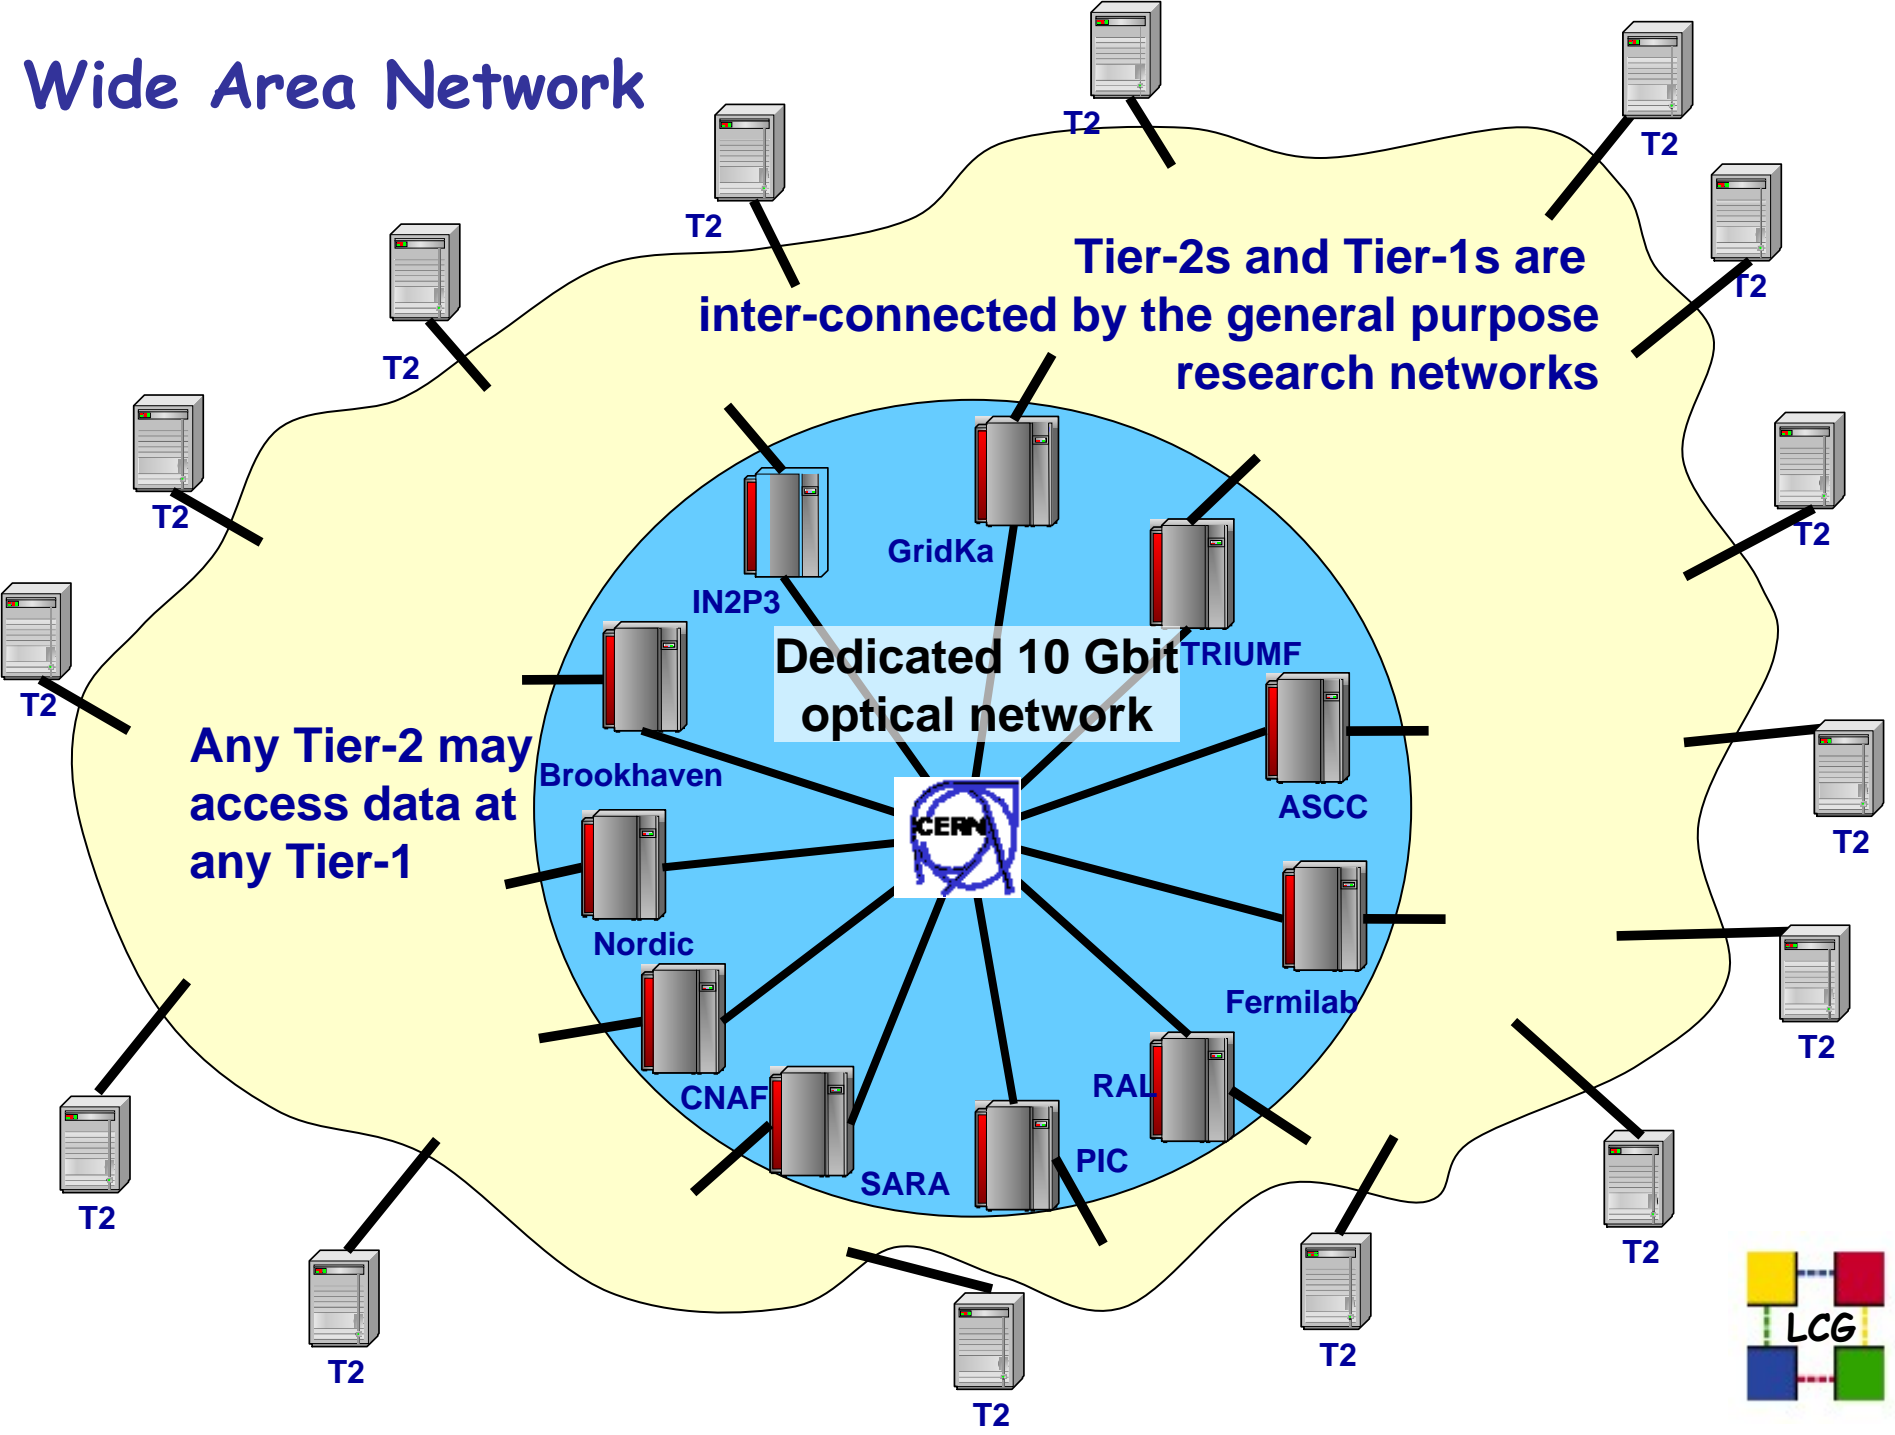


Wide Area Networking

- There is a special working group for bringing together the Tier-1 centres, the national and regional research networks to plan for LHC wide area networking
 - Responsibilities
 - Overall architecture and implementation schedule
 - Agree operational responsibilities as necessary
 - CERN/Tier-1 high performance connections (mostly at 10Gbps)
 - Ensure adequate Tier-1/Tier-2 connectivity
 - Active participation of NRENs
 - Abilene/ESNET/LHCnet (USA), ASnet (TW), Canarie (CDN), DFN (D), GARR (I), NorduGrid, RedIris (E), RENATER (F), SURFNET (NL), UKERNA,
- and the European backbone network GEANT



Wide Area Network







Grid Activity


- Continuing increase in usage of the EGEE and OSG grids
- All sites reporting accounting data (CERN, Tier-1, -2, -3)
- Increase in past 17 months - 5 X number of jobs
- 3.5 X cpu usage

- lhcb
- cms
- atlas
- alice





September 2007 - CPU Usage CERN, Tier-1s, Tier-2s



- > 80% of CPU Usage is external to CERN






CERN data export 2007



Daily Report

(VO-wise Data Transfer From CERN-CIC To All Sites)

Revert Source/Dest Site(s)


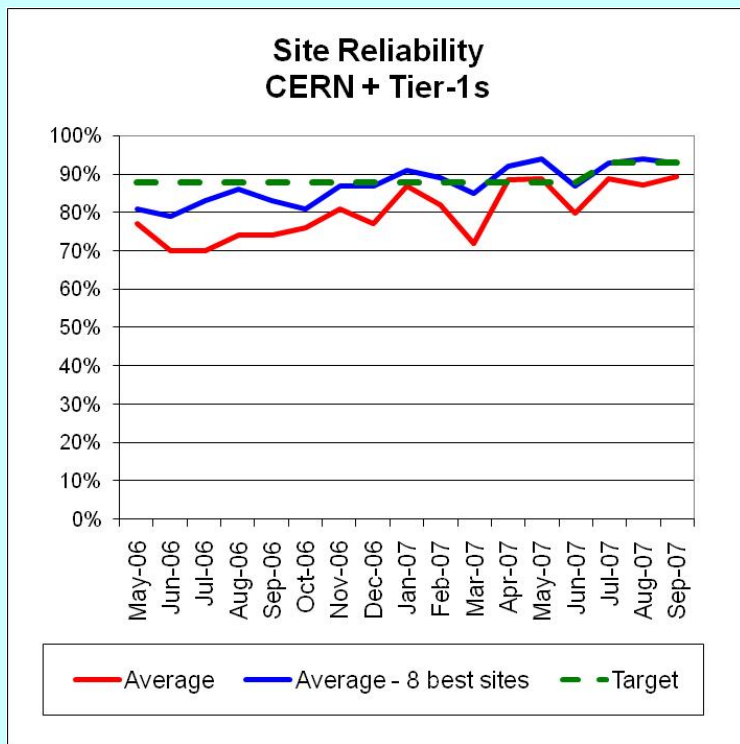


- Data distribution from CERN to Tier-1 sites
 - The target rate was achieved last year under test conditions
 - This year under more realistic experiment testing, reaching ~70% of the target peak rate





Site Reliability




Targets – CERN + Tier-1s

	Before July	July 07	Dec 07	Avg. last 3 months
Each site	88%	91%	93%	89%
8 best sites	88%	93%	95%	93%




Ramp-up Needed for Startup

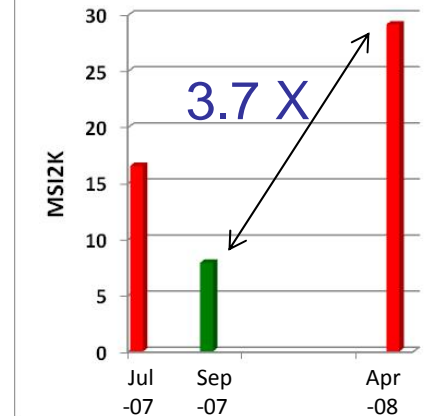
CERN - CPU Installed




Tier-1s - CPU Installed



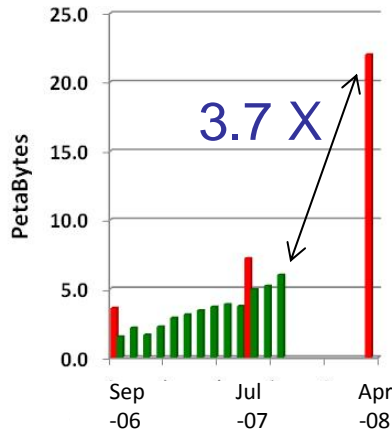
Tier-2s - CPU Usage Ramp-Up



CERN - Disk Installed



Tier-1s - Disk Installed



target usage
usage

pledge
installed



CERN openlab



CERN
openlab

www.cern.ch/openlab

PARTNERS



invent



ORACLE®

CONTRIBUTORS

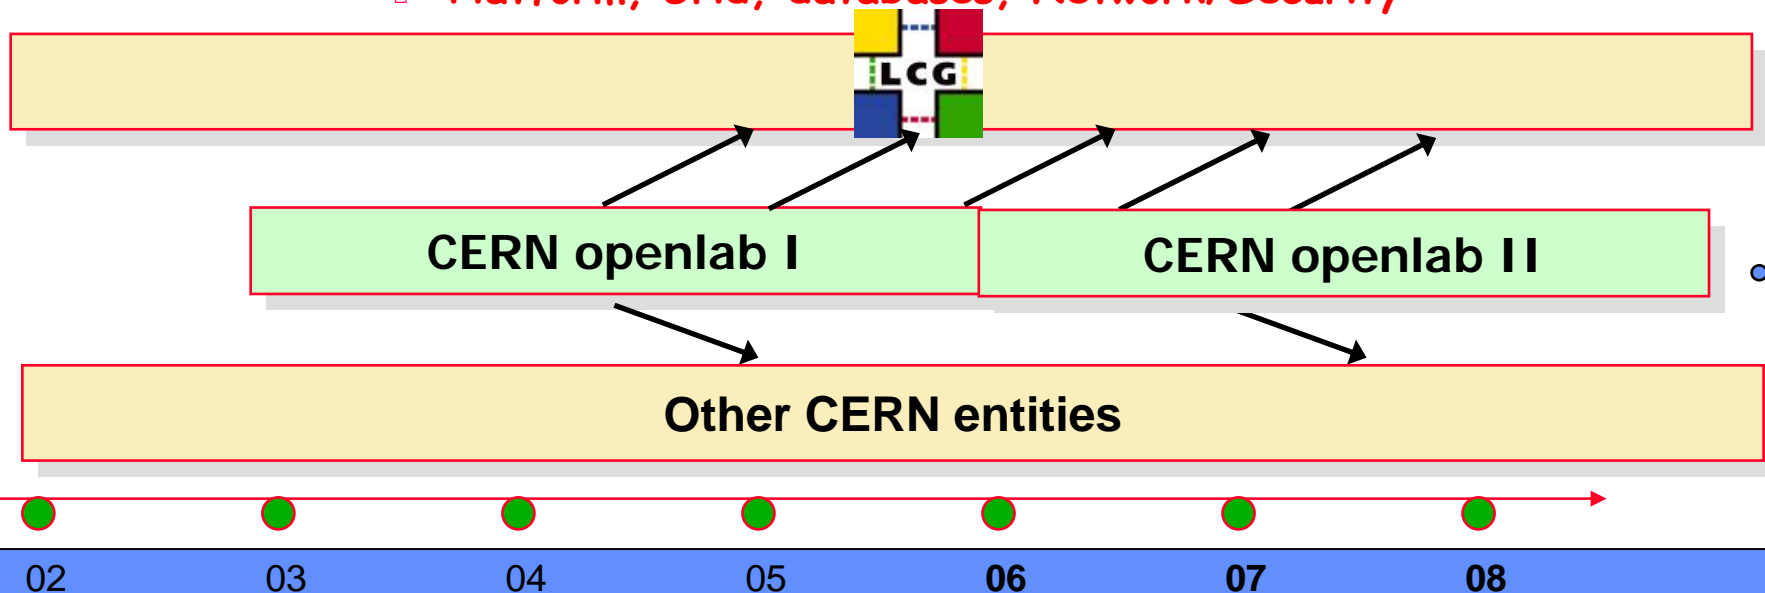


STONESOFT
Real World Business Security™
EDS



CERN openlab

- CERN-IT department's main R&D focus
- Framework for collaboration with industry
- Evaluation, integration, validation
 - of cutting-edge technologies that can serve the LHC Computing Grid (LCG)
- Sequence of 3-year agreements
 - 2003 - 2005: the "opencluster" project
 - 2006 - 2008: openlab Phase II with new projects:
 - Platform, Grid, databases, Network/Security





The Worldwide LHC Computing Grid

- The LHC physics data analysis service distributed across the world
 - CERN, 11 large *Tier-1* centres, ~ 140 active *Tier-2* centres
- Status in May 2007
 - Established the 10 Gigabit/sec optical network that interlinks CERN and the Tier-1 centres
 - Demonstrated data distribution from CERN to the Tier-1 centres at 1.3 GByte/sec - the rate that will be needed in 2008
 - ATLAS and CMS can each transfer 1 PetaByte of data per month between their computing centres
 - Running ~2 million jobs each month across the grid
 - The distributed grid operation, set up during 2005, has reached maturity, with responsibility shared across 7 sites in Europe, the US and Asia
 - End-user analysis tools enabling "real physicists" to profit from this worldwide data-intensive computing environment

