
LHC Computing
R.J.Cashmore
Director of Research,CERN

LHC Experiments

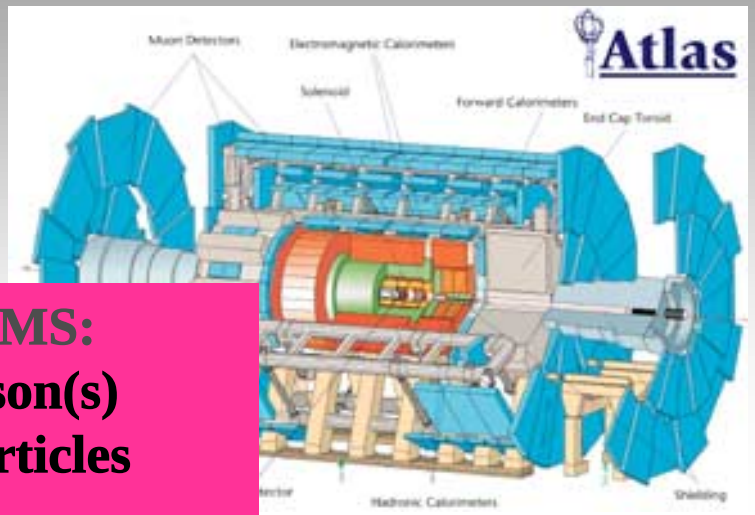

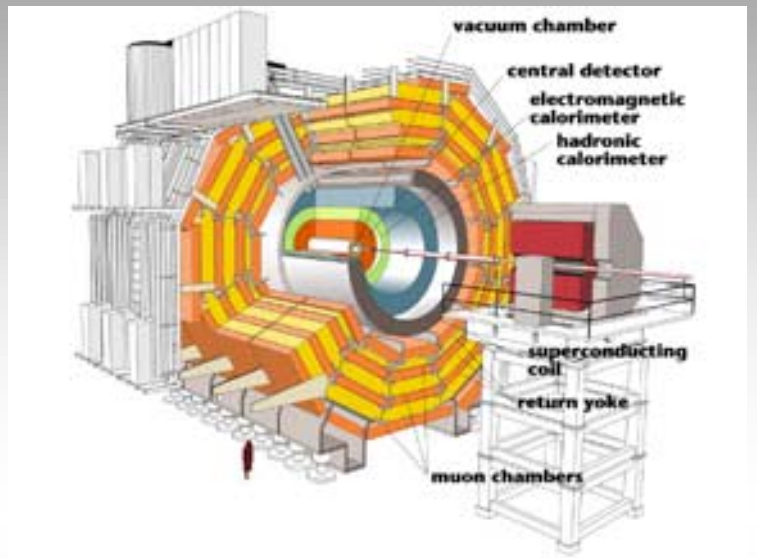

ATLAS, CMS:

- Higgs boson(s)

- SUSY particles

- …??

ALICE:

Quark Gluon Plasma

LHC-B:

- CP violation in B

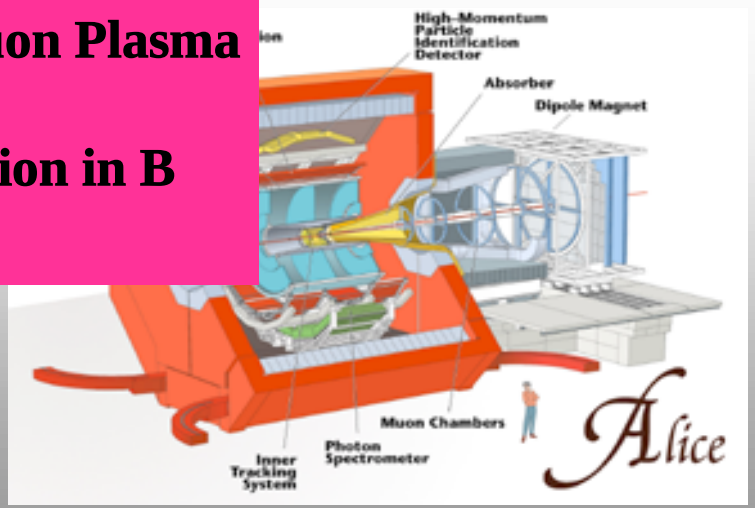

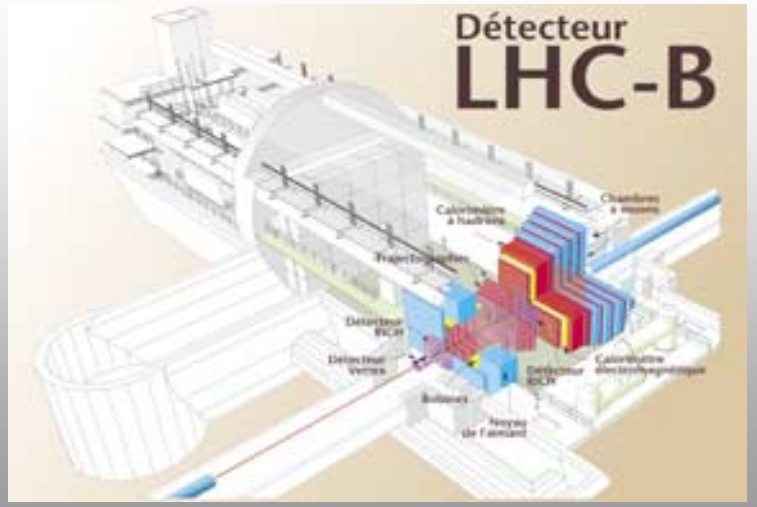

On-line System

- Large variety of triggers and thresholds: select physics à la carte
- Multi-level trigger
- Filter out less interesting
- **Online reduction 10^7**
- **Keep highly selected events**

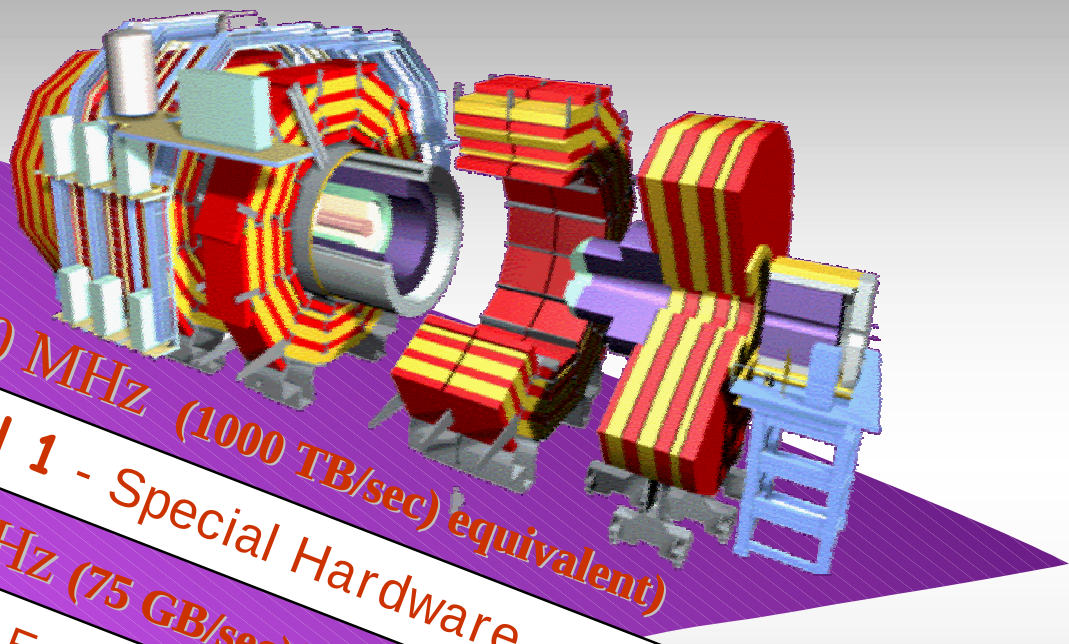

40 MHz (1000 TB/sec) equivalent)

Level 1 - Special Hardware

75 KHz (75 GB/sec)fully digitised

Level 2 - Embedded Processors/Farm

5 KHz (5 GB/sec)

Level 3 – Farm of commodity CPU

100 Hz (100 MB/sec)

Data Recording & Offline Analysis

Digital telephone 1-2 KB/sec

The Large Hadron Collider Project
4 detectors

ATLAS

CMS

Storage –

Raw recording rate 0.1 – 1 GByte/sec

Accumulating at 5-8 PetaBytes/year

10 PetaBytes of disk

Processing –

200,000 of today's fastest PCs

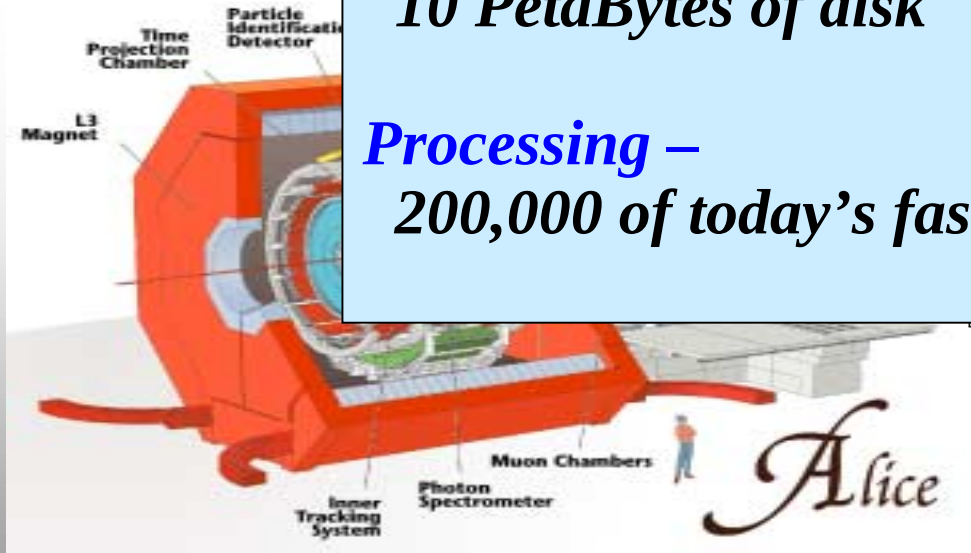

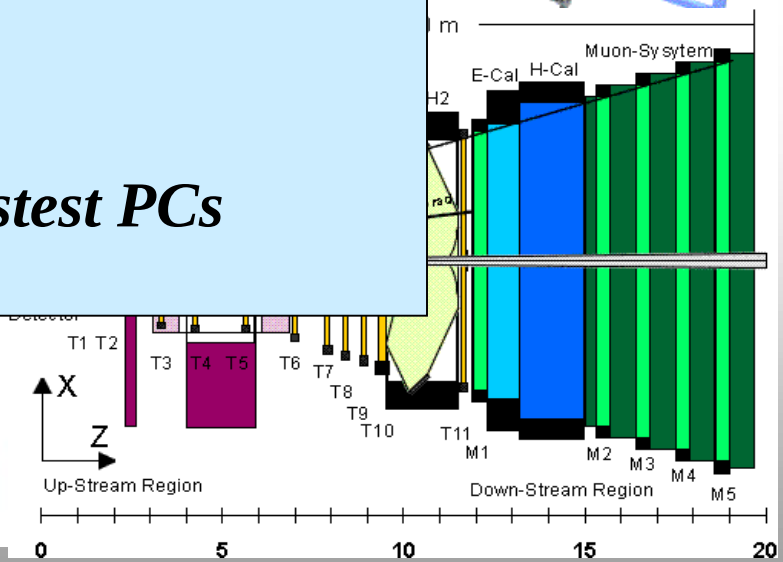

LHC Computing Review 2000 / 1

Parameter	Unit	ALICE p-p	ALICE Pb-Pb	ATLAS	CMS	LHCb	TOTAL	ATLAS (**)
# assumed Tier1 not at CERN		4		6	5	5		6
# assumed Tier2 not at CERN***					25			
Event recording rate	Hz	100	50	100	100	200		270
RAW Event size	MB	1	25	1	1	0.125		2
REC/ESD Event size	MB	0.1	2.5	0.5	0.5	0.1		0.5
AOD Event size	kB	10	250	10	10	20		10
TAG Event size	kB	1	10	0.1	1	1		0.1
Running time per year	M seconds	10	1	10	10	10		10
Events/year	Giga	1	0.05	1	1	2		2.7
Storage for real data	PB	1.2	1.5	2.0	1.7	0.45	6.9	8.1
RAW SIM Event size	MB	0.5	600	2	2	0.2		2
REC/ESD SIM Event size	MB	0.1	5	0.5	0.4	0.1		0.5
Events SIM/year	Giga	0.1	0.0001	0.12	0.5	1.2		0.12
Number of reconst. passes	Nb	2		2-3	2	2-3		2-3
Storage for simul. data	PB	0.1	0.1	1.5	1.2	0.36	3.2	1.5
Storage for calibration	PB	0.0	0.0	0.4	0.01	0.01	0.4	0.4
Tape storage at CERN T0+T1	PB (10**15 B)	3.23		2.86	4.17	1.22	11.5	9.00
Tape storage at each Tier1 (Avg.)		}0.37		}1.26	1.02	}0.32	}3.0	}1.80
Tape storage at each Tier2 (Avg.)***					0.05			
Total tape storage / year		**4.7**		**10.4**	**10.5**	**2.8**	**28.5**	**19.8**
Disk storage at CERN T0+T1	PB	0.53		0.31	1.14	0.33	2.3	0.41
Disk storage at each Tier1 (Avg.)		}0.27		}0.26	0.44	}0.15	}1.1	}0.36
Disk storage at each Tier2 (Avg.)***					0.10			
Total disk storage		**1.6**		**1.9**	**5.9**	**1.1**	**10.4**	**2.57**
Time to reconstruct 1 event	k SI-95 sec	0.4	100	0.64	3	0.25		0.64
Time to simulate 1 event	k SI-95 sec	3	2250	3	5	1.5		3
CPU for 1 rec. pass/y (real data)	k SI-95	20	250	200	434	50		385
CPU for 1 SIM pass/y (sim+rec)	k SI-95	19	269	30	200	660		30
CPU reconstruction, calib.	k SI-95	65	525	251	1040	50	1931	435
CPU simulation		19	269	30	587	660	1564	30
CPU analysis		880		1479	1280	215	3854	1479
Total CPU at CERN T0+T1	k SI-95	824		506	820	225	2375	690
Total CPU each Tier1 (Avg.)		}234		}209	204	}140	}787	}209
Total CPU each Tier2 (Avg.)***					43			
Total CPU		**1758**		**1760**	**2907**	**925**	**7349**	**1944**
WAN, Bandwidths								
Tier0 - Tier1 link, 1 expt.	Mbps	1500		1500	1500	310	4810	1500
Tier1 - Tier2 link		622		622	622			622

Further results of the Computing Review

- The Review recommends the multi-tier hierarchical model as one key element of the LHC computing model with the majority of the resources not based at CERN.

- The Review sees about equal efforts between Tier O at CERN, Tier 1's and lower level Tiers down to the desktops in the participating countries:

Tier0/Σ (Tier 1)/Σ (all Tier 2, etc) =1/1/1

- The Review recommends that all experiments perform "Data Challenges" of increasing size and complexity until LHC start-up:
  - to build up and test the distributed computing infrastructure
  - (to use the emerging structure for running experiments)

Summary of Computing Capacity Required for all LHC Experiments in 2007

source: CERN/LHCC/2001-004 - Report of the LHC Computing Review - 20 February 2001

(ATLAS with 270Hz trigger)

	CERN			Regional	Grand
	Tier 0	Tier 1	Total	Centres	Total
Processing (K SI95)	1,727	832	2,559	4,974	7,533
Disk (PB)	1.2	1.2	2.4	8.7	11.1
Magnetic tape (PB)	16.3	1.2	17.6	20.3	37.9

- **Worldwide distributed computing system**
- **Small fraction c**
- **ESD analysis – regional centres**
  - how to use the res
  - establishing and maintaining a uniform physics
  - development & management costs

Importance of cost containment

- components & architecture
- utilisation efficiency
- maintenance, capacity evolution

Evolution of Computing Capacity at CERN

LHC

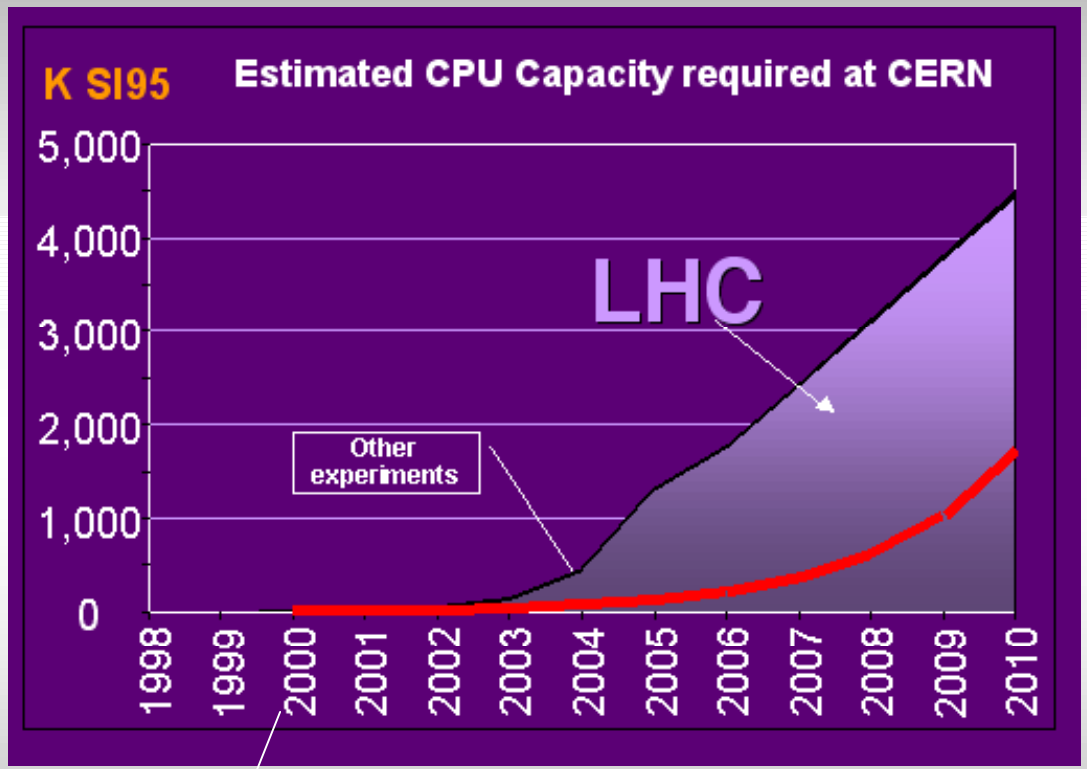

Jan 2000:
3.5K SI95

Moore's law –
some measure of the capacity technology advances provide for a constant number of processors or investment (based on 2000)

Summary of Additional Resources needed

Summary - Additional Resources needed

	R&D Phase (Phase 1)				First Production System (Phase 2)			Maint-enance	Total R&D 2001-04 (Phase 1)	Total First System 2005-07 (Phase 2)
year	*2001*	*2002*	*2003*	*2004*	*2005*	*2006*	*2007*	*2008*		
Services required at CERN										
Additional personnel (person-years)	**16**	**41**	**42**	**50**	**50**	**50**	**46**	**21**		
Cost if employed as CERN staff (MCHF)	2.4	6.2	6.3	7.5	7.5	7.5	6.9	3.2	*22.4*	*21.9*
Additional materials (MCHF)	2.1	6.6	10.1	10.7	30	33.4	32.4	22.6	*29.5*	*95.8*
Service funding required at CERN (MCHF)	**4.5**	**12.8**	**16.4**	**18.2**	**37.5**	**40.9**	**39.3**	**25.8**	*51.9*	*117.7*
and in addition										
Interface of experiments' Core Software to common Infrastructure										
Additional s/w professionals (person-years)	**6**	**6**	**6**	**6**	**6**	**6**				

Cost Estimates for the CERN Tier 0 + Tier 1 Centre for LHC

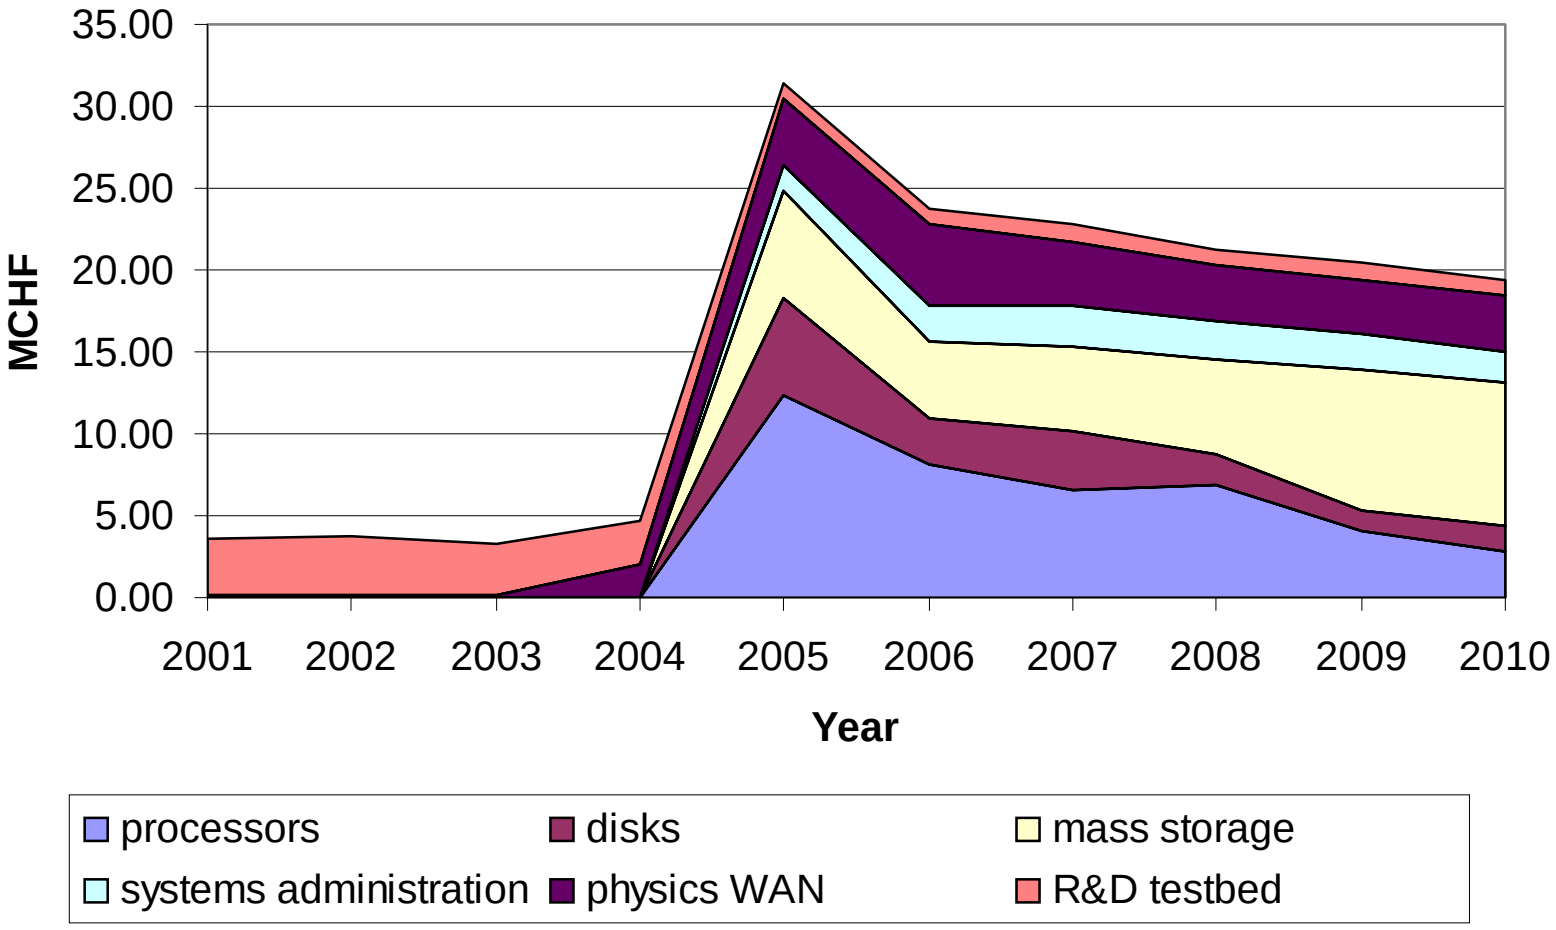

Core Software teams of the Experiments(RRB-matter)

- **Core software teams need additional resources**
- **Collaborating institutes and their funding agencies to provide for these persons;**
  - **software agreements are being discussed**

Required human resources to write the Core Software

		----- resources required (person-years) -----				
year		*2001*	*2002*	*2003*	*2004*	*2005*
collaboration						
ALICE		17.5	16.5	17.0	17.5	16.5
ATLAS		36.0	35.0	30.0	28.0	29.0
CMS		27.0	31.0	33.0	33.0	33.0
LHCb		25.0	24.0	23.0	22.0	21.0
Totals	64.0	**105.5**	**106.5**	**103**	**100.5**	**99.5**
Estimated Shortfall		**41.5**	**42.5**	**39.0**	**36.5**	**35.5**

- **Some of the work needs to be done at CERN**

CERN's Users in the World

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

Five Emerging Models of Networked Computing From *The Grid*

- **Distributed Computing**
  - || synchronous processing
- **High-Throughput Computing**
  - || asynchronous processing
- **On-Demand Computing**
  - || dynamic resources
- **Data-Intensive Computing**
  - || databases
- **Collaborative Computing**
  - || scientists

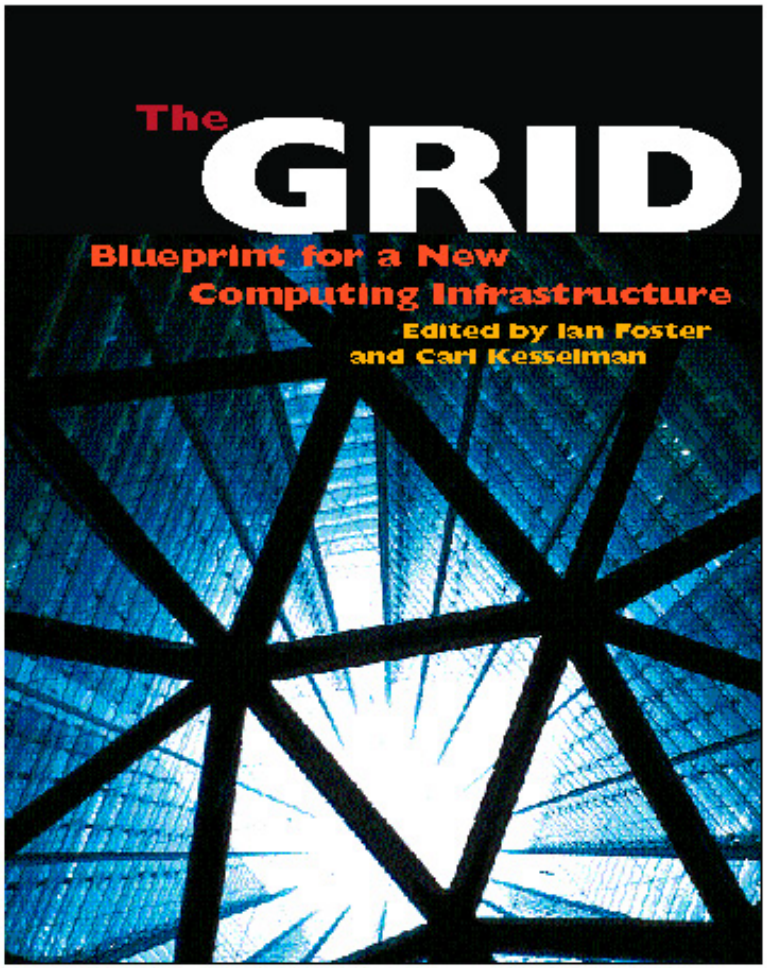

The MONARC Multi-Tier Model (1999)

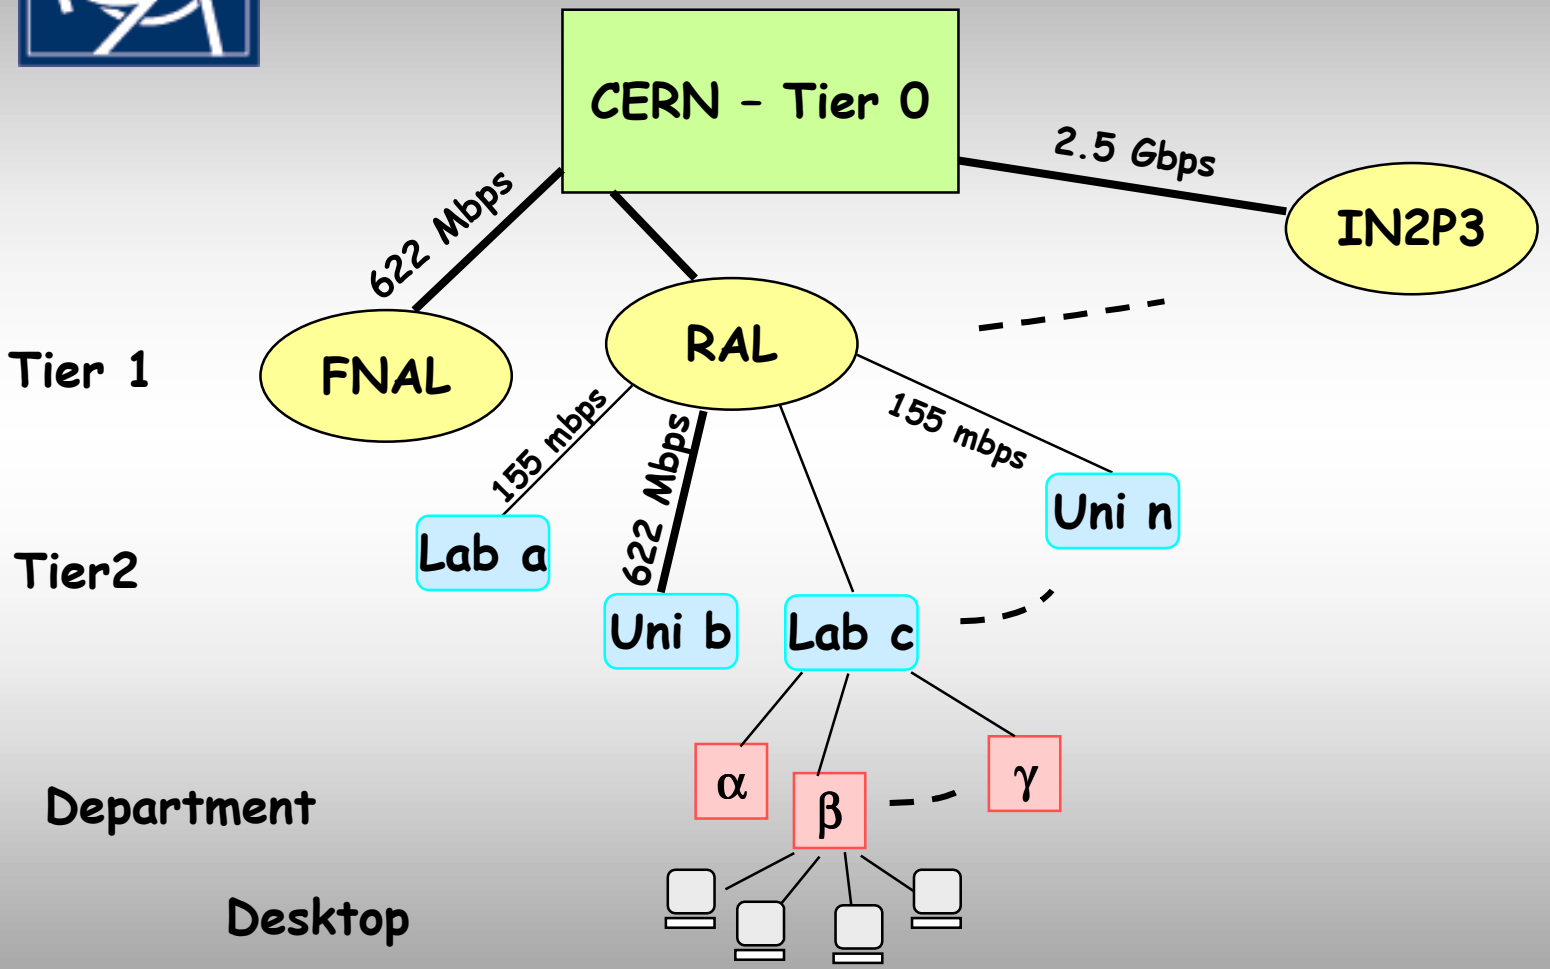

MONARC report: http://home.cern.ch/~barone/monarc/RCArchitecture.html

les.robertson@cern.ch

The opportunity of Grid technology

LHC Computing Model

2002 - evolving

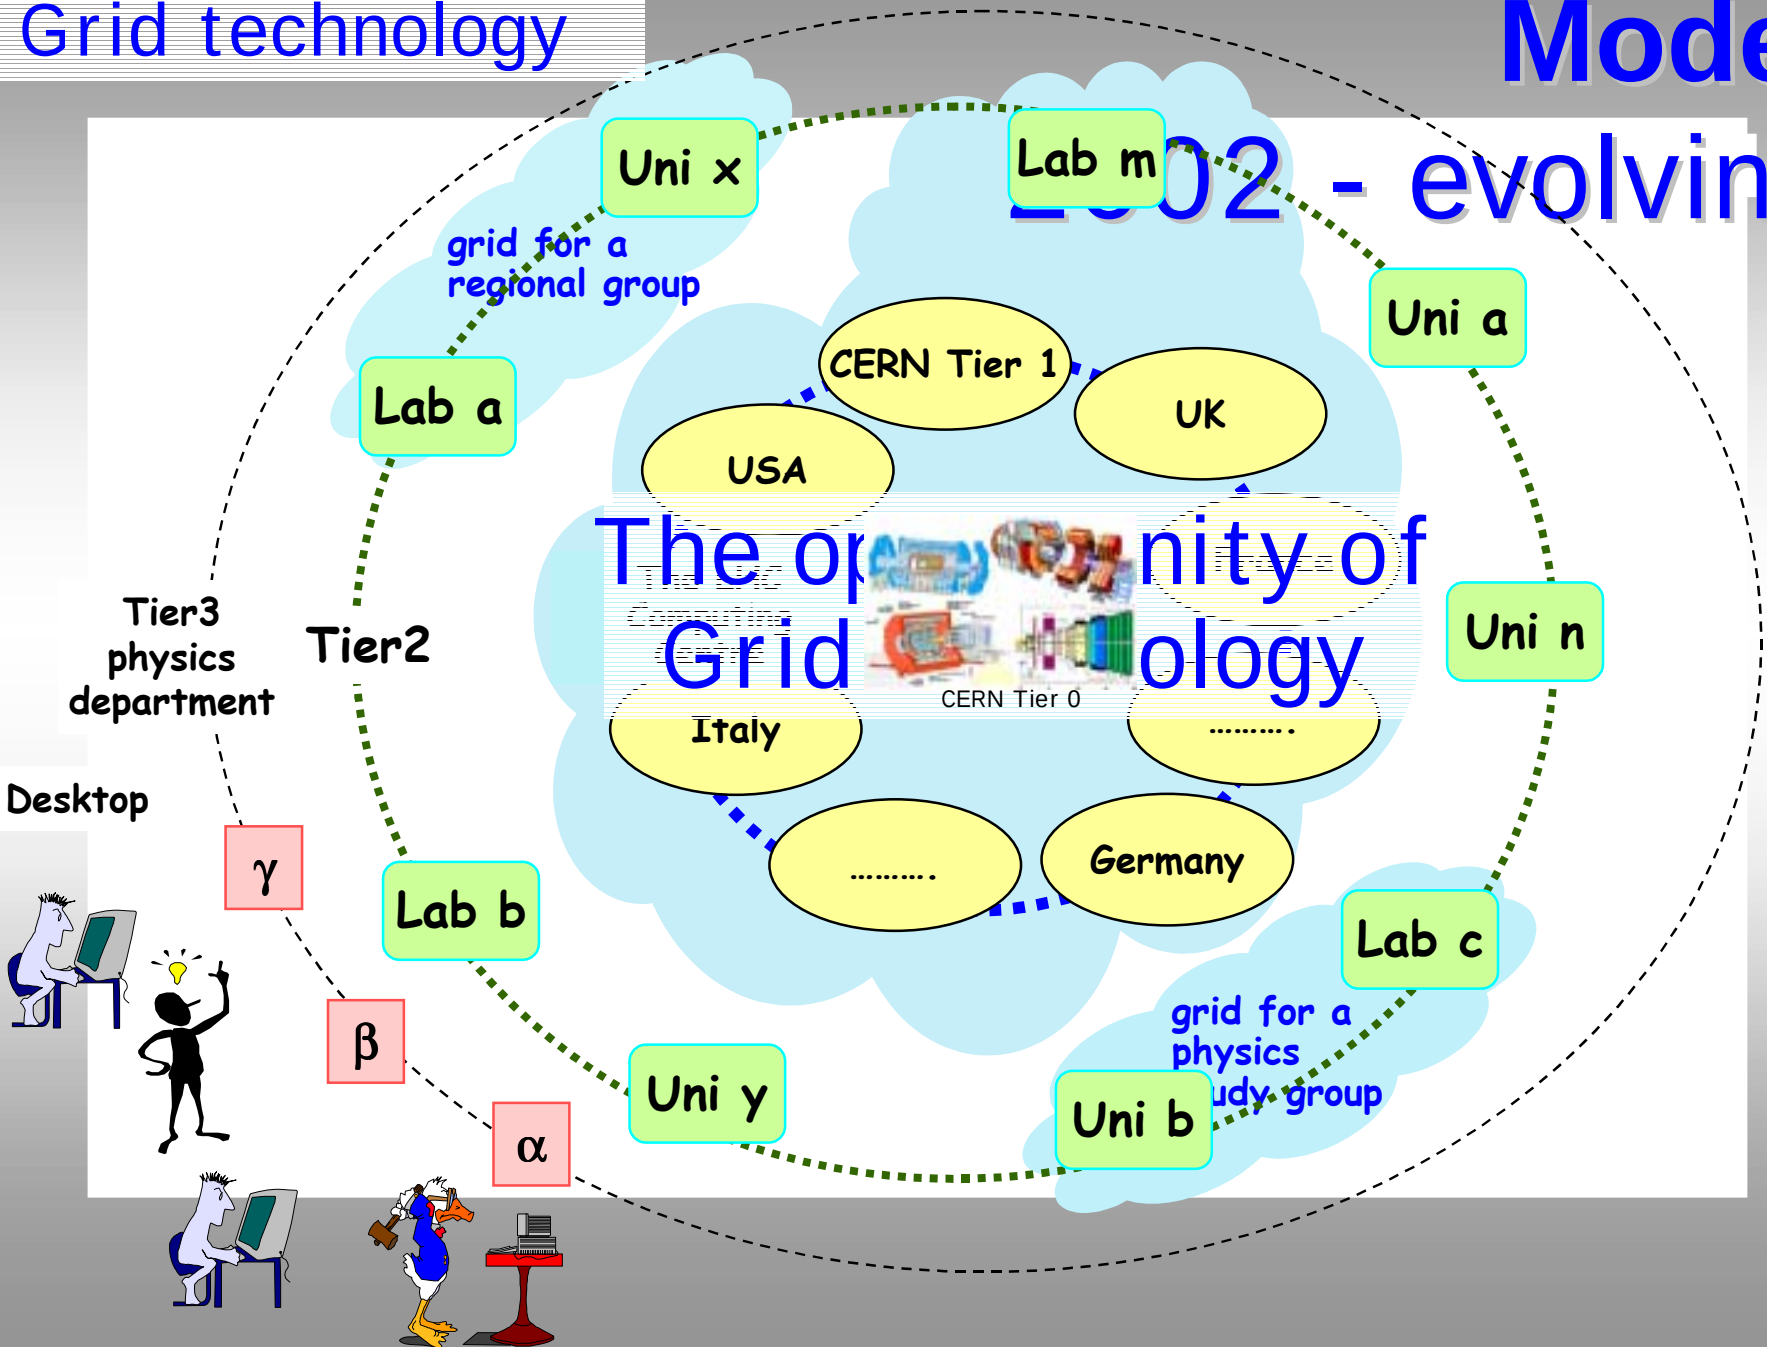

les.robertson@cern.ch

The DataGRID Projec

PPARC

www.eu-datagrid.or

Trans-Atlantic Testbeds

DataTAG – EU funding to enable Datagrid & US projects to build a common testbed

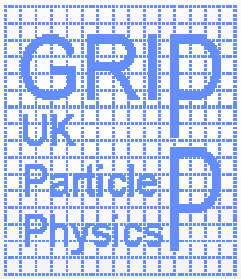

The Data Grid Project -Summary

- European dimension
  - EC funding 3 years, ~10M Euro
  - Closely coupled to several national initiatives
- Multi-science
- Technology leverage –
  - Globus, Condor, HEP farming & MSS, Monarc, INFN-Grid, Géant
- Emphasis –
  - Data – Scaling - Reliability
  - Rapid deployment of working prototypes - production quality
  - Collaboration with other European and US projects
- Status –
  - Started 1 January 2001
  - Testbed 1 in operation now
- Open –
  - Open-source and communication
  - Global GRID Forum
  - Industry and Research Forum

LHCGrid Organisation

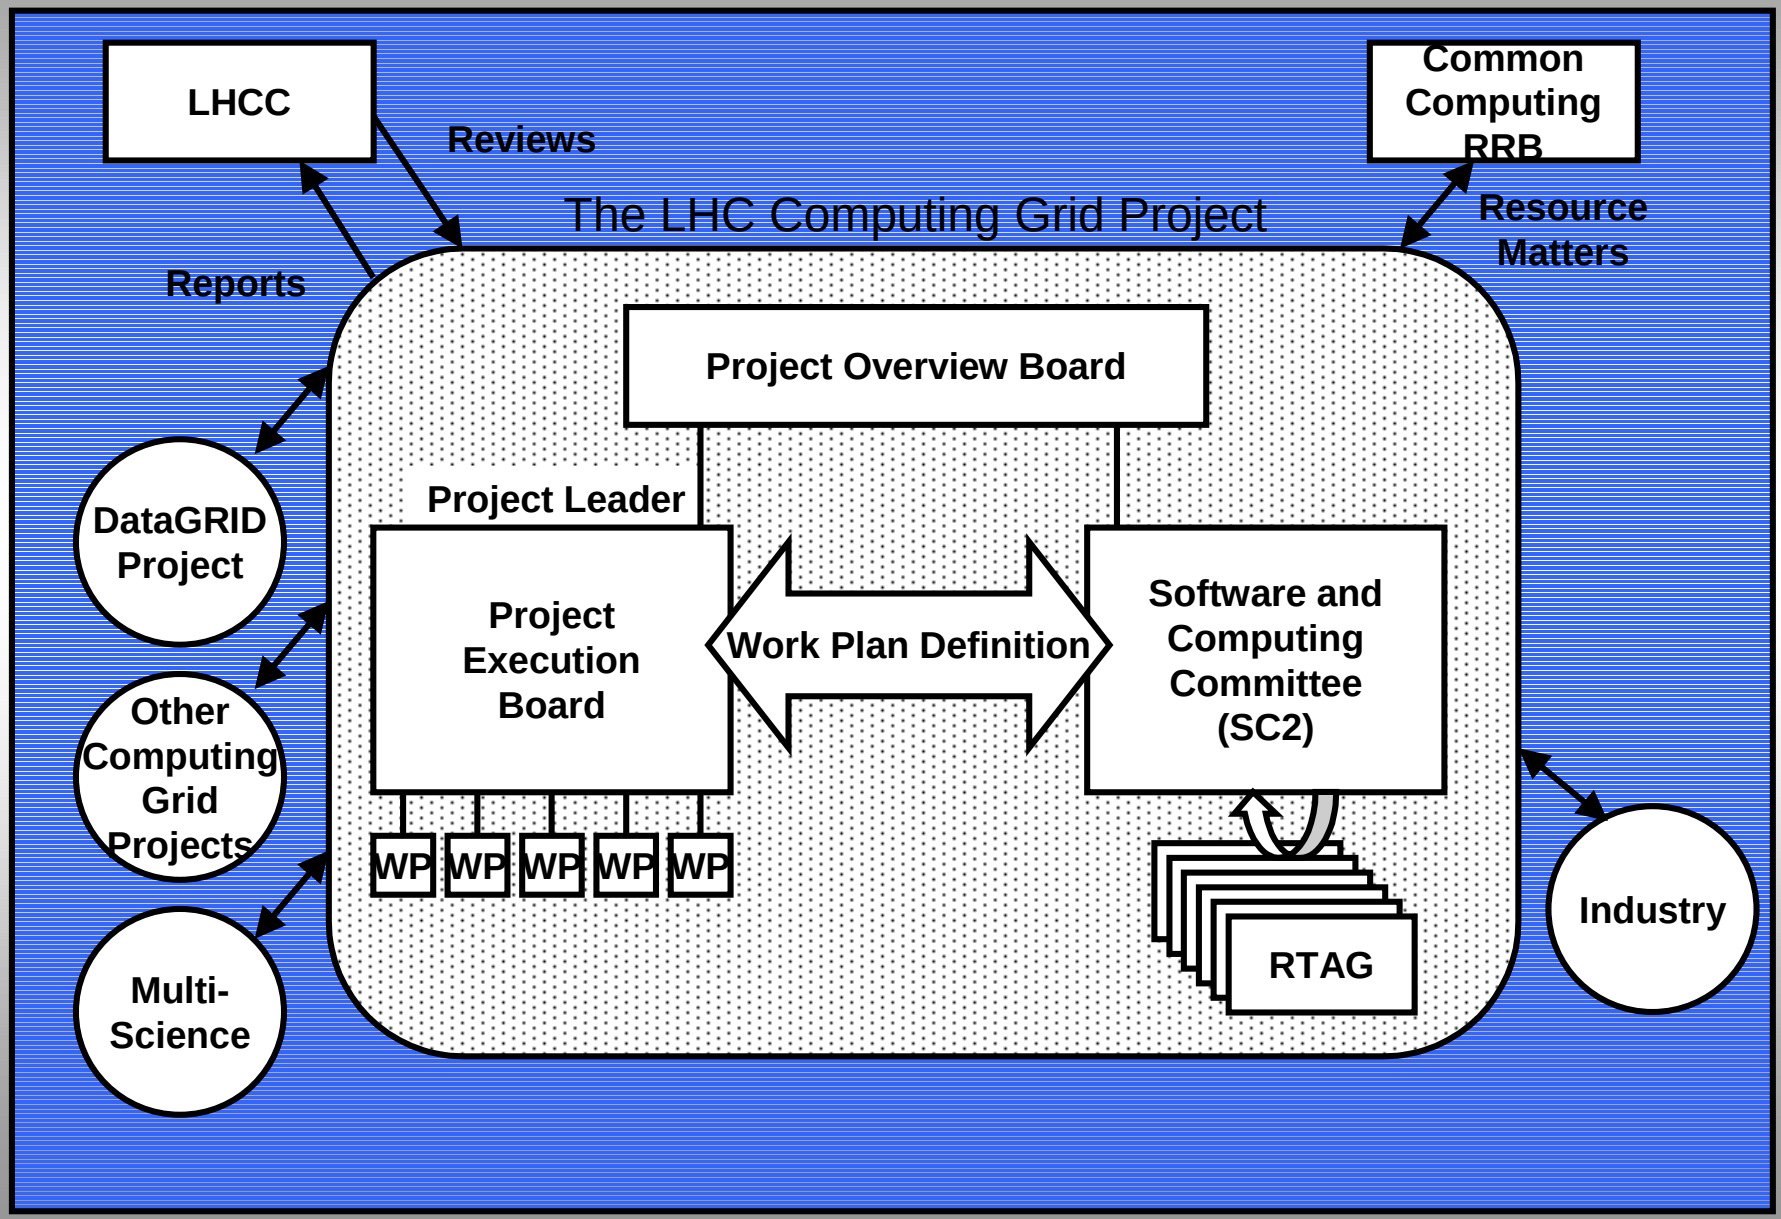

Distribution of CERN users, May 1, 2001

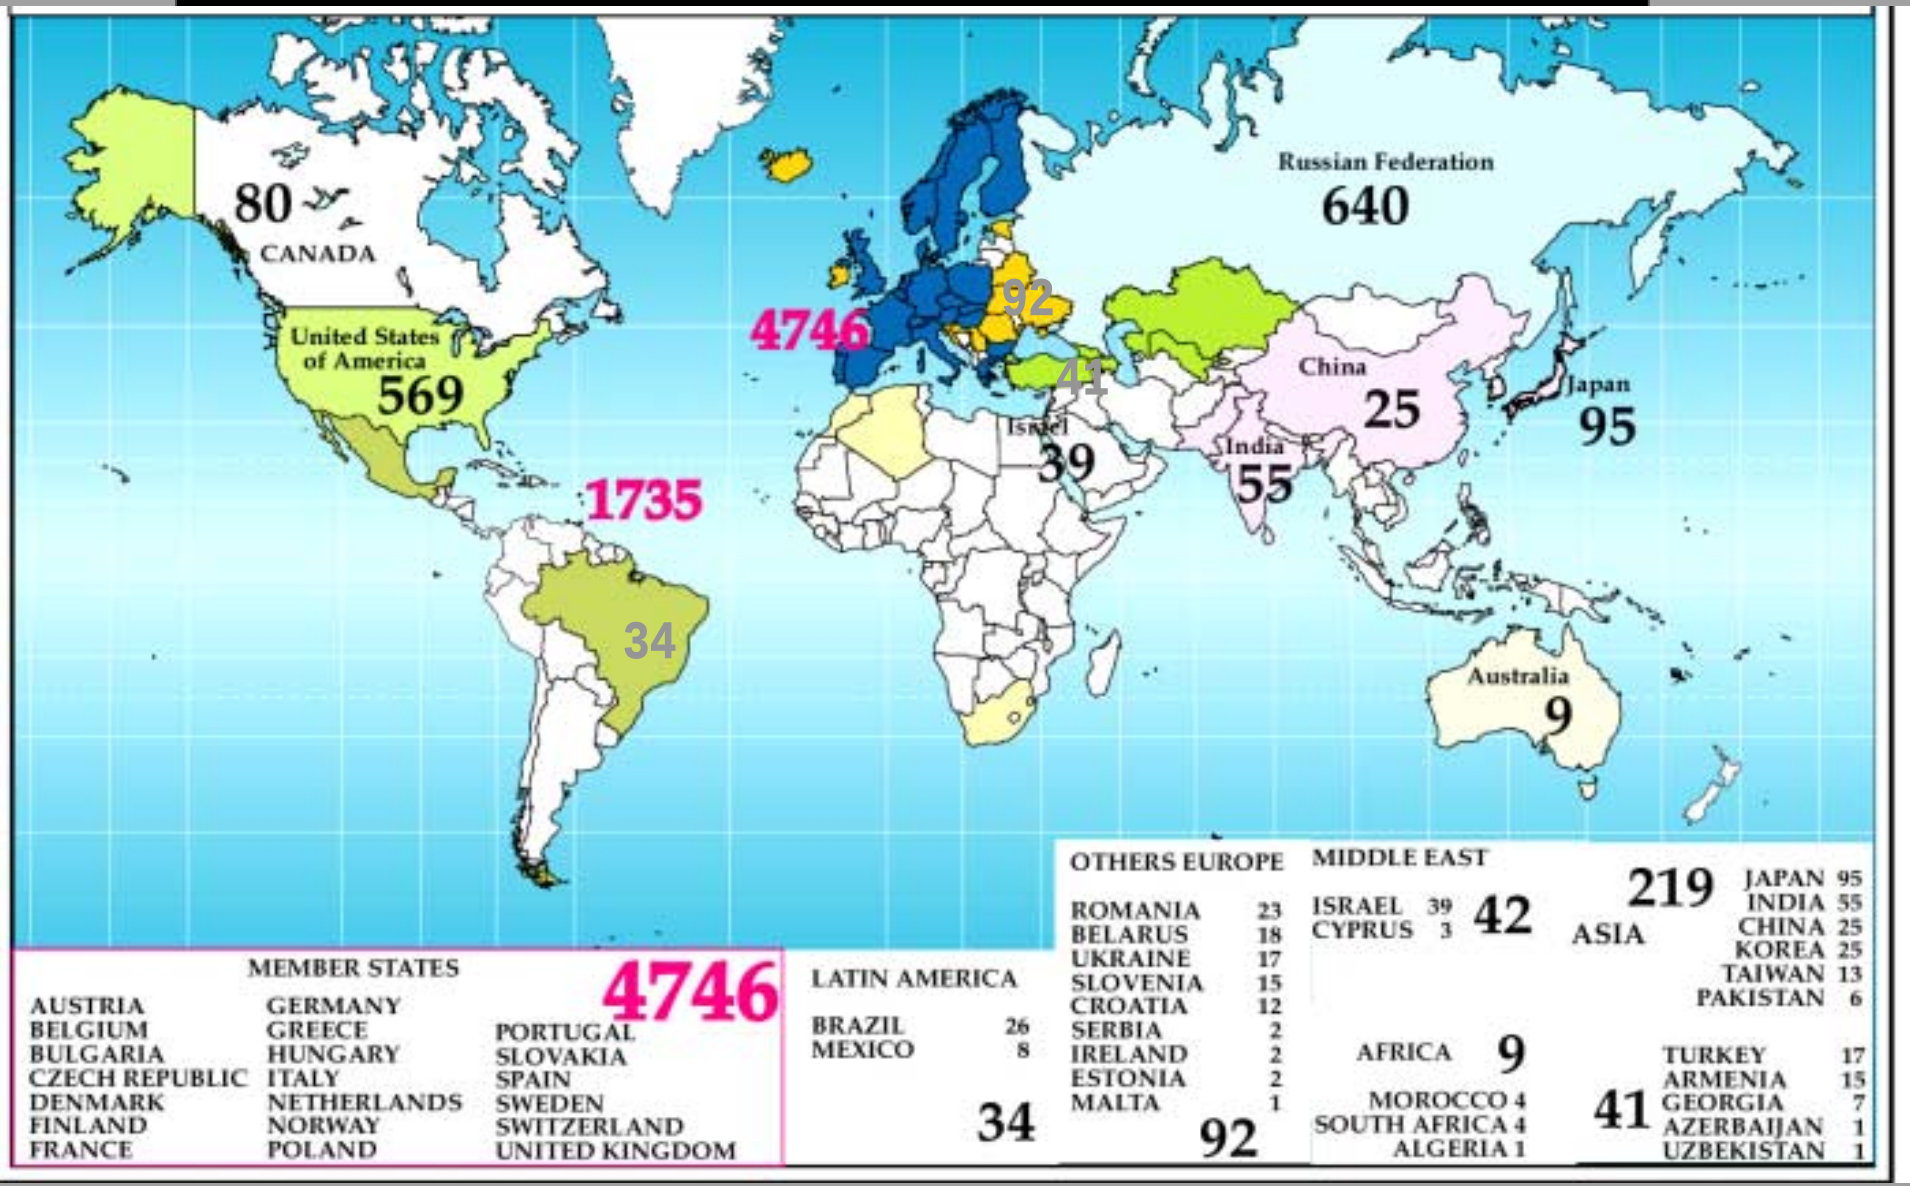

MEMBER STATES — **4746**

AUSTRIA	GERMANY	PORTUGAL
BELGIUM	GREECE	SLOVAKIA
BULGARIA	HUNGARY	SPAIN
CZECH REPUBLIC	ITALY	SWEDEN
DENMARK	NETHERLANDS	SWITZERLAND
FINLAND	NORWAY	UNITED KINGDOM
FRANCE	POLAND	

LATIN AMERICA — **34**

BRAZIL	26
MEXICO	8

OTHERS EUROPE — **92**

ROMANIA	23
BELARUS	18
UKRAINE	17
SLOVENIA	15
CROATIA	12
SERBIA	2
IRELAND	2
ESTONIA	2
MALTA	1

MIDDLE EAST — **42**

ISRAEL	39
CYPRUS	3

AFRICA — **9**

MOROCCO	4
SOUTH AFRICA	4
ALGERIA	1

ASIA — **219**

JAPAN	95
INDIA	55
CHINA	25
KOREA	25
TAIWAN	13
PAKISTAN	6

41

TURKEY	17
ARMENIA	15
GEORGIA	7
AZERBAIJAN	1
UZBEKISTAN	1

Summary

- GRID approach accepted
- Phase 1 of LHC GRID underway
  - Will determine Phase 2
  - World wide participation
- Phase 2 will be deployed in 2006
  - Important for resources but also for wide particpation