LHC Computing Grid Project – LCG

UK DTI Visit to CERN

29 June 2004

Les Robertson – LCG Project Leader
CERN – European Organisation for Nuclear Research
Geneva, Switzerland

les.robertson@cern.ch

last update: 29/06/2004

- Characteristics of HEP Computing
- Preparing the Computing Environment for LHC

The Atlas Detector

- The ATLAS collaboration is
  - ~2000 physicists from ..
  - ~ 150 universities and labs
  - from ~ 34 countries
  - distributed resources
  - remote development
- The ATLAS detector is
  - 26m long,
  - stands 20m high,
  - weighs 7000 tons
  - has 200 million read-out channels

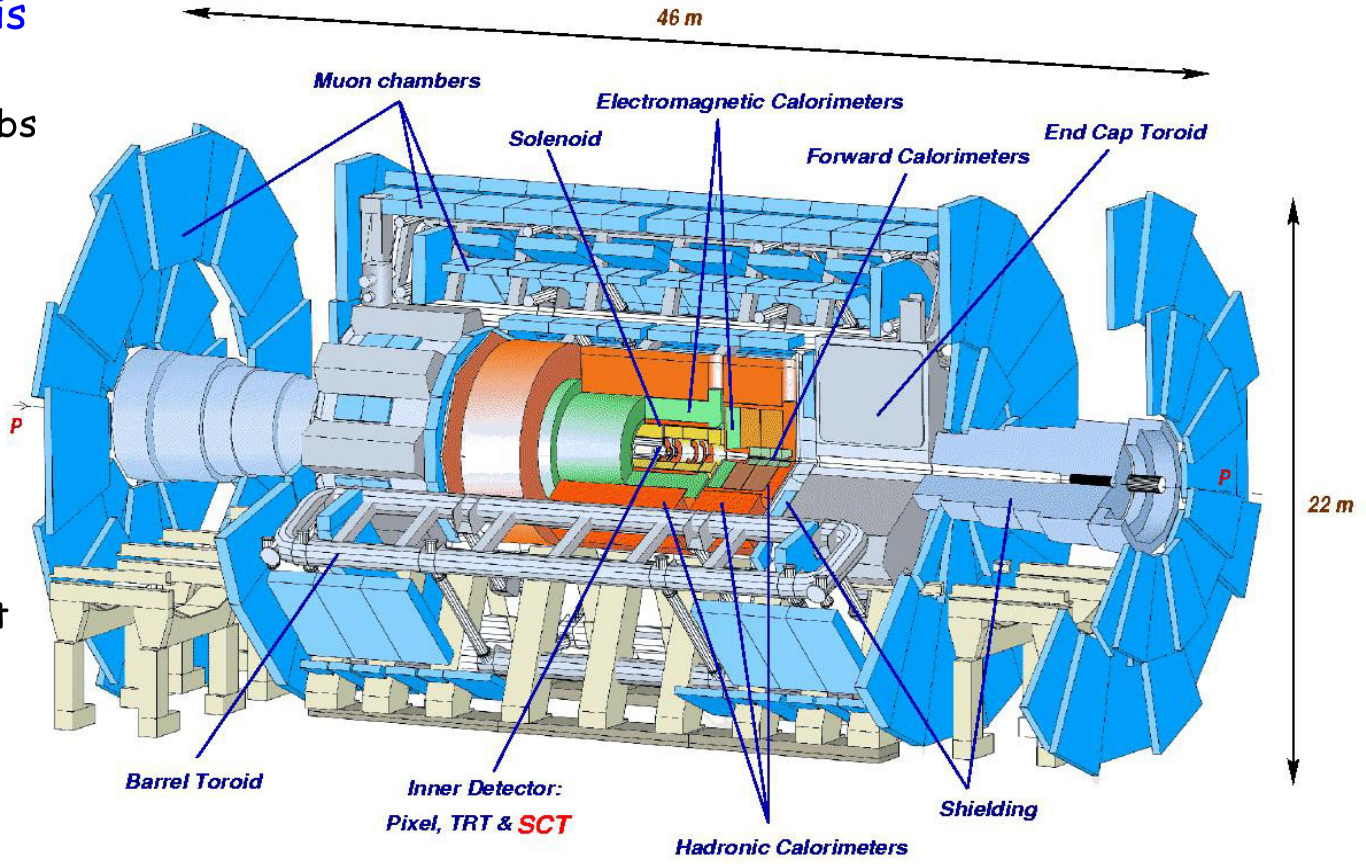

Data Handling and Computation for Physics Analysis

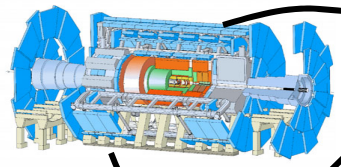

detector

event filter (selection & reconstruction)

reconstruction

raw data

event reprocessing

event summary data

processed data

batch physics analysis

analysis

analysis objects (extracted by physics topic)

interactive physics analysis

event simulation

simulation

Data Hierarchy

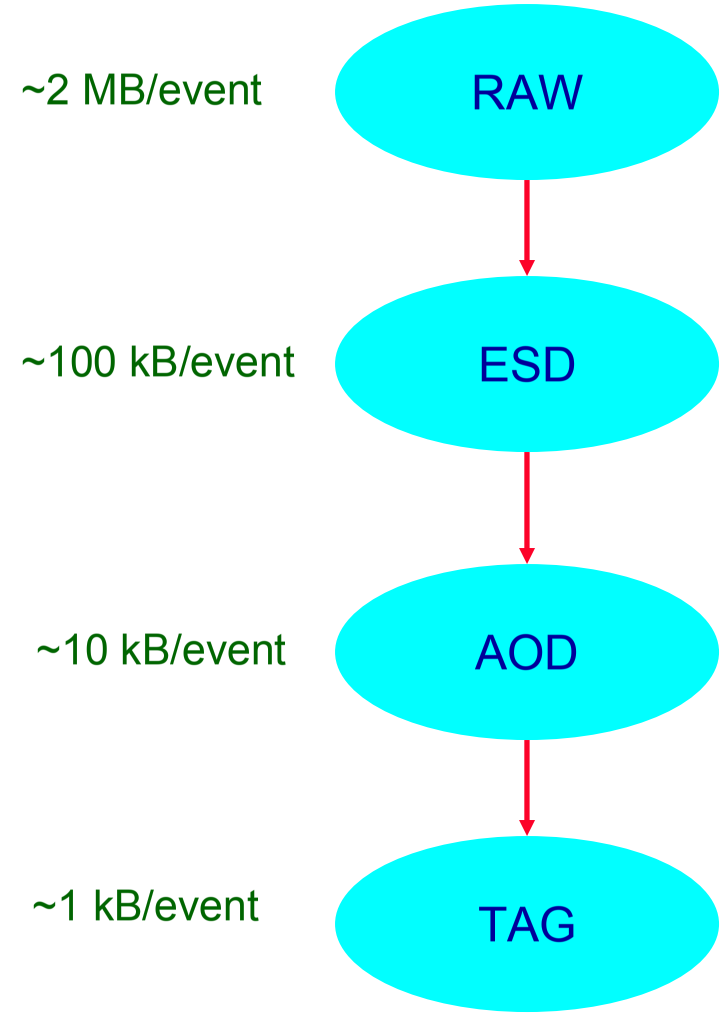

Event Data

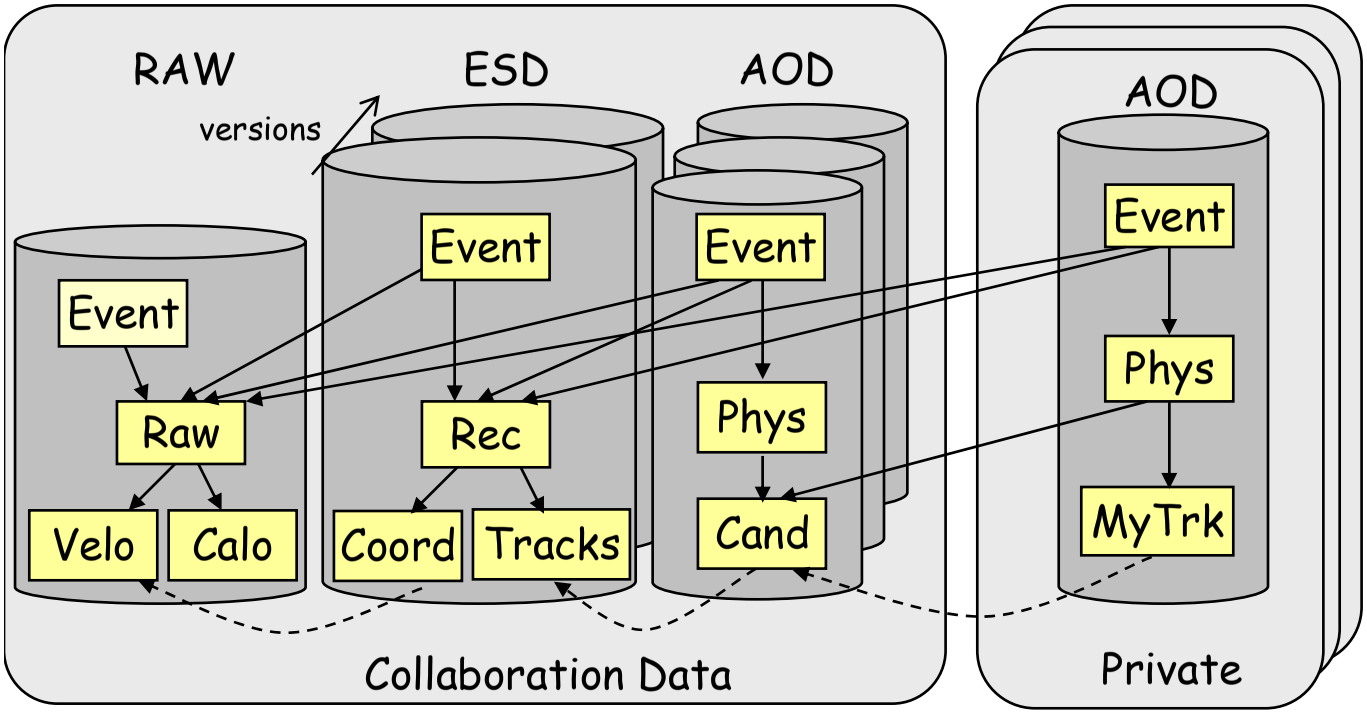

- Complex data models
  - ~500 structure types
- References to describe relationships between event objects
  - unidirectional
- Need to support transparent navigation
- Need ultimate resolution on selected events
  - need to run specialised algorithms
  - work interactively
- Not affordable if uncontrolled

Summary of Computing Characteristics

- small records – and most of the data is read-only
  - versions rather than updates
- meta-data (few %) in databases
- Modest I/O rates - few MB/sec per fast processor
- Modest floating point requirement - SPECint performance

Good fit for clusters of PCs

- very large aggregate requirements – computation, data, i/o
  - need to martial the resources of many computer centres
- chaotic workload –
  - batch & interactive
  - research environment
    → physics analysis by collective iterative discovery
    - unpredictable data access
    - no practical limit to the requirements

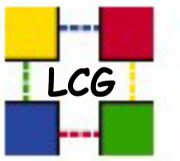

LHC Computing Grid Project

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

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

Software Engineers

Service Providers

Grid →The virtual LHC Computing Centre

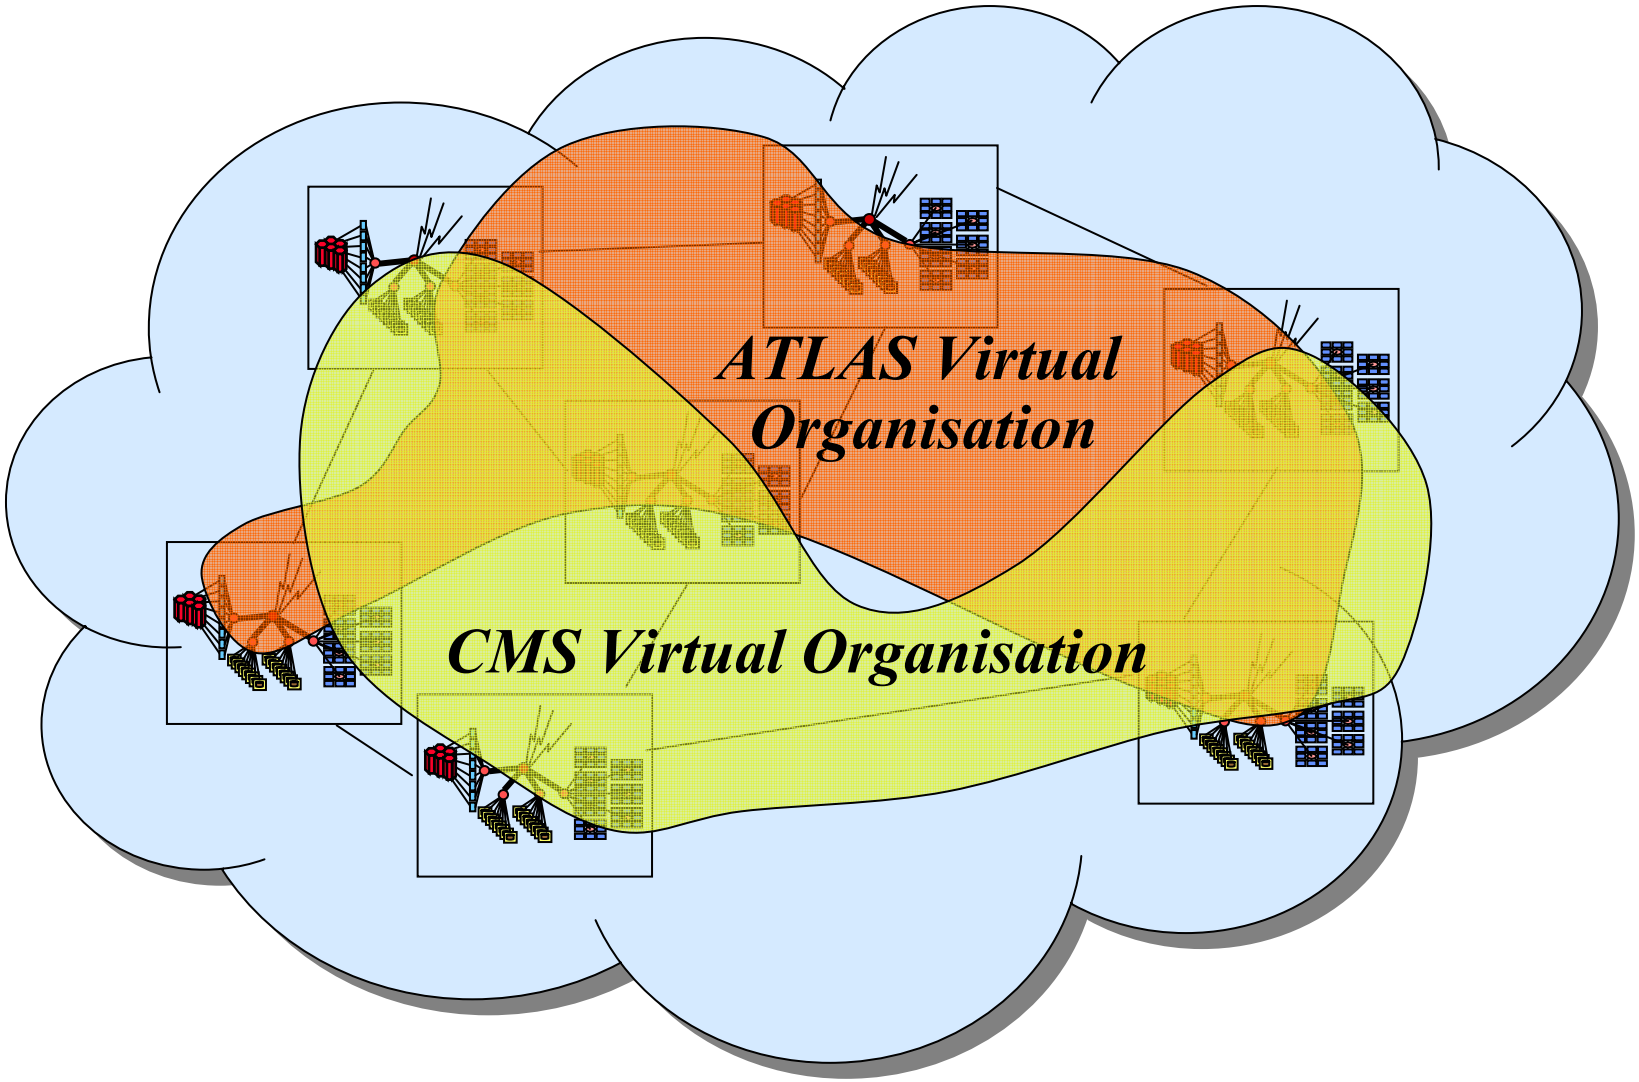

desktops portables

small centres

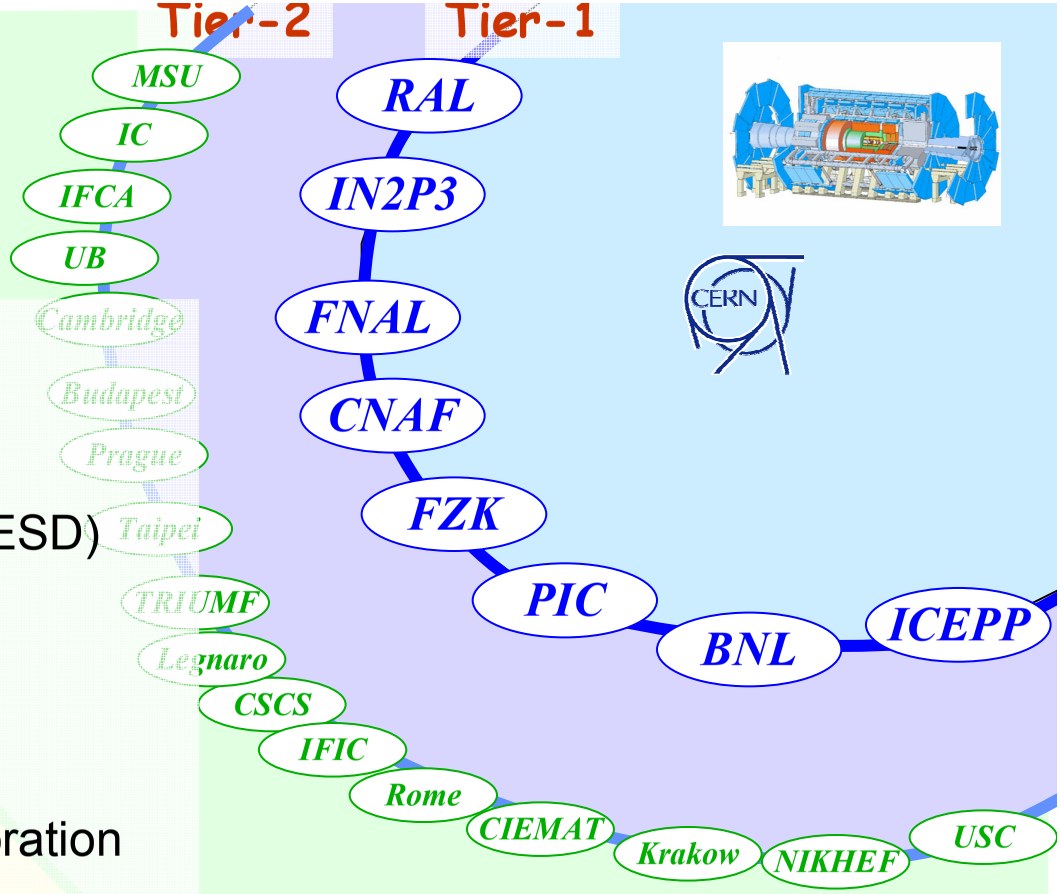

LHC Computing Model (simplified!!)

- Tier-0 – the accelerator centre
  - Filter → *raw data*
  - Reconstruction → *summary data* (ESD)
  - Record *raw data* and *ESD*
  - Distribute *raw* and *ESD* to Tier-1
- Tier-1 –
  - **Managed Mass Storage** – permanent storage *raw*, *ESD*, calibration data, meta-data, analysis data and databases
    → **grid-enabled data service**
  - Data-heavy analysis
  - Re-processing raw → ESD
  - National, regional support

"online" to the data acquisition process high availability, long-term commitment

- Tier-2 –
  - Well-managed disk storage – grid-enabled
  - Simulation
  - End-user analysis – batch and interactive
  - High performance parallel analysis (PROOF)

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)	45	20	18
Other large centres (Tier 2)	40	12	5
Totals	**105**	**37**	**43**

*** Current fast processor ~1K SI2000*

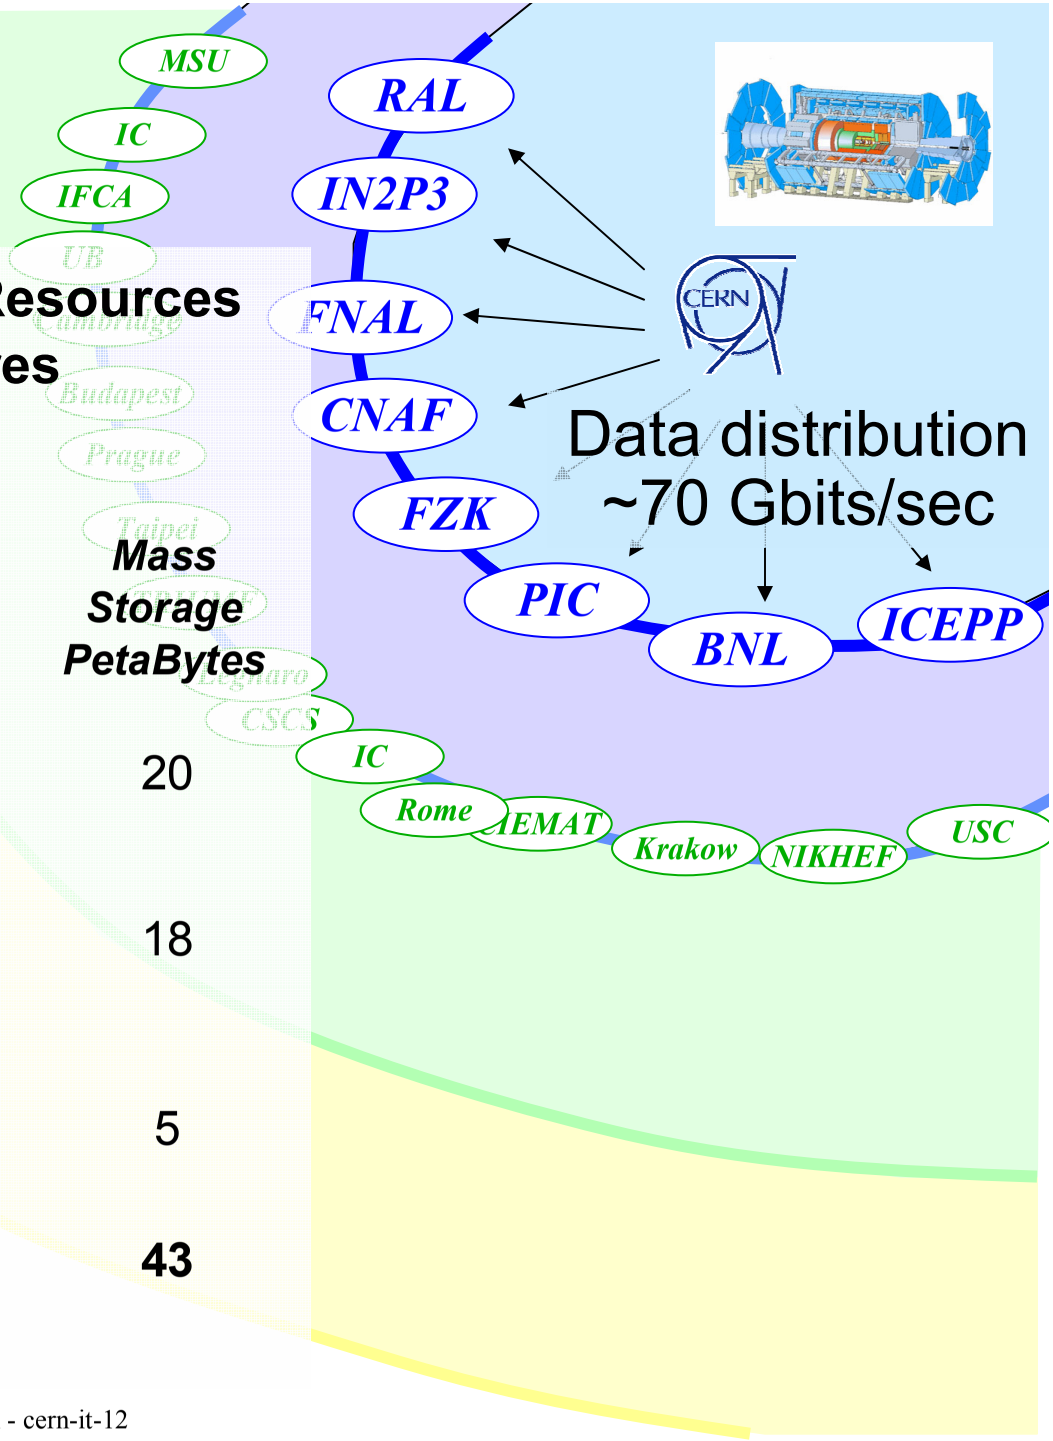

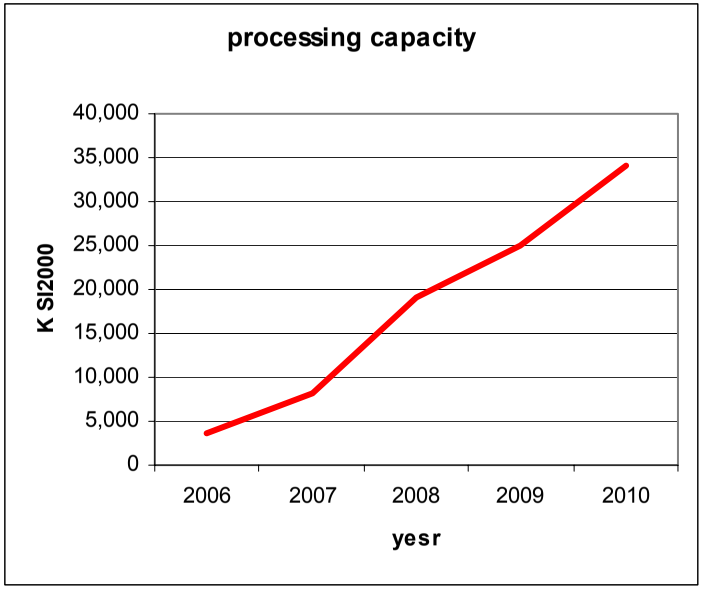

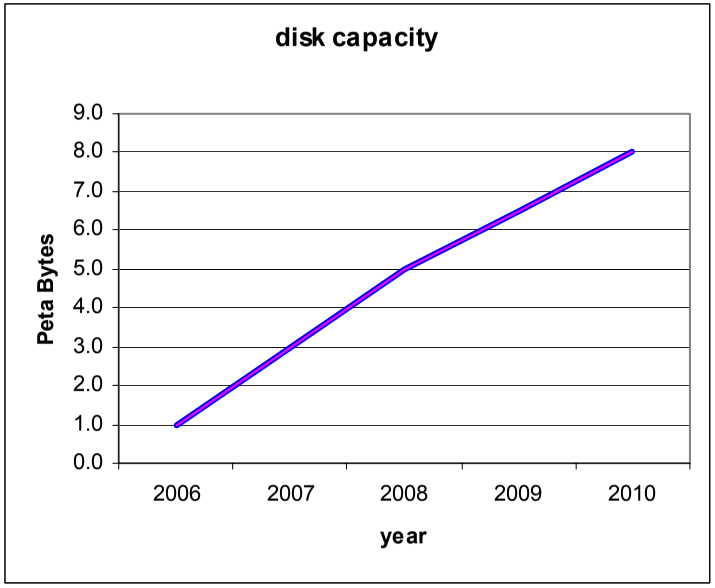

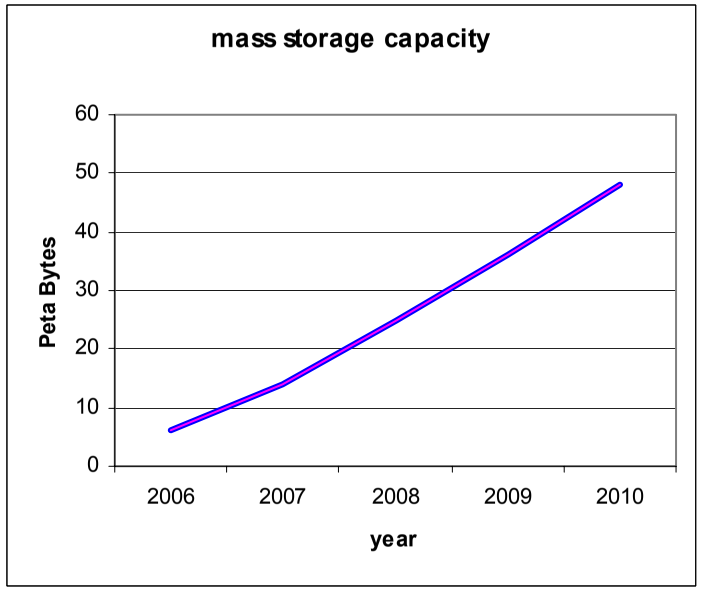

Planned evolution of capacity at CERN for LHC experiments

Estimated Resources Required by LHC Experiments in 2008

First full year of data - 2008

Experiment / Resource	Alice	Atlas	CMS	LHCb	Sum
CERN Tier 0 + Tier 1					
Disk PetaBytes	0.5	2.0	1.8	0.3	**5**
Mass Storage PetaBytes	2.3	7.6	9.2	1.0	**20**
*Processing M SI2000***	5.6	5.4	5.7	2.7	**19**
Sum of resources at all Tier1 centres					
Expected number of centres	3	6	6	5	
Disk PetaBytes	3.0	6.8	8.7	1.3	**20**
Mass Storage PetaBytes	3.6	7.2	6.6	0.4	**18**
*Processing M SI2000***	9.1	13.6	12.6	9.5	**45**
Sum of resources at all Tier2 centres					
Expected number of centres	16	24	25	15	
Disk PetaBytes	3.0	3.8	5.0	0.6	**12**
Mass Storage PetaBytes	0.0	1.6	2.9	0.0	**5**
*Processing M SI2000***	7.2	8.4	7.5	16.4	**40**

*** Current fast processor ~1K SI2000*

Estimates prepared as input to the MoU Task Force
Computing models under active development

Estimated Network Bandwidth in 2008 between Tier-1s and Tier-0

Gbits/sec

Tier-1 Centre	*RAL*	*Fermilab*	*Brookhaven*	*Karlsruhe*	*IN2P3*	*CNAF*	*PIC (Barcelona)*	*T0 Total*
Estimated T1 Bandwidth Needed **	6.9	1.7	4.2	6.9	6.9	6.9	4.4	38
Assumed Bandwidth Provisioned	10	10	10	10	10	10	10	70

** *Does not include traffic between Tier-1s and with Tier-2s*

Assumes that all Tier-1s share the load equally

LHC Grid Service

http://goc.grid-support.ac.uk/gppmonWorld/gppmon_maps/lcg2.html

> 60 sites
> 4K cpus

In process of being connected:

IHEP, Beijing
TIFR, Mumbai
NCP, Islamabad

H-P Puerto Rico

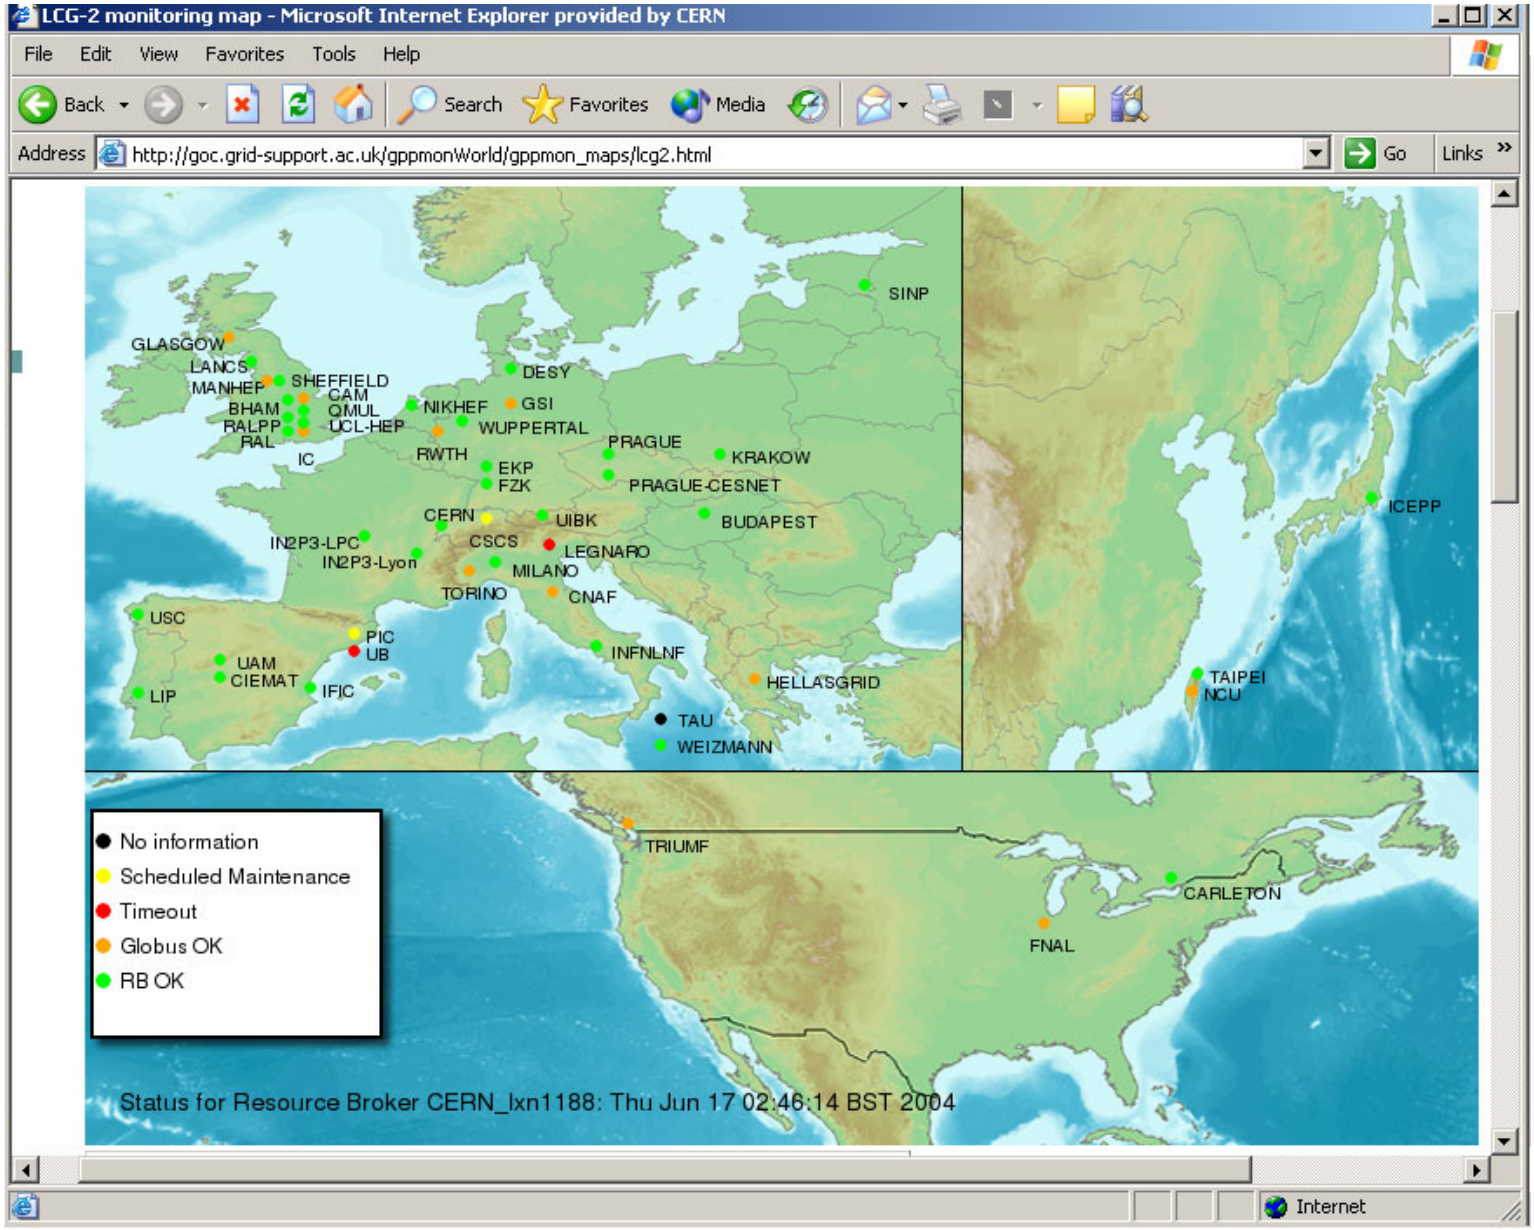

last update 29/06/2

Sites in LCG-2 – 4 June 2004

Austria	U-Innsbruck
Canada	Triumf
	Alberta
	Carleton
	Montreal
	Toronto
Czech Republic	Prague-FZU
	Prague-CESNET
France	CC-IN2P3
	Clermont-Ferrand
Germany	FZK
	Aachen
	DESY
	Wuppertal
Greece	HellasGrid
Hungary	Budapest (KFK)
Israel	Tel-Aviv
	Weizmann

Italy	CNAF
	Frascati
	Legnaro
	Milano
	Napoli
	Roma
	Torino
Japan	Tokyo University (ICEPP)
Netherlands	NIKHEF
Pakistan	NCP
Poland	CYFRONET (Krakow)
Portugal	LIP
Russia	SINP-Moscow
	JINR-Dubna
Spain	PIC
	Univ.Autonoma de Madrid
	Univ.Santiago de Compostela
	Univ. Barcelona
	IFCA Santander
	CIEMAT Madrid
	IFIC Valencia

Switzerland	CERN
	CSCS
Taiwan	ASCC
	NCU
UK	RAL
	RAL-Physics
	Universities:
	Birmingham
	Cambridge
	Glasgow
	Imperial
	Lancaster
	Manchester
	Queen Mary
	Sheffield
	UCL
US	BNL
	FNAL

In preparation/certification –

- **India – Tata Institute, Mumbai**
- **China – IHEP, Beijing**
- **New Zealand**
- **HP – Puerto Rico (and other HP sites)**

The LCG-2 service for the 2004 Data Challenges

LCG-2 target
– the 2004 "*LHC Data Challenges*"

- Large-scale tests of the experiments' computing models, processing chains, grid technology readiness, operating infrastructure
- ALICE and CMS data challenges started at the beginning of March
- LHCb and ATLAS – started at the beginning of May
- The big challenge for this year –
  - data –
  - file catalogue,
  - replica management,
  - database access,
  - integrating mass storage

Grid Operations Centre at RAL

User Support Centre at FZK

Planning for a second operations & support centre in Taipei

The Challenges of Grid C

Keeping the expectations in line with reality!

- Middleware maturity –
  - Service quality – standards ←→ innovation
  - a second (or third) generation of middleware will be needed before LHC starts
- Grid operation –
  - There is not yet any significant experience to build on
- Data migration
  - Replication, caching, security, policies
- Collaboration
  - From independent computer centres to a federation
  - Consensus rather than control
- International research networking
- Security – our biggest risk
- Standards
  - In the short-term there will many grids and middleware implementations

LCG-2 and Next Generation Middleware

LCG-2

focus on production, large-scale data handling

- The service for the 2004 data challenges
- Provides experience on operating and managing a global grid service
- Strong development programme driven by data challenge experience

- Evolves to LCG-3 as components progressively replaced with new middleware

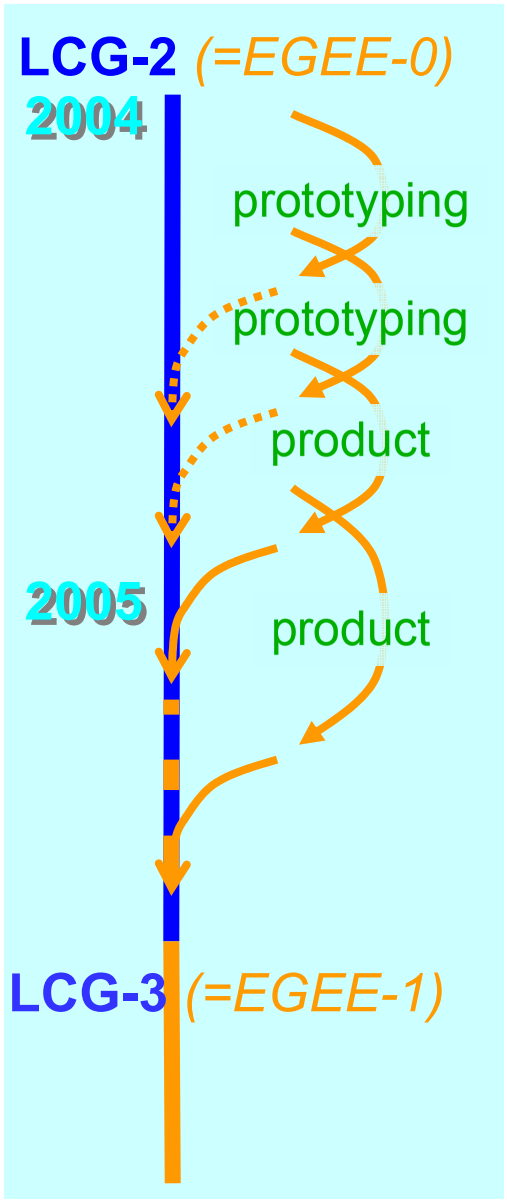

Next generation middleware

focus on analysis

- Developed by EGEE project in collaboration with VDT (US)
- LHC applications and users closely involved in prototyping & development (ARDA project)
- Short development cycles
- Completed components integrated in LCG-2

Enabling Grids for E-science in Europe

Preparing for 2007

- 2003 – has demonstrated event production
- In 2004 we must show that we can also handle the data – even if the computing model is very simple

  -- This is a key goal of the 2004 Data Challenges
- Target for end of this year –
  - Basic model demonstrated using current grid middleware
  - All Tier-1s and ~25% of Tier-2s operating a reliable service
  - Validate security model, understand storage model
  - Clear idea of the performance, scaling, operations and management issues

LCG Timeline

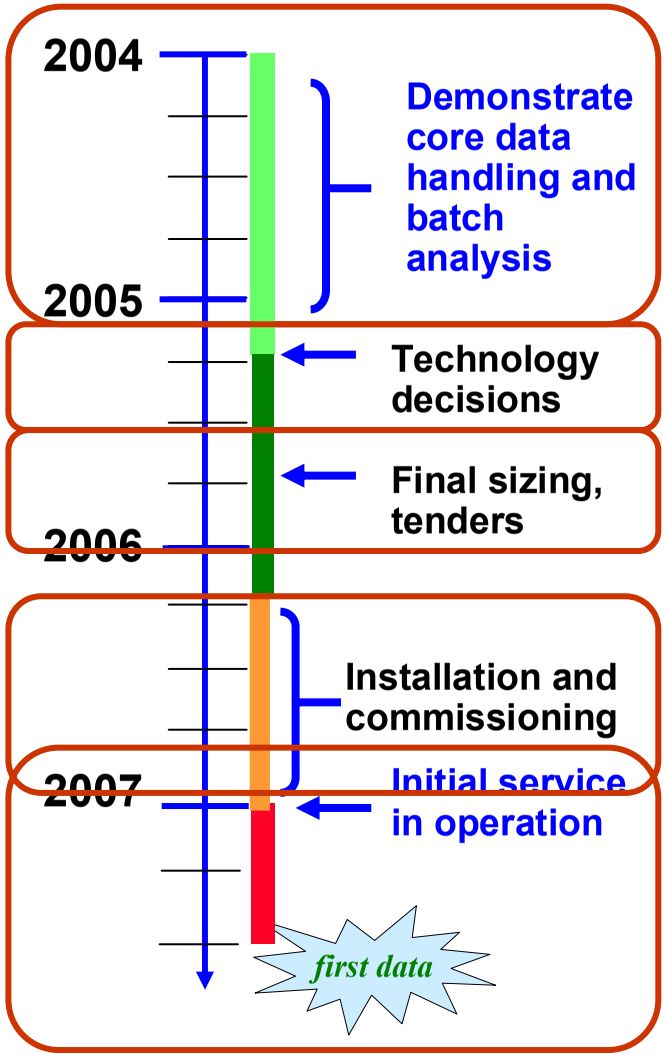

Summary

- LHC computing –
  - Geographically distributed
  - Independent regional centres
  - **Data intensive**
- LCG-2 –
  - Reliable environment for batch work
  - Already more than 60 sites, global reach
  - Just beginning to get experience with data handling .. and grid operation and management
  - An early example of a working grid
- Tight timescale to ramp up for LHC