The LHC Computing Review

S. Bethke

Max-Planck-Institute of Physics

Munich

CERN/LHCC/2001-004

CERN/RRB-D 2001-3

Original: English

22 February 2001

ORGANISATION EUROPÉENNE POUR LA RECHERCHE NUCLÉAIRE

CERN EUROPEAN ORGANIZATION FOR NUCLEAR RESEARCH

REPORT OF THE STEERING GROUP*

OF THE LHC COMPUTING REVIEW

http://lhc-computing-review-public.web.cern.ch

MEMBERSHIP OF THE REVIEW

Steering Group

Members:	S. Bethke (MPI Munich)	Chair	
	H.F. Hoffmann (CERN)	CERN Director for Sc. Computing	
	D. Jacobs (CERN)	Secretary	
	M. Calvetti (INFN Florence)	Chair of the Mgmt and Resources Panel	
	M. Kasemann (FNAL)	Chair of the Software Project Panel	
	D. Linglin (CC-IN2P3 / CNRS)	Chair of the Computing Panel	
In Attendance:		**Representative**	**Alternate**
	IT Division	M. Delfino (CERN)	L. Robertson (CERN)
	ALICE	F. Carminati (CERN)	K. Safarik (CERN)
	ATLAS	N. McCubbin (RAL)	G. Poulard (CERN)
	CMS	M. Pimia (CERN)	H. Newman (CALTECH)
	LHCb	J. Harvey (CERN)	M. Cattaneo (CERN)
Observers:	R. Cashmore (CERN)	CERN Director for collider programmes	
	J. Engelen (NIKHEF)	LHCC Chairman	

Worldwide Analysis / Computing Model Panel

	D. Lenglin (CC-IN2P3 / CNRS)	Chair	
	F. Gagliari (CERN)	Secretary	
Expt. Reps.:		**Representative**	**Alternate**
	ALICE	A. Masoni (INFN Rome)	A. Sandoval (GSI Darmstadt)
	ATLAS	A. Putzer (U. Heidelberg)	L. Perini (U. Milan)
	CMS	H. Newman (CALTECH)	W. Jank (CERN)
	LHCb	F. Harris (U. Oxford)	M. Schmelling (MPI Heidelberg)
Experts:	Y. Morita (KEK)	C. Michau (UREC-STIC / CNRS)	

Software Project Panel

	M. Kasemann (FNAL)	Chair	
	A. Pfeiffer (CERN)	Secretary and CERN-IT representative	
Expt. Reps.:		**Representative**	**Alternate**
	ALICE	R. Brun (CERN)	A. Morsch (CERN)
	ATLAS	D. Barberis (U. Genoa)	M. Bosman (U.A. Barcelona)
	CMS	L. Taylor (Northeastern U.)	T. Todorov (IN2P3 Strasbourg)
	LHCb	P. Mato (CERN)	O. Callot (LAL Orsay)
Experts:	V. White (FNAL)		

Management and Resources Panel

	M. Calvetti (INFN Florence)	Chair	
	M. Lamanna (INFN Trieste and CERN)	Secretary	
Expt. Reps.:		**Representative**	**Alternate**
	ALICE	P. Vande Vyvre (CERN)	K Safarik (CERN)
	ATLAS	J. Huth (Harvard)	H. Meinhard (CERN)
	CMS	P. Capiluppi (INFN Bologna)	I. Willers (CERN)
	LHCb	J. Harvey (CERN)	J.P. Dufey (CERN)
Experts:	F. Etienne (IN2P3 Marseille)	J. Gordon (RAL)	L. Robertson (CERN)
	F. Ruggieri (INFN Bari)	T. Wenaus (BNL)	K. Woller (DESY)
	G. Wormser (IN2P3 Paris)		

The LHC Computing Review

the challenge:

- 4 experiments; 50-200 Hz data taking rate
- raw event size: 0.12 / 1 / 1-25 MB (LHCb / ATLAS-CMS / ALICE)
- total raw data storage: 7 PB/a ($7 \cdot 10^{15}$ Bytes per year)
- total sim. Data storage: 3.2 PB/a
- world-wide* tape storage: 28.5 PB/a (40 million CD-Rom's)
- world-wide* disk storage: 10.4 PB/a (100k disks à 100 GB)
- world-wide* CPU capacity: 7350 k SI-95 (360k today's PCs)
- WAN bandwidth (Tier-0/-1): 1500 Mbps (1 experiment)
  (5000 Mbps when serving all 4 exp.'s)

* *all Tier-0, Tier-1 and Tier-2 computing centres, excl. Tier-3 and -4*

The LHC Computing Review

(Executive Summary – the LHC computing model)

1. review accepts scale of resource requirements of exp.s

2. recommend distributed, hierarchical model à la MONARC

   Tier-0: at CERN; raw data storage; reconstruction; …
   Tier-1: regional/supranational; analysis, MC generation, storage, …
   Tier-2: national/intranational;
   Tier-3: institutional;
   Tier-4: end-user workstations

3. GRID technology to be used (efficient resource usage, rapid turnaround)

4. need well-supported Research Networking of 1.5–3 Gbps (for each experiment), at affordable costs, by 2006.

Computing Resources[13] planned by the four LHC Experiments in 2007 (*)

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

(*) or the first full year with design luminosity

(**) further estimates envisaged by ATLAS, see Chapter 5.1.2.1 for details.

(***) for all except CMS, the Tier1 and Tier2 needs are merged together.

(all numbers include realistic usage efficiency factors)

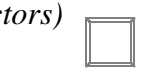

The LHC Computing Model

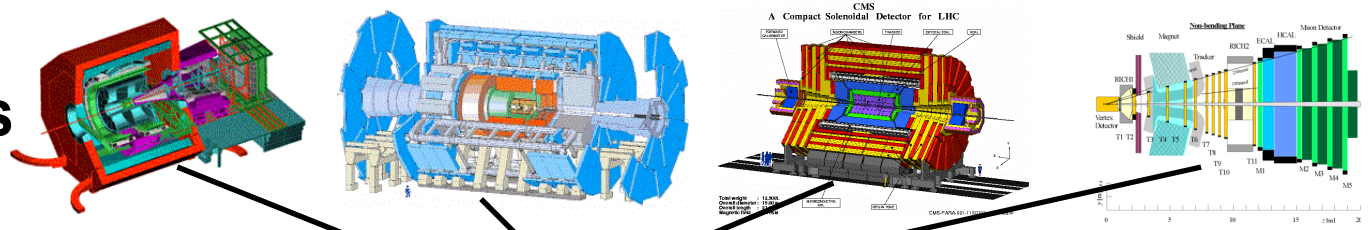

Experiments

Tier - 0 CERN

Tier - 1 Region 1 | Region 2 | Region 3 | Region 4 | CERN

Tier - 2 Country 1 | Country 2 | Country 3 | Country 4 | Country 5 | ...

Tier - 3 Lab 1 | Uni 2 | Uni 3 | Lab 4 | Uni 5 | Lab 6 | ...

Tier - 4 ...

grid middleware

The LHC Computing Review

(Executive Summary – Software)

5. recommend joint efforts and common projects between experiments and CERN-IT; support for widely used products

6. perform data challenges of increasing size and complexity

7. CERN should sponsor transition to OO programming

8. identified areas of concern:
   - limited maturity of current planning and resource estimates
   - insufficient development and support of simulation packages
   - insufficient support and future evolution of analysis tools

Examples of Common software products in use by LHC experiments

					X=yes			used by	
Project/Product		ALICE	ATLAS	LHCb	CMS	IT	as	HEP outside LHC	non-HEP
GEANT4	Detector Simulation	X	X		X	X	developer	X	X
	package written by GEANT		X	X	X	X	maintainer	X	X
	collaboration	X	X	X	X		user	X	X
GEANT3	Detector Simulation						developer	X	
	package written in Fortran	X			X	X	maintainer	X	
		X	X	X	X		user	X	
FLUKA	MC for radiation studies						developer	X	
							maintainer	X	
		X	X	X	X		user	X	
Event Generators	Many - including Pythia,						developer	X	
	Herwig, QQ,etc.						maintainer	X	
		X	X	X	X		user	X	
Objectivity DB +	Commercial Object				X	X	"developer"	X	X
tools/knowledge	Database				X	X	administrator	X	X
	tools+knowledge		X		X	X	user	X	X
ROOT	ROOT objects streamed	X					developer	X	X
persistency,CINT	to files for either data	X					maintainer	X	X
and file format	or conditions	X		X			user	X	X
Mass Storage	HPSS, Castor, other...					Castor	developer		
System						X	administrator		
		X	X	X	X	X	user		
Relational DB	ORACLE or MySQL						developer	X	X
for data handling						X	administrator	X	X
		X		X			user	X	X
ANAPHE	Replacement for CERNLIB				X	X	developer	X	
	several commercial and					X	maintainer	X	
	many HEP packages		X	X	X		user	X	

The LHC Computing Review

(Executive Summary – Management and Resources)

9. Current cost estimates based on forecast evolution of price and performance of computer hardware

10. hardware costs of initial set-up of LHC distributed computer centres (Tier-0 to -2): 240 MCHF

    CERN-based Tier-0+1 centre: about 1/3 of total.

    Significant uncertainties due to performance of LHC, detectors, triggers, backgrounds ...

11. investment for initial system to be spent in 2005, 2006 and 2007, in ~ equal portions (assuming LHC start-up in 2006 and reach of design luminosity in 2007)

12. major concern: core software teams severely understaffed

CPU price evolution (PASTA report)

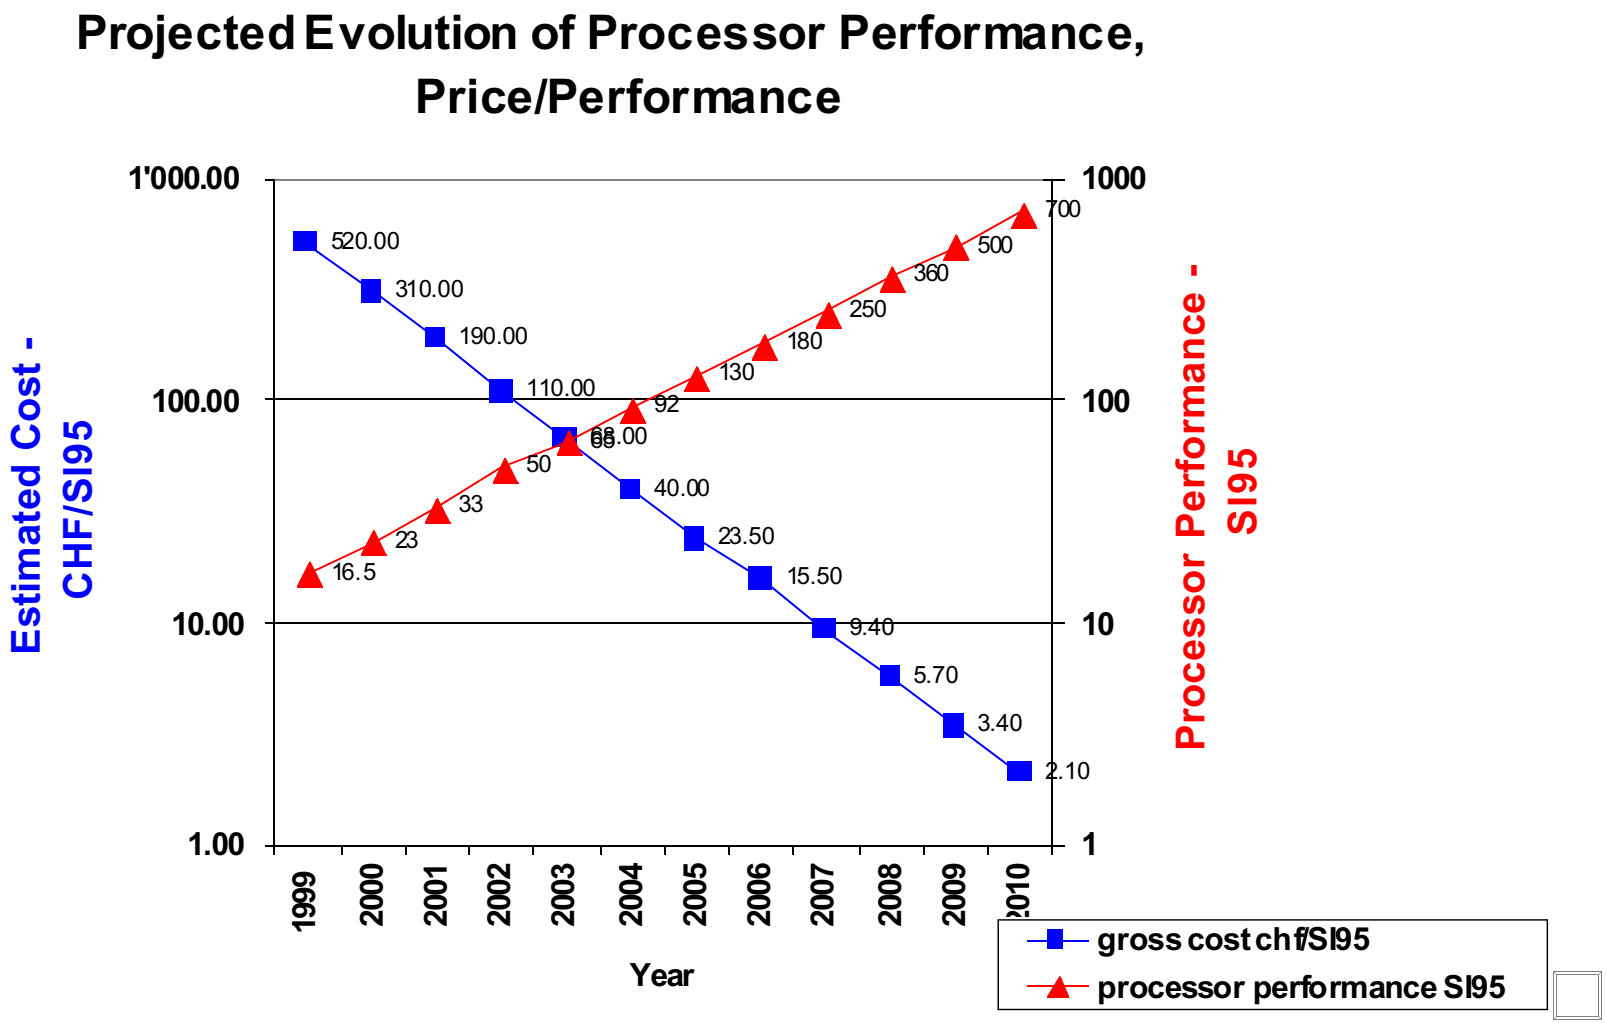

Comparison of recent CERN PC purchases with PASTA predictions

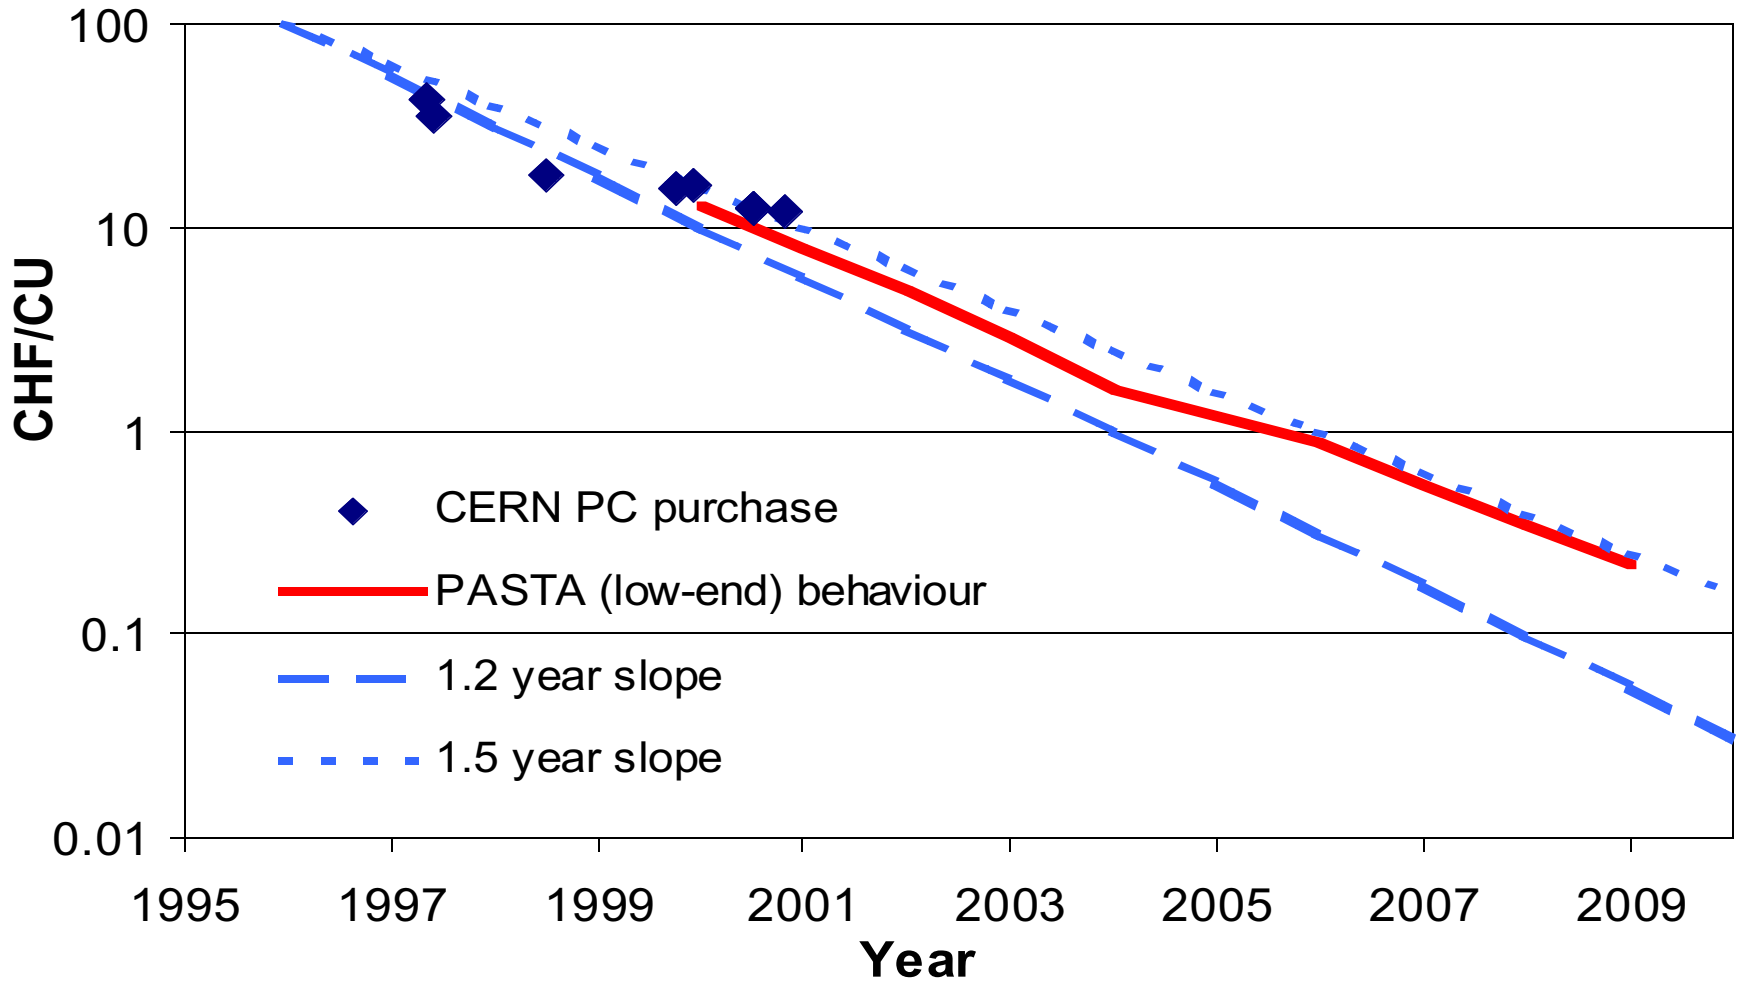

Experience from the LEP Era

- from *first estimates in 1983* (6 yrs before run-start)
  to *available power in 2000* (at final shut-down):
  increase in CPU power by factor 1000 ! (Moore's law: $[1.5]^{17} = 1000$)

- CERN computing capacity evolution
  'at constant budget':

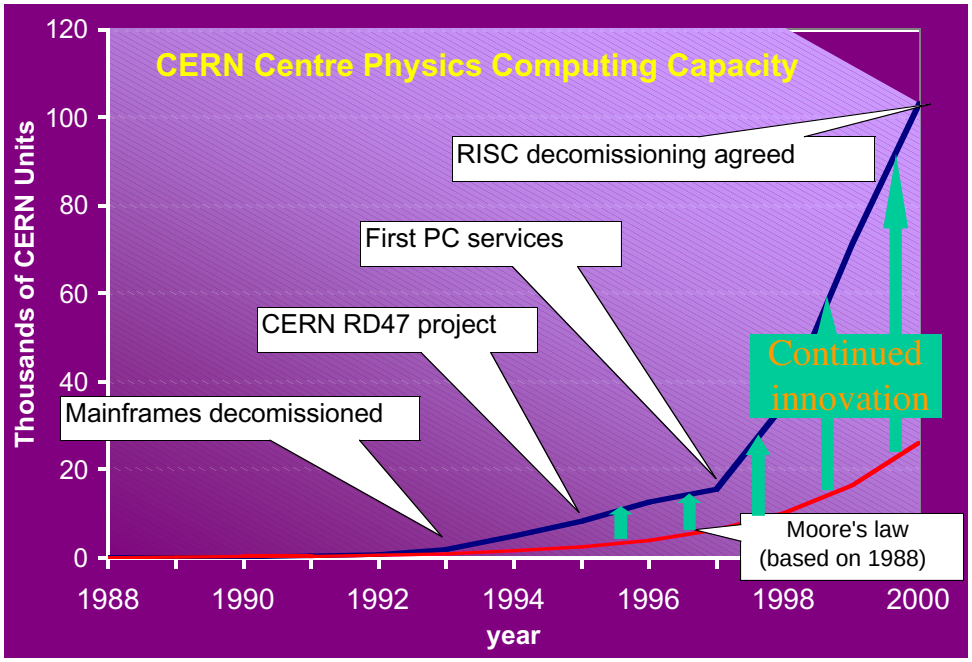

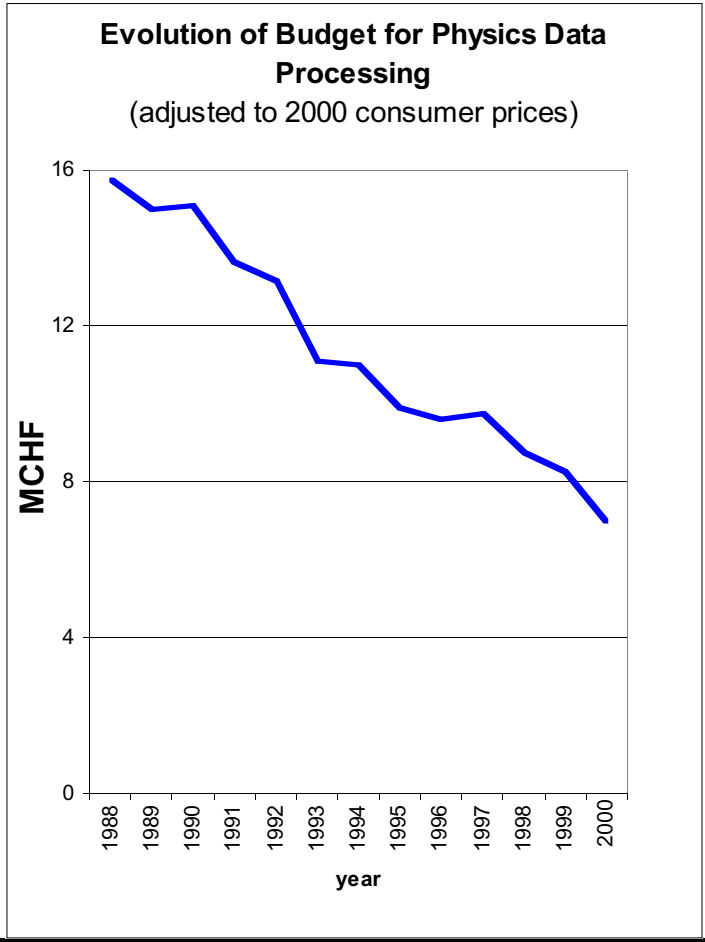

Disk price evolution (PASTA report)

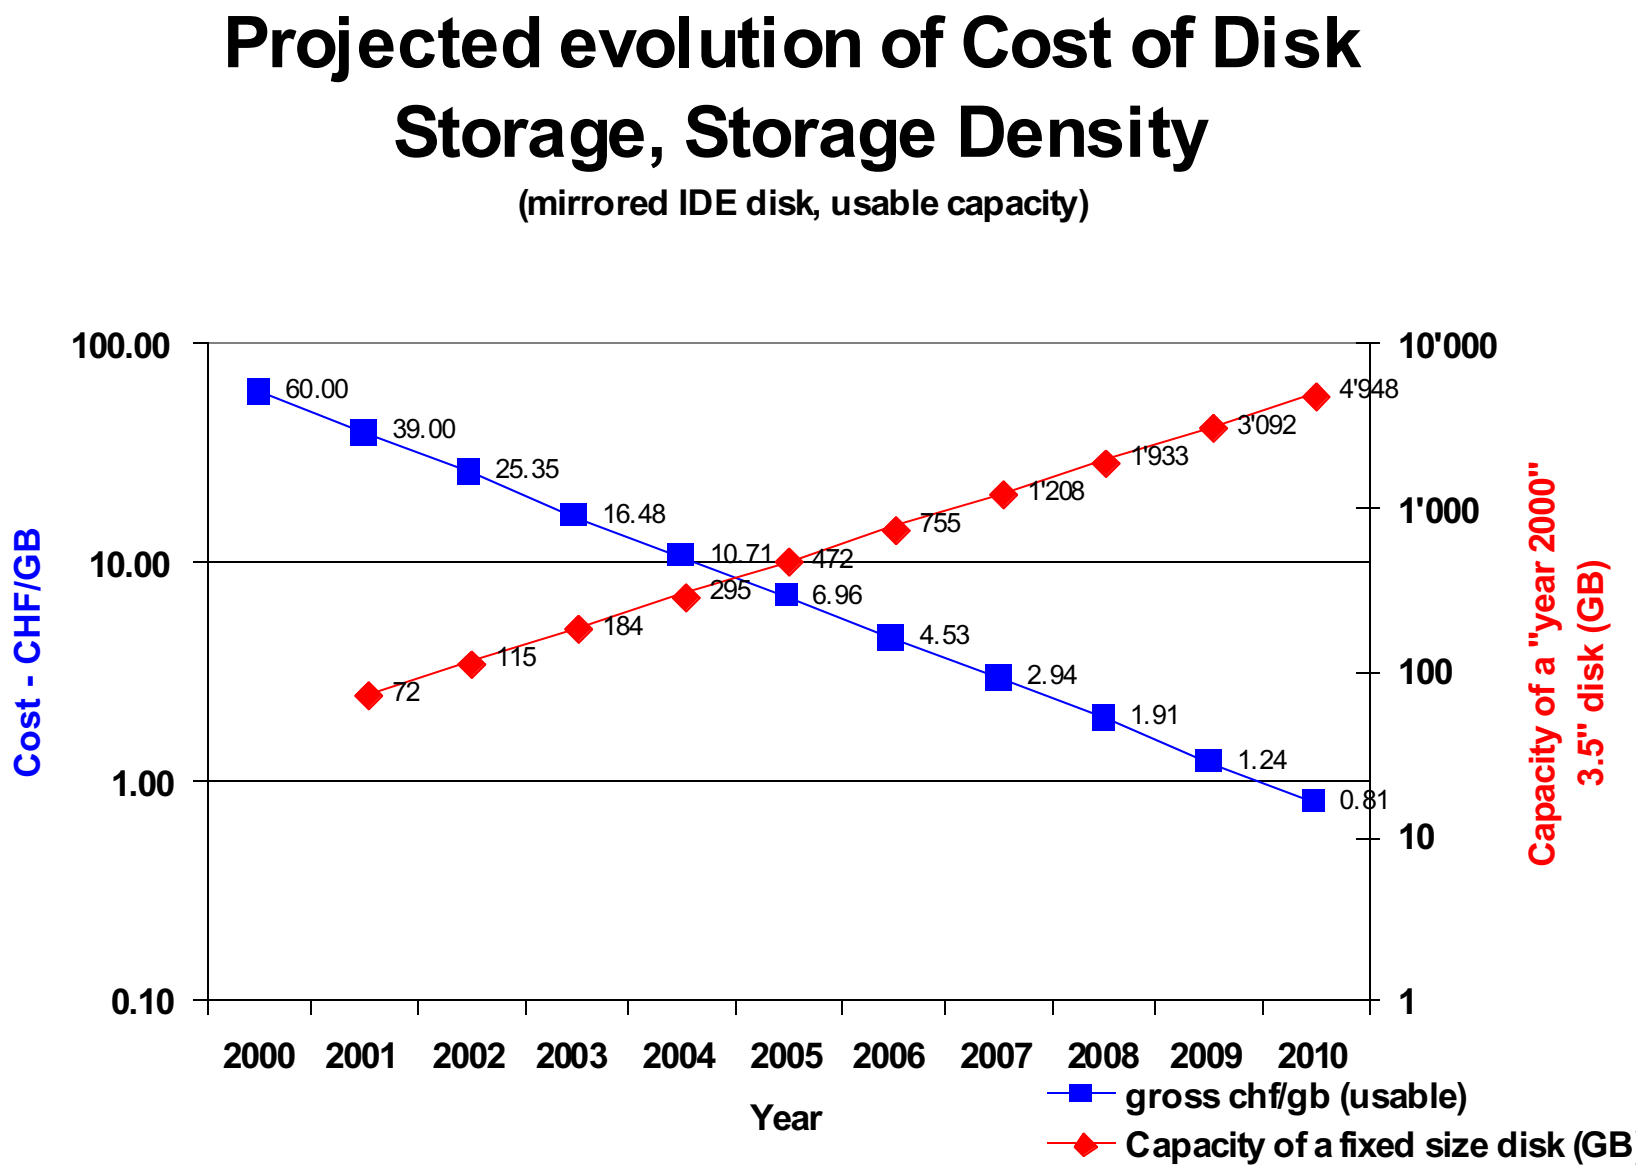

Required human resources (FTEs) to write the Core Software

Year	2000 have(missing)	2001	2002	2003	2004	2005
ALICE	12(5)	17.5	16.5	17.0	17.5	16.5
ATLAS[1]	23(8)	36	35	30	28	29
CMS	15(10)	27	31	33	33	33
LHCb	14(5)	25	24	23	22	21
Totals	**64(28)**	**105.5**	**106.5**	**103**	**100.5**	**99.5**

[1]CORE software includes everything except the algorithmic part of the reconstruction software, the simulation and physics analysis. Human resources attributable to the GRID are not included.

The LHC Computing Review

(Executive Summary – Management and Resources)

13. planned reduction of CERN-IT staff: incompatible with CERN-based LHC computing system and software support

14. M&O of LHC computing system: rolling replacement within constant budget: requires ~ 1/3 of initial investment per year (~ 80 MCHF world-wide) - includes steady evolution of capacity

15. set-up of a common prototype as joint project (experiments, CERN-IT, major regional centres), reaching ~50% of overall computing structure of 1 LHC experiment by ~2003/4

16. set up agreement about construction & cost sharing of prototype - now!

The LHC Computing Review

(Executive Summary – general recommendations)

17. Set up LHC Software and Computing Steering Committee (SC2) composed of highest level software and computing management in experiments, CERN-IT and regional centres to oversee deployment of entire LHC hierarchical system

18. SC2 establishes Technical Assessment Groups (TAG's) to prepare and initiate certain tasks and projects

19. Each collaboration must prepare a MoU for LHC computing describing funding and responsibilities for hard- and software, human resources etc. IMoU's by end of 2001

The LHC Computing Review: Summary

- emphasize utmost importance of proper funding, development, timely realisation and maintenance of computing for success of LHC !
- unprecedented challenge to international HEP and IT communities.
- follow distributed, hierarchical computing model à la MONARC.
- CERN-based Tier0+1 for all experiments: ca. 1/3 of total system.
- initial hardware costs of overall LHC Tier0-2 system: ~ 240 MCHF.
- maintenance and operation: rolling replacements; ~ 80 MCHF / yr.
- software development teams and CERN-IT severely understaffed!
- importance of support for commonly used software products.
- need for data challenges and common computing prototype.
- setup & sign MoU's: - prototype - software - LHC computing.
- setup SC2 committee to oversee LHC computing project.

The End

of LHC Computing Review

The Start

of LHC Computing Project