



User Facility Sub-Project Hardware and Material Costs

SUMMARY

This is a costing estimate of the User Facility subproject of the CMS S&C project. The first part of the document describes the scope of the User Facility project. Pricing of the various parts comes mainly from CDF/D0 Run 2 experience at Fermilab. The price evolution comes from Moore's laws with slopes described in the text and fixed at endpoints from Fermilab measurements. Costs for R&D, networking and CMS Detector Construction phase support are also included.

THE CMS DATA ANALYSIS MODEL

The LHC distributed data analysis model has been described in [1]. Here we will summarize from this document some important terms and describe the size estimate for the Tier 1 and Tier 2 Regional Centers.

Definition of Terms

Raw Event

This refers to a raw data event as written by the online system. The size of the raw data event is about 1 MB.

Raw Simulated Event

This refers to a data event as simulated by the Monte Carlo and reconstruction packages. The size of the raw simulated event is about 2 MB .

AOD (Analysis Object Data)

refers to objects which facilitate analysis, but which by construction are not larger than 10 Kbytes. So, if an analysis requires information that does not fit in this limited size, this analysis should access other objects (possibly larger in size). A more traditional term for a collection of such objects is "micro-dst" or "ntuple".

Calibration Data

refers to diverse kinds of information generally acquired online (which may be strictly calibration constants, or monitoring data) as well as special runs taken for calibration purposes. Cosmic ray runs are included in this category. Collections of normal event data used for alignment and calibration studies are not classified as calibration data.

Data Caching

refers to the capability of holding a copy of frequently accessed data on rapidly accessible storage medium, under algorithmic control to minimize data turnaround and maximize user response.

ESD (Event Summary Data)

refers to physics objects of about 400 Kbytes. A more traditional term for this kind of object would be "mini-dst".

Tags

refer to very small objects (1 kB) that identify (tag) an event by its physics signature. A Tag could, for example, be a set of 96 bits each tagging a given physics channel or it could be a set of 10 words with some packed overall information for an event. In this document the Tag is NOT a global object identifier (i.e. a set of bits identifying uniquely an object in a database). A more traditional term would be "nano-dst" or "ntuple" The border between AOD and Tags is well defined in size but less well defined in functionality.

Data Analysis Tasks

The offline software of each experiment is required to perform the following tasks:

- data reconstruction, which may include more than a single step, such as preprocessing, reduction and streaming. Some of these steps might be done online;
- Monte Carlo production, which includes event generation, detector simulation and reconstruction;
- offline (re)calibration;
- successive data reconstruction; and
- physics analysis.

Basic Assumptions

We assume that there exists one "central site", CERN, which is able to provide all the various services (but not with enough total capacity to do all analysis related tasks).

The following steps happen at the central site only:

- Online data acquisition and storage
- Possible data pre-processing before first reconstruction
- First data reconstruction

Other production steps (calibration data storage, creation of ESD/AOD/tags) are shared between CERN and the Regional Centers.

The central site holds:

- a complete archive of all raw data
- a master copy of the calibration data (including geometry, gains, etc.)
- a complete copy of all ESD, AOD, and Tags preferably online

For a typical large LHC experiment, the data-taking estimate is:

- 1 PB raw data per year per experiment
- 10^9 events (1 MB each) per year per experiment
- 100 days of data taking (i.e. 10^7 events per day per experiment)

Size of the CMS T1 Regional Center at Fermilab

Total CPU needs at the Regional Centers

Here we break down by function the total CPU needs at the Tier 1 and Tier 2 centers in order to come up with a sizing estimate.

Simulation:

We estimate from physics requirements that US-CMS will need to simulate 10^8 events/year and each event takes about 5k Si95-seconds. We use the metric 2×10^7 seconds/year to give us an overall efficiency factor of 65% for CPU usage. This gives a total simulation CPU need of 25k Si95's. Most of this is done at Tier II sites, however some provision is made for smaller simulations at the Tier 1 Regional Center.

Reconstruction of simulated events:

We will need to reconstruct the 10^8 simulated events/year. Each event takes 3000 Si95-sec to reconstruct and using 65% efficiency this gives us 15k Si95 total.

Re-Reconstruction:

At the Tier 1 Regional Center we anticipate re-reconstructing 10% of the total raw data twice during the year, for a total of 10^8 events. Each event takes about 3k Si95-secs to reconstruct and using the 65% efficiency we get 30k Si95's total.

Event Selection:

We assume we will need to process the ESD samples to select and tag events. The ESD sample is about 10% of the full data sample and the processing will be done by 10 physics groups 6 times/year. The data selection job on the ESD takes 50 Si95-secs/event and using a 65% efficiency we get 15k Si95's needed for this task.

Physics Object Creation:

This is a scheduled task that reads ESD data and recreates the AOD physics objects with better calibrations and algorithms. Here we assume that 10 physics groups will want to do this each month and that each production covers 10^7 events. Using a 65% efficiency factor we get 12k Si95's needed for this task.

Derived Physics Data (ntuple) Creation:

These are scheduled jobs reading AOD samples and preparing tag/ntuples for detailed analysis by individual physicists. We estimate that 10 groups will read 10^7 events 20 times per year in this activity, and that each event requires 50 Si95-secs to process. Using the 65% efficiency gives us 5k Si95's required for this task.

Individual Analysis:

This task covers chaotic jobs from individual physicists reading AOD and DPD samples to prepare their own private DPD's for further analyses. Here we assume 200 physicists look at 10^7 events 20 times/year and that the analysis on each event takes 20 Si95. Because the jobs are unpredictable, we assume the CPU efficiency is half of the scheduled jobs (i.e. 33% efficient) providing capacity to handle the relatively common peaks normally seen in CPU usage. This gives us a requirement of 80k Si95s at Tier 1. We assume the Tier 2 centers will support about 20% of the physicists supported at Tier 1, so each Tier 2 center gets 16k Si95s.

The table below summarizes the amount of CPU needed at the Tier 1 and Tier 2 Regional Centers.

	Tier 1 Regional Center	Tier 2 Regional Center (each)
Task	Total CPU	Total CPU
Simulation	5k Si95	6k Si95
Reconstruction of Simulation	5k Si95	6k Si95
Re-reconstruction of raw data	30k Si95	
Event Selection	15k Si95	
Physics Object Creation	12k Si95	
Derived Physics Data Creation	5k Si95	
Individual Analysis	80k Si95	16k Si95
Data Servers (0.2×CPU)	15k Si95	
Total	167k Si95	28k Si95

Total Online Storage needs at the Regional Centers

Below we list the data we believe needs to be kept online. Because of performance issues or redundancy needs, we also assign an efficiency to disk depending on type of data stored on it. Most disk gets a 50% efficiency factor. The mass storage data cache disk is much more straightforward to manage and gets a 75% efficiency factor.

Simulated ESD/AOD

This refers to the ESD and AOD objects created by the simulation. Although most of the simulation happens at Tier 2 sites, the Tier 1 Regional Center will keep all of the ESD and AOD from the simulation on disk. 10^8 events X 400kB/event yields 20 TB for the simulated ESD. Having 2 versions of the AOD dataset gives us an additional 10 TB in AOD.

Calibration Data

The exact size and number of calibration events is not yet known, however we are reserving 10 TB in online disk space for calibration data. CMS is currently defining calibrations and the size of this data.

ESD data

We want to keep online at the Tier 1 center two versions of the ESD which means $400\text{kB} \times 10^7 \text{ events} \times 2 = 80 \text{ TB}$.

AOD, TAG, etc.

These are the AOD and TAG data sets coming from physicist analysis. We imagine each AOD set to be on the order of 5TB and each TAG set being on the order of 0.5 TB.

Mass Storage Cache

This is disk used in front of the Mass Storage system to cache active data. It is estimated at about 20% of the 1 PB active data store.

AOD/Raw/ESD

This is a global estimate for online disk at a typical Tier 2 center.

User Disk

This is disk used by individual physicists or small workgroups to store their TAG or DPD event samples. For Tier 1 we estimate we need about 50 TB for this purpose and at Tier 2 sites we estimate about 25 TB.

	Tier 1 Regional Center		Tier 2 Regional Center (each)	
Task	Total Disk	Efficiency	Total Disk	efficiency
Simulated ESD/AOD	30 TB	50%		
Calibration Data	10 TB	50%		
2 ESD Versions	80 TB	50%		
AOD, TAG, etc.	20 TB	50%		
Mass Storage Cache	200 TB	75%		
AOD/raw/ESD			50 TB	50%
User Disk	50 TB	50%	25 TB	50%
Total	390 TB	650 TB	75 TB	150 TB

Total Offline Storage needs at the Regional Centers

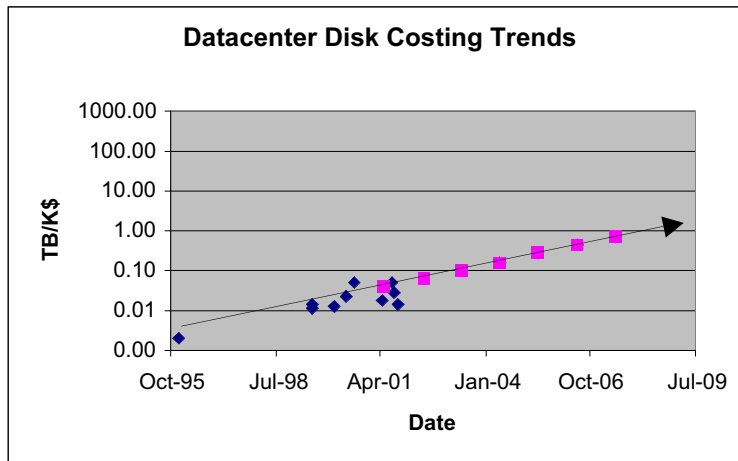
BASIS OF THE REGIONAL CENTER COST ESTIMATE

Assumptions of the Cost Estimates

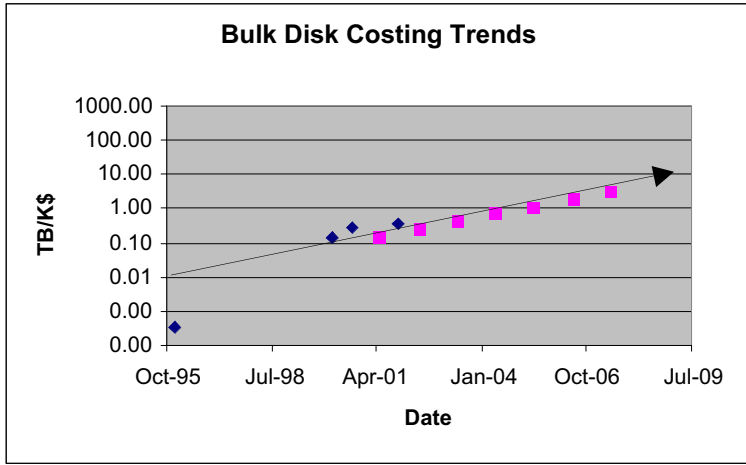
Hardware costs for computing come from FNAL experience over the last few years in purchasing for Run 2 computing. We take the slopes of the extrapolations as measured at CERN [2] using Femilab data over the past few years as boundary conditions. Because much of the hardware is not going to be purchased for several years, and because of the dramatic drop in computing costs over time, an extrapolation is used. The assumptions for cost evolution are given below.

Disk Storage

There are two types of disk: bulk commodity disks and datacenter quality disks. The plots below show Run 2 disk purchases for both bulk commodity disks and datacenter quality disks. In both cases, the slope of the predicted prices are the assumption that disks halve in cost every 1.4 years. The pink points (>2001) are extrapolated to the middle of each calendar year (i.e. June). For the datacenter quality disks, we have inputted some smallish (1 TB) purchase of cheap RAID. Even so these disks seem not to be following the predicted curve very well – in fact they are more expensive than the prediction.



For the bulk disks, we have plotted the prices of cheap IDE disks (i.e. from buy.com) without any infrastructure costs.



The extrapolated points in the above figures are listed in the table below.

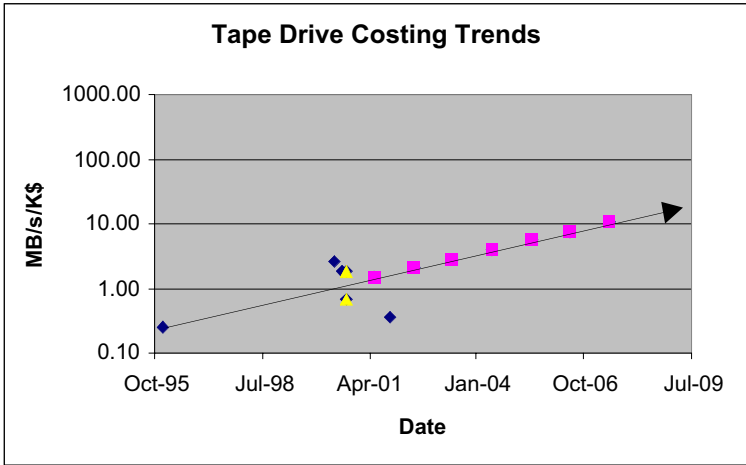
Year	2002	2003	2004	2005	2006	2007
Bulk Disk	\$2.5/GB	\$1.50/GB	\$0.90/GB	\$0.56/GB	\$0.34/GB	\$0.21/GB
Datacenter Disk	\$16.20/GB	\$9.90/GB	\$6/GB	\$3.70/GB	\$2.25/GB	\$1.40/GB

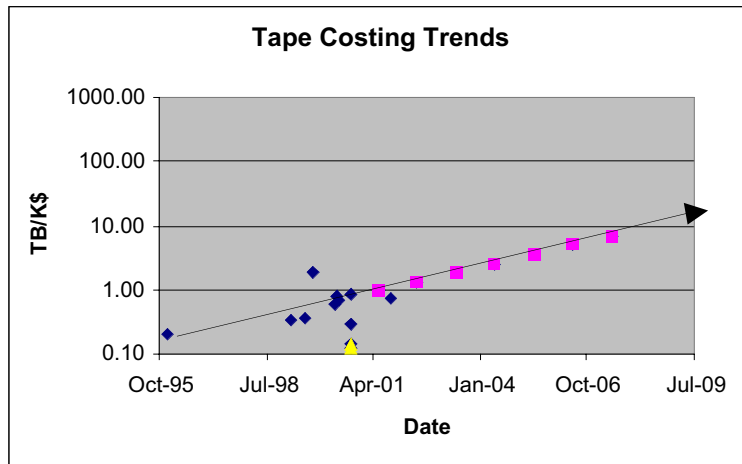
Table 1: Extrapolated \$/GB for bulk (commodity) disk and for data center quality disk. The extrapolated dates are for mid calendar year (June).

Tape Storage

Here the assumptions are that robotics cost \$75/slot, constant in time, and that tape density doubles every 2.1 years, implying that the robotics cost $\frac{1}{2}$ every 2.1 years. The cost of tape media is assumed to $\frac{1}{2}$ every 2.1 years. Tape drive costs/MB/sec are also expected to $\frac{1}{2}$ every 2.1 years.

In the following plots, the pink (>2001) points are extrapolated, and the yellow points are for DVD media and drives. The blue points are traditional tape media and drives. The extrapolated points are for the middle of the calendar year (June).





Year	2002	2003	2004	2005	2006	2007
Tape Drives	\$487/MB/sec	\$350/MB/sec	\$252/MB/sec	\$181/MB/sec	\$130/MB/sec	\$94/MB/sec
Media Costs	\$760/TB	\$540/TB	\$390/TB	\$280/TB	\$200/TB	\$145/TB

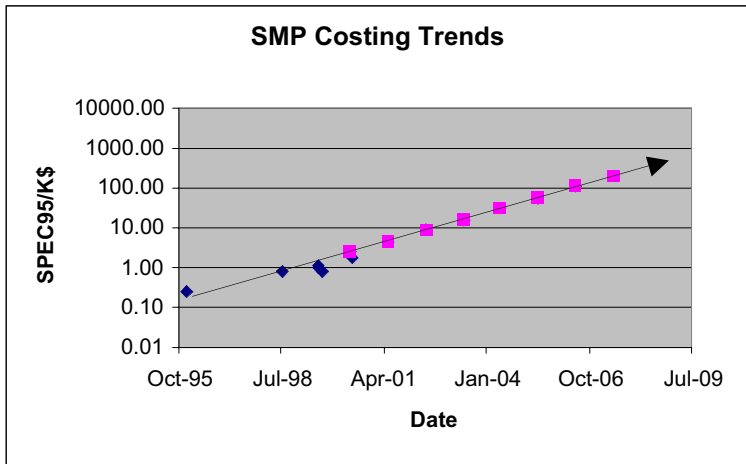
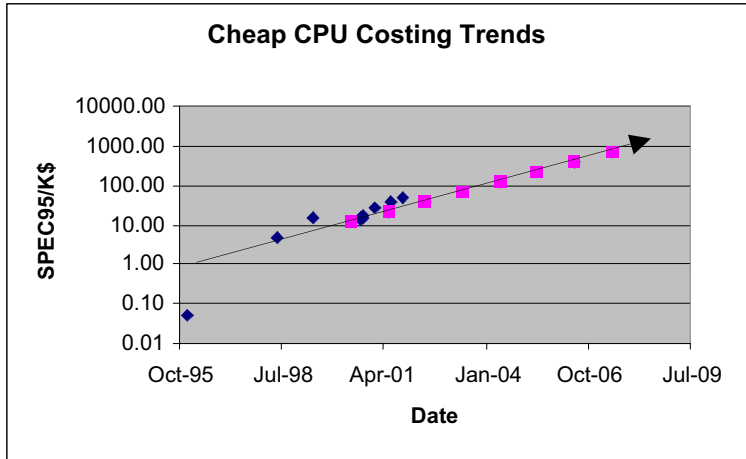
Table 2: Costs of tape media and drives. The pink points are for extrapolations to the middle of the calendar year (June). See text for more detail.

CPU

Commodity CPU costs are estimated by assuming the performance of the CPU doubles every 1.2 years and the price stays the same. The starting number is from Run 2 farm purchases in June 1999, where 2060 Si95's were bought for \$132,000 or 15.61 Si/\$1000. The canonical dual CPU box was ~30 Si95's and cost \$2000.

SMP CPU costs are estimated by assuming the performance of the CPU doubles every 1.1 years and the price stays the same. The starting point is for an SGI general use machine bought in 1998.

In each case the extrapolation is to the middle of the calendar year (June).



Year	2002	2003	2004	2005	2006	2007
Commodity CPU	\$27/Si95	\$15/Si95	\$8.40/Si95	\$5/Si95	\$3/Si95	\$1.5/Si95
SMP CPU	\$116/Si95	\$61/Si95	\$33/Si95	\$18/Si95	\$10/Si95	\$5/Si95

Table 3: CPU costs.

HARDWARE COST BY WBS ITEM FOR THE USER FACILITY PROJECT

Here we detail by WBS number the hardware costs for the User Facility.

1.1 Tier 1 Regional Center

1.1.1 Development and Test systems

This item covers test systems needed to test both hardware and software used in the User Facility.

1.1.1.1 Modular Test Stands

These are hardware test stands typically on expert desktops for testing new hardware such as disks, tape drives, tapes and software drivers. The numbers here are a guess based on a test stand cost of typically \$5000 and assuming we will need four of them each year.

2005	2006	2007	Total	2008 (ops)
\$20,000	\$20,000	\$20,000	\$60,000	\$15,000

1.1.1.2 Integration Test Stands

This is the cost for a fully integral test stand to test production software and hardware. Most likely this will be the fully functional prototype regional center hardware with some upgrades. The cost here is some incremental costs to upgrade or retrofit the prototype regional center with any new technologies necessary.

2005	2006	2007	Total	2008 (ops)
\$10,000	\$10,000	\$20,000	\$40,000	\$10,000

1.1.2 Simulation & Reconstruction systems

Assume the configuration from the Hoffman Review Panel 3. For costing we are staging by buying 10% in 2004, 30% in 2005 and 60% in 2006. Total here is installed capability in 2006.

Task	CPU (Si95)
Simulation	5,000
Reconstruction of Simulation	5,000
Re-reconstruction of raw data	30,000
Event Selection	15,000
Physics Object Creation	12,000
Derived Physics Data Creation	5,000
Total	72,000

The table below uses the extrapolated commodity computing costs, the total reconstruction/simulation CPU as detailed in the table above, and assumes a 10%, 30%, 60% staging over the years 2004-2006.

2005	2006	2007	Total	Operations
\$60,500	\$102,000	\$115,000	\$277,000	\$92,000

1.1.2.1 Disk

There must be a disk pool to support the simulation work. This is not a large pool and is not used to store data. 1 TB should be fine. Disk is replaced every three years starting in 2008. These are data center disks. The efficiency of the disk is assumed to be 80% to cover RAID or need for increased performance.

2005	2006	2007	Total	Operations
\$7,500	\$9,100	\$8,400	\$25,000	\$8,300

1.1.2.2 Hardware Infrastructure Costs

This is the infrastructure and other costs for the facility (racks, cables, power, cooling, consoles, etc.) Using combined experience with Run 2 and more recent totally integrated purchases, we put this at 20% of the total hardware costs. Because we used the more expensive data center disks for this, we didn't include the disk cost in this overhead.

2005	2006	2007	Total	Operations
\$12,000	\$20,000	\$23,000	\$55,000	\$18,400

1.1.3 Analysis Systems

The table below uses an average between SMP and commodity computing extrapolated prices to cost the 80,000 Si95's needed for user analysis as detailed above. We use a balance of SMP and commodity to hedge against cluster issues.

	2005	2006	2007	Total	2008 (operations)
CPU (Si95)	8,000	24,000	48,000	80,000	
Cost	\$165,000	\$267,000	\$288,000	\$720,000	\$240,000

1.1.3.1 Hardware Infrastructure Costs

Because building a large analysis cluster has a lot of unknowns, especially (but not limited to) the commodity computing clustering issues, we put a hardware infrastructure overhead of 25% the total hardware costs for the analysis cluster.

2005	2006	2007	Total	Operations
\$41,000	\$67,000	\$72,000	\$180,000	\$60,000

Note the networking costs are separated out in 1.6.1.

1.1.4 Data Access and Distribution Servers

This covers multiple AMS servers to serve the CMS database to/from the analysis & simulation systems. We use for a measurement the CPU usage for Objectivity which has been measured at ~ 20% overhead by CMS and BaBar. This overhead is calculated for all the CPU before any efficiency factors. Because the measurement included efficiencies, we don't add any additional efficiencies to this CPU, however we assume these servers are expensive SMPs. The operations costs assume a replacement every 3 years.

	2005	2006	2007	Total	2008 (operations)
Si95's	1500	4500	9000	15000	
cost	\$494,000	\$263,000	\$140,000	\$897,000	\$300,000

1.1.5 Calibration and Non-Event Database Servers

Here we assume we need about 5% of the reconstruction Si95 and that this is also a high-end server class machine. The operations costs assume a replacement every 3 years.

	2005	2006	2007	Total	2008 (operations)
Si95's	285	855	1710	2850	
cost	\$119,000	\$63,000	\$34,000	\$215,000	\$72,000

1.1.6 Data Storage Systems

This includes disk and tape. The costs below are just for the disks and/or tapes, not for the server machines to serve the data (which is covered in 1.2.4).

Tape estimates

Task	Tape (TB)
Simulated event data	200
General ESD's	200
Simulated ESD+AOD for analysis	40
Raw Event Data	50
Reprocessed and Revised ESDs	500
AOD, TAG, Calibration, etc	30
Total	1,020

Tape is assumed to be 100% efficient.

Disk estimates

Task	Disk (TB)	Efficiency (%)	Total Disk
Simulated ESD+AOD for Calibration Data	30	50	60
ESD for Local Analyses AOD, TAG, etc	10	50	20
Mass Storage Cache	80	50	160
User Disk	20	50	40
	200	75	270
	50	50	100
Total	390		650

1.1.6.1 Data Storage Costs

In the below table, the online disk was priced by assuming datacenter disk, due to recent experiences with commodity disks.

Robotics cost was placed at \$75/slot, which is the current cost for robotics. This price is not expected to go down (it's basically a cost for a mechanical system), however the number of slots needed goes down depending on tape density. We assume we will purchase a robot that can use commodity drives and tape media.

Robotic maintenance is currently \$4.80/slot and is not expected to change.

	2005	2006	2007	Total	2008 (Operations)
Disk	487,500	891,500	1,086,700	2,465,700	821,900
Media	78,000	168,000	241,900	487,900	162,633
# Robotic Slots	7,000	3,000	0	10,000	
Robotics cost	525,000	225,000	0	750,000	150,000
Robotics Maintenance	33,600	48,000	48,000	129,600	48,000
# Tape Drives	15	15	15	45	11
Tape Drives	375,000	375,000	375,000	1,125,000	275,000
Total	1,499,100	1,707,500	1,751,600	4,958,200	1,457,533

1.1.7 Data Import/Export

1.1.7.1 Tape import/export

We assume 100 tapes in/out per year. This is very uncertain. Will networking be good enough? Will we need more data import/export facility? This is a very modest sized system. We are assuming we can use an AML/J currently available

throughout the lifetime of the experiment. Tapes are assumed to cost \$75 each. Year 1 (2004) will start with 50 tapes/year and we go to 100 tapes/year after that. Tapes in year 2004 will hold on the order of 400 GB by our extrapolations which would mean exporting ~500 TB/year by tape.

Task	2005	2006	2007	Total	2008 (ops)
Tapes	\$3,750	\$7,500	\$7,500	\$18,750	\$7,500
Maintenance of AML/J	\$5,000	\$5,000	\$5,000	\$15,000	\$5,000
Refitting AML/J, drives, etc.	\$15,000	0	\$15,000	\$30,000	\$10,000
Total	\$23,750	\$12,500	\$27,500	\$63,750	\$22,500

1.1.7.2 Tape import/export

There are no explicit hardware costs related to this item.

1.1.8 Data and Physics Analysis Software

There are no explicit hardware costs associated with this item.

1.1.9 System Administration

There are no explicit hardware costs associated with this item.

1.2 System and User Support

1.2.1 Documentation - document generation and access tools.

There are no explicit hardware costs associated with this item.

1.2.2 Collaborative Tools

1.2.2.1 Video conferencing

For video conferencing we need a couple of video conferencing rooms and some personal video conferencing equipment. We assume that personal video conferencing equipment (such as cameras on desktop pc's, etc) is something that comes with the encumbered salary. (i.e. VRVS is "free" with a desktop.)

In order to support the ~ 100 FNAL resident users in 2006 and the myriad of visitors, we estimate we will need 3 video conference rooms

	2002	2003	2004	2005	2006	2007	Total	2008 (operations)
Video rooms	1 room	1 room		1			3	upgrades
costs	\$20,000	\$20,000		\$20,000			\$60,000	\$10,000

1.2.2.2 Web Servers

We assume a web server is a reasonably fancy desktop with lots of disk.

	2002	2003	2004	2005	2006	2007	Total	20078(operations)
Web Server	1	0	0	2	0	0	0	2/3
Costs	\$5,000			\$10,000			\$15,000	\$3,300

1.2.2.3 Scheduling Tools

There are no explicit hardware costs associated with this item.

1.2.2.4 Other

There are no explicit hardware costs associated with this item.

1.2.3 Software Development Environment

This would cover licensing costs. So far, CERN has bought CMS wide licenses for all CMS computing as in, for example, objectivity. Certainly licenses for software development is important, and for software professionals the license fee is contained in their total encumbered salaries.

One can imagine licenses such as lsf licenses, etc., or licenses for any special clustering software for analysis or commodity computing. This is included in the infrastructure overhead.

1.2.4 System Administration, Monitoring and Problem Resolution Software

There are no explicit hardware costs associated with this item.

1.2.5 User Help Desk

There are no explicit hardware costs associated with this item.

1.2.6 Training Office

There are no explicit hardware costs associated with this item. Although a training office requires infrastructure, we are planning to use existing or future Fermilab facilities for this.

1.2.7 Support for Tier 2 Centers

There are no explicit hardware costs associated with this item.

1.2.8 Computer Security

There are no explicit hardware costs associated with this item.

1.3 Operations and Infrastructure

1.3.1 Software license maintenance

There are no explicit hardware costs associated with this item.

1.3.2 Software license maintenance

There are no explicit hardware costs associated with this item.

1.3.3 Infrastructure support

There are no explicit hardware costs associated with this item.

1.4 Tier 2 Regional Centers

The Tier 2 hardware configurations were chosen so as to minimize equipment and support costs and provide more flexibility than is possible at a large Tier 1 facility (though less than would be found at an average university site). Certain tradeoffs in cost, performance and mission were considered. For example, Tier 2 CPUs will be less expensive than those used at Tier 1, but they will still cost more than the cheapest commodity processors in order to provide high performance, excellent reliability and simple maintenance.

The same considerations of performance, reliability and maintenance require other items such as high-speed LAN switches, large capacity and high-throughput data servers, uninterruptible power, dense rack mounted systems, software debugging tools, load balancing software, monitoring / management tools, and maintenance contracts. Most centers will not have Hierarchical Storage Managers (HSM) for transparent tape-disk handling because of the personnel overhead. However, one or possibly two centers might be able to take advantage of an existing HSM and its support personnel and not incur extra costs. Similarly, since the centers will be built from inexpensive CPUs, they will have more processing power than strictly needed for simulation in order to allow some of them to perform more simulations than average while other centers serve other needs.

The costing philosophy for Tier 2 hardware is that an initial purchase is made in some year, then 1/3 of that funding is used to purchase hardware for each following year, effectively replacing that hardware every three years. Other items such as maintenance and software licenses must be paid every year.

The funding profile for the Tier 2 facilities is documented elsewhere, and a summary is listing in the summary table at the end of this document.

1.5 Networking

Non-FTE networking costs fall into two general categories:

- hardware costs for network devices that support CMS systems at Fermilab;
- recurring costs for services associated with CMS network support, principally costs associated with any upgrade of Fermilab off-site network facilities in support of wide-area CMS collaboration.

The estimated non-FTE costs for networking support of CMS at Fermilab are listed below, broken down according CMS WBS item.

1.5.1 Onsite network infrastructure

A very high performance LAN switch to serve as the core Tier 1 CMS network switch fabric is tentatively planned for procurement in 2003. In 2001 & 2002, CMS system network support at Fermilab will be provided on existing switch fabric, although additional switch modules will need to be procured to provide switch ports for the expanding number of CMS systems on-site.

The CMS analysis, data management, and data storage systems, based in the FCC, are assumed to have gigabit network connections. Network support cost estimates are based on \$500/port in 2003, \$450/port in 2004, and \$400/port in 2005. The per system cost estimate includes the cost of the system's gigabit switch port, plus a small amount toward support the general network infrastructure and basic network services used by CMS systems at Fermilab.

CMS desktop network connections at Fermilab are assumed to be 100Mb/s. Desktop support cost estimates are based on \$250/port. The desktop cost estimate includes the cost of the switch port, a small amount toward support the general network infrastructure and basic network services used by CMS systems at Fermilab.

Year	2002	2003	2004	2005	2006	2007	Total	2008
CMS Core Switch Fabric	0	\$30,000	\$20,000	\$20,000	\$20,000	\$20,000	\$110,000	\$10,000
FCC networking	\$50,000	\$10,000	\$10,000	\$23,600	\$39,600	\$38,850	\$172,050	\$64,017
Desktop support	\$2,500	\$6,250	\$6,250	\$5,000	\$3,750	\$3,750	\$27,500	\$9,167
Total	\$52,500	\$46,250	\$36,250	\$48,600	\$63,350	\$62,600	\$278,800	\$83,184

1.5.2 Off-Site Networking

Network support for off-site CMS collaboration is generally assumed to be provided by shared wide-area HEP network facilities, including the LHC trans-Atlantic link, ESnet, and the Internet-2 networks. The Fermilab ESnet link, currently OC3, is assumed to be OC12 by 2002, as part of general upgrade of Fermilab wide-area network capabilities. Additional bandwidth capacity to meet CMS Tier 0/Tier 1 and Tier 1/Tier 2 bandwidth needs is projected as a recurring cost, starting in 2005. A one-time hardware cost for network equipment to support the additional bandwidth is projected for 2005. Modest hardware costs are projected for procurement of servers to support network-level functions of data grid services, starting in 2005.

Year	2002	2003	2004	2005	2006	2007	Total	2008(ops)
Off-site access hardware	0	0	0	30000	0	0	30000	10000
Leased line costs	0	0	0	250000	250000	250000		250000
Grid network support servers	0	0	0	10000	5000	5000	20000	5000
Total	0	0	0	290000	255000	255000	50000	265000

1.5.3 Network Security

Hardware costs for CMS LAN firewall protection, intrusion detection, and related network-level security devices are projected in 2005, when the core Tier 1 CMS switch is procured. Upgrade and enhancement of those devices is projected through 2007. Modest hardware costs are projected for procurement of servers to support security functions of data grid services, starting in 2005.

Year	2002	2003	2004	2005	2006	2007	Total	2008(ops)
Network security Hardware	0	0	0	20000	10000	10000	40000	10000
Grid security support servers	0	0	0	10000	5000	5000	20000	5000
Total	0	0	0	30000	15000	15000	60000	15000

1.5.4 Network Performance Monitoring

A wire-rate gigabit LAN analyzer and associated packet capture probes would be procured in 2003, with modest upgrades or enhancements in the following years. Modest hardware costs are projected for procurement of servers to support wide-area performance monitoring functions of data grid services, starting in 2003.

Year	2002	2003	2004	2005	2006	2007	Total	2008 (ops)
Monitoring tools & hardware	0	0	0	\$40,000	\$10,000	\$10,000	\$60,000	\$20,000
Grid WAN perf. Monitor servers	0	0	0	\$10,000	\$5,000	\$5,000	\$20,000	\$5,000
Total	0	0	0	\$50,000	\$15,000	\$15,000	\$80,000	\$25,000

1.6 Computing and Software R&D

1.6.1 Distributed Data and Computing Test Beds

This covers a test bed to study distributed data access and data access monitoring. Results from this will be used to validate the MONARC simulation tool. The testbed will support distributed data handling software and monitoring as well as studying data access patterns.

This does not require a lot of computing. A large size server with good network connectivity (e.g. Gb Ethernet) is fine. This machine must be for the exclusive use of these tests in order to monitor everything that goes on with them, however. We have purchased a DELL 4cpu linux box in 2000 we will use for this, but we believe we may need a real server class machine, which we plan to purchase in 2001. The dell machine cost \$25k with _ TB h/w raid. We are assuming in 2001 a server class machine with order TB of disk will cost about \$30k.

Year	2002	2003	2004	2005	2006	2007	Total	2008(ops)
Boxes	1	1	1	0	0	0	3	0
Cost	50,000	30,000	30,000	0	0	0	\$110,000	0

1.6.1.1 Hardware Testbed system design

There are no explicit hardware costs associated with this item.

1.6.1.2 Hardware testbed system

There are no explicit hardware costs associated with this item.

1.6.1.3 Distributed data management

There are no explicit hardware costs associated with this item.

1.6.1.4 Distributed Data Access

There are no explicit hardware costs associated with this item.

1.6.1.5 Monitoring Tools

There are no explicit hardware costs associated with this item.

1.6.2 R&D in Computing Hardware for T1/T2 Computing

The cost in this item comes from computing costs for investigating different kinds of computing technology. This is assumed to be bought during 2000-2002 as new things come in.

1.6.2.1 Technology Tracking and Testing

There are no hardware costs associated with this item.

1.6.2.2 Technology Investigation and Deployment

1.6.2.2.1 Data Intensive Computing

Data intensive computing refers to computing used for analysis needed a very high I/O. These machines must have excellent inter-connectivity and lots of I/O capability. Much of this can probably be achieved with a linux cluster, however we may need a more expensive SMP –type solution. The costs here are for a linux cluster + integration costs.

Year	2002	2003	2004	2005	2006	2007	Total	2007(ops)
Linux Boxes	15	15	20	0	0	0	50	0
Cost	\$75,000	\$75,000	\$100,000	0	0	0	\$250,000	0

1.6.2.2.2 CPU Intensive Computing

These requirements are driven by MC production milestones and the fully functional 5% RC milestone. We buy some each year in order to continually evaluate new boxes as they come out.

Year	2002	2003	2004	2005	2006	2007	Total	2008(ops)
Linux boxes	15	15	20	0	0	0	50	0
Cost	\$75,000	\$75,000	\$100,000	0	0	0	\$250,000	0

1.6.2.2.3 Data Access and Distribution Servers

Here we don't know if linux machines will work or if we need higher end servers. Experience from CMS MC production suggests one server class machine for each 20 cpus in order to get the I/O for the MC production. We need lots of disk as well. The cost is for larger SMP servers, but we have bought already some smaller linux servers in order to compare.

Year	2002	2003	2004	2005	2006	2007	Total	2008(ops)
Data server boxes	1	1	1	0	0	0	5	0
Disk (TB)	10	15	15	0	0	0	24	0
Hardware Cost	\$373,030	\$358,195	\$346,421	0	0	0	\$1,077,646	0

1.6.3 Prototype Fully Functional Regional Center

The fully functional regional center should be 5% of the final tier 1 center. It should come online middle of FY 2004. Listed here is an approximate cost for the additional incremental costs to finish it in addition to the R&D systems that can be used as parts.

Year	2002	2003	2004	2005	2006	2007	Total	2008(ops)
Total	0	0	\$50,000	0	0	0	\$50,000	0

1.6.4 Simulation Software

There are no explicit hardware costs associated with this item.

1.6.5 Simulation Software

There are no explicit hardware costs associated with this item.

1.6.6 Simulation Software

There are no explicit hardware costs associated with this item.

1.6.7 Data Import/Export R&D

These are incremental costs for R&D for tape input/output. We assume we can use the same AML/J we are using to support the construction project.

Year	2002	2003	2004	2005	2006	2007	Total	2008(ops)
# tapes	60	60	60	0	0	0	180	0
Tape costs	\$4,200	\$4,200	\$4,200	0	0	0	\$12,600	0
Drive cost	\$25,000	\$25,000	0	0	0	0	\$50,000	0
Total	\$29,200	\$29,200	\$4,200	0	0	0	\$62,600	0

1.7 Detector Construction Phase Computing

1.7.1 Computing

We assume we can support the detector construction project computing on existing central FNAL systems as well as CMS systems we already have, and the R&D systems as we replace them. We add no new costs here.

1.7.2 Data Storage

This item covers data storage to support physicists doing detector or physics studies during the construction project.

Year	2002	2003	2004	2005	2006	2007	Total	2008(ops)
Tape costs	\$17,500	\$15,000	\$12,500	0	0	0	\$45,000	0
Drive cost	\$50,000	\$25,000	\$25,000	0	0	0	\$100,000	0
Maintenance	\$4,500	\$6,000	\$7,500	0	0	0	\$18,000	0
Total	\$72,000	\$46,000	\$45,000	0	0	0	\$163,000	0

1.7.3 Network based file interchange

There are no hardware costs associated with this item.

1.7.4 Tape import/export system

This covers the cost of refitting the AML/J, exporting 50 tapes in 2002 and 100 tapes in 2003 and the maintenance on the library and tape drives. We originally planned to do this in 2001, but had no funding or FTE's. This project has subsequently been dropped and we plan to use only network file exchange for support of the construction project during the R&D phase..

1.8 Support for FNAL based computing

1.8.1 Desktop systems

1.8.1.1 Installation

There are no explicit hardware costs associated with this item.

1.8.1.1.1 System specification, design and installation

There are no explicit hardware costs associated with this item. FNAL costs are in the total encumbered salaries of computing professionals, or come from the base program support of scientists. Outside users who come to FNAL are expected to bring their own desktops (or laptops), which must conform to the FNAL standard.

1.8.1.1.2 Host server systems

We are assuming we support 100 desktops, 50 FNAL, 50 outside users by the year 2005. This ramps up from 10 in 2000 to 100 in 2005 and onwards. We are not buying the actual desktops but we are buying disk servers to support them. Run 2 experience suggests 1 server class machine for every 20 desktops.

	2002	2003	2004	2005	2006	2007	Total	2008 (operations)
#desiktops (integrated)	10	15	20	30	60	100	100	100
# new servers		1		1	1	2	4(+1)	1
costs	\$10,000	\$25,000	\$15,000	0	\$25,000	0	\$75,000	\$12,500

1.8.1.1.3 Backup systems

This item covers backups for user areas, etc. I am assuming a reasonable system is about \$10k/50 machines.

	2002	2003	2004	2005	2006	2007	Total	2008(operations)
#desiktops (integrated)	10	15	25	40	60	100	100	100
# new backup systems	0	1	0	0	0		1	0.25
costs	0	\$37,000	0	0	0		\$37,000	\$9,250

1.8.1.2 Support

There are no explicit hardware costs associated with this item.

1.8.1.3 User data storage and archiving

This is a backup and archiving system for user data from their desktops. Because the server class machines will hold a lot of data disk, we should not back it up automatically but allow users to backup and archive what they want. The minimum would be a small stacker on each server class machine, which should cost on the order of \$10k/machine.

	2002	2003	2004	2005	2006	2007	Total	2008 (operations)
costs	0	0	\$85,000	\$10,000	\$10,000	\$10,000	\$115,000	\$10,000

1.8.2 Remote control room

Year	2002	2003	2004	2005	2006	2007	Total	2008(operations)
costs	0	\$20,000	\$10,000	\$10,000	\$10,000	\$10,000	\$60,000	\$10,000

TOTAL HARDWARE COST FOR THE USER FACILITY PROJECT

Here we summarize the total cost for the User Facility Project. Each of the WBS items (except for 1.4) also has an additional overhead of 16% which is charged by Fermilab.

WBS #	2002	2003	2004	2005	2006	2007	Total	2008 (Ops)
1.1 T1 Regional Center	0	0	0	2,866,049	2,984,127	2,937,831	8,788,007	2,646,796
1.2 System Support	29,000	23,200	0	34,800	0	0	87,000	15,428
1.3 O&M	0	0	0	0	0	0	0	0
1.4 T2 Regional Center	232,000	240,000	870,000	1,870,000	1,500,000	1,750,000	6,462,000	1,250,000
1.5 T1 Networking	60,900	53,650	42,050	511,560	461,854	527,568	1,657,582	484,861
1.6 Computing R&D	511,489	472,203	492,080	0	0	0	1,475,772	0
1.7 Current Support	83,520	53,360	52,200	0	0	0	189,080	0
1.8 Local Comp. Supp.	11,600	95,120	127,600	23,200	52,200	23,200	332,920	48,430
Total	928,509	937,533	1,583,930	5,305,609	4,998,181	5,238,599	18,992,361	4,445,515

BIBLIOGRAPHY

- [1] *Regional Centres for LHC Computing*, Report of the MONARC Architecture Group