DØ Computing and Software Operations and Plan

The DØ Collaboration

September 2004

Contents

Chapter 1

Introduction

This document presents a status report on DØ Computing and Software operation as well as the current evolution plan for the next several years. It includes scope and cost estimates for hardware and software upgrades, with a detailed equipment spending description for the first few years.

At the end of the Summer of 2004, DØ's computing operation is running rather well. Reconstruction has kept up with the accelerator's excellent performance and the detector's high data-taking efficiency, data handling (SAM) works extremely well, remote MC production has reached the million events per week mark, databases are satisfying most of our needs, and analysis cpu power has generally been sufficient (with the exception of a short period before the Winter conferences, but *not* before the Summer conferences). In addition to that, a significant fraction of the DST reprocessing using our p14 reconstruction release was performed off-site. It is worth noting that some of the tools used by DØ (most notably SAMGrid and Runjob) have proven so useful that they are being adopted by other experiments .

However, maintaining this state will require significant effort since the dataset is expected to roughly double on a yearly basis, while resources are expected to increasingly be diverted to the LHC experiments. To maintain adequate computing resources in the future, it is therefore important that DØ's computing model evolve with the developing LHC model. This will allow us to use shared resources rather than facing an unavoidable decline. Our plan is therefore to gradually migrate all of our computing uses to GRID-compatible tools, without sacrificing DØ-specific needs or user-friendly interfaces like dØtools. The first step in this direction was the reprocessing of p14 DSTs at remote sites, which we plan to follow up with p17 reprocessing of raw data, with less manual intervention and increased GRID-compatibility. In the meantime, the Fermilab reconstruction farm is migrating to using JIM rather than custom scripts, making it very similar to remote resources. We anticipate that we will pursue this approach for analysis computing as well, which implies conversion of our central analysis resource (*CAB*) to using JIM at a later stage.

1.1 Managerial Structure

From a management point of view, DØ's Computing and Core Software is organized in two main branches: production and analysis. This structure has been chosen to reflect the two main categories of computing use cases of the experiment and presents the advantage of appearing "natural" as viewed from the user's standpoint. A potential drawback of this structure is that we are striving to develop tools that can be used seamlessly in production as well as analysis, and this is poorly expressed. The people in charge of coordinating development of these tools are very well aware of this goal however. The organization chart is shown in Fig. 1.1.

At the top management level, the Computing Planning Board serves a function which can best be described as intermediate between an advisory council and an executive board, the latter mainly with regard to remote computing issues.

The production branch has 8 components. The integration of JIM, dØtools, dØrte, Runjob and SAM is a key aspect of our computing model, since in the future we plan to use this set of tools both for production and for submission of analysis jobs by users. This is described in detail in Chapter 7, Section 7.2. The data processing infrastructure entails both keeping track of our processing power and performance, and management of all associated needs, like bookkeeping etc. The production of simulated events is now handled nearly exclusively using JIM, and the transition towards job submission by non-experts has started. Both data and MC processing are treated in Chapter 3. On the code development infrastructure and builds front, we have made two significant changes: the transition towards relatively stable operation of the experiment has led us to slow down the release pace. Instead of weekly builds with a goal of having executables that link, we now cut a new release when a significant number of improvements are available, and declare it complete based on executable performance. The second major change is the introduction of shared libraries to make code development in the offline framework significantly faster, and the introduction of so-called "OS-compatibility", which allows us to run our executables independently of the linux flavor or version of the machine.

The online system is transitioning from a set of core machines running OSF to a more distributed and expandable linux-only architecture. This is described in Chapter 3. Following the Computing Division's reorganization, and the fact that SAM is now used by both Run II collaborations, with expressions of interest from others, SAM is not a part of DØ's organization anymore, but we have a SAM liaison. Detailed information on the current status of SAMGrid can be found in Chapter 7. It is sufficient to say here that since DØ is so crucially dependent on SAMGrid's continued success, we work in close collaboration with the SAMGrid development group and are continuously trying to recruit manpower for that effort, with some success.

The experiment's database requirements are quite extensive, and are reasonably well covered, although continuous development remains needed to keep up with new detector capabilities. This is the subject of Chapter 4. The last but not the least component of the production branch concerns offline shifts, key in early detection of problems. Effort

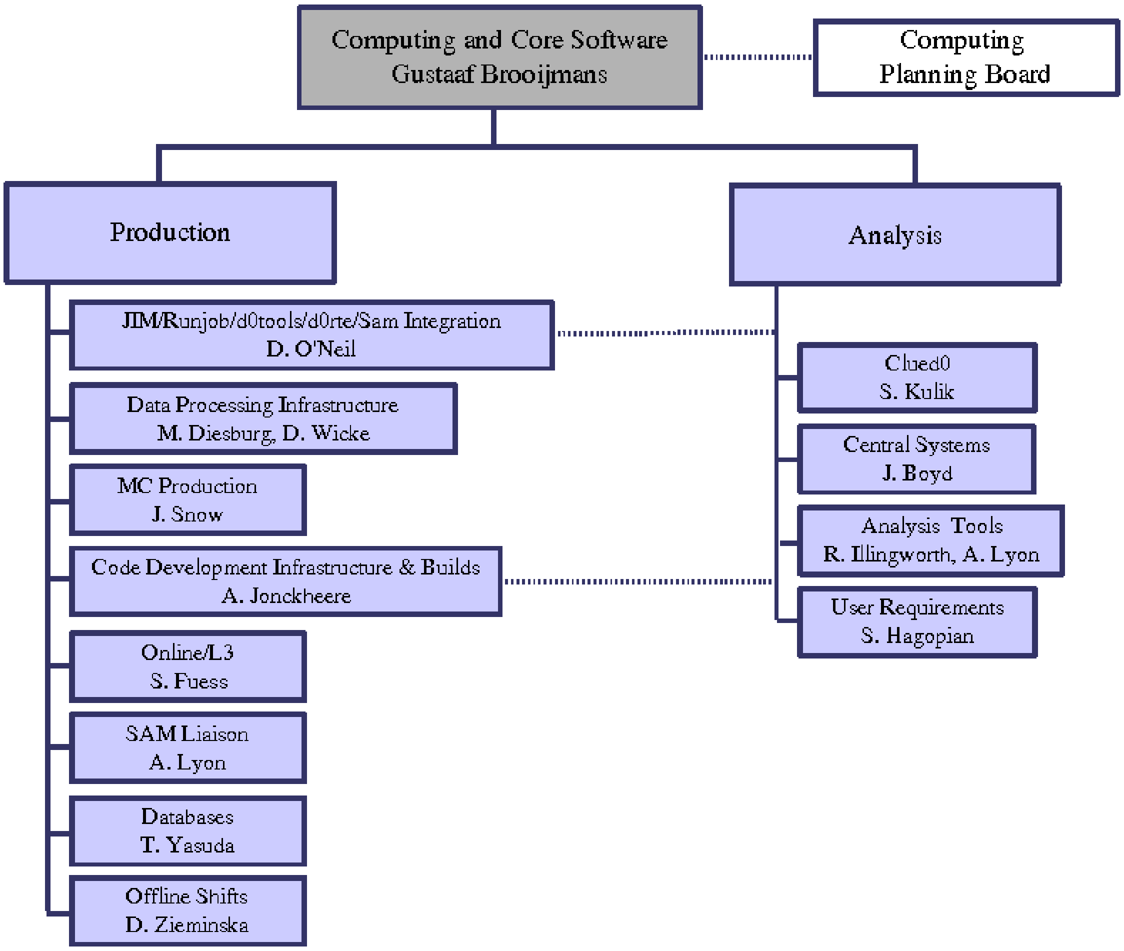

Figure 1.1: DØ Computing and Core Software Organization Chart.

continues to make it easier for offline shifters to not only identify failures, but also diagnose their precise cause and take corrective action.

The analysis branch contains 4 components. The first two concern our analysis computing infrastructure, one being our desktop cluster which offers both interactive and batch capabilities, and the other our central batch systems. More details on these can be found in Chapter 3. The analysis tools encompass things like dØtools, the user-friendly interface to running executables, as well as maintenance and development of SAM tools, etc. The "User Requirements" component is new, and is aimed at improving the general interactions of the users with the computing and algorithms experts. This task covers identification of outdated documentation, prodding of people to update and/or improve documentation, acting as a contact point for users to point out poor or non-existing documentation or hard-to-use software, and organizing of software tutorials during collaboration meetings. Progress on this front is being made.

Two items that are not part of the DØ organization chart but which will be discussed in this report are data storage, in Chapter 5, and networking in Chapter 6,

Chapter 2

Computing Needs

2.1 Production Executables

In the previous years, the simulation and reconstruction has seen a number of successive production releases, often several a year. As the algorithms have stabilized, the efforts have shifted towards a better understanding of the detector and the differences between real data and simulated data. These efforts demand long and complex studies and for this reason no new production release was introduced this year. We are currently converging towards production release p17 in which all major components of simulation (generators, detector simulation, trigger simulation), reconstruction and Level 3 trigger reconstruction will be certified.

2.1.1 Reconstruction

At the reconstruction level, this next production release will see major improvements in the treatment of the calorimeter, as significant progress has been made in understanding the behavior of the calorimeter electronics and of the calorimeter calibration. In addition, bad channels identified offline or online will be suppressed and an algorithm designed to further reduce the effect of calorimeter noise will be used.

Last year, we reported on the delivery of the thumbnail data tier (most compact data tier) which is essentially the sole input for analysis. It was decided to add the calorimeter data to the thumbnail. This has proven to be essential for the rapid development of algorithms (e.g. new electron and jet identification tools, new photon identification) and for the understanding of the calorimeter (data quality, calibration). This also allowed us to make two passes of partial reconstructions ("fixing") based on the corrections of known calorimeter electronics problems. The first pass provided the dataset used for the 2004 summer conferences, while the second pass, still ongoing, runs newer algorithms that will be run in the p17 releases. To build on this success, we decided that as of p17 clusters from the tracking devices and the preshower information would also be included. This is mainly based on the physics groups' interest in being able to re-run the tracking algorithms with a

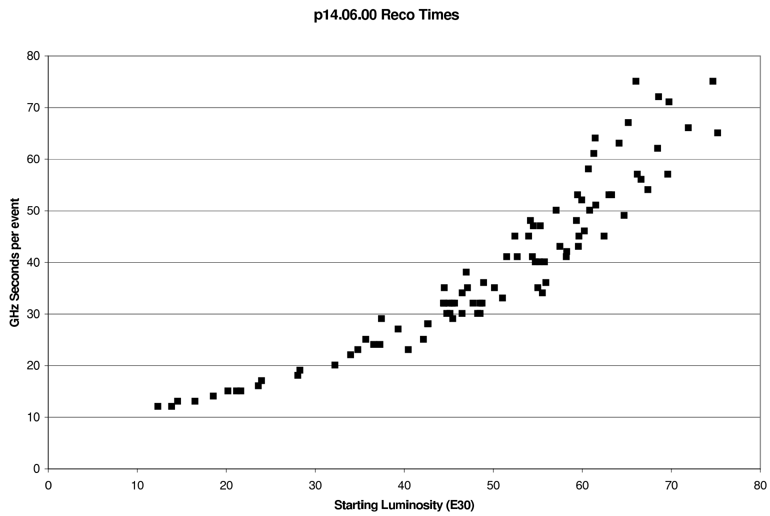

Figure 2.1: Average reconstruction time per event as a function of initial instantaneous luminosity for the run.

lower p_T cut on selected events, but since these enlarged thumbnails will allow efficient re-tracking, we think that keeping this additional information will allow for rapid development of algorithms and feedback on calibration.

2.1.2 Simulation

It is essential for the DØ physics program to provide a realistic simulation of the detector. First, an improved description of the material within the tracking volume has been implemented in the simulation (and in the reconstruction). Second, studies have shown that overlaying real data events over generated signal events to simulate the effect of pileup and multiple interactions significantly improves the agreement between data and simulation of the tracking performances. These capabilities are now available for large scale Monte-Carlo production.

2.1.3 Prospects

In the next year, the reconstruction and simulation (as well as most other algorithm software) will undergo significant upgrades. Experience has shown that with instantaneous luminosities around $10^{32}cm^{-2}s^{-1}$, the tracking part of the reconstruction is reaching its limitation in terms of speed, as can be seen in Figure 2.1. A significant effort is ongoing with help from the Computing Division to address this problem from the computing point of view (speed of the current algorithm) and the algorithm and physics points of view (tuning of algorithm versus physics trade off). Moreover, the simulation and the reconstruction will be modified to support the detector upgrades scheduled for the coming years, in particular the addition of the Silicon Layer 0, and the trigger and Central Fiber Tracker electronics upgrades.

2.2 Data Analysis

2.2.1 Postprocessing

After the data has been reconstructed on the farm, the Common Samples Group (CSG) runs it through the so-called *fixing* and *skimming* processes. The former applies corrections for problems that were discovered after the production release was cut, and basically consists of unpacking the thumbnail, correcting some calorimeter cell energies, re-running jet, electron and missing E_T algorithms and repacking the thumbnails. The latter selects events based on reconstructed physics objects and writes them to corresponding streams. One or more of these "skims" forms the base data sample for the vast majority of the physics analyses.

2.2.2 Data Format

Up to now, different physics groups and individuals have taken different approaches to using the thumbnails. Some have decided to work within the framework, using the thumbnails directly. Since the unpacking of the thumbnails depends on so much of the DØ code, however, linking of the executables required very significant amounts of memory, and remained slow even on well-equipped machines. The recently introduced possibility to used shared libraries will improve this situation dramatically. Since the relinking process was so slow, and physics analysis inherently requires frequent code changes, most of the users have chosen to convert the thumbnails to root format and work on those. There exist multiple root formats however, used by different groups. Producing these, and keeping the data on disk has proven to be a strain on analysis resources, and in May 2004 the Data Format Working group was formed to study the issue and formulate a recommendation for improvement. Their report is available as DØ Note 4473. In summary, they recommend creating a root-tree based on the existing TMBTree, a root tree format containing a large fraction of the thumbnail information. This would be enhanced to satisfy all analyzers' requirements, and would be centrally produced and stored in SAM.

Since the boundary conditions are that *d0reco* produces thumbnails and plots are produced in root, it is expected that we will have two main groups, one working directly on the thumbnails, and another working on the TMBTrees. This will most likely reduce the resource requirement quite significantly.

2.2.3 Resources

As mentioned in the introduction, the computing power available at Fermilab for data analysis has proven to be adequate for most of the year, with the exception of the period preceding the Winter conferences. As the experiment matures and analyses converge toward publication, we expect the analysis cpu requirements to further even out over the course of the year. The dataset is growing however, and disk and cpu demand is likely to increase in proportion. While we expect to be able to keep supplying the necessary resources in the short term at least, we are working towards enabling the usage of external resources for analysis as well as production. To achieve this without introducing a significant learning curve, the necessary tools and interfaces are being built into *DØtools*, with which the majority of DØ users is already familiar. This is described in detail in Chapter 7.

Chapter 3

Computing Systems

3.1 Reconstruction Farm

3.1.1 Current Status

The current DØ reconstruction farm consists of 368 dual processor worker nodes, 19 dual processsor input staging nodes and an 8 processor SGI Origin system which is used as a server machine and output buffer. The worker nodes are a mix of 1 GHz PIII, 1.67 GHz Athlon and 2.6 GHz Xeon class machines. The total compute power of the system is approximately 1225 GHz.

The Origin server node, the input staging nodes, and 128 of the worker nodes are housed at FCC in the second floor computer room. The remaining 240 worker nodes were moved to the computer facility at the New Muon Lab (NML) in June of this year in an effort to avoid power overload problems at FCC. Operation of the farm with the workers distributed in two locations has so far caused no difficulties, however, stability of power at NML has been somewhat lower than we would like. We expect to move these nodes to the new High Density Computing Facility (HDCF) at Wideband when that installation is complete.

3.1.2 Performance

Since the Fall of 2003, the farm has been running the p14.06 version of the reconstruction code, which has proven to be the most robust version of *d0reco* that we have run. Excluding files not processed because the raw data themselves were corrupted, the completion rate for reconstructing data on the farm is in excess of 99.97%, that is there is less than 0.03% data loss that could be recovered by correcting problems in the reconstruction code. There are no losses from the farm operation itself, in other words, with the exception of the 0.03% noted above, all raw data that can be processed are eventually reconstructed and the output stored to tape.

Until the completion of the p14 reprocessing pass we have needed all available cycles on the farm to keep up with new data and to complete the reprocessing. During that time the farm has operated at better than 80% efficiency. Since early January, when the reprocessing finished, we have used roughly 60% of the available cycles on the farm with available cycles and demand almost ideally matched since January. Note that is a usage statistic and not an efficiency. In that period, we have been able to have raw data processed and available to users within two days of the time it was collected. The small excess of available cycles has also allowed us to quickly accommodate special processing requests for test runs and do some additional offline processing that would otherwise have to be done on the analysis systems (TMB fixing, for instance).

The match of resources to demand is dependent on the average execution time per event. Figure 2.1 shows the average execution time per event for a series of runs as a function of the initial luminosity of the run. On its busiest days the DØ detector records about 3.5 million physics events that go to the reconstruction farm. Assuming an 80% operating efficiency for the farm and 1225 GHz of compute capacity we can see from the plot that the farm can keep up with a detector running at 100% duty cycle if the average initial luminosity is no more than $\approx 35 \times 10^{30}\ cm^{-2}s^{-1}$. As the average initial luminosity grows beyond that, we become dependent on the limited duty cycle of the accelerator and detector to keep up with data acquisition. With an accelerator and detector duty cycle of 33%, we could survive up to average initial luminosities of $\approx 75 \times 10^{30}\ cm^{-2}s^{-1}$ under the pessimistic assumption that the reconstruction code execution time cannot be improved.

3.1.3 Planned Expansion

Current plans are to add 80 dual processor nodes to the reconstruction farm for the 2004 fiscal year, with the expectation that these nodes will be 3 GHz processors. That will bring the total capacity of the farm to 1700GHz. With the additional processors we can probably survive up to luminosities of $\approx 85 \times 10^{30} cm^{-2}s^{-1}$ with a 33% accelerator duty cycle. Note that HDCF will need to be ready for occupancy by the end of October in order for the new nodes to be online by the beginning of the 2005 data taking period.

3.2 Central Analysis Systems

We successfully reduced DØmino from 128 processors to 64 in the last year. We recently started deploying linux sam cache servers (7 have been deployed) and will start pulling project disk off of DØmino soon. The new linux based interactive login pool is being tested and appears to be ready. We will release it very soon for users well ahead of the planned retirement of DØmino to have time to solve any issues that arise. There are still many remaining services DØmino is supplying and we will be working hard to move those functions onto linux machines through the rest of this year in preparation for DØmino's retirement (which is planned for the end of CY 2004).

We currently have 389 linux nodes, in a cluster called CAB, set up for batch analysis jobs using Torque (PBS) as the batch system. We are adding another 200 machines this year, which will be located in the new High Density Computing Facility and hooked to the rest of the DØ network through a 10 Gbit network link. The link between the DØ network and the external network will also be upgraded to a 10 Gbit link this year. More details on the network infrastructure are given in chapter 6.

In addition to DØmino, we currently have 21 TB of linux SAM cache disk on 6 servers and will continue moving the disk off of DØmino onto the linux boxes. DØmino currently has 14 TB of SAM cache. We have also purchased over 100 TB of cache disk this year and it will be deployed as soon as it is available.

We have 17 TB of linux based project disk and about 10 TB of project and user disk on DØmino. We have ordered approximately 28 TB of project disk that will be attached to linux servers. All of these disk servers will be in FCC attached to the DØ network by at least gigabit connections.

We are looking at options to move off of our IRIX based home directory server. Since we were recently joined with the CDF department, a solution similar to their Network Appliance home space would be ideal. This would allow us to have one type of home directory server for both experiments.

We are also investigating the possibility of leveraging the Tibbs backup system being deployed by the Computing Division to use as our main backup solution. This would allow us to eliminate the maintenance and administrative cost of having our own backup software and hardware.

3.3 ClueDØ

3.3.1 Introduction

ClueDØ is a cluster of desktop Linux PCs at Fermilab, supported by the institutes of the collaboration. It is the main interactive development and analysis platform for physicists at DØ. ClueDØ appears as a single integrated cluster to the user. Services like user management, batch system, web server etc. are provided centrally. Every machine is accessible to everybody and contributes to the batch processing as long as the machine fullfills some minimal CPU and memory requirements.

Most machines are desktop nodes, but over the past two years institutes have begun to add rack-mounted nodes as large disk servers (usually based on inexpensive IDE RAID arrays). These servers are located on the second floor of the DØ Assembly Building, where racks, power and Gigabit Ethernet connectivity has been made available. This is now the preferred way for institutes to buy disk space that is used for specific analyses or physics groups.

At the moment ClueDØ consists of 365 nodes. More than 880 users have an account on ClueDØ. While in the past access to ClueDØ was based on contributing to the cluster in the form of hardware, this requirement has been dropped in the meantime and every DØ member can have an account on ClueDØ. Fermilab is providing hardware for running several of the central services (racks plus about 16 machines).

3.3.2 Batch

Only machines with a CPU faster than 400 MHz and some minimal memory requirement (256 MB) are kept in the batch system. At the moment there are 264 active nodes, many of them dual-cpu, which provide 358 batch slots.

The typical limit for jobs is the memory, not CPU power, so usually there are not more than 250 simultaneous batch jobs running. At peak times before conferences there have been more than 1000 jobs in queues. A fair-share scheduling algorithm, based on the "number of GHz" an institution contributed to the cluster is used to prioritize jobs.

3.3.3 Disk Space

Each machine has disk space reserved for batch jobs (5-10 GB) and SAM cache (10-20 GB). It also exports a local scratch disk to the whole cluster. These sizes range from 5 GB to more than 300 GB, with a total of probably a few TB. Scratch disks are writable by everybody, although owners of the machines are supposed to be contacted before large amount of data are stored there.

The large IDE raid arrays are exported separately and appear as different volumes under /rooms/... Currently there are 69 rooms on 50 servers which together add up to approximately 64 TB of disk space. The institution which owns the server can decide how the space should be used. Typically one part is reserved for members of the institute, while another part is "donated" to a physics group of their choice.

All home directories are mounted from DØ2ka and are backed up. UPS products and software releases are mounted from central servers. A user has direct access to DØmino home directories and to project disks on the new Linux file servers.

No backup is provided for any of the scratch disks or large disk servers. A failure of a scratch disk usually involves data loss. Failure of more than one disk in the RAID arrays, which weren't replaced quickly enough, has led to loss of data in the past as well.

3.3.4 System Administration

System administration is done by members of the DØ collaboration There is a core team of about 7-8 people, plus one "local administrator" per institution, if available. Most of these people can only contribute a small fraction of their time (10-20 %), which can be claimed as service task for the collaboration. Recruitment of new people has been difficult at times, but this could always be remediated by application of light pressure by management.

No service beyond Monday to Friday, 9-5 is guaranteed, although in practice problems are solved any time a sysadmin is around or can fix it remotely – a big advantage of a user administered cluster. Given the size and age of the cluster (the first incarnation existed in 2000), hardware problems occur almost every day. They are supposed to be fixed by the owning institute instead of the ClueDØ administrators however.

After the recent upgrade to a newer Linux distribution, there are no more single points of failure left which can bring down the whole cluster. Downtime of DØ2ka, which provides

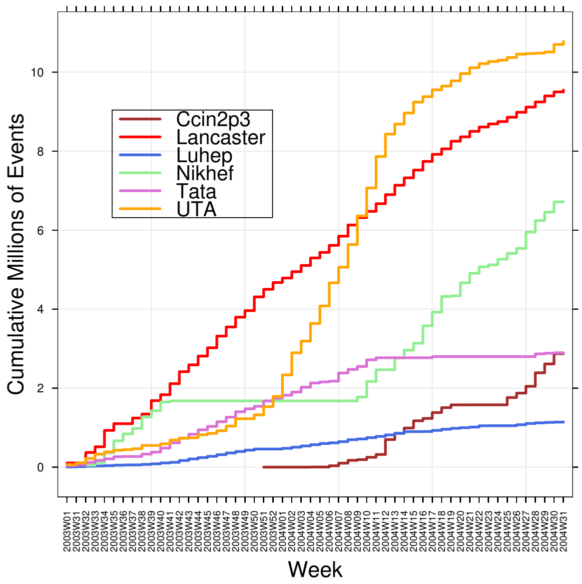

Figure 3.1: Cumulative number of MC events produced at various remote farms.

the home and release areas, makes the cluster relatively useless, however.

3.4 Remote Computing Resources

3.4.1 Monte-Carlo Production

The total number of Monte Carlo events produced between August 2003 and August 2004 was 37 Million events, with the breakdown between sites listed in Table 3.1. As can be seen from Fig. 3.1, the site at the University of Texas, Arlington expanded significantly in capacity in 2004. A small test site in Sao Paulo, Brazil became available (Sprace) which is expected to expand in 2005, and we are now adding a site in China. The University of Oklahoma is also commissioning a site.

To date, all of the DØ Monte Carlo production has taken place at remote sites, and has been the foundation for building tools that facilate site independent distributed computing. The original SAM tests of remote data access and storage were done with MC data, the development and use of mc_runjob facilitated the reprocessing using remote resources. In the past year, we have deployed the job manager and monitoring functionality of SAMGrid for MC production in preparation of deploying it for reconstruction and analysis.

The JIM functionality was orginally tested on resources at the University of Wisconsin, and the installation tested on a SAMGrid test bed at FNAL. JIM is being used for MC production at six sites including IN2P3, Lancaster, Manchester, University of Wisconsin, the DØ Southern Analysis Region (DØSAR) sites (using locally developed control software

Group	Events	Files
CCIN2P3	2874445	6253
Lancaster	9547122	37929
Latech	216500	419
Luhep	1147000	45886
Manchester	219902	916
Nikhef	6719849	26887
Ouhep	737000	2948
Prague	583447	4273
Sprace	631500	1813
Tata	2898500	5600
UTA	10779063	20171
Wisconsin	247102	1420
Wuppertal	202000	808
Total	36803430	155323

Table 3.1: Monte-Carlo production by group.

known as mcfarm) and Prague.

3.4.2 Results From p14 Reprocessing Round

The initial planning for the last round of reprocessing was begun in June of 2003 at the DØ Beaune Workshop. The goal was to have the the entire DØ data set processed with the latest p14 version of the reconstruction program in time for the Winter conferences. This put a cutoff date for completion of the reprocessing at around January 1st, 2004. The dataset to be processed consisted of all data taken between April 20, 2002 and the final cutoff date. This was projected to consist of 500M-600M events.

The amount of processing power needed to accomplish this task exceeded what was expected to be available at Fermilab in this time period. Consequently it was decided to move as much of the reprocessing as possible to remote sites. In the end, 7 computing centers in 6 countries participated in the reprocessing.

Some changes to normal production procedures were necessary to simplify the processing task at remote sites. Reconstruction from raw data requires numerous accesses to calibration data bases located at Fermilab. In order to avoid these accesses it was decided to do the reconstruction from p13 DST output rather than from raw data. This output already has calibration constants applied and removes the need for remote access to the Fermilab databases. Furthermore, the final step of processing involves merging the small thumbnail output files into larger files suitable for efficient storage in SAM. This merging process requires multiple accesses to the SAM database. It was therefore decided to ship the unmerged thumbnails to the DØ central analysis system at Fermilab and do the final merging and storage to SAM at FNAL rather than deal with excessive DB access over the wide area network.

Debugging and certification of the p14 code proceeded through the summer with a final version ready for use at remote sites in early November of 2003. Each remote site was required to run the same certification data sample before proceeding with production. Go ahead was given only when the sites produced identical results on the sample. Most sites were ready for production by the middle of November.

Production reprocessing began in earnest in the middle of November and was complete by January 5th, 2004. The final data sample consisted of 524M events, of which approximatively 200M were initially processed from raw data at Fermilab with p14. Another roughly 200M done at Fermilab had been previously reconstructed with older versions. The total number of events reprocessed at remote sites was about 100M, which represented approximatively 25TB of data moved between FNAL and the remote institutions. The full reprocessing was completed in 7 weeks after certification and required an average of 0.5 FTE at each participating site to run the production.

3.4.3 Prospects

Following the successful remote p14 reprocessing in autumn of 2003 and the successful production of MC using SAMGrid earlier this year, SAMGrid is increasingly becoming the default for production computing. SAMGrid installations have been established in the Czech republic, France, the UK and the USA, with installations underway in Canada, Germany and Holland, as well as elsewhere. Considerable activity is also underway to make SAMGrid interoperable with LCG and other grids, giving DØ access to the extensive shared LHC resources.

The next data reprocessing, using release p17, is scheduled to start in late fall 2004, using SAMGrid. The sample to be reprocessed will be approximately ten times larger than the p14 sample reprocessed remotely. Significant resources in Europe and the USA totalling 1.5 THz equivalent have been identified for this activity. This would leave an estimated 1 THz available for MC production in Europe, India, South America and the USA. The reprocessing sites will be free for MC production once the p17 reprocessing is completed, though since many are shared resources the fraction available to DØ may well be reduced outside the selected p17 reprocessing slot. The use of a single standard set of automated tools for all DØ production computing will enable sites to rapidly switch from data reprocessing to MC, enabling the most efficient use of resources.

Increasingly the resources are shared, both with other HEP experiments, most significantly LHC, and by the wider community. It is thus imperative that the SAMGrid software can run on these shared systems. Additionally structured planning is required to minimize conflict with other large production users, e.g. LHC data challenges.

3.5 Online Systems

The DØ Online computing system provides the functionalities of data acquisition (DAQ), data quality monitoring, and detector control and monitoring. The configuration of the

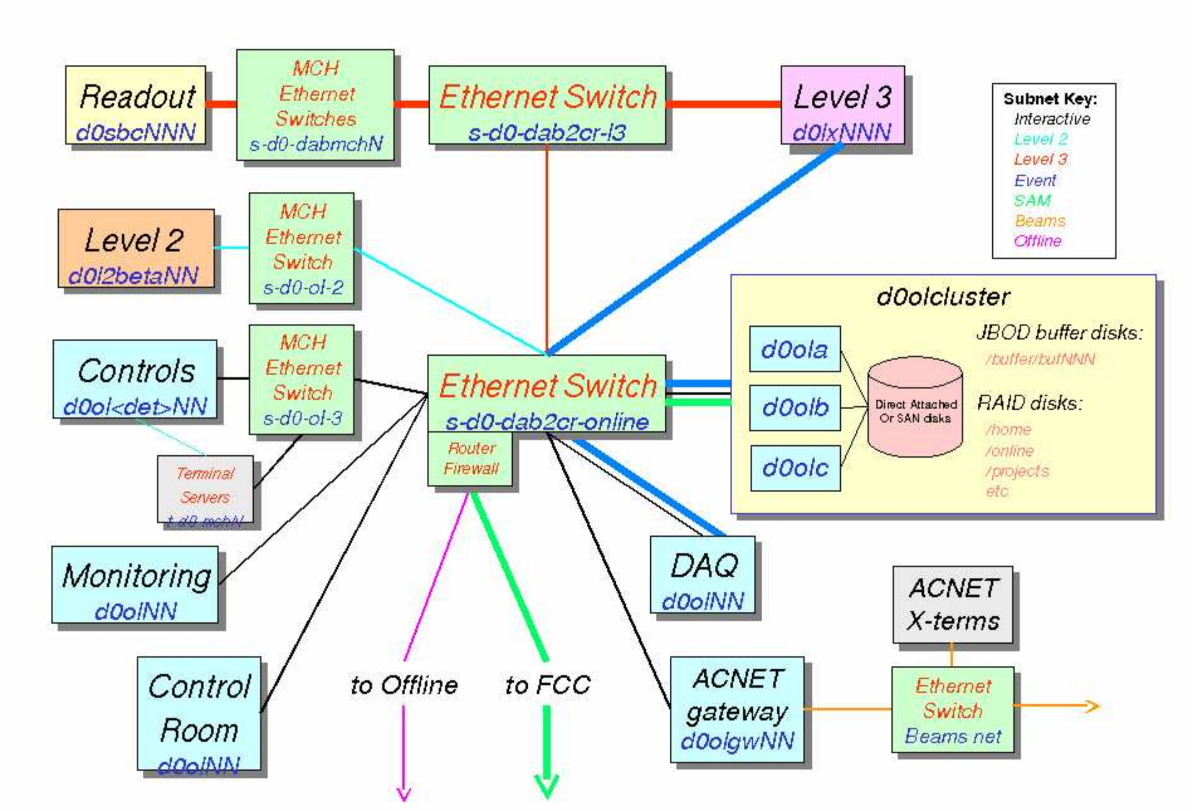

Figure 1. DØ Online computing elements

Figure 3.2: Online system configuration.

Online system is illustrated in Figure 3.2. The principal computing elements are:

- Network infrastructure
  Two Cisco Catalyst 6509 switches provide the backbone of the network. The Level 3 switch (s-d0-dab2cr-l3) is dedicated to the event data flow and connects the readout crate processors to the Level 3 farm nodes. The Online 6509 (s-d0-dab2cr-online) contains a router module, and subdivides the remaining Online network into functional virtual LANs. This switch supports the nodes responsible for Data Logging, event data monitoring, control room activities, detector slow controls and monitoring, and general detector operations. A 1 Gb connection from the Online 6509 to a switch in FCC provides the path for raw event data to be sent to the enstore/SAM system.

- DAQ components
  The DØ event data path from electronic readout crates to final logging to tape is completely Ethernet based. Each of the approximately 64 readout crates contains a single board computer (SBC) running Linux. The SBCs send their event fragments to one of the 114 Level 3 dual-cpu nodes, where a software filter selects events satisfying trigger requirements. The Level 3 nodes send accepted events through a DAQ Collector node to a Data Logger running on the d0olcluster Host system. Events are accumulated on buffer disks, and then complete files are transferred over the network to FCC to be stored in the enstore/SAM system.

- Data monitoring components
  Event data are replicated at the DAQ Collector nodes and passed to Monitoring nodes for analysis. There are approximately 20 distinct monitoring processes and associated nodes, which provide feedback to shift operators in the Control Room.

- Detector control and monitoring components
  DØ utilizes the EPICS [1] control system software, running on approximately 110 embedded Controls processors in both VME readout and detector monitoring crates. Various interactive and background applications are used for detector operations.

- Other components
  A cluster of 3 AlphaServers provides file services and ORACLE database services. These same servers support the event Data Logger. Additional Online components include approximately 10 Level 2 processors, which participate in the Trigger decision but are not part of the event data path. There are also a pair of dual-homed Linux nodes which act as a non-routing interface between the Online and Accelerator networks.

The DØ RunIIb upgrade includes several improvements to the Online system. An expansion/upgrade of Level 3 with 96 new dual-processor nodes will be performed in 2005. Approximately 36 of the Control system embedded processors will be replaced in 2005 by newer models. And finally, the AlphaServer Host systems will be replaced in late 2004 by a set of Linux clusters. These improvements are funded as part of the RunIIb upgrade project.

This last item of the upgrade noted above has a potential impact on DØ Offline computing needs. To date, the raw data acquisition (or level 3 accept) rate has been limited by the capacity of the Host systems to receive data from Level 3, buffer it to disk, and transfer it to the FCC enstore/SAM systems. Trigger configurations were set to match these hardware limits, which were also consistent with the limits set by the ability to reconstruct and analyze the data. With the transition from the older (purchased in 1997/1998/2000) AlphaServers to more modern and dedicated DAQ Host servers, the hardware will be able to support a much higher rate. The current limit of approximately 50 Hz (* 250 kBytes/event = 100 Mbits/sec) is expected to easily surpass 100 Hz for a single output stream, and can be increased up to the network limit (currently 1 Gbit/sec) using multiple output streams with minimal incremental cost (perhaps $5000 per 100 Hz). The effective removal of any rate limitation at the Online system means that there will be pressure to loosen trigger restrictions and record data at the highest possible rate. The level 3 accept rate will then exclusively need to be set by a policy which takes into account the downstream Offline reconstruction and analysis capabilities.

Chapter 4

Databases

4.1 Offline Calibration Database

The calibration constants for the SMT, CFT, and CPS are served to the reconstruction processes on the Fermilab reconstruction farm and to the jobs that users run for their own purposes via the DAN client-server system based on the CORBA technology and developed by the Computing Division and DØ. This system can also deliver calibration constants for the muon subsystems to *d0reco*, but at this time the muon group has chosen not to use this facility.

The database has been in use since June, 2003 and is fully reliable. There is one dedicated linux node for running the calibration DB servers for the farm and another for the non-farm users. Since these nodes are not supported 24/7, there is a failover node and a procedure in case of troubles with the server nodes. Several server node failures have occurred over the past year due to malfunctioning cpu cooling fans, as well as one failure due to bad memory. All of the fans have now been replaced. These nodes are two year old machines which will need to be replaced in the near future.

The DØ farm runs 600 processes and requests a new set of calibration constants on average every 100 seconds. It takes about 90 seconds of wall-clock time for a *d0reco* process to receive one whole set from the servers, including the time for marshalling data on the servers and unmarshalling them on the client. The cpu usage of the farm server node is about 5-10 % and the network usage out of the server node is about 30 %. This implies that up to a three-fold increase in the number of the farm nodes can be handled without upgrading the network interfaces.

To prepare for p17 remote processing, tests have been performed for the remote proxy servers and the problems found during the tests have been solved. These tests were performed at GridKa, Imperial College, IN2P3 and Westgrid; these groups are now running the proxy servers routinely. We believe that the system is ready for serving calibration constants at the remote institutions for large scale data re-processing from raw.

One major component that has been missing in the calibration DB server-client system is the calorimeter. The client and server code has been completed recently and is now included in the test release of *d0reco*. It will be included in the imminent p17 release.

4.2 Luminosity Database

The information required to calculate luminosities delivered to the DØ detector has been kept in the form of ascii files. Our goal is to transfer all of the existing information to an Oracle database to profit from the better tools available with a modern database system. At this time, we plan to move this database into production before the end of FY04.

The schema design has gone through several iterations. We believe the current design is sufficient to proceed to large scale testing and have loaded the database with six weeks' worth of data to test possible use-cases. Our next step is to add another 6 weeks of data and do similar tests. If everything goes well, we will then load the remaining data into the database.

The web-based luminosity report system has successfully been converted to access the database instead of reading ascii files, however this still needs to be optimized for better performance. In parallel, we are working on a system to load data directly or indirectly from the luminosity data acquisition system into the database.

To access the luminosity database from the analysis programs, we are developing a client-server system similar to the calibration system. The database server for this has been generated using the same code generation system as the calibration system, and the client code called lm_access has been converted to use the CORBA protocol instead of reading the ascii files. At this time, the new lm_access performs as fast as or faster than the old system.

4.3 Trigger Database

The trigger database has been in use since the start of Run II. The database server is similar to the SAM servers and contains many modules with some business logic in it. The server modules all need extensions and improvements however, some of which may require modifications in the base server code. We also need to address the performance of the web interface for the data entry. The database group in the CD has agreed to provide help in this project, and we have now produced a detailed, prioritized description of the work that needs to be completed. We expect to make significant progress on this front before the end of CY04.

Chapter 5

Data Storage

For the analysis systems, we currently have a 14 TB disk cache on DØmino, the legacy SMP Origin 2000 IRIX machine. We have deployed 21 TB of IDE disk on 6 linux fileservers as SAM cache for the central analysis system, CAB, and 64 TB is available on ClueDØ to support desktop analysis. In addition, the local worker node disks there can be used as cache space.

The fileservers used as SAM cache are deployed as *pass through* in which requests for files from tape are first staged to the fileservers and then distributed to the worker nodes. While this mode is somewhat slower for tape access than directly staging to the worker nodes and leads to duplication of some of the files on disk, our experience is that cache management should be made as automatic as possible. The alternative mode would be populating the fileservers by hand with the frequently used samples and staging directly from tape to worker nodes for files not resident in the large cache. This is effectively how the combination of DØmino cache and CAB was configured. Our experience using the DØmino cache with CAB showed decreasing numbers of transfers from DØmino as a function of time, as fewer people ran analysis jobs directly on DØmino, or tended to use DØmino for staging raw samples, and thus the cache was not populated with samples for general analysis. Until the fileservers were deployed, in busy weeks, up to 2/3 of the files would be staged from tape to CAB, an indication that the cache formed by the worker node disk was insufficient, leading to cache and tape thrashing.

Fig. 5.1 shows the number of gigabytes consumed per month on the FNAL analysis stations, with Fig. 5.2 showing the same in terms of number of events. Typical monthly consumption on CAB (and CABSRV1 which serves part of the CAB nodes) is 150 TB, with local disk cache of roughly 12 TB. We do not have complete statistics for the relative transfers from enstore, DØmino, or intra-station for CAB or CABSRV1 for 2003/2004, or cache utilization statistics for CAB or CABSRV1. We are planning on using dCache, developed by the FNAL Computing Division and DESY, as a rate adapting buffer to the tape drive system.

In addition to SAM cache and dCache, we supply the physics and analysis groups with project space. This space is intended for highly specialized samples, code development,

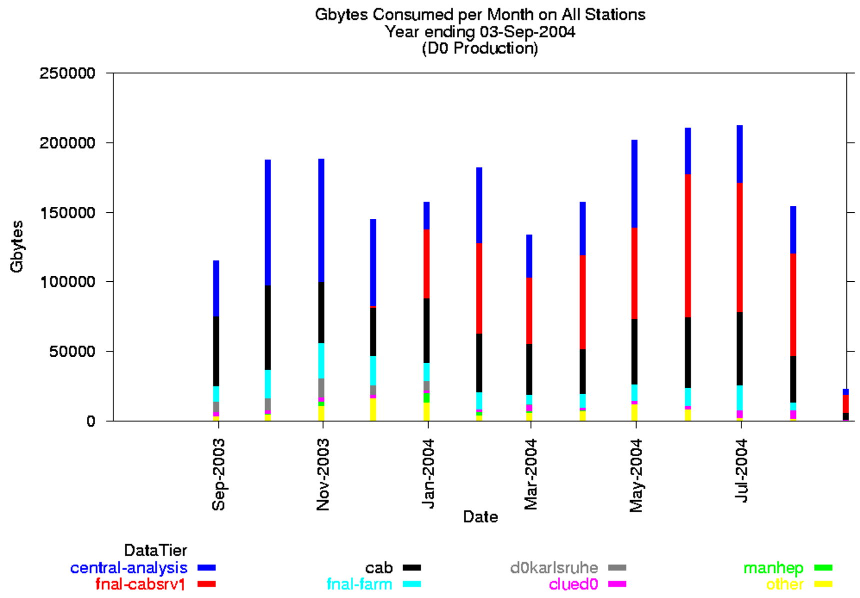

Figure 5.1: Gigabytes consumed per month on various stations, including central-analysis (DØmino), CAB (and CABSRV1 which serves part of the CAB nodes), ClueDØ (the desktop cluster), DØKarlsruhe and other remote stations.

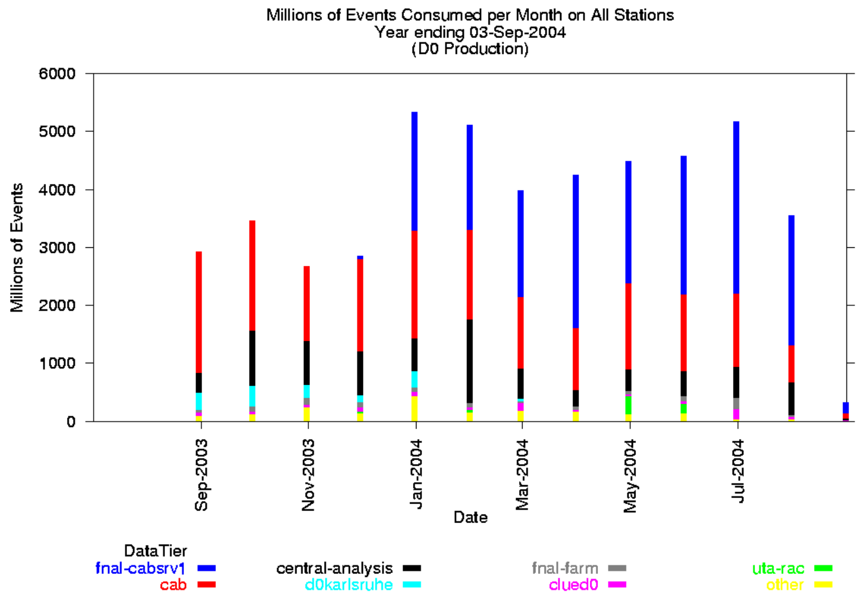

Figure 5.2: Events consumed per month on various stations, including central-analysis (DØmino), CAB (and CABSRV1 which serves part of the CAB nodes), ClueDØ (the desktop cluster), DØKarlsruhe and other remote nodes.

	Written Tapes	Data Volume	Number of Drives
9940a	3826	209 TB	10
9940b	2037	365 TB	20
LTOI	1998	178 TB	8
LTOII	298	54 TB	5 (+10)

Table 5.1: Snapshot of the 4 active DØ tape libraries. 10 LTO II drives are on order.

documentation and other group specific needs. For budgeting purposes, we assume that the amount of project disk needed is one quarter of the amount of space estimated for disk resident analysis samples.

5.1 Robotic Storage

The DØ data management system relies heavily on Hierarchical Storage Management (HSM) systems for archival storage. The principal HSM used by SAM is Enstore, developed at Fermilab by the Integrated Systems Department of the Fermilab Computing Division (now within CCF), and largely influenced by DØ requirements.

Currently DØ is writing different types of 200 GB cartridges in three different robots. We use 2 STK Powderhorn silos (each with 5500 tape slots) currently writing with 9940b drives, and 1 tower (3500 tape slots) in an ADIC AML2 (a 3 tower system shared with SDSS who are using DLT media) writing with LTO II drives. Table 5.1 shows the current use in the various libraries (the amount of data on particular type of tape), which also shows the amount of legacy data, as DØ has not migrated completely to the 200 GB cartridges. The DØ experiment has adequate tape slots, and plans to use 200 GB cartridges through FY2005, however in 2006, we will need to replace the 200 GB drives with 500 GB drives and media and copy the older data to new media.

We are considering turning on the second arm of the ADIC robot to ensure adequate mount capacity, and might need to activate the third tower. While we have found it useful to preserve the capability to use commodity drives, the AML2 is aging technology in an aging system, and it may be less costly and operationally more robust to replace the robot in 2006.

The projected technology is 500 GB cartridges in 2005. The SuperAIT drives are available and currently in use at FNAL for system backups. The LTOIIIs and 9940cs are projected to be available for beta-testing within the next six months. Putting new types of drives into production is typically a six month to 1 year process: a few drives are purchased and first evaluated using standard read/write tests, and depending on the level of problems seen and fixed, the drives are put into service for an experiment using a non-critical data format. With smooth operation, more drives are purchased, and if necessary or desirable, other data formats are directed to that library, and legacy data can be migrated to the new media. In this way, we build up confidence in the drives with minimal risk; however, it is a labor intensive and time consuming process. The last four drive evaluations have gone well, with those drives going into production service for use for experimental data. However, prior to that, 3 successive drive evaluations did not go well (Mammoth I, Mammoth II, AIT2). For this reason, DØ would like to see a testing program with superAIT drives start soon, and encourage participation pursuing a beta testing program for 9940cs and LTOIIIs. The use of superAIT drives for system backups does not adequately model the "write once/read many" use of experimental physics data.

5.2 Physics considerations

In order to project the storage needs for the DØ experiment, we assume a data collection rate, a size for the various formats, and that all of the formats scale with the raw data rate. Whenever possible, we use our current experience as a guide. We assume that the disk cache under SAM control should be large enough to hold the most frequently accessed data samples with 100 % over-capacity for other samples.

The anticipated plan for DØ data formats represents a change from past use. The DST format is being discarded in favor of adding tracking system hits to the tmb+ data format, which itself corresponds to the thumbnail format augmented with calorimeter cell information, originally planned as a short term convenience for debugging. This new format is called the tmb++. In addition a new centrally supported root-based user format is being developed, with the intention that this format serves the bulk of the user analysis needs. However, some groups have communities that maintain that they will not switch to that format, preferring to continue work with either the tmb or with their own format. This implies that the final step of the data processing is available in both centrally supported formats and that the SAM cache is sized to hold both.

DØ is planning a reprocessing of the full data set with p17. Past experience with processing and reprocessing on DØ indicates that well need to reprocess the full data set once per year, and run a fixing pass on the reconstructed data twice per year, with a tmb++ output format. After that, the data are skimmed into samples for physics group use, with tmb++ output and root analysis format outputs. Current experience indicates that the ratio of the output to input of the skimming step is 1:1.

We assume that the data collection rate remains constant in 2005 at 16 Hz, and increases to 30 Hz after the 2005 shutdown. Folding in data sizes and scaling factors (Table 5.2), Tables 5.3 and 5.4 show the estimated amount of data for tape and disk storage. As a note, the disk estimate is specifically in support of FNAL-based analysis. Additional factors are applied to account for the fact that this is cache space and to provide dCache to serve offsite analysis, MC generation and secondary reconstruction. In 2005, due to having two primary analysis formats, and an increase in the size of the tmb, an adjustment is added.

5.3 Cost Projections

The cost of tape is estimated in Table 5.5, which includes the cost of migrating to new media in 2006. As a note, starting in 2006, it would cost \$20K per year to duplicate the

category	size	2004	2005	2006	2007	2008	2009
raw event	0.25 MB	1 (0.01)	1 (0.01)	1 (0.01)	1 (0.01)	1 (0.01)	1 (0.01)
raw/RECO	0.5 MB	0 (0)	0 (0)	0 (0)	0 (0)	0 (0)	0 (0)
data DST	0.15 MB	2 (0)	0 (0)	0 (0)	0 (0)	0 (0)	0 (0)
data tmb+	0.02 MB	3 (1)	0 (0)	0 (0)	0 (0)	0 (0)	0 (0)
data tmb++	0.04 MB	0 (0)	4 (1)	6 (1)	8 (1)	10 (1)	4 (1)
data roottuple	0.03 MB	0 (0)	4 (1)	6 (1)	8 (1)	10 (1)	6 (1)
user formats	0.02 MB	0.5 (1)	0.5 (1)	0.5 (1)	0.5 (1)	0.5 (1)	0.5 (1)
MC D0Gstar	0.7 MB	0.1 (0)	0.1 (0)	0.1 (0)	0.1 (0)	0.1 (0)	0.1 (0)
MC D0Sim	0.3 MB	0 (0)	0 (0)	0 (0)	0 (0)	0 (0)	0 (0)
MC DST	0.3 MB	0.2 (0)	1 (0)	1 (0)	1 (0)	1 (0)	1 (0)
MC tmb	0.05 MB	1 (0.1)	3 (0.1)	3 (0.1)	3 (0.1)	3 (0.1)	3 (0.1)
PMCS MC	0.02 MB	0 (0)	2 (0.5)	2 (0.5)	2 (0.5)	2 (0.5)	2 (0.5)

Table 5.2: Data sizes and tape (disk) factors for various tiers.

	Current	2005	2006	2007	2008
raw event	236.57	119.37	223.81	261.11	261.11
data DST	283.88	0.00	0.00	0.00	0.00
data tmb+	56.78	0.00	0.00	0.00	0.00
data tmb++	0.00	76.39	114.59	152.79	190.99
data roottuple	0.00	57.30	85.94	114.59	143.24
user format	9.46	4.77	4.77	4.77	4.77
MC D0Gstar	66.24	33.42	33.42	33.42	33.42
MC D0Sim	0.00	0.00	0.00	0.00	0.00
MC DST	56.78	143.24	143.24	143.24	143.24
MC tmb	47.31	71.62	71.62	71.62	71.62
PMCS MC	0.00	19.10	19.10	19.10	19.10
Yearly storage (TB)	757	525	697	801	867
total storage (TB)	757	1,282	1,979	2,779	3,647
MC Yearly (TB)	170	267	267	267	267
MC Total (TB)	170	438	705	972	1,240

Table 5.3: Data accumulation on tape (TB).

	Current	2005	2006	2007	2008
raw event	2.37	1.19	2.24	2.61	2.61
data DST	0.00	0.00	0.00	0.00	0.00
data tmb	18.93	0.00	0.00	0.00	0.00
data tmb++	0.00	19.10	35.81	41.78	41.78
data roottuple	0.00	14.32	26.86	31.33	31.33
user format	18.93	9.55	17.90	20.89	20.89
MC D0Gstar	0.00	0.00	0.00	0.00	0.00
MC D0Sim	0.00	0.00	0.00	0.00	0.00
MC DST	0.00	0.00	0.00	0.00	0.00
MC TMB	4.73	2.39	4.48	5.22	5.22
PMCS MC	0.00	4.77	8.95	10.44	10.44
Yearly storage (TB)	45	51	96	112	112
adjusted for format change in 2005	0	43	0	0	0
Yearly adjusted storage (TB)	45	95	96	112	112
total storage (TB)	45	140	236	348	460
MC Yearly (TB)	5	7	13	16	16
MC Total (TB)	5	12	25	41	57

Table 5.4: Data accumulation on disk (TB).

raw data. The old raw data tapes from the 9940a and 9940b libraries could continue to be maintained to provide a complete duplicate of the raw data at nominal cost.

For project disk, we intend to use an Infortrend chassis with 16 400GB disks (4.8 TB useable capacity with 1 global hot spare and 3 parity drives), attached to a 2U server at a cost of $16k each. For SAM cache, we will purchase 24 bay systems defined as 24 hot swap IDE drive bays with 400GB drives in a 5U chassis with server components (8 TB useable with 4 parity drives and no global spare) at $17k each. For out-year cost projections, we assume that the cost of the system remains constant with an 18 month capacity doubling. The associated costs for file servers are shown in Table 5.6. The 2004 number is projected from the spreadsheet and is in reasonable agreement with the 2004 pending requisitions for 100 TB SAM and dCache, and 17 TB project space.

Year	2004	2005	2006	2007	2008
Data Volume (TB)	757	525	697	801	867
Data to Rewrite (TB)	0	0	1,282	0	0
Tape Volume (Tapes)	4543	3152	1672	1922	2082
Replacements (Tapes)	0	0	2565	0	0
Purchase (Tapes)	4543	3152	4237	1922	2082
Cost	$363,440	$252,160	$338,960	$153,760	$166,560
cost to duplicate raw data	$20,000	$20,000	$20,000	$20,000	$20,000

Table 5.5: Tape cost estimate. The 2004 number shown is not the 2004 increment, but the total amount of data DØ currently has on tape.

Year	2004	2005	2006	2007	2008
IDE	$208,000	$288,000	$288,000	$240,000	$240,000
Infortrend	$68,000	$85,000	$68,000	$85,000	$68,000
Total	$276,000	$373,000	$356,000	$325,000	$308,000

Table 5.6: Disk cost estimate.

Chapter 6

Networking

6.1 Infrastructure

The principal elements of the current network infrastructure supporting DØ Computing are noted in Figure 6.1. These elements and the functions they provide are:

- CAS 6513 (s-d0-fcc1w-cas)
  This Cisco Catalyst 6513 currently acts as the central switch supporting DØ. All of the other large switches connect to the CAS switch. A large number of nodes on FCC1 are also connected to this switch: 61 ports for enstore mover nodes, 185 ports for CAB nodes, 12 ports for database servers, 8 ports for SAM servers, and 7 ports for DØmino. Including the switch interconnects, 276 of the total of 384 ports on the 6513 are in use.

- CAB 6509 (s-d0-fcc2w-cab)
  This Cisco Catalyst 6509 supports CAB nodes on FCC2. There are currently 224 worker, 20 server, and 4 mover nodes attached. Including interconnects, 250 of 296 available ports are in use.

- NML 4006 (s-d0-nml-farm-1)
  This Cisco Catalyst 4006 supports Farm nodes in the New Muon Lab. There are 240 worker nodes attached. Including interconnects, 241 of 242 total ports are in use.

- FCC 4006 (s-d0-fcc2w-farm-2)
  This Cisco Catalyst 4006 supports Farm nodes on FCC2. There are 150 worker nodes attached. Including interconnects, 152 of 202 total ports are in use.

- DAB 6509 (s-d0-dab2cr-offline)
  This Cisco Catalyst 6509 supports all Offline computing at the DØ site, including ClueDØ servers and desktops, and Windows servers and desktops. A number of smaller subsidiary switches are located throughout the DØ Assembly Building (DAB), Outback, and Portakamp office areas. Recently a number of ClueDØ file servers have been located on DAB2, and have direct connections to Gb ports on this switch.

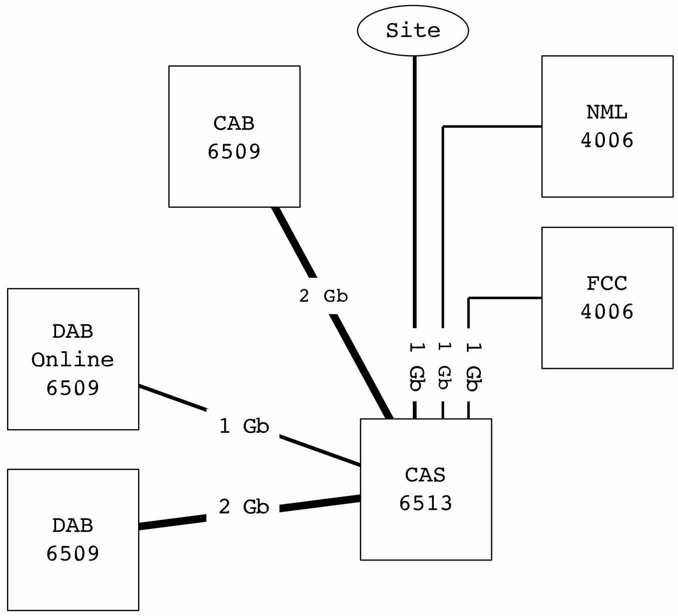

Figure 6.1: Principal elements of the network infrastructure.

- DAB Online 6509 (s-d0-dab2cr-online)
  This Cisco Catalyst 6509, along with a companion 6509 and smaller subsidiary switches, serve the Online and DAQ systems. The functions and utilization of the Online network components is described elsewhere.

The large switches enumerated above are interconnected by single or aggregated dual gigabit connections.

6.2 Utilization

Figure 6.2 is a snapshot of the utilization of these interconnects in the period from roughly mid-July to mid-August 2004. The plots show the average and peak bandwidth both in to and out of the switches connected to the CAS 6513. The sense of the in/out is relative to the ports on the subsidiary switches. There are a number of features visible:

- The DAQ/Online generates, during accelerator operations, an average of 60-80 Mbits per second of data, with peaks approaching 160-180 Mbits per second.

- The reconstruction farms ingest data at an average rate similar to that produced by the DAQ. There are higher peaks in traffic observed, as much as 300-400 Mbits per second. As the reconstruction generates an output, there is nearly as much traffic outgoing from the farms.

- The principal analysis site, CAB, ingests data at several hundred Mbits per second, with peaks above 1 Gbit per second. Outgoing traffic is approximately 50% that of incoming.

- The analysis nodes at the DØ complex appear to sink and source approximately equal amounts of data, at an average of about 100 Mbits per second.

Figure 6.3 is the extension of Figure 6.2 to the network external to Fermilab. This figure shows the network traffic at the points between the DØ CAS 6513 and ESNET. The plots are made from the point of view of the switches closest to DØ that is Outgoing traffic is that destined for off site. There are a number of switch elements between the DØ network and the external ESNET. At the FCC Core router (s-s-hub-fcc) the various Fermilab workgroups are attached. External traffic traverses the Border router (r-s-bdr) and the DMZ switch (s-x-dmz-fcc) before reaching ESNET (or eventually the Starlight network). The external DØ traffic appears to be at the 100 Mbits per second level, but the aggregate Fermilab external traffic uses a sizable fraction of the 622 Mbits per second ESNET connection. Since the latter could become a limitation for remote raw data processing, DØ welcomes the planned increase of bandwidth between Fermilab and the outside world.

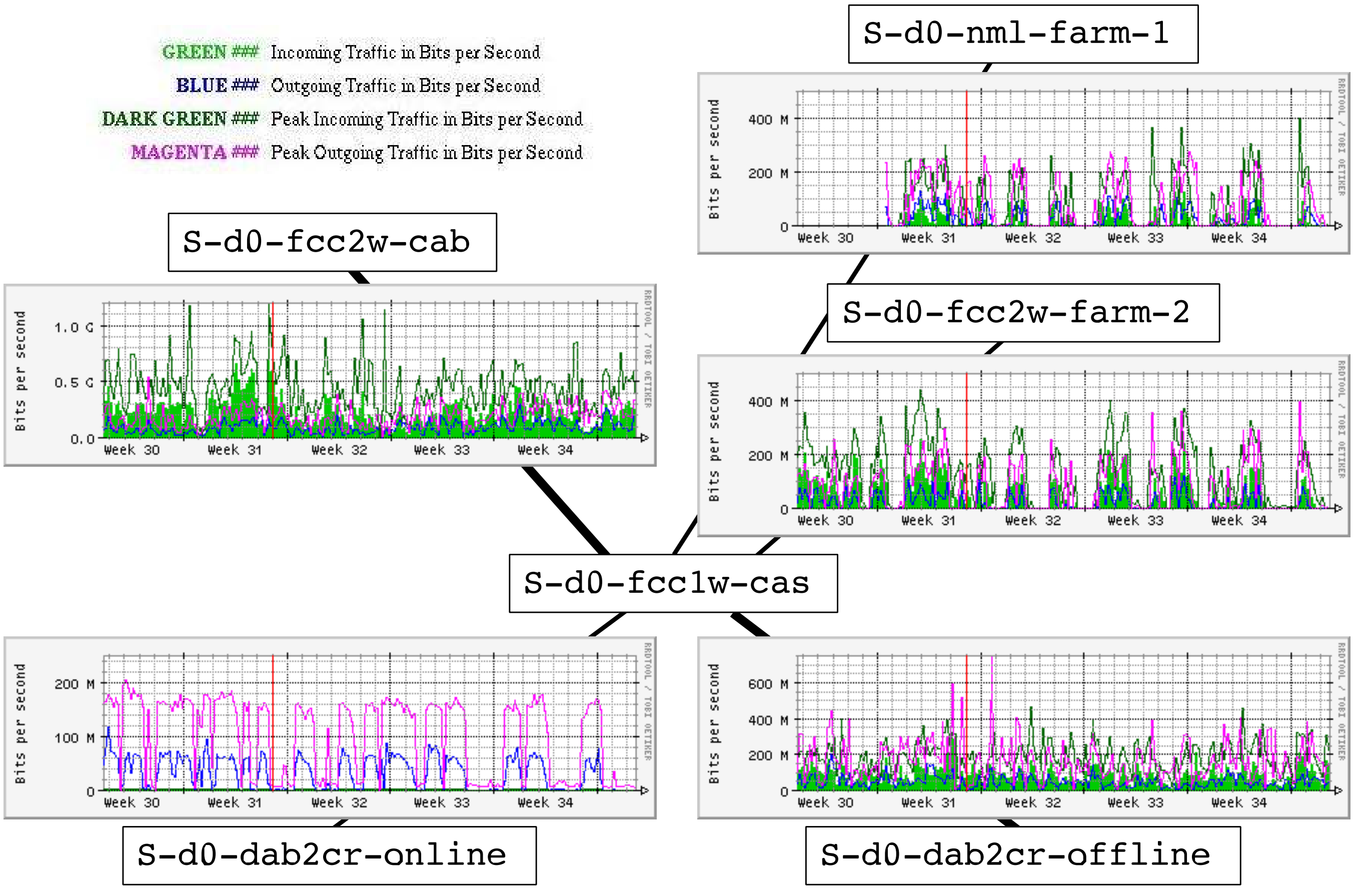

Figure 6.2: Snapshot of utilization of switch interconnects.

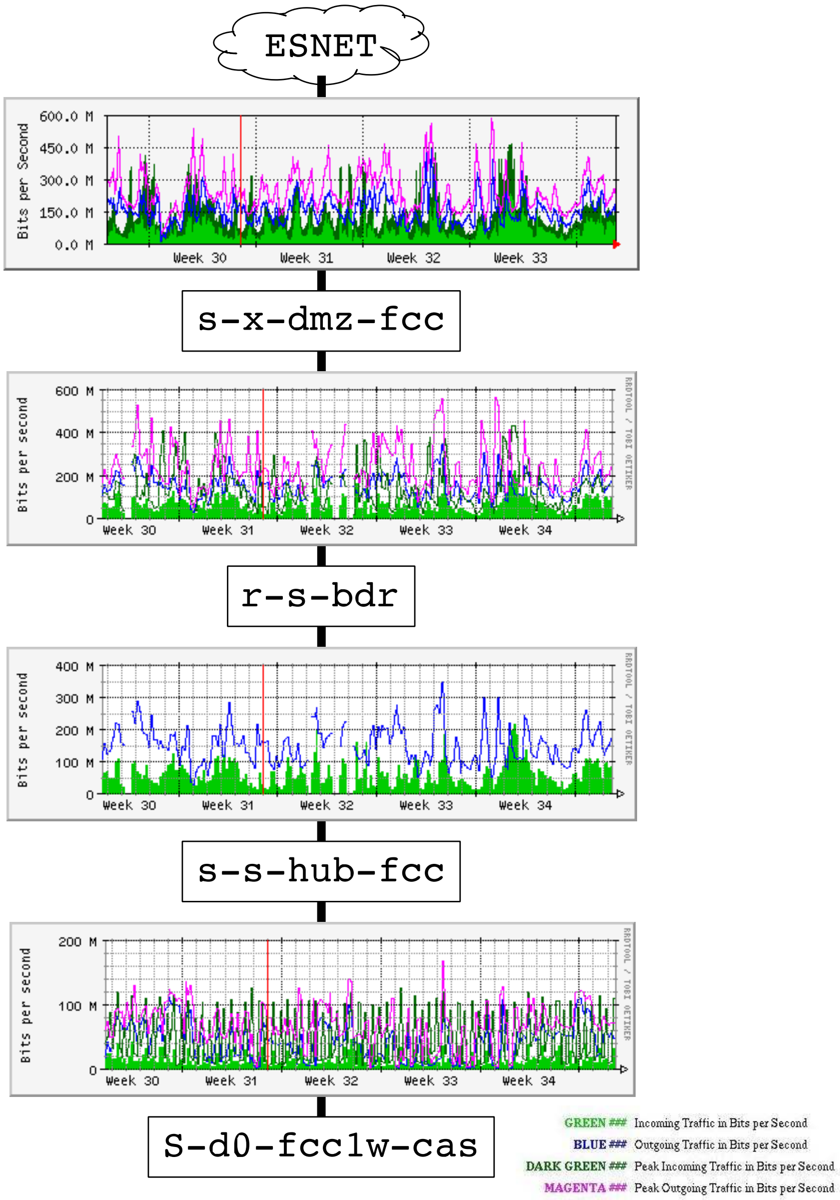

Figure 6.3: Snapshot of traffic from the main DØ switch to ESNET.

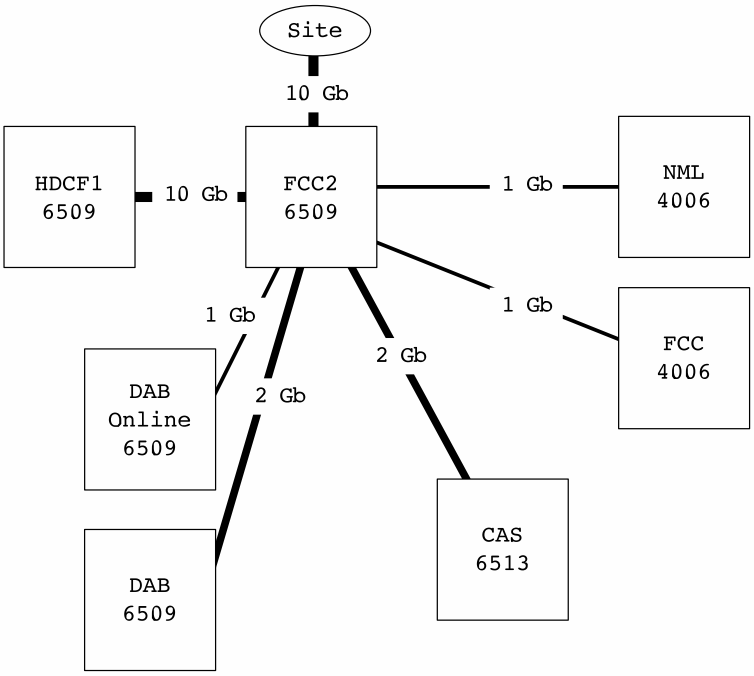

Figure 6.4: Reconfigured DØ network.

6.3 Prospects

Upgrades to the network are necessary for several reasons: a physical relocation of nodes, a greater number of supported nodes, higher bandwidth requirements, and improved flexibility of configuration. Figure 6.4 indicates a proposed reconfiguration of the DØ network which is in progress for completion at the end of FY04, though may also require FY05 funds. The principal components of the upgrade are:

- The current CAB 6509 will be upgraded with a new Supervisor module, giving it a 720 Gb/s backplane and the capability of supporting newer copper Gb and 10 Gb modules. This switch, now renamed the FCC2 6509, becomes the center hub to which other DØ switches are attached. The existing FCC2 CAB worker, server, and mover nodes remain attached to this switch.
- A new Cisco Catalyst 6509, named the HDCF1 6509, will be purchased to support CAB and Reconstruction nodes in the new High Density Computing Facility (the old Wideband Hall). This switch will be outfitted with six 48-port copper Gb blades and a 4-port 10 Gb blade, providing a total of 292 ports. It is expected to initially support 200 CAB and 80 Reconstruction nodes. The HDCF1 6509 will be uplinked to the FCC2 6509 with a 10 Gb connection.
- The 10 Gb module in the FCC2 6509 will support an uplink to the FCC core router and then on to the external network. This will facilitate high bandwidth utilization of the Starlight link for data transfer to/from remote sites.

In FY05, to support additional CAB and Reconstruction nodes to be placed at HDCF, an additional 6509 chassis, supervisor, and copper Gb blades will need to be purchased. The exact number of ports required is dependent upon the number of nodes to be purchased. Future network enhancements also include the upgrade of the supervisor module in the CAS 6513 switch to allow copper Gb connections. Finally, a 10 Gb blade would be added to the CAB 6509 to improve the connectivity of the ClueDØ servers to the SAM/enstore servers.

Chapter 7

Data Handling, Resource Management and Tools

DØ uses the SAMGrid data handling system at all levels in its online and offline systems, to move files, manage resources, and handle process bookkeeping. SAMGrid is supplied by the Fermilab Computing Division via a joint project with participation from the Run II Computing and Analysis Department, the CCF Department, and the CEPA Department, as well as external effort from the GridPP project and other collaboration effort. The experiment supplies tools that interface with SAMGrid as well as with simple filelists for small tasks. At this point in time, these tools are being revisited in order to provide more commonality and maintainability.

7.1 SAMGrid Architecture and Technology

SAMGrid is a system built around the sequential delivery of files which are organized into flexible, user-defined datasets, based on comprehensive descriptive metadata stored in a centrally maintained Oracle database at Fermilab (Run IIs Tier 0 center). The system stores processing information in addition to file metadata. Its services are primarily implemented in a set of servers written in Python, communicating with each other via CORBA. It includes C++ and Python user APIs and a test harness. The Grid job submission tools are built on Condor and Globus tools in the Virtual Data Toolkit.

7.1.1 SAMGrid Status

SAMGrid is in production with local job submission for central (i.e. local to Fermilab) analysis and production activities, remote analysis and production activities, production and user-generated file storage. It is also in production with remote job submission for simulation activities. Detailed status reports from February 04, and a more recent update, are available at [2].

Station	Number of files	Billions of Events	PetaBytes of data
fnal-analysis	10,828,901	213	6.1
fnal-farm	346,621	1	0.2
remote	40,634,502	28	1.5

Table 7.1: Data consumed from 8/1/03 - 8/1/04.

7.1.2 SAMGrid Plans

SAMGrid development projects for the next year, for DØ, are shown on slide 13 of the second hyperlink in Ref [2]. SAM deployment plans for DØ are shown on slide 6. The other slides in this talk note the increasing commitments of the SAMGrid team to other experiments.

7.1.3 SAMGrid Effort

SAMGrid development effort is experiencing a short term shortfall, due to CD reassignments. At the same time, commitments to other experiments are rising and the Grid deployments for DØ are entering an integration/production phase, which requires large support commitments plus a sufficient reserve of development effort to respond to problems arising in that phase. At the same time, the project is expected by both its stakeholders and the Computing Division to move rapidly in the directions of interoperability with various Grid middleware environments, technology changes driven by the need to consolidate effort with other Grid projects in CD, and feature enhancements driven less by experiment than by facility requirements (e.g., tape migration activities, site security enhancements). Some of the effort reduction is expected to be recovered in other reorganization moves. The long term roadmap of SAMGrid for DØ is to continue to supply the experiment with a performant and robust data handling system able to respond to the changing Grid environment and hence able to supply the necessary flexibility in the quest for resources that will grow in importance as Run II datasets multiply at the same time as Run II resource infusions decline.

7.1.4 SAMGrid Performance

The SAM data handling system records a multitude of information on system performance into its main database. Various monitoring tools also run to keep us abreast on short term operations and efficiency of the system. We are actively improving SAM's performance as well as our monitoring tools.

Table 7.1 shows consumed data statistics for the past year (see also Figs 5.1 and 5.2. Fnal-analysis indicates the Fermilab stations devoted to analysis, fnal-farm is the Fermilab reconstruction farm, and remote indicates remote stations away from Fermilab. There are four analysis farms at Fermilab: central-analysis resides on the soon to be decommissioned SGI SMP machine and is mainly used for picking individual events, ClueDØ is the DØ

Station	Number of Projects Run
cab	26,615
central-analysis	16,326
dØppdg-wisconsin-2	12,316
fnal-cabsrv1	12,180
cluedØ	12,094
manhep	11,317
uta-rac	10,973
fnal-farm	8,034
ccin2p3-analysis	7,206
dØppdg-wisconsin	5,734

Table 7.2: Top ten active stations from 8/1/03 - 8/1/04.

Station	% unconsumed	Total files
central-analysis	0.32	2,323,621
fnal-cabsrv1	0.74	3,418,477
cluedØ	1.11	1,777,718
cab	1.21	3,398,645
dØppdg-wisconsin-2	2.15	20,632,432
dØkarlsruhe	2.29	557,680
westgrid-sfu	2.33	30,493
manhep	3.26	91,75,570
fnal-farm	3.81	360,340

Table 7.3: Unconsumed file rates for 8/1/03 - 8/1/04.

desktop cluster mainly used for small jobs and testing, cab and cabsrv1 are independent large farms for running analysis jobs. There are forty-two remote stations around the world.

133,126 projects were run on the SAM stations in the past year (8/1/03-8/1/04). The top ten most active SAM stations are shown in Table 7.2. The Wisconsin, UTA (uta-rac), Manchester, UK (manhep), and Lyon, France (ccin2p3) stations are remote.

SAM keeps track of how many files a project left unconsumed. Causes for unconsumed files are usually disk errors, temporary tape problems, and the user's application crashing before the current file is completely processed. The unconsumed file rates for several SAM stations are shown in Table 7.3 (total files include both consumed and unconsumed files). It is interesting to note that the failure rate is quite a bit lower for central-analysis than for other Fermilab stations. This is likely because the other stations are linux clusters with inexpensive IDE disks having a copy failure rate of about 1%. In fact due to this rate the CRC of files must be checked after a copy to ensure no corruption occurred.

7.2 Tools

In addition to data handling and grid infrastructures DØ must develop tools to manage DØ-specific applications and their interactions with this infrastructure. This includes applications such as reconstruction software aimed at a production environment as well as user applications. This section describes the toolset currently under development to meet these needs.

7.2.1 OS Compatibility, Run Time Environment and Sandboxing

In order to submit DØ applications to grid resources, the binary executable, runtime libraries and configuration files must be packaged for distribution and the user must have some reasonable assurance that the package will actually run in the remote environment despite "local" differences. To address these needs DØ has developed an OS-compatibility system, a developer-level system to publish lists of required files (dØrte) and a sandboxing mechanism to package the user environment and pass it to SAMGrid.

Care must be taken to ensure that DØ executables will run in a variety of operating system environments. Even within Fermilab this is a concern as, in general, the online and analysis farms do not run the same OS as the desktop cluster (ClueDØ) on which the code is developed and initially tested. Once DØ applications are regularly submitted through grid interfaces it is to our advantage that our applications can run on a variety of linux-based systems. To this end the DØ release managers have implemented an OS compatibility system consisting of UPS/UPD products which can accompany the DØ release. This system provides enough system-level code with each release to allow executables to run on many flavors and versions of linux. This system is already in use in production release p17.

The DØ runtime environment (dØrte) provides developers with a straightforward mechanism to publish a list of required libraries and runtime files for their component or executable. This is already in use to create the Monte Carlo (MC) tarball distributed to remote production farms. It's ease of use at the developer level is essential to keeping responsibility for runtime files at the right level: component authors must be the responsible for providing the list of files they need.

In considering how best to integrate dØrte with the rest of the DØ tools (eg. dØrunjob) it was felt that, while the publication mechanism was suitable for integration the distribution mechanism was not. In this context the DØSandbox product was created. This is a python library which uses an XML-based description file to find the needed files for an executable in a particular software release. This allowed the creation of a heirarchy of environments starting from the general DØ release and ending with a small "user" tarball which can be quickly and easily distributed to many nodes in a farm. This sandboxing was implemented in a very general way and was considered a useful development for the runjob project. Aspects of this product have therefore been integrated into runjob. Sandboxing of DØ executables is now being tested using this new infrastructure.

7.2.2 DØtools, DØRunjob and Runjob

Historically people in DØ have used either of two independent packages to run their executable:

DØtools

DØtools is a set of scripts (mainly written in bash) developed primarily by users to run DØ executables. Its traditional use was code testing and debugging with single, locally-built executables. However, it has been organically extended to include job submission to various batch systems (eg. ClueDØ and CAB), transportation of user tarballs (i.e. sandboxing) and some awareness of useful SAM features and interfaces. DØtools has even been adapted to a few remote sites in DØ and has worked in those environments. While users appreciate its ease-of-use (it is commandline-driven), the code itself has grown to the point where it is unmaintainable and cannot be extended to a grid environment.

MCRunjob/Runjob/DØRunjob

MCRunjob is a package traditionally used for production activities on DØ MC production farms. Its strength is workflow management: it has the ability to chain together various DØ executables to, for example, use the output of the MC simulation as the input to the reconstruction program. While it was originally referred to as MCRunjob it has since been rewritten, generalized and adopted as a Fermilab Computing Division project known simply as Runjob. The DØ-specific interface to Runjob is referred to as DØRunjob.

The independent development of these two interfaces to DØ executables is not sustainable. While each has played a valuable role in the past we cannot afford to extend each of them into the grid era independently. Further, no longterm maintenance plan exists for dØtools and simple maintenance of such a complex mix of scripts is onerous. For these reasons the experiment has decided to rewrite dØtools as a commandline user interface to DØRunjob. DØtools is currently being re-implemented in python and will in essence become a macro-generator for DØRunjob while maintaining the existing dØtools UI. This addresses user concerns about the user interface to DØRunjob (users strongly dislike a macro language approach to simple tasks) while allowing support for grid tools to be concentrated in Runjob/DØRunjob. This also alleviates dØtools support concerns as it becomes a thin layer on top of DØRunjob. Support for Runjob and DØRunjob is provided by the Fermilab Computing Division and UK personnel respectively. The timeline for this merger is driven by the availability of the new Runjob/DØRunjob product this fall. This new version of the software will contain both the needed interfaces for dØtools and the new sandboxing mechanisms described earlier.

7.2.3 Summary

Support for future production and user operation on grid resources has prompted a careful consideration and rationalization of tools to package, transport and execute DØ code on remote installations. DØrte will be used by individual developers to publish lists of runtime requirements for their components. DØtools will then provide a friendly user interface to the DØRunjob/Runjob suite. DØRunjob can provide the needed sandboxing tools and workflow management needed by the user and also handle interfacing to the SAMGrid infrastructure.

Chapter 8

Budget Summary

We present two different concepts for budgeting. The first concept is "yearly cost" and is targeted towards estimating the best use of the FNAL equipment budget. The second presentation is the "virtual center" that would be required to meet all of DØ 's computing needs.

8.1 Yearly Cost

There are four major costing areas used for equipment budget planning; worker nodes for processing tasks, disk and associated servers, components for the mass storage and infrastructure hardware needed to support the computing and user access. To the greatest extent possible, we use our knowledge about the past use of the system and past hardware costing trends to estimate future need and equipment cost.

In our experience, the data collection rate (seen in Table 8.1), and total number of events (Table 8.2) is a useful estimator for scaling usage as a function of time. We assume an increased rate after the 2006 shutdown, as planned for Run IIb. For outyear projections, given in Table 8.3, we assume that the hardware cost per unit remains constant, while scaling up performance based on the history of the past four years of purchases. The estimate of computing cycles needed for a task (reconstruction, Monte Carlo generation, user analysis) will depend on the efficiency at which a facility can run a task, the rate at which the task needs to be accomplished and the processing time per event. We also allow for the possibility to add some contingency to the estimate.

For the FNAL farm, we assume that the facility keeps up with data collection, and use

	2005	2006	2007	2008	2009
Average event rate	16	30	35	35	35
Geant MC rate	1.60	3	4	3	4
PMCS MC rate	8	8	8	8	8

Table 8.1: Data rates in Hz.

	Current	2005	2006	2007	2008
event rate	1.00E+09	5.05E+08	9.46E+08	1.10E+09	1.10E+09
total events		1.50E+09	2.45E+09	3.55E+09	4.66E+09

Table 8.2: Data sample size.

year	"Nominal GHz"	SpecInt	GHZ
1999		238	0.5
2000		336	1
2001	1.1	475	1.4
2002		672	2
2003	2.6Hz	950	2.8
2004	3GHz	1343	3.9
2005	4GHz	1899	5.6
2006	6GHz	2685	7.9
2007	10GHz	3797	11.3
2008	15GHz	5370	15.9
2009		7594	22.6
2010		10739	31.9

Table 8.3: CPU evolution expectations.

a luminosity profile and measured time per event to calculate the mean time per event, which for 2005 is expected to be 55 seconds per event on a 1 GHz machine, using the current p14 timing as an estimator. The same estimates yield 100 GHz*sec per event in 2006, although we intend to improve the reco speed performance, and therefore assume 80 GHz*sec per event. We calculate the total number of cycles needed to keep up and subtract off the existing system, and then estimate the number of nodes to purchase. Table 8.4 details this process step by step. The "network cost" is an estimate based on number of nodes purchased; however, it doesn't include all costs. The same procedure is followed for the FNAL based analysis system, using the average weekly data consumption and past system usage as estimators. On the analysis system, the usage pattern is less steady than on the farm, however, the average usage scales with the number of events collected and reconstructed. Mechanically, we maintain a series of linked Excel spreadsheets that contain the assumptions, and calculations. Current versions of these spreadsheets are available at Ref. [3]. For the FNAL farm estimate, in addition to the primary reconstruction, an estimate for the number of computing cycles needed for the post-production fixing was made, and found to be within the contingency.

The estimates for disk are detailed in Chapter 5.

Mass storage estimates focus on the amount of tape to buy and insuring that we have enough slots in the existing robots to accommodate them. We do not yet have a good data rate driven model for estimating the number of drives to purchase based on fundemental

Year	2004	2005	2006	2007	2008
Average Rate	16	16	30	35	35
efficiency	70%	80%	70%	70%	70%
contingency	20%	20%	20%	20%	20%
Reco time	30	55	80	80	80
Required CPU		628320	1958400	2284800	2284800
Existing system	0	344947	381649	1302087	1791625
Nodes to purchase	160	92	361	159	57
Node Cost	$320,000	$202,147	$795,244	$350,400	$124,339
Networking Cost	$25,000	$0	$75,000	$25,000	$0
Number of Nodes at FNAL	476	308	573	613	577

Table 8.4: Primary Reconstruction Cost Estimate.

parameters, and typically use the capacity of the system at peak as an indicator that more are needed (or not).

The other large spending category is infrastructure, which covers equipment for database support, interactive use, networking, and I/O support for the farm and analysis cluster. Based on the estimated number of nodes and fileservers to purchase at FNAL, we make a crude estimate of the networking costs, which is refined and then added to the infrastructure total. For other items, we look at the aging elements of the system and work out strategies for replacement. As most of the initial systems were purchased at roughly the same time, many of them are now due to be replaced. We have been budgeting for the past several years to replace various elements; however, in the face of other budget demands, the infrastructure updates are frequently postponed. This is to be expected – the nominal equipment computing budget for the Run II experiments was $2M per year, but starting in 2003 the budget was reduced to $1.5M, and in 2004 to $1.4M with additional calls to support the new High Density Computing Facility. The details of the infrastructure costs are shown in Table 8.5.

As a point of interest, comparisons between the estimates for 2003 and 2004 ($2M and $1.5M guidance) and actual purchases are shown in Table 8.6. In general, it can be seen that within a constant budget, DØ has typically chosen to purchase worker nodes and postpone infrastructure upgrades to meet the immediate physics needs of the experiment. This clearly cannot be maintained.

The current cost estimates for 2005, with outyear projections for 2006-2008 are shown in Table 8.7.

8.2 Virtual Center Value

Our calculation of the value of the Virtual Center is designed to recognize that non-FNAL computing contributions can be used to offset contributions to the DØ operating fund. The value of the center is calculated by assuming that all necessary nodes and fileservers (and

		2005
DØmino	code master nfs	$20,000
	7 2us for login	$28,000
	NIS slave 4 2U	$12,000
	san hardware for SAM cache	$15,000
	dØworld(Dell 2650)	$15,000
DØmino total		$90,000
DØ2ka	Netapps nfs server appliance	$100,000
	production linux boxes NIS server, disk management	$15,000
DØ2ka total		$115,000
Support for new nodes		$115,000
10 Gb Uplink		$40,000
3 48 port Gb Blades		$30,000
Supervisor Module		$17,000
DØora2 Disk		$42,000
DØbbin Replacement		$48,000
Total		$497,000

Table 8.5: Summary of the projected infrastructure costs for networking, the replacement of the aging SMP machines, and database disk.

	2003 Projected	Projected(2)	Purchased	2004 Projected	Purchased
FNAL Analysis CPU	$505,400	$500,000	$470,400	$339,000	$277,000
FNAL Reconstruction	$200,000	$40,000	$200,000	$83,000	$370,000
File Servers/disk	$262,000	$200,000	$111,000	$490,000	$350,000
Mass Storage	$460,000	$285,000	$280,000	$230,000	$254,700
Infrastructure	$640,000	$500,000	$244,000	$290,000	$140,000
FNAL Total	$2,067,400	$1,525,000	$1,305,400	$1,432,000	$1,391,700

Table 8.6: Projected and spent budgets for FY2003 and FY2004.

	Purchase 2005	Purchase 2006	Purchase 2007	Purchase 2008
FNAL Analysis CPU	$417,132	$534,926	$444,304	$385,174
FNAL Reconstruction	$486,194	$623,851	$335,834	$191,522
File Servers/disk	$373,000	$356,000	$325,000	$308,000
Mass Storage	$40,000	$600,000	$300,000	$100,000
Infrastructure	$497,000	$200,000	$200,000	$200,000
FNAL Total	$1,813,326	$2,314,777	$1,605,138	$1,184,697

Table 8.7: Projected budgets for FY2005 through FY2008.

Estimated Value	2005	2006	2007	2008
FNAL Analysis CPU	$417,132	$534,926	$817,048	$738,631
FNAL Reconstruction	$820,089	$1,087,730	$773,295	$543,752
File Servers/disk	$560,000	$688,000	$528,000	$560,000
Mass Storage	$1,182,000	$1,201,000	$1,501,000	$1,501,000
FNAL Infrastructure	$0	$0	$0	$0
MC	$128,353	$170,152	$160,390	$85,056
Reprocessing	$1,792,632	$3,317,845	$3,245,506	$2,940,120
Virtual Center Total	$4,900,207	$6,999,653	$7,025,239	$6,368,560

Table 8.8: Virtual center cost estimates for FY2005 through FY2008.

the crude networking estimate mentioned above) to do all production activities (including reprocessing and MC production) and FNAL analysis would be purchased in the current year. In this way, for past fiscal years, we can calculate a cash value based on a country's fractional contribution to a given activity as measured by the number of events delivered [4].

For future years, we estimate the value in a way similar to that described above. Currently, analysis at remote facilities is not considered to be part of the virtual center; however, that is likely to change once SAMGrid is deployed. Note that the concept of "value" is one that we are still developing. The estimates of the infrastructure value are difficult to make for systems that are purchased and function for many years at time. As an example, clearly, the "value" of DØmino, the Origin 2000 and fiber channel disk, is not the purchase cost of $1.25M, but neither is it obvious that its value is the replacement cost in linux compute nodes and IDE fileservers. We do not currently assign a value for infrastructure. For the mass storage system, we take the value of the silos to be the (used) purchase cost as they cost about the same, assign the AML2 a comparable value, and pro-rate the value of the drives based on capacity and speed. As an example, a 9940b drive is assigned replacement cost for 2005, and 1/2 that cost in 2006 when the 9940c drives will be available. The virtual center cost estimate is given in Table 8.8.

Bibliography

[1] http://www.aps.anl.gov/epics/

[2] http://wwwserver2.fnal.gov/cfdocs/projectsdb/projdetail.cfm?ProjectID=122 (scroll down to see the 4 talks from Feb 18);

http://computing.fnal.gov/docdb/documents/0003/000314/001/ SAM%20Project%20Status%20Jun%2004.pdf (updated project report from 6/16/04).

[3] http://d0server1.fnal.gov/projects/Computing/Reviews/Sept2004/Index.html.

[4] http://d0server1.fnal.gov/projects/Spokes/FinancialCommittee/scrutiny_march04.doc