COCOTIME Requested IT Services for 2002

Version 2.6

17th December, 2001

Editor Tim Smith

Contents

1 Changes

COCOTIME	NA49 missed out 1TB disk server	+25
	CAST wanted 1TB usable not a 1TB server, so extra server	+12
	Oracle 9iRAC extra networking	+20
	LHCb networking (completed estimation)	+40
	Serco rescaling correction to include piquet costs also	-68
	LTO moved to LHC prototype	-30
	CUT: Disk contingency	-70
	CUT: CPU contingency	-100
	CUT: SUNDEV	-60
	CUT: Import/Export CD writer	-10
	CUT: PC remote access	-100
	CUT: GPFS	-100
		-441
LHC Prototype	Added 200k for LTO robot and drives (tender in progress)	+200
	LTO moved from general tape infrastructure	+30
	Serco rescaling correction to include piquet costs also	-18
		+212

DELPHI would like another disk server not mentioned in request
OPAL, need to replace 2TB of disks from shift6

2 Financial Summary

This paper presents a plan to fulfil the requests of the collaborations, as expressed in their replies to the COCOTIME questionnaire, for computing facilities in 2002. In the case of the larger experiments, the details of the requests have been discussed with the experiments and referees, and for the experiments with smaller requests, the IT link man clarified items individually. Table 1 gives an overview of the estimated costs for 2002 if all of the requests are satisfied. The equivalent costs for 2001 are given for comparison.

Table 1 — LHC Project View of Costs

		2001	2002
Direct Physics			
	LHC experiments	647	866
	non-LHC experiments	1,367	1,360
Infrastructure			
	LHC experiments	2,538	2,100
	non-LHC experiments	1,525	1,260
Sub-Totals		**6,077**	**5,586**
LHC Prototype			
	Prototype	1,282	3,649
Totals		**7,359**	**9,235**

A number of measures have been applied in order to respond to the cuts imposed as a result of the review of Cost to Completion of the LHC. These include:

- The elimination of all maintenance from RISC hardware
- The anticipated retirement of RISC machines for ATLAS, ALICE and PARC by the beginning of 2002 and of the LEP experiments by the middle of 2002
- The scaling down of the systems administration costs to fit within a budget of 2MCHF

In addition the group operational expenses are no longer included in the paper since the services covered by the paper now span a large number of groups, whose budgets are only partially covered by this paper.

3 Introduction

The year 2001 was when we have been consolidating our services onto candidate technologies for the LHC. A key part of this is the RISC decommissioning program which is concentrating services on Linux and Solaris. As part of this, the AIX batch service RSBATCH was terminated as was the HP simulation service CSF. Other services have been downsized or restricted to specific communities (like LEP) requiring continued access for a limited time. The HiPPI network was phased out in favour of gigabit Ethernet. The IBM and COMPAQ tape robots were stopped to concentrate on the large STK robotic installations.

This paper presents an acquisition proposal to continue this strategy, concentrating on Linux CPU and Disk servers. More tables are included in the appendices to expose the cost calculations, and shown breakdowns of expenses per activity. In addition the infrastructure costs are detailed in order to give more accurate estimates, and as a consequence explicit contingency lines have been added to the infrastructure budget in order to accommodate the usual extra experiment requests during the year for unforeseen needs.

3.1 Capacity ratings

All previous planning papers have used 'CERN Units' (CU) for capacity ratings. This year we wished to take the first steps away from local measurements and use international standards instead. To this end, the SPEC CINT95 (Si95) was chosen, from the SPEC organisation (http://www.spec.org). In fact the better choice for the longer term would have been the SPEC CINT2000 (Si2k), since the Si95 was retired in June 2000. However only the most recent processor families have been bench marked on the Si2k scale, not the ones in service in the computer centre, and so a matrix of comparisons and scalings will have to be established to consistently quote everything in Si2k. Instead Si95 was chosen, and for the questionnaire, a simple scaling factor of 10 was proposed to be used from CU, which unfortunately turned out to be too inaccurate.

The factor of 10 between CU and Si95 has been propagated through many documents in the past years (for instance the MONARC planning papers). However it now seems that this number was based on direct comparisons between Si92 and CU combined with scaling factors between Si92 and early Si95 numbers. Measurements made on recent PC models find a factor of nearer 5. Hence in the COCOTIME questionnaire, the 800 MHz PIII quoted as 16 (the 160 CU from LSF divided by 10), should read 32. This 'CERN measurement' compares more favourably with the actual value quoted on the SPEC site of 37.6.

So for the purposes of this paper, as a stepping stone to using Si2k in the longer term, processor ratings are given in true Si95 from the SPEC site. Previous years requests for CU-hrs have been scaled using a factor of 5, and this years requests have been scaled up by 2 since they erroneous factor of 10 not 5 was quoted in the questionnaire!

4 Interactive Services

The main interactive platform for 2002 will continue to be Linux, with a smaller Solaris cluster, SUNDEV, offered as an alternative development and verification platform. The public interactive service configurations during 2001 are detailed in Table 2 together with figures which give an impression of the average usage. RSPLUS will not form part of the planning for 2002 since it will be retired at the end of 2001 and so is not shown. It should be noted that most dedicated work group services were merged into the PLUS services during 2001 and the plan is to merge the remainder during 2002. The only service remaining on maintenance during 2002 will be SUNDEV.

Table 2 — Interactive Service Configuration and usage

Service	Configuration	Monthly Users	logged-in at peaks	active in last 5 minutes average	active in last 5 minutes Peaks
LXPLUS	51 dual CPU PCs PIII@550-800 MHz				
SUNDEV[1]	9 dual CPU SUN 220R/280R servers				
HPPLUS[2]	5 dual CPU HP J2240 servers				
DXPLUS[2]	5 dual CPU DEC PWS500 servers				

Notes:

1. SUNDEV is reserved primarily for LHC experiments, as a development facility
2. HPPLUS and DXPLUS are reserved primarily for LEP experiments

4.1 LXPLUS expansion

The set of responses to the questionnaires indicate that the total number of users requiring access to LXPLUS will go up by 10% and that on average the fraction of active users will rise by 25%. Since some of the load will be reduced by the creation of an LXBUILD cluster (described below), we estimate that an increase of 10% capacity seems reasonable. The current capacity having taken into account additions from merging in the remaining CMS and ALICE dedicated interactive facilities, is 3000 Si95. So a 10% increase represents 4 x 1GHz machines.

4.2 SUNDEV expansion

The questionnaire responses indicate a very large increase in the Solaris developer community in 2002, with 80% more users requiring access and 55% more active users. Due to the large uncertainty in these numbers the suggested strategy would be to expand by 20% in 1H2002 and optionally another 20% in 2H2002. The current capacity is estimated to be 400 Si95. So a 20% increase represents 2 x 900 MHz machines

4.3 Build services

In general the usage of dual processor Linux boxes as interactive nodes supporting up to 50 simultaneous sessions seems to have been successful. However it is found that compilations which consume a lot of resources for a considerable time are becoming more prevalent and affect other users of a node. Therefore we propose to create a new build facility without interactive users, on which such compilations can be launched. Compilations will be launched using the LSF 'lsrun' command. The configuration of the machines will be optimised for the purpose with larger memory, as well as larger AFS cache. A 5 node cluster is proposed as the proof of principle for 2002.

4.4 Summary

The detailed breakdown of the 2002 Public Interactive Services budget is presented in Table 3. The PC server costing includes infrastructure estimates as detailed in Appendix I.

Table 3 — Public Interactive Services Budget

Item		kCHF
SUNDEV Maintenance		10
LXPLUS expansion		9
LXBUILD acquisition		25
SUNDEV expansion		60
Systems administration		130
Miscellaneous purchases and operations		20
	Total	254

5 Batch Services

5.1 Overview

The merging of dedicated Linux services into a combined single large LXBATCH facility has proved very successful operationally. It has greatly reduced the complexity, bringing benefits through the whole chain of operators, system administrators and service managers. The first stage of this merge has been completed for the majority of experiments, which now share common configurations, but are still logically separated by soft partitions in the batch queuing software. As expected, statistics show that such static partitions which have been sized for the peak load tend to be under utilised on average. This is where the full benefit of the harmonisation of configurations can be realised by sharing of resources between experiments. However such sharing is not trivial due to the extra requirement that experiments should have guaranteed access to their lent resources when they need them again.

During 2001 we were able to prototype and demonstrate the benefits of sharing by setting up an experiment production queue underneath 33% of the dedicated experiment nodes. This has been well used by MC production managers and has been demonstrated to increase the overall usage of LXBATCH resources. However this approach can not address the entire problem, since such queues are not appropriate for general user jobs, only large productions. The reason is that in order to guarantee the response for the allocated experiment, jobs running in the shared_prod queue are suspended when experiment jobs are submitted. This can leave jobs stranded for some considerable time until the load again drops.

The next steps in 2002 will be to develop a way of running the batch system fully in 'fairshare' mode, without logical partitions but still guaranteeing the correct share of the resources when required as well as a reasonable response to regain the resources after a period of under-utilisation by an experiment. We are investigating how to implement this scheme together with the company who supply the batch system. Additional ways to optimise the resources will also be investigated like over-subscription of CPUs and better compiler technology

5.2 Interpretation of Utilisation Statistics

Subjectively the batch system utilisation is in general high and the queues long. However objectively the statistics gathered so far do not report the full picture, and so need some careful interpretation. Figure 1 presents a summary of the utilisation of LXBATCH, shown to be around 50% before the creation of the shared_prod queues, and around 60% afterwards.

Focusing on the public partition (the experiment partitions are treated elsewhere), the overall CPU utilisation broken down by queue is shown in Figure 2. There it can be seen that 27% of the CPU time does not go on running jobs. The fact that the machines are non the less constantly busy is shown in Figure 3 which presents the queue length (always non-zero) and the running job activity. Initial investigations of where the remaining CPU is being used has come up with two causes, but further investigation is needed before concluding on their relative importance. Firstly the batch statistics do not include waiting for remote IO since this does not appear in the batch accounting of the local machines. This can be a significant contribution in our distributed architecture. Secondly processes run outside the batch system are not included, notably system processes (run under root) such as monitoring and alarm scripts.

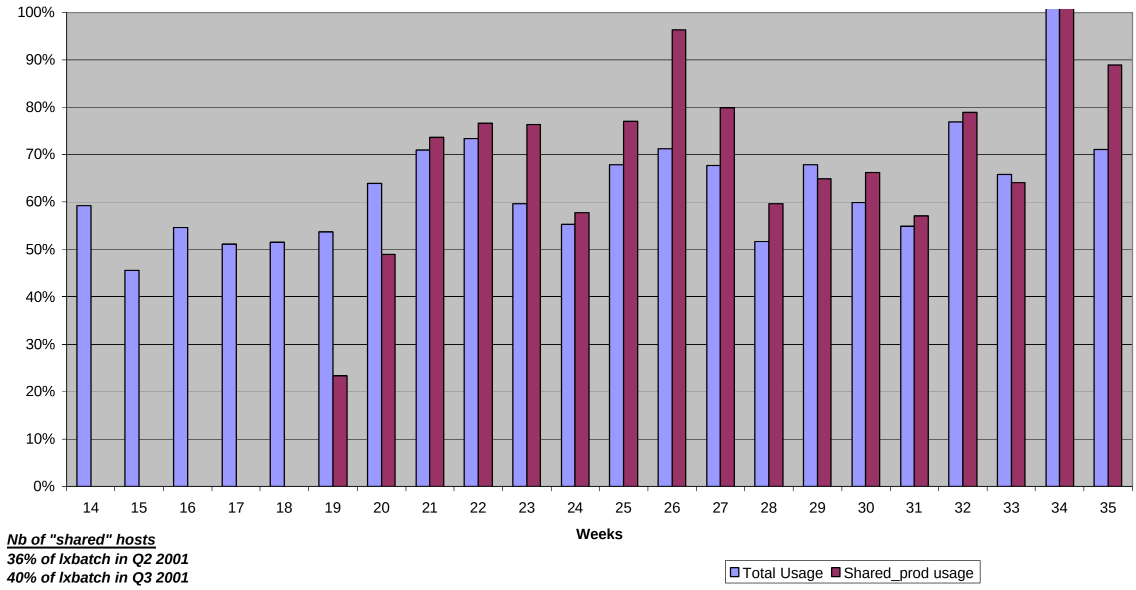

Figure 1 — LXBATCH and Shared_prod usage in Q2/Q3 2001

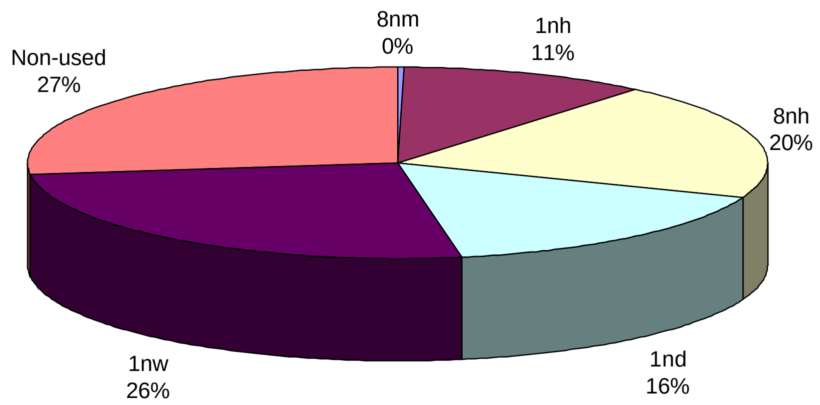

Figure 2 — CPU time per queue (September 2001)

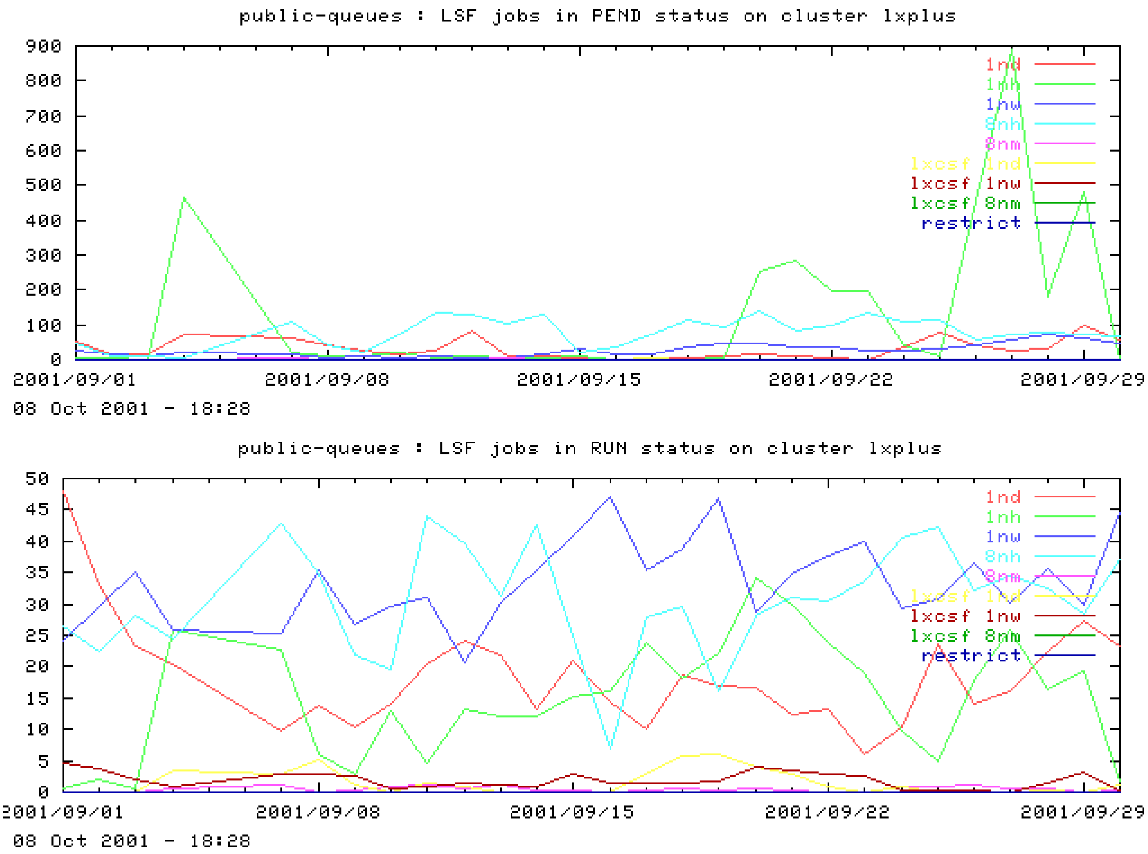

Figure 3 — LXBATCH queues: Pending and Running jobs

In conclusion, the initial analysis of LXBATCH usage shows the benefits of sharing resources, but that the overall utilisation can not be expected to reach 100%. Better accounting of full resource usage is also needed in order to perform better optimisation of disk to CPU server capacity, and to be able to judge the benefits of different schemes for mixing jobs in any future over-subscribing scheme.

5.3 Requests for 2002

Given that all batch requests are for Linux capacity, and that we operate in a mode of shared resources, together with the fact that the PCSF facility is no longer a separate entity, but is simply now a set of queues in LXBATCH, we now chose to combine the batch and simulation requests from all experiments into a combined request for LXBATCH expansion. We have thus almost come full circle back to the mainframe scheduling days. So in Table 4 we present a line for each experiment showing the total available capacity within or still outside LXBATCH combined with the requests for simulation and batch capacity in 2001 plus the increment requested in 2002.

Table 4 - Batch Requests

Experiment	2001 Capacity (kSi95hr)				2002 Capacity (kSi95hr)			
	Installed in LXBATCH on 2001/10/01	To be merged into LXBATCH after 2001/10/01	Simulation Request	Public Batch Request	Simulation Increase Requested	Public Batch Increase Requested	Total Increase	Total 2001 Capacity
Atlas	12,230	0	0	0	0	3,520	3,520	12,230
CMS	4,472	10,722	400	0	8,000	26,400	34,400	15,594
LHCb	6,710	0	3,600	0	14,400	2,240	16,640	10,310
Alice	3,844	2,522	200	0	0	4,800	4,800	6,566
Aleph	12,440	0	1,200	0	0	0	0	13,640
Delphi	9,644	0	400	600	6,400	0	6,400	10,644
L3	0	0	440	0	0	400	400	440
L3C	20,478	0	0	0	0	0	0	20,478
Opal	0	0	400	1,600	1,800	1,760	3,560	2,000
NA48	0	0	0	0	0	0	0	0
NA48-1	0	12,908	0	0	0	16,000	16,000	12,908
Compass	39,766	0	0	400	0	0	0	40,166
NA49	0	15,554	0	0	0	4,000	4,000	15,554
NA45	0	17,766	0	0	0	0	0	17,766
NA60	0	0	0	200	0	720	720	200
Chorus	0	0	40	2,000	0	0	0	2,040
WA98	0	0	4	56	0	0	0	60
Nomad	0	0	0	0	0	0	0	0
Opera	4,752	0	0	0	1,200	0	1,200	4,752
Dirac	0	0	2	4	1,600	240	1,840	6
TOTEM	0	0	8	8	0	0	0	16
Harp	0	0	0	0	0	0	0	0
Ntof	1,120	0	1,140	3,500	800	600	1,400	5,760
Cast	0	0	0	0	1,120	560	1,680	0
PARC	3,214	0	0	0	0	0	0	3,214
SL/AP	11,182	0	0	0	0	0	0	11,182
Public	21,666	0	0	0	0	2,166	2,166	21,666
CSF	7,828	0	0	0	0	0	0	7,828
Totals	159,346	59,472	7,834	8,368	27,320	63,406	98,726	235,020

As can be seen in the two summary columns at the right, the total increment represents 42% of the installed capacity, which could be satisfied with 122 x 1GHz machines.

The detailed breakdown of the 2002 Batch Service budget is presented in Table 5.

Table 5 — Batch Service Budget

Item		KCHF
Expansion of Linux capacity by 11,300 Si95		283
NA48 private funding		-23
Systems administration		1029
Miscellaneous purchases and operations		10
	Total	1299

6 The LHC Prototype Facility

6.1 Overview

In 2001 the LHC prototype facility consisted of 2 farms, LXSHARE for data challenges and production, and TestBed for EU DataGrid work. The reason for this distinction was that the GLOBUS software which forms the heart of EU DataGrid work was not integrated into the AFS environment of CERN general services and so the TestBed was set up for GLOBUS developments, and LXSHARE with the standard LXBATCH environment. For 2002 it is hoped that this distinction can be eliminated by 2H2002, by which time GLOBUS could be put on the entire LXSHARE farm, in order to satisfy the request of LHCb.

6.2 Predicted growth

The predicted growth of the LHC prototype facility as presented to the CERN Council in June 2001 (in paper CERN/2379Rev) is shown in Table 7. These numbers are taken as the target growth over the coming years in order to steadily advance the scale to that of the LHC. Such a growth will not only allow the experiments to do software scalability tests, but also the computer centre to scale the acquisitions, installation and management as well.

Table 6 — Foreseen Prototype Growth

		2001	2002	2003	2004	2001 October
CPU Capacity						
	Number of dual CPU systems	200	400	800	1,200	218
	Estimated Si95	13,000	33,000	85,000	158,000	12,000
Disk Storage						
	Capacity (TB of usable disk)	16	44	120	320	9
	Numbers of disks	200	400	800	1,600	190
	Performance (GB/s I/O rate)	1	5	?	?	<1
Tape Storage						
	Capacity (TB)	30	120	350	600	-
	Performance (MB/s I/O rate)	200	350	500	800	-

6.3 CPU servers

The ensemble of requests for LXSHARE capacity represent a much more complicated initial schedule than presented in last years paper. This scheduling is shown in Figure 4 where it can be seen that the requests peak at the end of each half year, since the experiment cycles tend to be in synchronism. This puts additional strain on the network, disk and tape infrastructure which must be sized for such peaks. In addition to the experiment requests, in 2002 IT division will also ask for 'machine development' time on the farm at regular intervals throughout the year in order to run realistic functionality and scalability tests of the fabric management software being developed. We assume COCOTIME will arrange the requests to fit in the 400 foreseen, and therefore price this scenario.

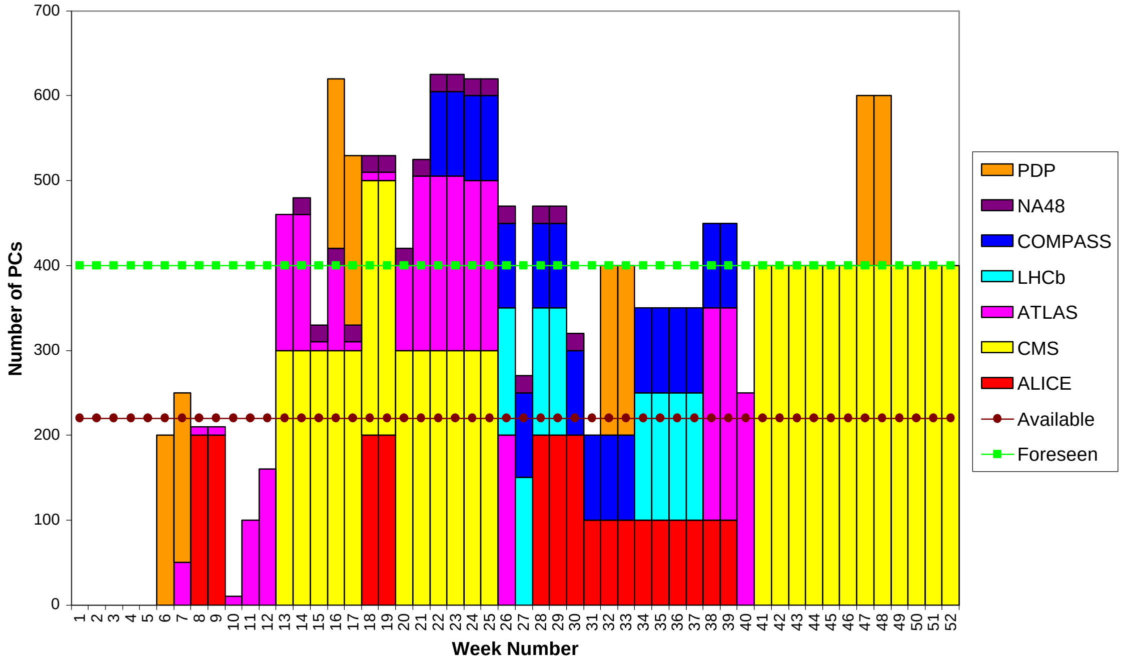

Figure 4 — Proposed LXSHARE scheduling for 2002

6.4 Disk servers

The foreseen disk server capacity for 2002 is 44TB. We already have 9 TB distributed over 11 servers, and we are in the process of acquiring 19 bare servers to increase the disk to network performance ratio. Therefore an additional 35 TB needs to be acquired in 2002. The configuration currently under tender is estimated to have 500GB of disk capacity per server, thus 70 nodes would be required. This, together with the existing servers is considered sufficient to answer the needs of the LHC experiment data challenges for 2002. These nodes are costed at 10.5kCHF each, and the additional gigabit infrastructure they require is costed at 128kCHF.

In addition to the dedicated disk servers, we propose to prototype an alternative architecture in 2002, having disks distributed across the CPU nodes. To do a realistic test in this configuration would require equipping 100 nodes with 200 GB of local disk space, costed at 1.5kCHF per server.

6.5 Objectivity servers

In order to provide addition lock, journal and federation server capacity during Data Challenges to experiments relying on Objectivity, 2 new fast reliable SUN servers would be required.

6.6 Tape drives, media and servers

In order to reach the foreseen 350 MB/s in 2002, it is estimated that 20 new drives should be dedicated to the LXSHARE facility. Anticipating the technology currently under tender (high quality, high reliability linear technology), each drive should be connected through Fibre Channel to a single tape server, and all tape servers should be on gigabit infrastructure. Costing estimates are based on 45kCHF drives, 3kCHF servers with 2kCHF fibre channel boards and 36kCHF for network infrastructure. In addition the maximum foreseen need for media is 1000 tapes for the ALICE data challenge.

Also in the area of tape drives we are starting an acquisition process for high quality, high reliability linear technology tape drives to satisfy the requirements of the LHC Computing Grid project for its first phase. This requires a sustained data recording rate to tape of 200MB/sec in 2002 rising to 500MB/sec in 2003. We believe it is too early to rely on the cheaper mid-range technology of LTO for this but will be acquiring a small robot with several LTO vendor tape drives to intensively test this new technology and also support import/export.

6.7 Summary

The detailed breakdown of the 2002 LHC Prototype budget is presented in Table 7.

Table 7 — LHC Prototype Budget

Item		KCHF
Expansion of LXSHARE (by 182) to 400 nodes		422
Disk servers		862
Distributed disks		150
Objectivity SUN servers		20
Systems administration		596
Tape drives and servers		1000
Tape media		150
LTO media and SuperDLT drive and media		30
LTO robot and drives		200
GigaBit Network Infrastructure		218
	Total	3649

7 Dedicated Services

7.1 Comments on the requests of the major experiments

ALICE Assume all HPs phased out by the end of 2001. Request for ftp server to be satisfied by WACDR service.

ATLAS Assume all HPs phased out by the end of 2001

CMS Remove 4 old weak suns without replacement. To split functionality from shift19,20, propose shift20 becomes disk server only, and production stager goes to new Linux server, shift19 goes to user disk server only and user stager, and buy a new standard lock server for user federation tests (in addition to the requested official user federation server). 6 Linux servers to be replaced on IT request by EIDE

LHCb During their run on the LXSHARE, require GLOBUS software on all machines. We will try and meet the need for build servers by the new LXBUILD central service.

ALEPH Once AFS space is available for the transfer of scratch functionality, and space is available for the transfer of working group disk functionality, shift50 should be able to be switched off.

NA48 50% of 20 extra 1GHz batch machines will be paid by outside lab funding as well as 50% of 4 extra disk servers

NA49 10 Gigabit Ethernet interfaces to be added to na49i interactive machines

7.2 Disk Servers

Since all interactive and batch acquisitions are now to be in the Public Services, the only remaining dedicated resources are the experiment disk servers. The requests for additional disk server capacity are summarised in Table 8. The costing includes infrastructure estimates as detailed in Appendix I.

Table 8 – Disk Server Requests

Experiment	2001 Capacity (TB)	Requested increase (TB)	Physical configuration	Costs kCHF
Atlas	2.8	4.5	9 x 1.0	111
CMS	4.7	3.5	7 x 1.0	86
LHCb	1.9	1	2 x 1.0	25
Alice	1.4	4.5	9 x 1.0	111
L3	4.5	0.75	2 x 1.0	25
Opal	5.3	2	4 x 1.0	50
NA48-1		2	4 x 1.0	50
Compass		2	4 x 1.0	50
NA49		2	4 x 1.0	50
Opera		0.75	2 x 1.0	25
Ntof		1	2 x 1.0	25
HARP		0.75	2 x 1.0	25
Cast		1	2 x 1.0	25

Totals	25.75	53 machines	658

7.3 Disk server replacements

The Linux SCSI disk servers we have in production for NA45, NA48 and COMPASS were some of the first models which we tried, and after a few years service, are now starting to cause operations troubles, suffering from numerous failures. We therefore would like to replace them by standard EIDE disk servers with mirrored disks. This would necessitate 3 x 500 GB for NA45, 6 x 500 GB for NA48 and 4 x 500 GB for COMPASS.

Additionally a couple of the earlier examples of EIDE disk servers we set up did not have mirrored disks. Again for operational reasons, although these have not yet caused too much trouble, we would like to make the systems more uniformly standard by adding extra servers to make equivalent mirrored capacity. This would entail 6 new servers for ALEPH, NA48 and NA49.

7.4 Objectivity Servers

CMS request an expansion, and PDP request a rationalisation, of their objectivity server setup, concentrating on SUN servers for the most critical lock, federation and journal files serving. Such a choice is not only well motivated within CERN, but has strong support from operational measurements in SLAC. The total would be 3 new fast reliable SUN servers.

7.5 Dedicated Services Summary

The detailed breakdown of the 2002 Dedicated Services budget, including maintenance and administration costs, is presented in Table 9.

Table 9 — Dedicated Services Budget

Item		kCHF
Disk Server acquisitions		658
NA48 private contribution		-25
Replacement for SCSI servers		161
Eliminating non-mirrored configurations		75
Objectivity SUN servers		30
SHIFT Maintenance		0
S/W licenses & support		25
Systems Administration		135
	Total	1059

8 Oracle Services

To satisfy the requests for production Oracle services from the 4 LHC experiments, we plan to acquire a Sun cluster, initally 2-node, but with the capability for further expansion. This today implies fibre-channel storage (350GB of mirrored storage). This cluster would run Oracle 8i – the same as all other existing production Oracle servers - and would host applications such as the ATLAS LAr calorimeter, ALICE book-keeping, CMS end-cap, CMS Cristal needs, LHCb conditions DB and book-keeping.

To meet CMS' request for an Oracle 9i test environment, we propose an additional cluster with 1TB of storage. A third cluster, Linux-based, is proposed to evaluate Oracle 9i's new clustering capabilities, in particular on commodity processors and storage. A ten-node cluster with at least 1TB of storage is required, together with a dedicated interconnect (GBit Ethernet) for cluster traffic, a 2nd for I/O, as well as a standard network connection for the rest.

Table 10 — Oracle Servers Budget

Item		kCHF
2 x 280R Sun Servers (8i cluster)		60
Storage for 8i		50
2 x 280R Sun Servers (9i cluster)		60
Storage for 9i		100
10 node 9iRAC cluster		50
	Total	320

9 Central Data Recording

High data rate experiments have dedicated CDR setups. Lower rate experiments and test beams are all catered for by a set of 4 Public CDR machines acquired and configured in 2001. These provide a reliable and robust 24 hour service, writing data into CASTOR. The requests for CDR services in 2002 are summarised in Table 11 and it is estimated that no expansion of the current service is necessary to cope with the declared load.

Table 11 - CDR Requests

Experiment	Location	Raw Data Rate (MB/s)	Time Profile
Atlas	H6/H8	5	April-October
CMS	H2A/H2B	4	April-October
	H4	5	April-October
	X5B/X5C	3	April-October
LHCb		Same as 2001	
Alice	T11 (T9/T10)	1	PS
	X5 (H4/H6)	1	SPS
NA48-1		30	April-July
Compass	SPS M2	35	May-September
	Test Beams	1	Spring
NA60	B918	5.5	May + Oct/Nov
Dirac		1	6 months
Harp	PS East Hall	5	April-July
Ntof	ISR (PS)	15	200 days
Cast	SPS	1	Continuous

10 Data Management

10.1 AFS Services

During 2002 we will complete the migration of personal scratch space off experiment dedicated local or NFS mounted disks onto Linux PCs configured with small RAID arrays and migrate the larger project scratch space to the Sun scratch servers, also now RAID. This will permit the final closedown of some of the old proprietary Unix servers.

We will attempt to improve performance by buying 2002 servers as dual CPU models with gigabit ethernet as both are limiting peak performance. The current servers could later be upgraded. On the software side we will build an OpenAFS server test cell to evaluate total migration to OpenAFS (Linux clients are already OpenAFS). This will be done in collaboration with other HEP sites. We will also continue to follow the evolution of alternative global file systems.

10.2 AFS Requests for 2002

The COCOTIME requests listed in Table 12 require an additional 1140 GB of AFS project space for 2002 (illustrated in Figure 5). We estimate that an additional 180 GB of AFS Home Directory space will be required in 2002 based on the growth pattern (see Figure 6). The current usage of Home Directory space is shown in Figure 7 (overall) and Figure 8 (per user).

Table 12 - AFS Project Space Requests

Experiment	Prior Capacity GB	Increase Requested GB
Atlas	400	200
CMS	275	150
LHCb	210	100
Alice	80	25
Aleph	82	20
Delphi	85	15
L3	53	50
Opal	225	100
NA48	150	0
NA48-1	0	30
Compass	60	60
NA49	25	4
NA45	4	0
NA60	16	0
Chorus	60	0
Nomad	10	0
Opera	60	200
Dirac	50	35
Harp	30	20
Ntof	40	80
Cast	0	50
Totals	**1915**	**1139**
IT Services	360	
IT Parc	150	
Geant4	110	

Accelerators	125

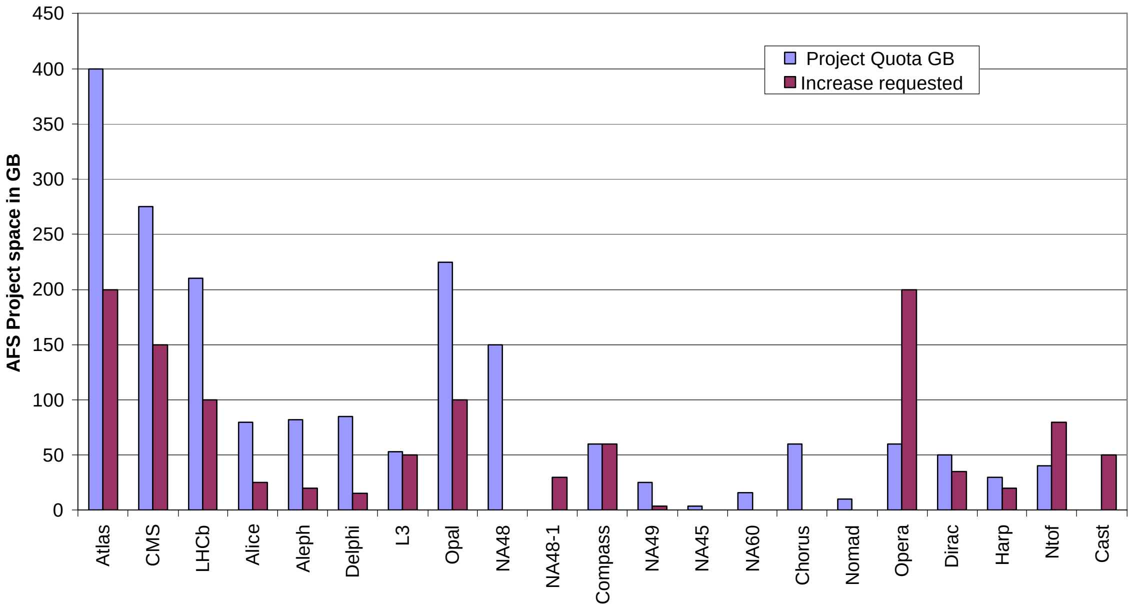

Figure 5 — AFS project space, requests and current usage

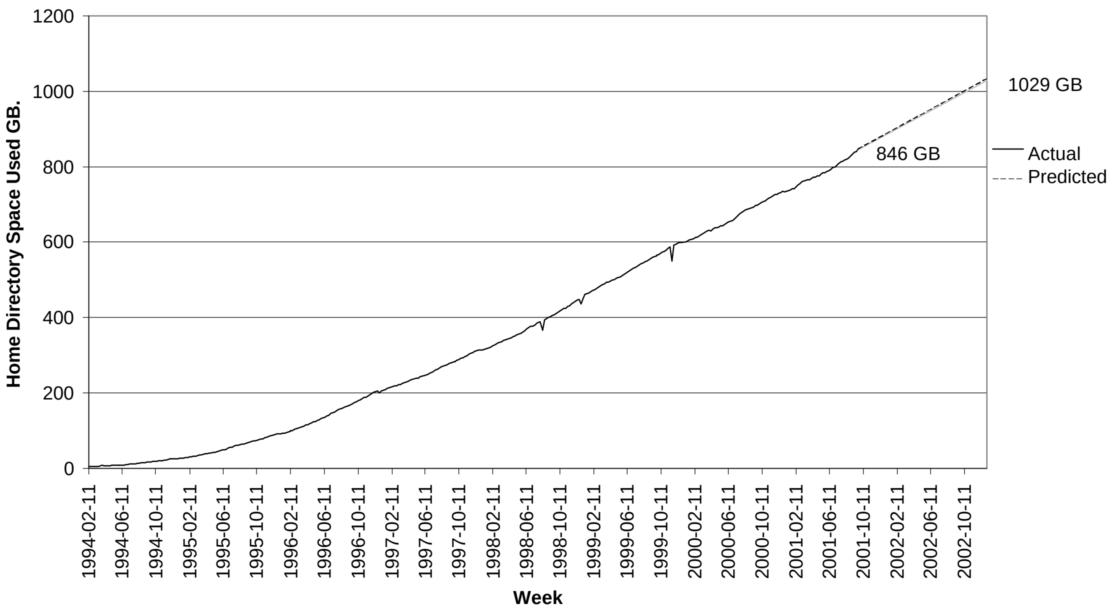

Figure 6 — Evolution of AFS Home Directory Usage since 1994 with Extrapolation to end-2002

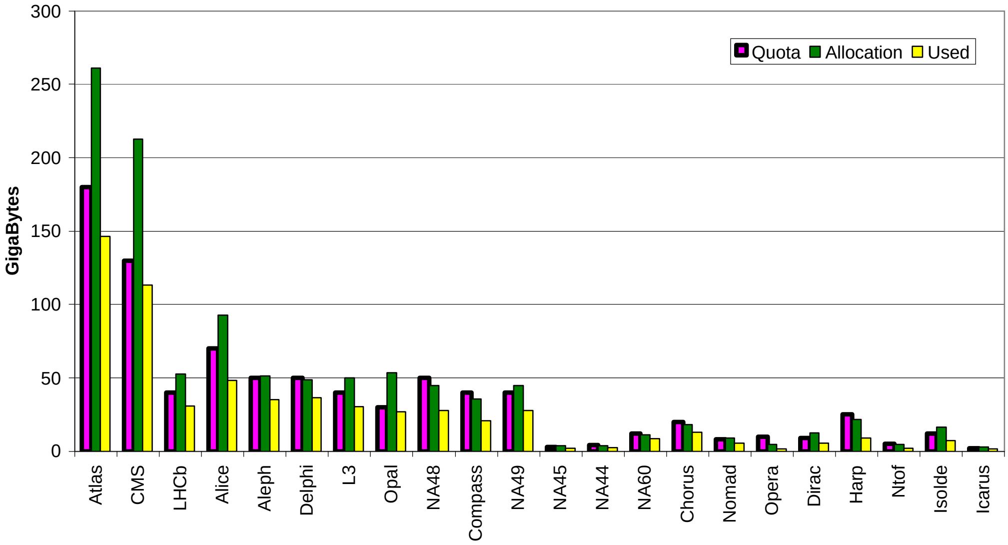

Figure 7 — Total AFS home directory space for major groups

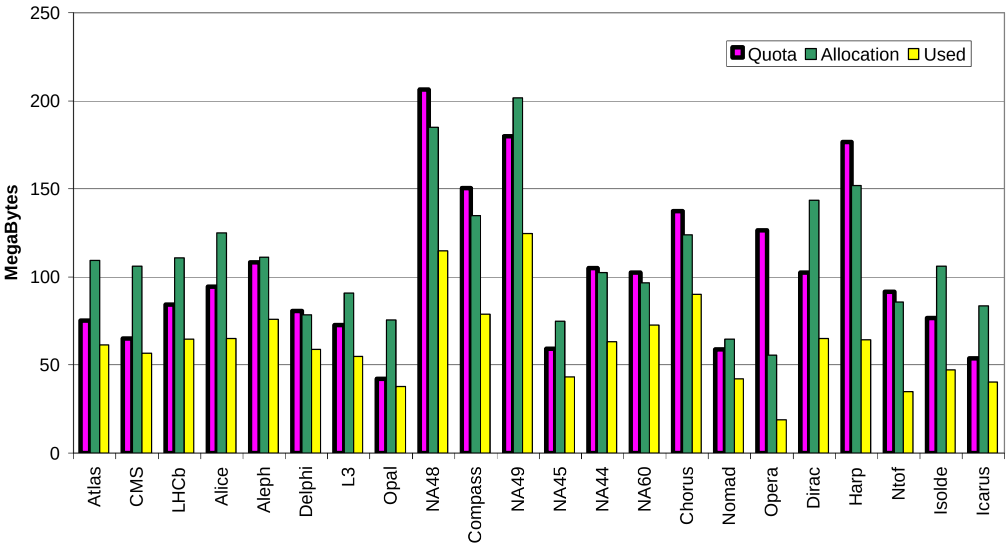

Figure 8 — Per-user AFS home directory space for major groups

10.3 Backup and Archive Services

IT does not backup the individual desktop expecting local disk to be used for scratch purposes only and permanent files to be either on the Windows servers, in AFS or under managed storage, either used from there or copied locally. We backup critical computer centre machines using Legato and group or experiment local servers using the IBM TSM product. TSM is also used for explicit user driven archiving (the 'pubarch' command). This separation is historical and we believe we can now decide on only one of these products. Much of the service is assured by outsourcing and we do not have the time or experience any more to decide how to proceed so have appointed a consultancy to advise us.

10.4 Managed storage

IT is adding a staff post to strengthen the CASTOR team and have put robustness and reliability as the highest priority until we are satisfied, followed by accounting, control and new functionality. Over the next few months we will be working with users and experiments to migrate remaining data from HPSS to CASTOR. Hence we budget for another 3 months of maintenance on the HPSS servers.

10.5 Import/Export service

Our strategy is to support a manageable number of low cost removable storage media devices for which there is sufficient demand. Most data will be sent over the network. We have already introduced a dedicated server for external ftp access to CASTOR (the wacdr001d machine) and budget on adding a second one. Next year we will try to add a CD-Rom reader/writer as requested by the HARP experiment. For this we estimate the PC and reliable burner will cost 10kCHF.

10.6 FATMEN/HEPDB

A replacement FATMEN/HEPDB server is now installed in the computer centre and the software is being migrated. Next year we will schedule data migration with the concerned experiments but those that can easily stop using FATMEN/HEPDB should do so. No extra expenditure is foreseen on this for 2002.

10.7 Data Management Summary

The detailed breakdown of the 2002 Data Management budget is presented in Table 13.

Table 13 — Data Management Budget

Item		kCHF
AFS Software Licenses		75
AFS Maintenance		135
AFS Project and Home Directory Space		250
Backup and Archive Services		211
Large User Files media		20
HPSS maintenance & operations		20
Public Objectivity servers maintenance		35
Import/Export Service		18
Systems administration		15
Maintenance		20
Miscellaneous		20
	Total	819

11 Tape Services

11.1 Tape Usage Statistics, October 2000 to September 2001

As can be seen from Figure 9 the most frequently mounted tape drive type changed twice over the past 12 months. Firstly at the end of 2000 the 9840s took over the lead from the Redwoods and then by 2Q2001 the 9940 became the most common.

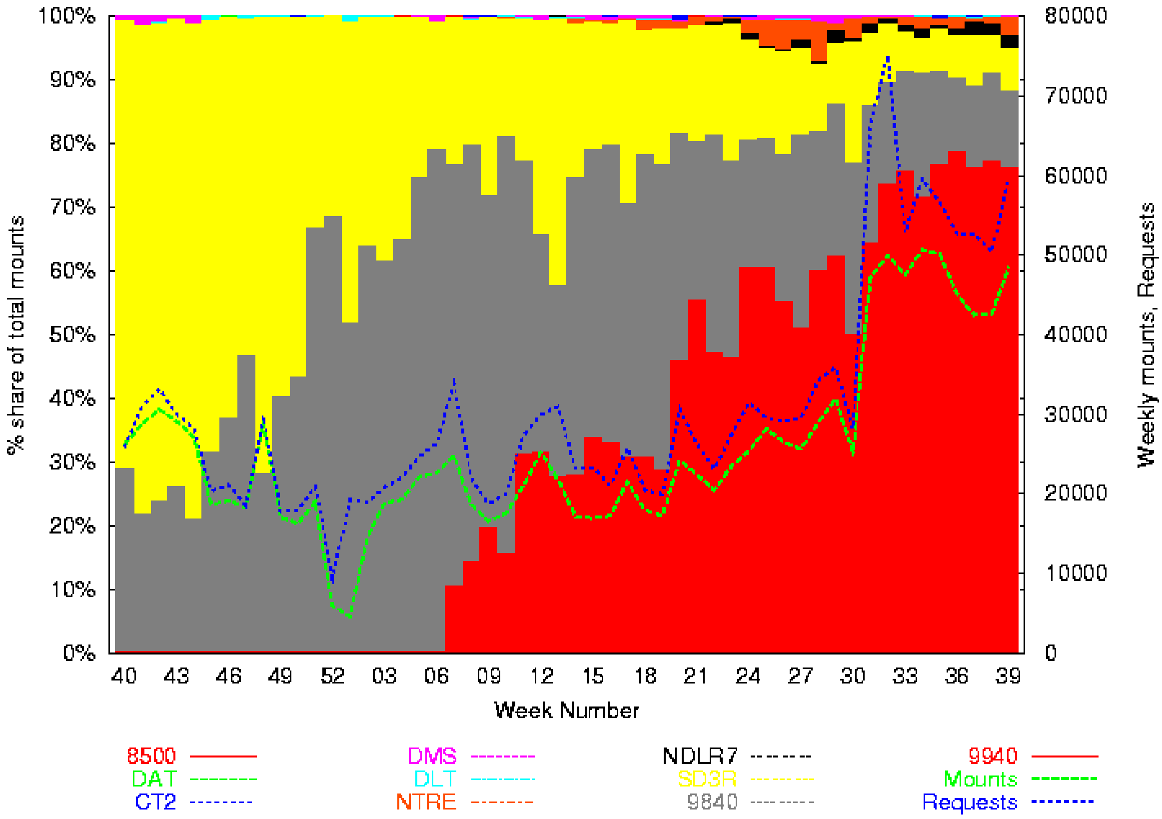

Figure 9 — Share of total tape mounts by drive type

A breakdown of tape mounts and data transfer volume by experiments is shown in Figure 10.

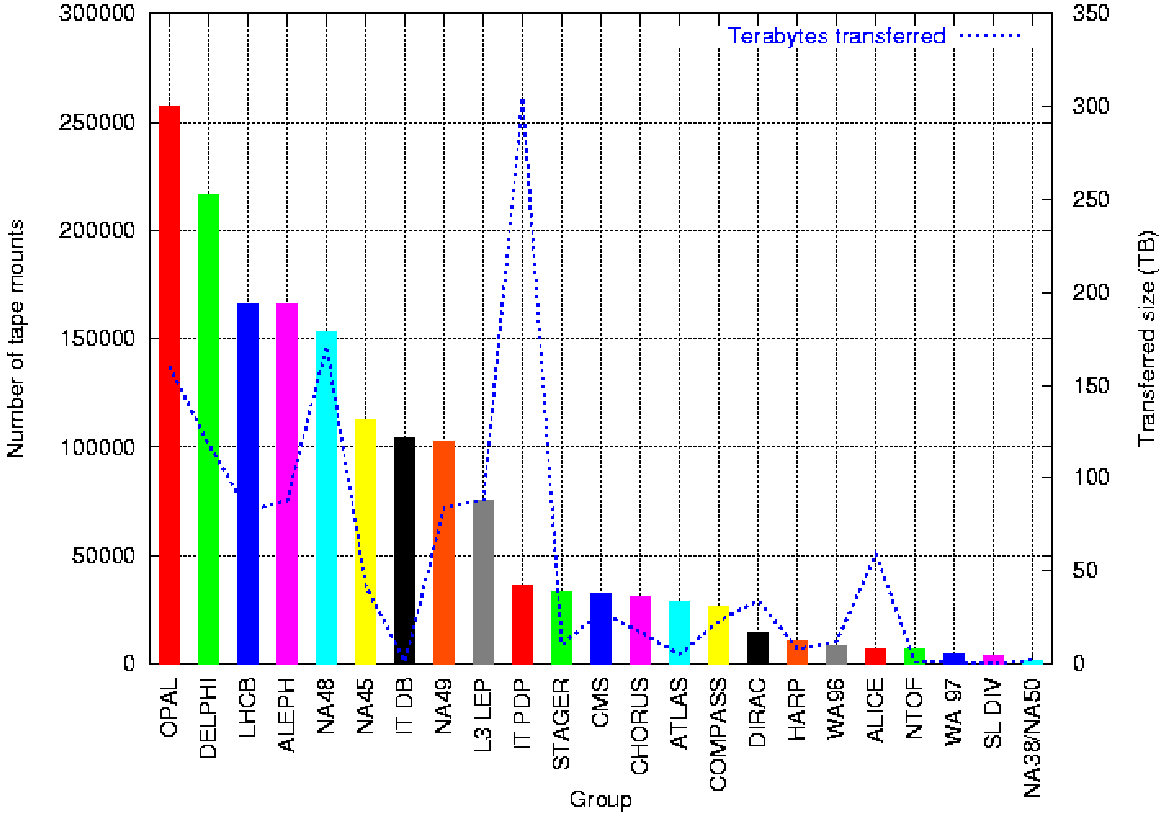

Figure 10 — Tape Mounts and Data Transferred by Group

During the course of 2001 the CASTOR service came into full production. The growth of data stored in CASTOR is presented in Figure 11 by data volume as well as number of files.

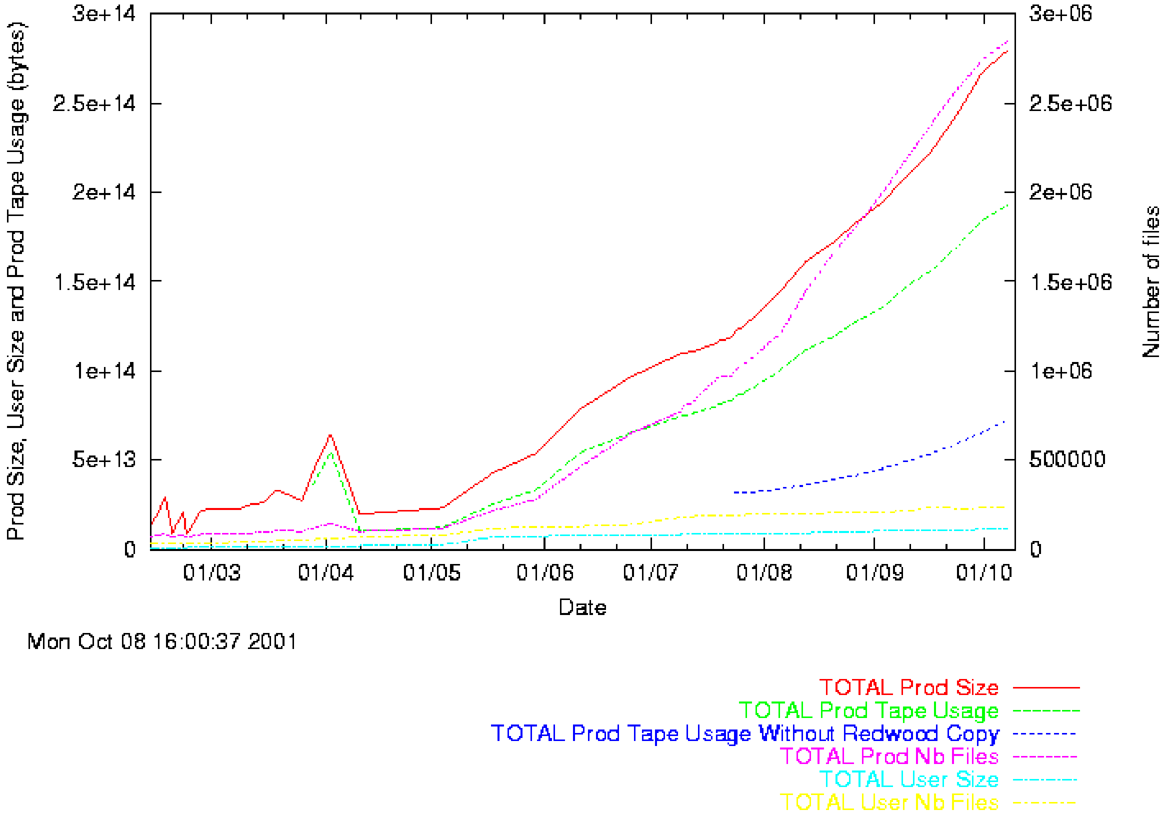

Figure 11 — Evolution of data stored in CASTOR

11.2 Tape Service plans for 2002

Two years ago we decided on a disaster avoidance scenario of having two physically separate tape robot installations and a new building, 613, has been prepared for this purpose. During October/November 2000 the STK silos in the tape vault have been moved to this new building and two new silos have been installed for next years tape capacity. We hence have two physically separate tape robot complexes each of 5 silos of 6000 slots.

At the end of 2000 we agreed a project with STK to replace the unreliable helical scan Redwood tape drives with the linear 9940 technology over two years. We now have 28 9940 drives and they have proved very successful. We have already reduced from 32 to 20 Redwood drives and will go down to 16 as part of the move to the new building. By the end of 2002 we plan to have only 4 Redwood drives and for them to be only lightly used. About half of the total number of Redwood cartridges have so far been copied to 9940, mostly as CASTOR files.

An important objective of this timing is to clear the vault to prepare it to be used for the LHC Computing Grid project. Many archive tapes are stored there, plus some DLT tapes, and experiments have already been requested to tell us what we can scrap or to where they should be moved. There will be a limited amount of space in IT for 'active' DLTs that can be cycled into the small 600 slot DLT robot. There will be no more manual DLT mounts and we anticipate only small numbers of manual mounts for 4mm and 8mm media.

Table 14 – Tertiary Storage Requests

Experiment	Prior Capacity TB	Increase Requested TB
Atlas		3
CMS		25
LHCb		10
Alice		17
Aleph		1
Delphi		
L3		5
Opal		5
NA48		
NA48-1		30
Compass		300
NA49		8
NA45		10
NA60		20
Chorus		
Nomad		
Opera		3
Dirac		20
Harp		30
Ntof		82
Cast		4
Totals		573

11.3 Tape Services Summary

The detailed breakdown of the 2002 Tape Services budget is presented in Table 15.

Table 15 — Tape Service Budget

Item	kCHF
Replacement media for Redwoods	150
Tape handling services	100
Tape operations	50
Drive, robot (Timberwolf) and server maintenance	480
S/W licenses & support	40
Total	820

12 PC R & D

To date only critical servers (such as disk or tape servers) have been equipped with the necessary infrastructure to access their consoles remotely. In 2002 we would like to prototype new fully functional low level infrastructures including consoles and remote reset capabilities across the centre (ie 2000 nodes) with a view to making this the standard for the LHC era.

All our PC acquisitions for production services to date have been based on the Intel Pentium Pro, II and III families. Other manufacturers offer competitive alternatives (such as AMD) and there are other 4 or 6 processor configurations all of which we would like to experiment with in the CERN environment to assess the advantages or disadvantages.

In 2001 we have started to acquire a GPFS test setup consisting of 4 servers to investigate the applicability of this technology in our intermediate file size storage domain. Assuming this to be as successful as it promises to be, we envisage and extension in 2002 to be able to run a proper pilot service.

The breakdown of the 2002 PC R&D budget is presented in Table 16.

Table 16 — PC R & D Budget

Item		kCHF
Remote PC access infrastructure prototyping		100
General PC R & D		100
GPFS pilot project		0
	Total	200

13 Infrastructure

Infrastructure includes the support facilities common to the different services such as the LSF license, service monitors, operator consoles and server contingencies. In addition there are central servers for ASIS and the Operating System installation servers.

The detailed breakdown of the 2002 Infrastructure budgets are presented in Table 17.

Table 17 — Infrastructure Budget

Item		kCHF
General servers		20
Systems administration (including piquet services)		95
Maintenance		20
Disk server contingency		30
CPU server contingency		100
OS install servers (production and development)		20
Replace ASIS FTP server		10
Software licenses and support		10
LSF maintenance		100
Other operations		50
Miscellaneous & contingency		30
	Total	485

14 Networking

The central networking requirements for 2002 are presented in Table 18.

Table 18 — Networking Budget

Item		KCHF
Additional SSR		120
Redundant Gigabit Ethernet infrastructure		50
	Total	170

The experiment requests expressed either in the COCOTIME 2002 questionnaire or the COCOTIME meetings with the experiments are summarised in Table 19 which also includes the comments and cost estimates. There were no extra network requests from other experiments

Table 19 – Experiment Networking Requests

	Request and Comments	kCHF
ALICE	Building 13 : 1 switch FastEthernet for daily library builds	4.5
	Wireless in meeting rooms in building 13 and 14 : Service not available for the moment.	0
ATLAS	Gigabit uplinks for DAQ in hall 887 for H6 and H8	25.4
	Gigabit uplink for lab in building 40. During COCOTIME meeting, it was pointed out that this type of exercise should be done in Computer Center	12.7
	Extension of the structured cabling in counting rooms in hall 887 (After review turned out not to be required)	0
CMS	No extra network requirements	0
LHCb	Structured cabling for hall 156	40
	72 FastEthernet connection and 2 GigabitEthernet connection (1 for uplink) in hall 156	24.5
ALEPH	No extra network requirements	0
DELPHI	No extra network requirements	0
L3	No extra network requirements	0
OPAL	Building 28 and around : 3 FastEthernet switches	15
NA60	24 FastEthernet connections plus GigabitEthernet uplink This request was first estimated at 10000 CHF. But we had to install more equipment which increased the cost up to 14700 CHF.	14.7
NTOF	10 Gigabit copper ports + 8 Gigabit optical ports + 1 FastEthernet switch This has been already installed and has been paid by IT/PDP.	14
CAST	FastEthernet switch in experimental area (imply installation of an optical fiber trunk)	9.4
		160.2

15 Licences

The experiment requests for licences are summarised in Table 20

Table 20 — Infrastructure Budget

ATLAS	NAG C, Objectivty
CMS	NAG C, Objectivity
LHCb	NAG C
DELPHI	GPHIGS, Licenced FORTRAN compilers (pgf77 currently)
OPAL	GPHIGS
NA45	MAFIA, Maxwell
NA48	Objectivity
COMPASS	Objectivity, Inusre / Sniff / Together / FrameMaker
HARP	NAG C

15.1 Comments

The requested NAG C licences are well within the expected profile agreed by Finance Committee.
The maintenance for GPHIGS licences will be continued as in previous years.
We believe the Objectivity requests are covered by the level of currently acquired licences.

Appendix I Assumptions

Table 21 — Assumptions used in Budget calculations

Item		CHF
PC CPU servers (1 GHz rated at 46 Si95 per processor)		2000
Average systems administration costs per CPU		1000
PC EIDE disk servers (1.0 TB raw, so 500 GB usable in mirrored configuration)		10500
Gigabit infrastructure for disk and tape servers		
	Share of 12 port switch populated with 7 (performance reasons)	1000
	NIC for PC	250
	Share of 2 uplinks (switch to Cabletron)	570
	Total	1820
Infrastructure for PC CPU servers		
	Network	200
	Racks, cables, etc	120
	Total	320

Appendix II Summaries by Experiment

Table 22 — Experiment Investment Summary

	LXBATCH	Disk Servers	Replace SCSI	Eliminate non-mirrored	Objectivity Servers	Totals
ALICE	14	111	0	0	0	125
ATLAS	9	111	0	0	0	120
CMS	100	86	0	0	30	216
LHCb	49	25	0	0	0	74
ALEPH	0	0	0	25	0	25
DELPHI	19	0	0	0	0	19
L3	0	25	0	0	0	25
OPAL	9	50	0	0	0	59
NA45	0	0	37	0	0	37
NA48	9	50	74	25	0	158
NA49	12	25	0	25	0	62
COMPASS	0	50	50	0	0	100
NA60	2	0	0	0	0	2
OPERA	2	25	0	0	0	27
NTOF	5	25	0	0	0	30
HARP	0	25	0	0	0	25
DIRAC	5	0	0	0	0	5
CAST	5	13	0	0	0	18
Totals	239	621	161	75	30	

Appendix III Summaries by Activity

Table 23 — Server Acquisition Summary

PC CPU Servers		**294**
	LXPLUS	9
	LXBATCH	103
	LHC Prototype	182
Disk Servers		**140**
	Dedicated	69
	Infrastructure	1
	LHC Prototype	70
SUN Servers		**12**
	SUNDEV	4
	Objectivity	5
	Oracle	3

Table 24 — Systems Administration Charges Summary

Section	Charge
PLUS	134
BATCH	1059
LHC Prototype	614
Dedicated	139
Data Management	25
Infrastructure	29
	2001

Appendix IV Processor and Disk Summaries

Table 25 — PC CPU Server Acquisition Summary

MHz	Number of Machines	Year of Acquisition
200	5	1997
300	70	1998
400	28	1998
450	115	1999
500	19	1999
550	121	1999
600	347	2000
800	267	2001
	944	

Table 26 — SCSI Disk Summary

Manufacturer	Model	Disk Capacity	#	Total Capacity (GB)	Year of Acquisition
IBM OEM	0664N1D	2	207	414	
DEC	RZ29B (C) DEC	4	24	96	
SEAGATE	ST15150W	4	23	92	
SGI	IBM DCHS04X	4.2	70	294	
SGI	IBM DFHSS4E	4.2	431	1810	
FUJITSU	MAB3091S	9	9	81	
IBM	DDRS39130	9	1	9	
SEAGATE	ST19171W	9	173	1557	
SEAGATE	ST410800W	9	276	2484	
COMPAQ	BD018122C0	18	6	108	
DEC	RZ1EF-CB (C) DEC	18	108	1944	
DEC	RZ2EA-LA (C) DEC	18	15	270	
SGI	IBM DGHS18Y	18	357	6426	
SEAGATE	ST423451W	21.6	159	3434	
HITACHI	DF400	36	5	180	
SEAGATE	ST136475LC	36	118	4248	
SEAGATE	ST150176LC	50	168	8400	2000
			2150	31,848	

Table 27 — EIDE Disk Summary

Manufacturer	Model	Disk Capacity	# of Disks	Total Capacity (GB)	Year of Acquisition
IBM	IBM-DJNA-352500	25.6	8	205	2000
IBM	IBM-DPTA 353750	37.5	20	750	2000
Maxtor	Maxtor 94098U6	41	20	820	2001
Maxtor	Maxtor 94098U8	41	6	246	2001
Seagate	ST150176LC	50	9	450	2001
IBM	IBM-DTLA 307075	71.6	752	53,843	2001
WDC	WD800BB-00BSA0	74.5	20	1,490	2001
WDC	WD800BB-00CCB0	74.5	196	14,602	2001
IBM	IBM-DTLA 307075	76.9	120	9,228	2001
Maxtor	Maxtor 98196H8	82	12	984	2001
			1163	82,618	

Appendix V CPU Server Configurations

Appendix VI Disk Server Configurations

Table 28 — Disk Server Configurations

	Host	Type	Manufacturer	Model	Capacity	#	Type Sum GB	Total GB
ALICE								2445.9
	aliceb01	SCSI	HPPBNSUP	ST136475LC	33.9	9	305.1	
	aliceb01	SCSI	SEAGATE	ST410800W	8.5	16	136.0	441.1
	alice001d	EIDE	IBM	IBM-DTLA 307075	71.6	14	1002.4	1002.4
	alice002d	EIDE	IBM	IBM-DTLA 307075	71.6	14	1002.4	1002.4
ATLAS								3569.5
	atlasd01	SCSI	SGI	IBM DGHS18Y	17.0	30	510.0	
	atlasd01	SCSI	DEC	RZ1EF-CB (C) DEC	17.0	7	119.0	
	atlasd01	SCSI	SEAGATE	ST136475LC	33.9	9	305.1	
	atlasd01	SCSI	SEAGATE	ST150176LC	46.6	10	466.0	
	atlasd01	SCSI	SEAGATE	ST19171W	8.5	13	110.5	
	atlasd01	SCSI	SEAGATE	ST410800W	8.5	44	374.0	1884.6
	atlobj01	SCSI	SEAGATE	ST423451W	21.6	4	86.4	86.4
	atlobj02	SCSI	HITACHI	DF400	33.3	5	166.5	166.5
	atlas001d	EIDE	IBM	IBM-DTLA 307075	71.6	20	1432.0	1432.0
CMS								6537.0
	cmsc01	SCSI	SGI	IBM DGHS18Y	17.0	36	612.0	
	cmsc01	SCSI	FUJITSU	MAB3091S SUN9.0G	8.4	5	42.0	
	cmsc01	SCSI	SEAGATE	ST136475LC	33.9	9	305.1	
	cmsc01	SCSI	SEAGATE	ST150176LC	46.6	5	233.0	1192.1
	shift19	SCSI	DEC	RZ1EF-CB (C) DEC	17.0	28	476.0	
	shift19	SCSI	SEAGATE	ST19171W	8.5	24	204.0	
	shift19	SCSI	SEAGATE	ST410800W	8.5	16	136.0	816.0
	shift20	SCSI	IBM	DDRS39130SUN9.0G	8.4	1	8.4	
	shift20	SCSI	FUJITSU	MAB3091S SUN9.0G	8.4	4	33.6	
	shift20	SCSI	SEAGATE	ST136475LC	33.9	9	305.1	
	shift20	SCSI	SEAGATE	ST150176LC	46.6	13	605.8	952.9
	cms001d	EIDE	WDC	WD800BB-00BSA0	74.5	20	1490.0	1490.0
	cms002d	EIDE	WDC	WD800BB-00CCB0	74.5	14	1043.0	1043.0
	cms003d	EIDE	WDC	WD800BB-00CCB0	74.5	14	1043.0	1043.0
LHCb								2942.8
	shd18	SCSI	SGI	IBM DGHS18Y	17.0	18	306.0	
	shd18	SCSI	SEAGATE	ST423451W	21.6	12	259.2	
	shd18	SCSI	SEAGATE	ST150176LC	46.6	8	372.8	938.0
	lhcb001d	EIDE	IBM	IBM-DTLA 307075	71.6	14	1002.4	1002.4
	lhcb002d	EIDE	IBM	IBM-DTLA 307075	71.6	14	1002.4	1002.4
ALEPH								6357.8
	shift50	SCSI	SGI	IBMA DFHSS4E 509	4.2	104	436.8	
	shift50	SCSI	DEC	RZ1EF-CB (C) DEC	17.0	59	1003.0	
	shift50	SCSI	DEC	RZ29B (C) DEC	4.0	18	72.0	
	shift50	SCSI	DEC	RZ2EA-LA (C) DEC	17.0	2	34.0	
	shift50	SCSI	SEAGATE	ST19171W	8.5	32	272.0	1817.8
	aleph001d	EIDE	IBM	IBM-DPTA 353750	37.5	20	750.0	750.0
	aleph002d	EIDE	Maxtor	Maxtor 94098U6	41.0	20	820.0	820.0
	aleph003d	EIDE	IBM	IBM-DTLA 307075	76.9	20	1538.0	1538.0
	aleph004d	EIDE	IBM	IBM-DTLA 307075	71.6	20	1432.0	1432.0
DELPHI								4618.2
	shd01	SCSI	SGI	IBM DGHS18Y	17.0	30	510.0	
	shd01	SCSI	SEAGATE	ST150176LC	46.6	7	326.2	
	shd01	SCSI	SEAGATE	ST19171W	8.5	56	476.0	
	shd01	SCSI	SEAGATE	ST410800W	8.5	52	442.0	1754.2
	delphi001d	EIDE	IBM	IBM-DTLA 307075	71.6	20	1432.0	1432.0
	delphi002d	EIDE	IBM	IBM-DTLA 307075	71.6	20	1432.0	1432.0
L3								5730.9

	shift3	SCSI	IBM OEM	0664N1D	1.9	97	184.3	
	shift3	SCSI	SGI	IBM DCHS04X	4.2	12	50.4	
	shift3	SCSI	SGI	IBM DFHSS4E	4.2	2	8.4	
	shift3	SCSI	SGI	IBM DGHS18Y	17.0	68	1156.0	
	shift3	SCSI	SGI	IBMA DFHSS4E 509	4.2	174	730.8	
	shift3	SCSI	SEAGATE	ST150176LC	46.6	10	466.0	
	shift3	SCSI	SEAGATE	ST15150W	4.0	23	92.0	
	shift3	SCSI	SEAGATE	ST410800W	8.5	16	136.0	
	shift3	SCSI	SEAGATE	ST423451W	21.6	20	432.0	3255.9
	l3001d	EIDE	IBM	IBM-DTLA 307075	71.6	20	1432.0	1432.0
	l3002d	EIDE	WDC	WD800BB-00CCB0	74.5	14	1043.0	1043.0
L3C								982.8
	l3c001d	EIDE	IBM	IBM-DJNA-352500	25.6	8	204.8	204.8
	l3c001d	EIDE	Maxtor	Maxtor 98196H8	82.0	4	328.0	328.0
	l3c001d	EIDE	Seagate	ST150176LC	50.0	9	450.0	450.0
OPAL								7509.5
	shift6	SCSI	IBM OEM	0664N1D	1.9	110	209.0	
	shift6	SCSI	SGI	IBM DCHS04X	4.2	58	243.6	
	shift6	SCSI	SGI	IBM DFHSS4E 506	4.2	1	4.2	
	shift6	SCSI	SGI	IBM DGHS18Y	17.0	19	323.0	
	shift6	SCSI	SGI	IBMA DFHSS4E 506	4.2	2	8.4	
	shift6	SCSI	SGI	IBMA DFHSS4E 509	4.2	148	621.6	
	shift6	SCSI	SEAGATE	ST136475LC	33.9	9	305.1	
	shift6	SCSI	SEAGATE	ST150176LC	46.6	18	838.8	
	shift6	SCSI	SEAGATE	ST19171W	8.5	8	68.0	
	shift6	SCSI	SEAGATE	ST410800W	8.5	48	408.0	3029.7
	opal001d	EIDE	IBM	IBM-DTLA 307075	71.6	20	1432.0	1432.0
	opal002d	EIDE	IBM	IBM-DTLA 307075	71.6	14	1002.4	1002.4
	opal003d	EIDE	IBM	IBM-DTLA 307075	71.6	14	1002.4	1002.4
	opal004d	EIDE	WDC	WD800BB-00CCB0	74.5	14	1043.0	1043.0
NA45								1513.6
	na45pc07	SCSI	SEAGATE	ST136475LC	34.4	4	137.6	137.6
	na45pc08	SCSI	SEAGATE	ST136475LC	34.4	5	172.0	172.0
	na45pc09	SCSI	SEAGATE	ST136475LC	34.4	4	137.6	137.6
	na45pc10	SCSI	SEAGATE	ST136475LC	34.4	5	172.0	172.0
	na45pc11	SCSI	SEAGATE	ST136475LC	34.4	4	137.6	
	na45pc11	SCSI	SEAGATE	ST150176LC	47.3	4	189.2	326.8
	na45pc12	SCSI	SEAGATE	ST136475LC	34.4	4	137.6	
	na45pc12	SCSI	SEAGATE	ST150176LC	47.3	4	189.2	326.8
	na45pc14	SCSI	DEC	RZ1EF-CB (C)	17.2	14	240.8	240.8
	na45001d	EIDE					0.0	0.0
	na45002d	EIDE					0.0	0.0
NA48								11759.0
	shd58	SCSI	SGI	IBM DGHS18Y	17.0	36	612.0	
	shd58	SCSI	SEAGATE	ST423451W	21.6	24	518.4	1130.4
	shd59	SCSI	SGI	IBM DGHS18Y	17.0	36	612.0	
	shd59	SCSI	SEAGATE	ST423451W	21.6	24	518.4	1130.4
	shd60	SCSI	SGI	IBM DGHS18Y	17.0	36	612.0	
	shd60	SCSI	SEAGATE	ST423451W	21.6	19	410.4	1022.4
	na48d001	SCSI	SEAGATE	ST136475LC	34.5	13	448.5	448.5
	na48d002	SCSI	SEAGATE	ST136475LC	34.5	18	621.0	621.0
	na48d003	SCSI	SEAGATE	ST150176LC	47.3	13	614.9	614.9
	na48d004	SCSI	SEAGATE	ST150176LC	47.3	9	425.7	425.7
	na48d005	SCSI	SEAGATE	ST150176LC	47.3	9	425.7	425.7
	na48d006	SCSI					0.0	0.0
	na48007d	EIDE	IBM	IBM-DTLA 307075	76.9	20	1538.0	1538.0
	na48010d	EIDE	IBM	IBM-DTLA 307075	76.9	20	1538.0	1538.0
	na48011d	EIDE	IBM	IBM-DTLA 307075	71.6	20	1432.0	1432.0
	na48012d	EIDE	IBM	IBM-DTLA 307075	71.6	20	1432.0	1432.0
COMPASS								4640.8
	ccf009d	SCSI	DEC	RZ2EA-LA (C)	17.2	6.0	103.2	103.2
	ccf010d	SCSI	DEC	RZ2EA-LA (C)	17.2	7.0	120.4	120.4

	ccf011d	SCSI	SEAGATE	ST150176LC	48.8	4.0	195.2	195.2
	ccf012d	SCSI	SEAGATE	ST150176LC	48.8	4.0	195.2	195.2
	ccf013d	SCSI	SEAGATE	ST150176LC	48.8	4.0	195.2	195.2
	ccf014d	SCSI	SEAGATE	ST150176LC	48.8	4.0	195.2	195.2
	ccf018d	SCSI	SEAGATE	ST423451W	21.9	4.0	87.6	87.6
	ccf019d	SCSI	SEAGATE	ST150176LC	48.8	4.0	195.2	195.2
	ccf020d	SCSI	SEAGATE	ST150176LC	48.8	4.0	195.2	195.2
	ccf021d	SCSI	SEAGATE	ST150176LC	48.8	4.0	195.2	195.2
	ccf022d	SCSI	SEAGATE	ST136475LC	35.5	8.0	284.0	284.0
	ccf023d	SCSI	SEAGATE	ST136475LC	35.5	8.0	284.0	284.0
	ccf024d	SCSI	SEAGATE	ST150176LC	48.8	4.0	195.2	195.2
	ccf025d	SCSI	SEAGATE	ST150176LC	48.8	4.0	195.2	195.2
	compass001d	EIDE	IBM	IBM-DTLA 307075	71.6	14	1002.4	1002.4
	compass002d	EIDE	IBM	IBM-DTLA 307075	71.6	14	1002.4	1002.4
CHORUS								2206.0
	shd14	SCSI	SEAGATE	ST150176LC	46.6	6	279.6	
	shd14	SCSI	SEAGATE	ST410800W	8.5	47	399.5	
	shd14	SCSI	SEAGATE	ST410800WSUN9.0G	8.5	1	8.5	
	shd14	SCSI	SEAGATE	ST423451W	21.6	4	86.4	774.0
	chorus001d	EIDE	IBM	IBM-DTLA 307075	71.6	20	1432.0	1432.0
NOMAD								641.6
	shd62	SCSI	SGI	IBM DGHS18Y	17.0	12	204.0	
	shd62	SCSI	DEC	RZ29B (C) DEC	4.0	6	24.0	
	shd62	SCSI	SEAGATE	ST19171W	8.5	8	68.0	
	shd62	SCSI	SEAGATE	ST423451W	21.6	16	345.6	641.6
NA49								5926.8
	shd09	SCSI	SGI	IBM DGHS18Y	17.0	36	612.0	
	shd09	SCSI	SEAGATE	ST150176LC	46.6	16	745.6	
	shd09	SCSI	SEAGATE	ST19171W	8.5	32	272.0	
	shd09	SCSI	SEAGATE	ST410800W	8.5	32	272.0	
	shd09	SCSI	SEAGATE	ST423451W	21.6	12	259.2	2160.8
	na49001d	EIDE	Maxtor	Maxtor 98196H8	82.0	8	656.0	656.0
	na49001d	EIDE	Maxtor	Maxtor 94098U8	41.0	6	246.0	246.0
	na49002d	EIDE	IBM	IBM-DTLA 307075	71.6	20	1432.0	1432.0
	na49003d	EIDE	IBM	IBM-DTLA 307075	71.6	20	1432.0	1432.0
HARP								5339.0
	harp001d	EIDE	IBM	IBM-DTLA 307075	71.6	20	1432.0	1432.0
	harp002d	EIDE	IBM	IBM-DTLA 307075	71.6	20	1432.0	1432.0
	harp003d	EIDE	IBM	IBM-DTLA 307075	71.6	20	1432.0	1432.0
	harp004d	EIDE	WDC	WD800BB-00CCB0	74.5	14	1043.0	1043.0
nToF								1432.0
	ntof001d	EIDE	IBM	IBM-DTLA 307075	71.6	20	1432.0	1432.0
NA60								1043.0
	na60001d	EIDE	WDC	WD800BB-00CCB0	74.5	14	1043.0	1043.0
OPERA								1432.0
	opera001d	EIDE	IBM	IBM-DTLA 307075	71.6	20	1432.0	1432.0
SLAP								1432.0
	slap001d	SCSI					0.0	0.0
	slap002d	EIDE	IBM	IBM-DTLA 307075	71.6	20	1432.0	1432.0
LHC Prototype								13807.0
	tbed001d	EIDE	IBM	IBM-DTLA 307075	71.6	20	1432.0	1432.0
	tbed002d	EIDE	IBM	IBM-DTLA 307075	71.6	20	1432.0	1432.0
	tbed003d	EIDE	IBM	IBM-DTLA 307075	71.6	20	1432.0	1432.0
	tbed004d	EIDE	IBM	IBM-DTLA 307075	71.6	20	1432.0	1432.0
	tbed005d	EIDE	IBM	IBM-DTLA 307075	71.6	20	1432.0	1432.0
	tbed006d	EIDE	IBM	IBM-DTLA 307075	71.6	20	1432.0	1432.0
	tbed007d	EIDE	WDC	WD800BB-00CCB0	74.5	14	1043.0	1043.0
	tbed008d	EIDE	WDC	WD800BB-00CCB0	74.5	14	1043.0	1043.0
	tbed009d	EIDE	WDC	WD800BB-00CCB0	74.5	14	1043.0	1043.0
	tbed010d	EIDE	WDC	WD800BB-00CCB0	74.5	14	1043.0	1043.0
	tbed011d	EIDE	WDC	WD800BB-00CCB0	74.5	14	1043.0	1043.0
Public Stage								6561.0

	shd17	SCSI	COMPAQ	BD018122C0	17.0	6	102.0	
	shd17	SCSI	SEAGATE	ST410800W	8.5	3	25.5	
	shd17	SCSI	SEAGATE	ST410800WSUN9.0G	8.5	1	8.5	
	shd17	SCSI	SEAGATE	ST423451W	21.6	20	432.0	466.0
	pub001d	EIDE	IBM	IBM-DTLA 307075	71.6	20	1432.0	1432.0
	pub002d	EIDE	IBM	IBM-DTLA 307075	71.6	20	1432.0	1432.0
	pub003d	EIDE	WDC	WD800BB-00CCB0	74.5	14	1043.0	1043.0
	pub004d	EIDE	WDC	WD800BB-00CCB0	74.5	14	1043.0	1043.0
	pub005d	EIDE	WDC	WD800BB-00CCB0	74.5	14	1043.0	1043.0
Infrastructure								4296.0
	migr001d	EIDE	IBM	IBM-DTLA 307075	71.6	20	1432.0	1432.0
	migr002d	EIDE	IBM	IBM-DTLA 307075	71.6	20	1432.0	1432.0
	migr003d	EIDE	IBM	IBM-DTLA 307075	71.6	20	1432.0	1432.0
	migr004d	EIDE	IBM	IBM-DTLA 307075	71.6	20	1432.0	1432.0
	pem001d	EIDE	IBM	IBM-DTLA 307075	71.6	20	1432.0	1432.0
	pubcdr001d	EIDE	IBM	IBM-DTLA 307075	76.9	20	1538.0	1538.0
	pubcdr002d	EIDE	IBM	IBM-DTLA 307075	76.9	20	1538.0	1538.0
	pubcdr003d	EIDE	IBM	IBM-DTLA 307075	76.9	20	1538.0	1538.0
	pubcdr004d	EIDE	IBM	IBM-DTLA 307075	71.6	20	1432.0	1432.0
	wacdr001d	EIDE	IBM	IBM-DTLA 307075	71.6	20	1432.0	1432.0

Table 29 — Summary of 2002 Cost Estimates

Summary of 1996-2001 Costs and 2002 Estimates

Materials and services only (no CERN personnel costs are included). All figures in kCHF

	←		Costs			→	Budget	Increase
	1996	1997	1998	1999	2000	2001	2002	2001-2002
Service						(est.)		
Batch Services[1]	209	130	245	109	177	178	1,299	1,121
Interactive Services[1]		363	410	133	86	398	254	-144
LHC Prototype					1,149	1,282	3,649	2,367
Simulation	75	120	210	152	102	122	0	-122
Dedicated Services[2]	3,126	3,385	3,663	3,232	2,858	2,155	1,059	-1,096
Oracle Services[6]							320	320
Data Management[4]		144	580	908	939	1,046	819	-227
Tape Services	2,105	2,189	2,492	2,569	1,736	1,440	820	-620
CORE Networking	316	167	291	157	100	47	170	123
CORE Infrastructure	626	420	480	459	315	338	485	147
Unix Support Services[5]					248	212	0	-212
R & D	60	267	156	335	238	141	200	59
Experiment Networking[6]							160	160
IBM Mainframe & SP2 financing	2,813	2,300						
COCOTIME Total	**9,330**	**9,485**	**8,527**	**8,054**	**7,948**	**7,359**	**9,235**	**1,876**

Notes:

1. Public Batch and Interactive services were partly financed by the Mainframe and SP2 Project in 1995-97
2. Work Group Servers included from 1997; CS2 & PIAF included in 1996-99
3. Systems administration costs introduced progressively from 1996
4. AFS Services included from 1999
5. Additions because of OSE joining PDP in January 2000
6. Oracle and Networking Services included for 2002