Available on CMS information server

CMS NOTE 2001/035

The Compact Muon Solenoid Experiment

CMS Note

Mailing address: CMS CERN, CH-1211 GENEVA 23, Switzerland

CMS Interim Memorandum of Understanding

The Costs and How They are Calculated

Version 3.1

Ian Willers
CERN, Geneva, Switzerland

Abstract

This document explains how we arrived at the costs that are to be used in the Interim Memorandum of Understanding, iMoU.

It integrates information taken from the Hoffmann review, PASTA and Lucas Taylor. I would like to thank Petr Moissenz for proof reading which helped me to improve the accuracy of the information in a short period of time. Also Claude Charlot for his valuable comments.

The formatting of this document may appear strange since it is a collection of EXCEL spreadsheets. My attempts to integrate this with WORD failed.

Status

This version (from 3.0 onwards) will be circulated to a wide audience. It does not cover some costs such as general infrastructure e.g. buildings or cooling. Wide area network costs are thought to be 20% before the experiment runs and building up to 40% thereafter.

Contents

Part I: The Tier0 and Tier1 Centre at CERN

The Disk and Tape Storage

The following table is taken from the Hoffmann report. The cache disk is now included inside the active tape costs.

	# of events	Ev-size MBytes	Tape/disk		Tbytes
			Active	Archive	Disk
Raw data	1.E+09	1	1000	1000	
Rec.Raw	1.E+09	0.5	500		200
Calibrat.				10	10
Simulation (repository)	5.E+08	2		1000	
Re-proc. ESD	1.E+09	0.5		200	200
Rec-simulation	5.E+08	0.4	40	200	30
Reprocessed ESD (Tier1)	2.E+08	0.5		100	40
Revised ESD	2.E+08	0.5		100	40
General AOD	1.E+09	0.01		10	10
Revised AOD	2.E+08	0.01		2	2
Local AOD, TAG, DPD	2.E+08			10	10
User Data					100
Total			1540	2632	642

Table 1 - CERN tape and disk storage

This represents the useable tape and disk storage. Efficiency will be introduced later.

CPU

These numbers are taken from the Hoffmann review. The figures do not include any efficiency (it does in the Hoffmann report).

Data Processing	# of events to Mass Stor.	CPU per event KSI95/ev.s	CPU total kSI95
Reconstruction	1.E+09	3	347
Reprocessing	1.E+09	3	included above
Selection	1.E+07 - 1.E+08	0.025	8
Analysis and DPD	1.E+07	0.01	100
Total			455

Part I: Tier1 not at CERN

The Disk and Tape Storage

The following table is taken from the Hoffmann report. The cache disk is now included inside the active tape costs.

	# of events	Ev-size MBytes	Tape/ disk Tbytes		
			Active	Archive	Disk
SIM.Out	1.E+08	2		200	
SIM.Rec.	1.E+08	0.4	40		30
Raw-sample	5.E+07	1	50		
Calibration				10	10
ESD	1.E+09	0.5	500		
Re-proc.ESD	2.E+08	0.5		100	40
Re-vised ESD	2.E+08	0.5		100	40
General AOD	1.E+09	0.01		10	10
Revised AOD	2.E+08	2.E+08		2	2
TAG	1.E+09	0.001		1	1
Local AOD, TAG, DPD	2.E+08			10	10
User Data					50
Total			590	433	193

Table 1 - Tier1 tape and disk storage

CPU

These figures are taken from the Hoffmann panel 3 report. They do not include any efficiency.

Data Processing	# of events to Mass Stor.	CPU per event KSI95/ev.s	CPU total kSI95
Simulation	2.5E+07	5	5
Rec-Simulation	2.5E+07	3	3
Re-Processing	2.E+08	3	39
Selection	1.E+07 - 1.E+08	0.025	8
Analysis and DPD	1.E+07	0.01	50
Total			105

Table 2 - Tier1 cpu

In the case of selection a range of number of events is given and the CPU per event is a maximum.

Part I: Tier2

Tape and Disk

We expect there to be about 25 Tier2 centres. The following numbers are extracted from the Hoffmann panel 3 report.

	# of events	Ev-size MBytes	Tape/ disk Tbytes		
			Active	Archive	Disk
Local cached data (real + simulated)					50
User Data				50	20
Total			0	50	70

Table 1 - Tape and disk for a Tier1

CPU

These numbers are from the Hoffmann panel 3 report. They do not include any efficiency.

Data Processing	# of events to Mass	CPU per event KSI95/ev.	CPU total kSI95
Simulation	50000000	5	10
Rec-Simulation	50000000	3	6
Analysis	10000000	0.01	10
Total			26

Part II: The PASTA Cost Evolution

Here is a copy of the table used in the Panel 3 report.

Year	CPU CHF/SI95	box size SI95/box	Disk CHF/GB	Autom.tape CHF/GB	Shelf tape CHF/GB	Tape I/O CHF/(MB/s)	SysAdm CHF/box
2000	310.0	46	60.00	3.70	2.50	7.25	2,500.00
2001	190.0	66	39.00	2.80	2.00	4.83	1,875.00
2002	110.0	100	25.30	2.40	1.60	4.83	1,667.00
2003	68.0	130	16.50	1.90	1.20	3.63	1,250.00
2004	40.0	184	10.70	1.50	1.00	2.90	1,000.00
2005	23.5	260	7.00	1.30	0.80	2.90	833.00
2006	15.5	360	4.50	1.00	0.70	1.81	652.00
2007	9.4	500	3.00	0.90	0.60	1.81	500.00
2008	5.7	720	1.90	0.86	0.55	1.81	500.00
2009	3.4	1000	1.20	0.81	0.50	1.81	500.00
2010	2.1	1400	0.80	0.81	0.50	1.81	500.00

Table 1 - PASTA performance/cost evolution

I will expand on these numbers in the appropriate sections.

Disks

The PASTA evolution table and corresponding chart:

Year	Disk CHF/GB
2000	60.00
2001	39.00
2002	25.30
2003	16.50
2004	10.70
2005	7.00
2006	4.50
2007	3.00
2008	1.90
2009	1.20
2010	0.80

Table 2 - PASTA disk cost evolution

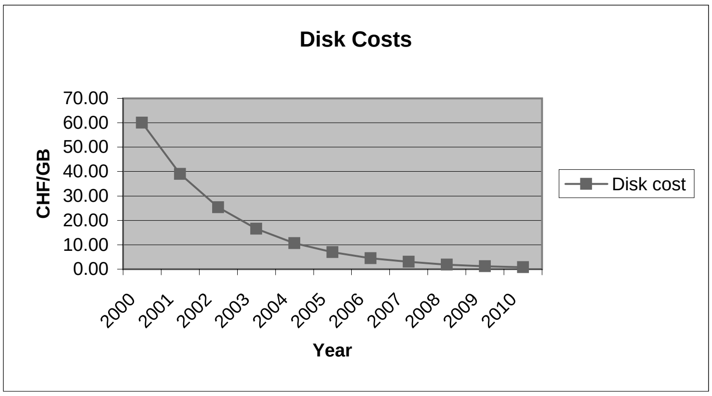

Tapes

The PASTA evolution table and corresponding chart:

Year	Autom.tape CHF/GB	Shelf tape CHF/GB
2000	3.70	2.50
2001	2.80	2.00
2002	2.40	1.60
2003	1.90	1.20
2004	1.50	1.00
2005	1.30	0.80
2006	1.00	0.70
2007	0.90	0.60
2008	0.86	0.55
2009	0.81	0.50
2010	0.81	0.50

Table 3 - PASTA tape cost evolution

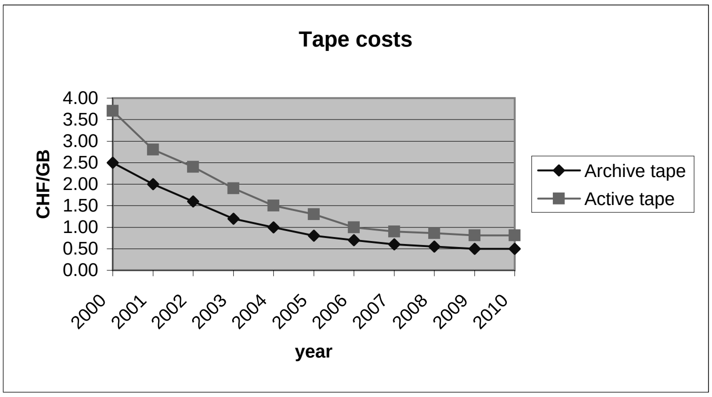

Please note that the "Automatic tape" referred to by PASTA is our "Active tape" and the "Shelf tape" is our "Archive tape".

We also have to give the infrastructure cost corresponding to the bandwidth required. This includes the cost of the drives, disk cache and networking. The PASTA evolution table and corresponding chart:

Year	Tape I/O CHF/(MB/s)
2000	7.25
2001	4.83
2002	4.83
2003	3.63
2004	2.90
2005	2.90
2006	1.81
2007	1.81
2008	1.81
2009	1.81
2010	1.81

Table 4 - PASTA tape I/O evolution

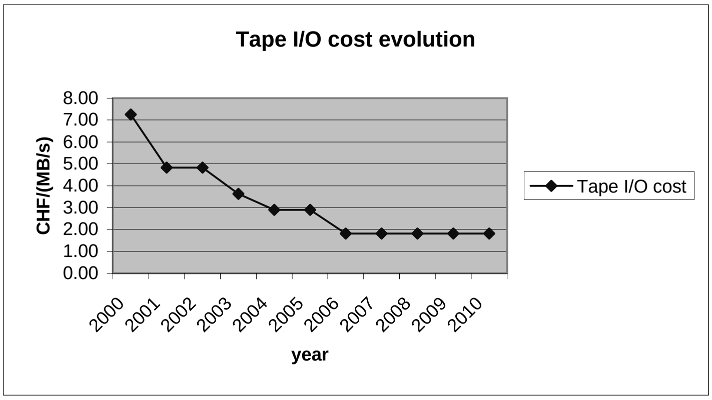

The table above seems to be incorrect by a factor of 1000 - either the units are KCHF/(MB/s) or the decimal points should be ',' to indicate 1,000's (Petr Moissenz).

CPU

The PASTA evolution table and corresponding chart:

Year	CPU CHF/SI95
2000	310
2001	190
2002	110
2003	68
2004	40
2005	24
2006	16
2007	9
2008	6
2009	3
2010	2

Table 5 PASTA cpu cost evolution

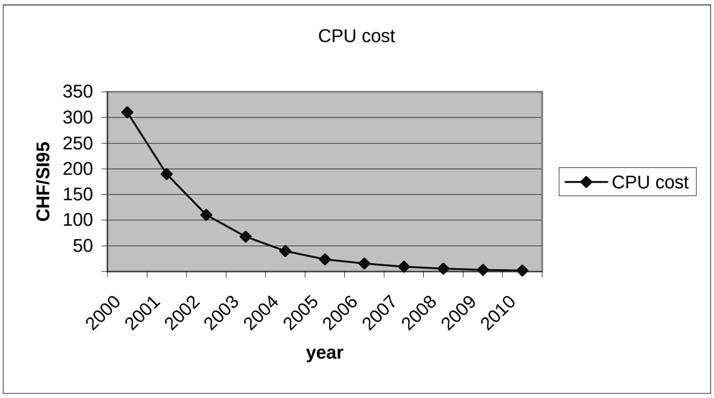

The PASTA report also gives the power per box. This can be used to calucate the complexity of a set up with respect to the number of boxes. The PASTA evolution table and corresponding chart:

Year	box size SI95/box
2000	46
2001	66
2002	100
2003	130
2004	184
2005	260
2006	360
2007	500
2008	720
2009	1000
2010	1400

Table 6 PASTA box size

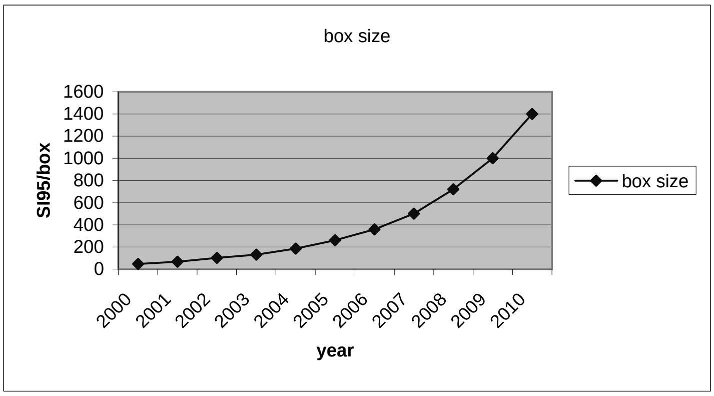

System Administration

This gives the cost of supporting a box for a period of one year. The PASTA evolution table and corresponding chart:

Year	SysAdm CHF/box
2000	2500
2001	1875
2002	1667
2003	1250
2004	1000
2005	833
2006	652
2007	500
2008	500
2009	500
2010	500

Table 7 PASTA System Administration

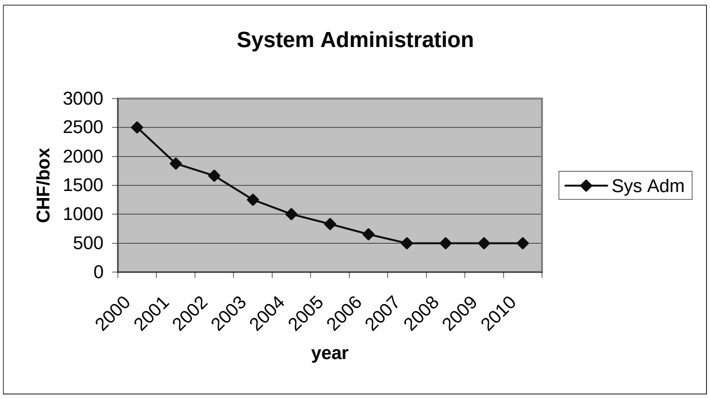

Part II: Efficiency

The following efficiency percentages, except for tape I/O, were established by Les Robertson and agreed by the Hoffmann review steering committee. We need to agree on the tape I/O figure.

Efficiency	
CPU scheduled	85%
CPU analysis	60%
Disk	70%
Active tape	100%
Archive tape	100%
Tape I/O	100%

Table 2 - Efficiency

Part III: Prototype Tier0 and Tier1 at CERN

The CERN prototype is that proportion of the CERN system needed by CMS or ATLAS. Here we use the CMS figures only. The size of the prototype is determined by the complexity of the system. The system is scaled by the number of boxes and the aim is to have a similar complexity as the system at CMS startup.

We first calculate the number of boxes required to build the CERN system (with efficiency included).

The cpu required for the final CERN system in 2007.	588	kSI95
The box size in 2007	500	SI95/box
The number of boxes in the CERN system	1176	boxes

We now calculate the cpu of the prototype based on this number of boxes.

The cpu power of the prototype in 2004	216	kSI95

We now scale all final system numbers with respect to the cpu power.

Scaling factor according to cpu power	37.00%

Now we decide what proportion of this should be in the prototype. The number of boxes will also be scaled by this amount.

Proportion of complexity in prototype	50%	588 boxes.

The final scaling factor	18.500%

We now divide the properties of the system by efficiency catagories.

	final	prototype	
CPU scheduled	408	75	kSI95
CPU analysis	180	33	kSI95
Disk	917	170	Tbytes
Active tape	1540	285	Tbytes
Archive tape	2632	487	Tbytes
Tape I/O	800	148	MB/s

The Hoffmann review said that the cost would be over three years.

Cost division.	2002	2003	2004
	25.00%	25.00%	50.00%

Now we calculate the ratio of the cost division over the cost per unit for each category, sum them up over the three years then divide the units required by this sum to give the total cost.

	2002	2003	2004	sum	cost (CHF)
CPU scheduled	0.002273	0.003676	0.0125	0.0184492	4,065,217
CPU analysis	0.002273	0.003676	0.0125	0.0184492	1,788,696
Disk	0.009881	0.015152	0.0714286	0.0964615	1,762,361
Active tape	0.104167	0.131579	0.3333333	0.5690789	500,809
Archive tape	0.15625	0.208333	0.5	0.8645833	563,277
Tape I/O	0.051728	0.068966	0.1724138	0.293107	504,935

Actual costs	2002	2003	2004	total
	2,296,324	2,296,324	4,592,648	9,185,295

WAN costs (20%)	459,265	459,265	918,530	1,837,059

Part III: Prototype Tier1

The Tier1 prototype is that proportion of the Tier1 system needed by CMS. The size of the prototype is determined by the complexity of the system. The system is scaled by the number of boxes and the aim is to have a similar complexity as the system at CMS startup.

We first calculate the number of boxes required to build the Tier1 system (with efficiency included).

The cpu required for the final Tier1 system in 2007.	152	kSI95
The box size in 2007	500	SI95/box
The number of boxes in the Tier1 system	304	boxes

We now calculate the cpu of the prototype based on this number of boxes.

The cpu power of the prototype in 2004	56	kSI95

We now scale all final system numbers with respect to the cpu power.

The scaling factor according to cpu power	37.00%

Now we decide what proportion of this should be in the prototype. The number of boxes will also be scaled by this amount.

Proportion of compexity in prototype	50.00%	152 boxes.

The final scaling factor	18.500%

We now divide the properties of the system by efficiency catagories.

	final	prototype	
CPU scheduled	55	10	kSI95
CPU analysis	97	18	kSI95
Disk	276	51	Tbytes
Active tape	590	109	Tbytes
Archive tape	433	80	Tbytes
Tape I/O	400	74	MB/s

The Hoffmann review said that the cost would be over three years.

Cost division.	2002	2003	2004
	25.00%	25.00%	50.00%

Now we calculate the ratio of the cost division over the cost per unit for each category, sum them up over the three years then divide the units required by this sum to give the total cost.

	2002	2003	2004	sum	cost
CPU scheduled	0.002273	0.003676	0.0125	0.0184492	542,029
CPU analysis	0.002273	0.003676	0.0125	0.0184492	975,652
Disk	0.009881	0.015152	0.0714286	0.0964615	528,708
Active tape	0.104167	0.131579	0.3333333	0.5690789	191,538
Archive tape	0.15625	0.208333	0.5	0.8645833	92,530
Tape I/O	0.051728	0.068966	0.1724138	0.293107	252,468

Actual costs	2002	2003	2004	total
	645,731	645,731	1,291,462	2,582,925

WAN costs 20%	129,146	129,146	258,292	516,585

Part III: Prototype Tier2

The Tier2 prototype is that proportion of the Tier2 system needed by CMS. The size of the prototype is determined by the complexity of the system. The system is scaled by the number of boxes and the aim is to have a similar complexity as the system at CMS startup.

We first calculate the number of boxes required to build the Tier2 system (with efficiency included).

The cpu required for the final Tier2 system in 2007.	36	kSI95
The box size in 2007	500	SI95/box
The number of boxes in the Tier2 system	72	boxes

We now calculate the cpu of the prototype based on this number of boxes.

The cpu power of the prototype in 2004	13	kSI95

We now scale all final system numbers with respect to the cpu power.

Scaling factor according to cpu power	36.00%

Now we decide what proportion of this should be in the prototype. The number of boxes will also be scaled by this amount.

Proportion of complexity in prototype	50.00%	36 boxes.

The scaling factor	18.00%

We now divide the properties of the system by efficiency catagories.

	final	prototype	
CPU scheduled	19	3	kSI95
CPU analysis	17	3	kSI95
Disk	100	18	Tbytes
Active tape	0	0	Tbytes
Archive tape	50	9	Tbytes
Tape I/O	100	18	MB/s

The Hoffmann review said that the cost would be over three years.

Cost division.	2002	2003	2004
	25.00%	25.00%	50.00%

Now we calculate the ratio of the cost division over the cost per unit for each category, sum them up over the three years then divide the units required by this sum to give the total cost.

	2002	2003	2004	sum	cost (CHF)
CPU scheduled	0.002273	0.003676	0.0125	0.0184492	162,609
CPU analysis	0.002273	0.003676	0.0125	0.0184492	162,609
Disk	0.009881	0.015152	0.0714286	0.0964615	186,603
Active tape	0.104167	0.131579	0.3333333	0.5690789	0
Archive tape	0.15625	0.208333	0.5	0.8645833	10,410
Tape I/O	0.051728	0.068966	0.1724138	0.293107	61,411

Actual costs	2002	2003	2004	total
	145,910	145,910	291,820	583,641

WAN costs 20%	29,182	29,182	58,364	116,728

Part IV: Tier0 and Tier1 at CERN

This calculates the cost of the CERN Tier0 plus Tier1.

	final	
CPU scheduled	347	kSI95
CPU analysis	108	kSI95
Disk	642	Tbytes
Active tape	1540	Tbytes
Archive tape	2632	Tbytes
Tape I/O	800	MB/s

We now apply the efficiency factors to the CERN system.

CPU scheduled	408	kSI95
CPU analysis	180	kSI95
Disk	917	Tbytes
Active tape	1540	Tbytes
Archive tape	2632	Tbytes
Tape I/O	800	MB/s

The Hoffmann review said that the cost would be over three years.

Cost division.	2005	2006	2007
	30.00%	30.00%	40.00%

Now we calculate the ratio of the cost division over the cost per unit for each category, sum them up over the three years then divide the units required by this sum to give the total cost.

	2005	2006	2007	sum	cost (CHF)
CPU scheduled	0.012766	0.019355	0.0425532	0.074674	5,463,750
CPU analysis	0.012766	0.019355	0.0425532	0.074674	2,410,478
Disk	0.042857	0.066667	0.1333333	0.2428571	3,775,882
Active tape	0.230769	0.3	0.4444444	0.9752137	1,579,141
Archive tape	0.375	0.428571	0.6666667	1.4702381	1,790,186
Tape I/O	0.103448	0.165563	0.2207506	0.4897617	1,633,447

Actual costs	2005	2006	2007	total
	4,995,865	4,995,865	6,661,154	16,652,885

WAN costs	20%	30%	40%	
	999,173	1,498,760	2,664,462	5,162,394

Part IV: Typical Tier1

This calculates the cost of a typical Tier1.

	final	
CPU scheduled	47	kSI95
CPU analysis	58	kSI95
Disk	193	Tbytes
Active tape	590	Tbytes
Archive tape	433	Tbytes
Tape I/O	400	MB/s

We now apply the efficiency factors to the Tier1 system.

CPU scheduled	55	kSI95
CPU analysis	97	kSI95
Disk	276	Tbytes
Active tape	590	Tbytes
Archive tape	433	Tbytes
Tape I/O	400	MB/s

The Hoffmann review said that the cost would be over three years.

Cost division.	2005	2006	2007
	30.00%	30.00%	40.00%

Now we calculate the ratio of the cost division over the cost per unit for each category, sum them up over the three years then divide the units required by this sum to give the total cost.

	2005	2006	2007	sum	cost (CHF)
CPU scheduled	0.012766	0.019355	0.0425532	0.074674	736,535
CPU analysis	0.012766	0.019355	0.0425532	0.074674	1,298,980
Disk	0.042857	0.066667	0.1333333	0.2428571	1,136,471
Active tape	0.230769	0.3	0.4444444	0.9752137	604,996
Archive tape	0.375	0.428571	0.6666667	1.4702381	294,510
Tape I/O	0.103448	0.165563	0.2207506	0.4897617	816,724

Actual costs	2005	2006	2007	total
	1,466,464	1,466,464	1,955,286	4,888,215

WAN costs	20%	30%	40%	
	293,293	439,939	782,114	1,515,347

Part IV: Typical Tier2

This calculates the cost of a typical Tier2.

	final	
CPU scheduled	16	kSI95
CPU analysis	10	kSI95
Disk	70	Tbytes
Active tape	0	Tbytes
Archive tape	50	Tbytes
Tape I/O	100	MB/s

We now apply the efficiency factors to the Tier2.

CPU scheduled	19	kSI95
CPU analysis	17	kSI95
Disk	100	Tbytes
Active tape	0	Tbytes
Archive tape	50	Tbytes
Tape I/O	100	MB/s

The Hoffmann review said that the cost would be over three years.

Cost division.	2005	2006	2007
	30.00%	30.00%	40.00%

Now we calculate the ratio of the cost division over the cost per unit for each category, sum them up over the three years then divide the units required by this sum to give the total cost.

	2005	2006	2007	sum	cost (CHF)
CPU scheduled	0.012766	0.019355	0.0425532	0.074674	254,439
CPU analysis	0.012766	0.019355	0.0425532	0.074674	227,656
Disk	0.042857	0.066667	0.1333333	0.2428571	411,765
Active tape	0.230769	0.3	0.4444444	0.9752137	0
Archive tape	0.375	0.428571	0.6666667	1.4702381	34,008
Tape I/O	0.103448	0.165563	0.2207506	0.4897617	204,181

Actual costs	2005	2006	2007	total
	339,615	339,615	452,820	1,132,049

WAN costs	20%	30%	40%	
	67,923	101,884	181,128	350,935

Part V - The Maintenance and Operation's figures

	Costbook kCHF	# Scient.	b/Costb.	b/Scient.	50/50 mix
Austria	3,900	15	0.86%	1.33%	1.09%
Belgium	5,000	25	1.10%	2.21%	1.66%
Bulgaria	600	17	0.13%	1.50%	0.82%
CERN	85,200	107	18.74%	9.46%	14.10%
China	4,765	33	1.05%	2.92%	1.98%
Croatia	280	4	0.06%	0.35%	0.21%
Cyprus	600	2	0.13%	0.18%	0.15%
Estonia	90	2	0.02%	0.18%	0.10%
Finland	5,000	12	1.10%	1.06%	1.08%
France-CEA	5,600	10	1.23%	0.88%	1.06%
France-IN2P3	19,700	49	4.33%	4.33%	4.33%
Germany	17,000	45	3.74%	3.98%	3.86%
Greece	5,000	15	1.10%	1.33%	1.21%
Hungary	1,000	20	0.22%	1.77%	0.99%
India	4,400	20	0.97%	1.77%	1.37%
Italy	55,000	190	12.10%	16.80%	14.45%
Korea	2,600	17	0.57%	1.50%	1.04%
Pakistan	1,000	7	0.22%	0.62%	0.42%
Poland	3,000	8	0.66%	0.71%	0.68%
Portugal	2,000	10	0.44%	0.88%	0.66%
RDMS-Russia	20,500	60	4.51%	5.31%	4.91%
RDMS-DMS	6,400	72	1.41%	6.37%	3.89%
Spain	6,000	34	1.32%	3.01%	2.16%
Switzerland-ETHZ/Univ.	78,500	36	17.27%	3.18%	10.23%
Switzerland-PSI	8,500	8	1.87%	0.71%	1.29%
Taiwan	2,230	6	0.49%	0.53%	0.51%
Turkey	1,000	7	0.22%	0.62%	0.42%
UK	9,100	38	2.00%	3.36%	2.68%
USA-DoE and USA-NSF	100,610	262	22.13%	23.17%	22.65%
Totals	454,575	1,131	100.00%	100.00%	100.00%

Part V: The Software Engineers

We assume that CERN takes a large proportion of the responsibility for offline software (25%). The remainder is devided amongst the countries according to the number of scientist.

These numbers are ordered putting CERN first and then followed by the largest contributers. The total number of software engineers is assumed to be 40.

Country	Number of CMS Scientists	Percentage of Software Engineers	Software Engineers FTE's
CERN (CMS groups)	107	25.0%	10.00
USA (DOE and NSF)	262	19.2%	7.68
Italy	190	13.9%	5.57
RDMS-DMS	72	5.3%	2.11
RDMS-Russia	60	4.4%	1.76
France (CEA+IN2P3)	59	4.3%	1.73
Germany	45	3.3%	1.32
UK	38	2.8%	1.11
Switzerland-ETHZ/Univ.	36	2.6%	1.05
Spain	34	2.5%	1.00
China	33	2.4%	0.97
Belgium	25	1.8%	0.73
India	20	1.5%	0.59
Hungary	20	1.5%	0.59
Korea	17	1.2%	0.50
Bulgaria	17	1.2%	0.50
Greece	15	1.1%	0.44
Austria	15	1.1%	0.44
Finland	12	0.9%	0.35
Portugal	10	0.7%	0.29
Switzerland-PSI	8	0.6%	0.23
Poland	8	0.6%	0.23
Turkey	7	0.5%	0.21
Pakistan	7	0.5%	0.21
Taiwan	6	0.4%	0.18
Croatia	4	0.3%	0.12
Cyprus	2	0.1%	0.06
Estonia	2	0.1%	0.06
Total	1,131	100%	40

Part VI: Cost Sharing

This sharing is proportional to the number of scientists. CERN is treated specially as the host laboratory and the other costs are in proportion.

These numbers are taken from the table used initially for the Maintenance and Operations MoU.

Country	Number of CMS Scientists	Percentage of CMS Scientists	CERN T0+T1	Tier1	Tier2
CERN (host lab)	107		1.00		
USA (DOE and NSF)	262	25.6%		1.3	6.4
Italy	190	18.6%		0.9	4.6
RDMS-DMS	72	7.0%		0.4	1.8
RDMS-Russia	60	5.9%		0.3	1.5
France (CEA+IN2P3)	59	5.8%		0.3	1.4
Germany	45	4.4%		0.2	1.1
Switzerland-ETHZ/Univ.	36	3.5%		0.2	0.9
UK	38	3.7%		0.2	0.9
Spain	34	3.3%		0.2	0.8
China	33	3.2%		0.2	0.8
Belgium	25	2.4%		0.1	0.6
India	20	2.0%		0.1	0.5
Hungary	20	2.0%		0.1	0.5
Korea	17	1.7%		0.08	0.4
Bulgaria	17	1.7%		0.08	0.4
Greece	15	1.5%		0.07	0.4
Austria	15	1.5%		0.07	0.4
Finland	12	1.2%		0.06	0.3
Portugal	10	1.0%		0.05	0.2
Switzerland-PSI	8	0.8%		0.04	0.2
Poland	8	0.8%		0.04	0.2
Turkey	7	0.7%		0.03	0.2
Pakistan	7	0.7%		0.03	0.2
Taiwan	6	0.6%		0.03	0.1
Croatia	4	0.4%		0.02	0.1
Cyprus	2	0.2%		0.01	0.05
Estonia	2	0.2%		0.01	0.05
Total	1,131	100%	1	5	25

Part VII CERN T0/T1 Staff

This represents the full CERN T0/T1 staff (FTEs) for the four LHC experiments. The portion attributed to CMS is roughly one third.

IT Staff Estimates for LHC Computing Support

Dedicated Physics Computing Services	
Physics applications services	
Simulation	5
Analysis & visualisation	8
Common libraries, tools and base support	8
Controls for physics	8
Tier-0 + Tier-1 centre	
Basic farm mgt/planning	8
Operation & support for LHC experiments	12
High bandwidth WAN for LHC	3
Direct support for non-LHC experiments	5
Total - dedicated support for physics	**57**
Specialised support for engineering	**18**
Infrastructure and shared services (including base support for physics and engineering)	
Data management	8
Desktop support	24
Campus networking	11
Controls infrastructure	5
Database services (relational and object)	13
External networking	6
Internet applications	12
User support, operation and infrastructure	12
Management, administration	21
Total - infrastructure and shared services	**112**
Total IT staff	**187**
Outsourced system administration	30
Total staff	**217**

Part VII - The Tier1 staff.

This table is based on the staff (FTEs) estimates for 2006 made in some detail by Fermilab. The other laboratories are scaled according to the number of physicists in the maintenance and operation memorandum of understanding.

	FNAL	IN2P3	Germany	INFN	Moscow cluster	RAL
Tier1 RC Hardware support	9.0	2.0	1.5	6.5	2.5	1.3
Tier1 Software Systems	11.0	2.5	1.9	8.0	3.0	1.6
System and User Support	6.0	1.4	1.0	4.4	1.6	0.9
Maintenance and operation	5.0	1.1	0.9	3.6	1.4	0.7
Support for Tier2 centres	2.0	0.5	0.3	1.5	0.5	0.3
Networking	2.0	0.5	0.3	1.5	0.5	0.3
Total	35	8	5.9	25.5	9.5	5.1

Part VII - The Tier2 staff

A typical Tier2 centre needs about one quarter of the staff of a typical Tier1.

Country	Staff
CERN (CMS groups)	3.6
USA (DOE and NSF)	8.8
Italy	6.4
RDMS-DMS	2.4
France (CEA+IN2P3)	2
Germany	1.5
Switzerland-ETHZ/Univ.	1.2
UK	1.3
Spain	1.1
China	1.1
Belgium	0.8
India	0.7
Hungary	0.7
Korea	0.6
Bulgaria	0.6
Greece	0.5
Austria	0.5
Finland	0.4
Portugal	0.3
Poland	0.3
Turkey	0.2
Pakistan	0.2
Taiwan	0.2
Croatia	0.1
Cyprus	0.1
Estonia	0.1

Part VIII: Installed capacity

We calculate the capacity to be installed each year of the various resources in the prototype and initial LHC systems

The CERN T0/T1	2002	2003	2004	2005	2006	2007	
CPU scheduled	9	15	51	70	106	233	kSI95
CPU analysis	4	7	22	31	47	103	kSI95
Disk	17	27	126	162	252	503	Tbytes
Active tape	52	66	167	364	474	702	Tbytes
Archive tape	88	117	282	671	767	1193	Tbytes
Tape I/O	26	35	87	169	270	361	MB/s
Number of CPU boxes	133	166	398	387	423	670	

Now we sum each year and the previous two years to get the installed useful capacity in any given year

The CERN T0/T1	2002	2003	2004	2005	2006	2007	
CPU scheduled	9	24	75	136	226	408	kSI95
CPU analysis	4	11	33	60	100	180	kSI95
Disk	17	44	170	314	539	917	Tbytes
Active tape	52	118	285	597	1005	1540	Tbytes
Archive tape	88	205	487	1070	1720	2632	Tbytes
Tape I/O	26	61	148	291	526	800	MB/s
Number of CPU boxes	133	299	696	950	1208	1480	

Repeat the same exercise for the Regional T1

A Single Regional T1	2002	2003	2004	2005	2006	2007	
CPU scheduled	1	2	7	9	14	31	kSI95
CPU analysis	2	4	12	17	25	55	kSI95
Disk	5	8	38	49	76	152	Tbytes
Active tape	20	25	64	140	181	269	Tbytes
Archive tape	14	19	46	110	126	196	Tbytes
Tape I/O	13	17	44	84	135	180	MB/s
Number of CPU boxes	34	43	103	100	109	173	

The number of such "effective" T1 centers	2	2	2	5	5	5

Now we take the product of the number of such centers and the sum of each year and of the previous two years to get the installed useful capacity in any given year

The Effective Sum of Regional T1's	2002	2003	2004	2005	2006	2007	
CPU scheduled	2	6	20	91	152	275	kSI95
CPU analysis	4	12	36	162	270	485	kSI95
Disk	10	26	102	472	811	1380	Tbytes
Active tape	40	90	218	1143	1925	2950	Tbytes
Archive tape	29	67	160	880	1415	2165	Tbytes
Tape I/O	26	61	148	727	1316	2000	MB/s
Number of CPU boxes	69	155	361	1230	1562	1913	

Repeat the same exercise for the Regional T2

A Single Regional T2	2002	2003	2004	2005	2006	2007	
CPU scheduled	0.4	0.6	2.0	3	5	11	kSI95
CPU analysis	0.4	0.6	2.0	3	4	10	kSI95
Disk	1.8	2.8	13.3	18	27	55	Tbytes
Active tape	0.0	0.0	0.0	0	0	0	Tbytes
Archive tape	1.6	2.2	5.2	13	15	23	Tbytes
Tape I/O	3.2	4.2	10.6	21	34	45	MB/s
Number of CPU boxes	7	9	22	24	26	41	

The number of such "effective" T2 centers	3	6	10	15	20	25

Now we take the product of the number of such centers and the sum of each year and of the previous two years to get the installed useful capacity in any given year

The Effective sum of Regional T2's	2002	2003	2004	2005	2006	2007	
CPU scheduled	1	6	30	88	204	475	kSI95
CPU analysis	1	2.108696	30	83	187	425	kSI95
Disk	6	28	180	507	1169	2500	Tbytes
Active tape	0	1	0	0	0	0	Tbytes
Archive tape	5	23	90	302	651	1250	Tbytes
Tape I/O	10	10.52995	180	539	1310	2500	MB/s
Number of CPU boxes	22	100	387	824	1434	2265	

Finally we sum the computing power available to CMS:

Total computing power available to CMS	2002	2003	2004	2005	2006	2007	
CPU scheduled	13	36	125	315	583	1158	kSI95
CPU analysis	10	24	99	305	556	1090	kSI95
Disk	33	99	452	1294	2519	4797	Tbytes
Active tape	92	209	503	1741	2930	4490	Tbytes
Archive tape	122	296	737	2252	3785	6047	Tbytes
Tape I/O	62	132	476	1557	3153	5300	MB/s
Number of CPU boxes	224	553	1444	3004	4204	5659	