CDF/DOC/COMP_UPG/PUBLIC/5822

A Tape Handling System for CDF in Run 2

Version 4
January 11, 2002

Edited by
Robert M. Harris
Fermilab Computing Division

Contributions from
Bonnie Alcorn, Jon Bakken, Andy Beretvas, Joe Boyd, Yen-Chu Chen, Rick Colombo, Glenn Cooper, Phil Demar, Dave Fagan, Richard Glosson, Paul Hubbard, Richard Jetton, Matthias Kasemann, Robert Kennedy, Donna Lamore, Mark Leininger, Dmitri Litvintsev, Phil Lutz, Kevin McFarland, Marc Mengel, Luciano Piccoli, Don Petravick, Bill Robertson, Miroslav Siket, David Tang, Rick Thies, Al Thomas, Jeff Tseng, Tony Vaiciulis, Margaret Votava, Terry Watts and Steve Wolbers

Abstract

We discuss the motivation, components and schedule of the tape handling system now being constructed for CDF in run 2. We require robust and supported tape handling for uninterrupted logging of raw and reconstructed data to tape and reliable staging of data to disk for user analysis. The system we are constructing consists of an STK tape robot housing network accessible STK 9940 tape drives operated via the Enstore tape handling software. The schedule consists of three stages, a prototype stage of limited bandwidth that is available now, a first stage with bandwidth adequate for the next 9 months scheduled to begin to be available in a month, and a second stage adequate for all of run 2A scheduled for roughly 9 months from now.

Contents

1 Acknowledgments

This work would not have been possible without the significant assistance of the following departments of the computing division: CSD, CDF, DCD, D0, FSS, ISD and division management. First we acknowledge Matthias Kasemann and Steve Wolbers in division management for their support of this direction. From ISD we thank Don Petravick, Jon Bakken, and Bonnie Alcorn for their hard work in constructing a proposal for CDF tape handling, keeping it on schedule, and their willingness to provide operational support. Much of the schedule section is taken directly from mail messages from Don and Jon. From DCD we thank David Tang, Donna Lamore, Phil Demar and Al Thomas for their support of networking. From FSS we thank Phil Lutz for coordinating power and space in FCC. From D0 we thank Joe Boyd and Dave Fagan for their help with setting up networking connections to our SGI similar to that done for D0. From CSD we thank Rick Thies, for coordinating the removal of tapes from the run 1 STK silo. From ODS we thank Margaret Votava and Luciano Piccoli and from OSS we thank Marc Mengel, for their help analyzing and patching the legacy tape handling system. From the CDF department we thank Kevin McFarland for the migration schedule, Steve Wolbers, Yen-Chu Chen and Miroslav Siket for the plan to use Enstore on the farms, Dmitri Litvintsev and Paul Hubbard for crucial development work on CDF software to implement Enstore at CDF, Robert Kennedy for reviewing the adaptation of CDF software to use Enstore, Rick Colombo for purchasing everything we needed, Richard Jetton and Glenn Cooper for numerous disk installations, and Eric Wicklund, Andy Beretvas, Bill Robertson and Glenn Cooper for the heroic work of emptying the Run 1 silo on schedule. Also from the CDF department we thank Mark Leininger, Paul Hubbard, Richard Glosson, and Dmitri Litvintsev for help with analyzing and patching the legacy tape handling system. From the CDF collaboration we thank Terry Watts and Jeff Tseng for oversight and planning on usage of Enstore, and Tony Vaiciulis for help with problems in the legacy system.

2 Introduction and Motivation

2.1 Informing Data Handling Review and Collaboration

This document is primarily a transcription of presentations made to the CDF Data Handling Reviews [1] on November 19 and December 19, 2001. The review was charged by the spokespeople and the reviewers were Joel Butler, Frank Chlebana, Henry Frisch, Young-Kee Kim, Pekka Sinervo (chair), Jeff Tseng (up to November 30) and Torre Wenaus. The Nov. 19 review approved this tape handling system up to stage 1, and expressed interest in our embryonic plans for stage 2. Although the reviewers, spokespeople and computing division are well informed of the course we are embarked on, the collaboration is not as well informed. This document is primarily intended to inform the collaboration. It is written to present the information available, without going into too much technical detail, and therefore may be both incomplete and imprecise. The editor takes full responsibility for any omissions or mistakes.

2.2 Existing Data Handling System

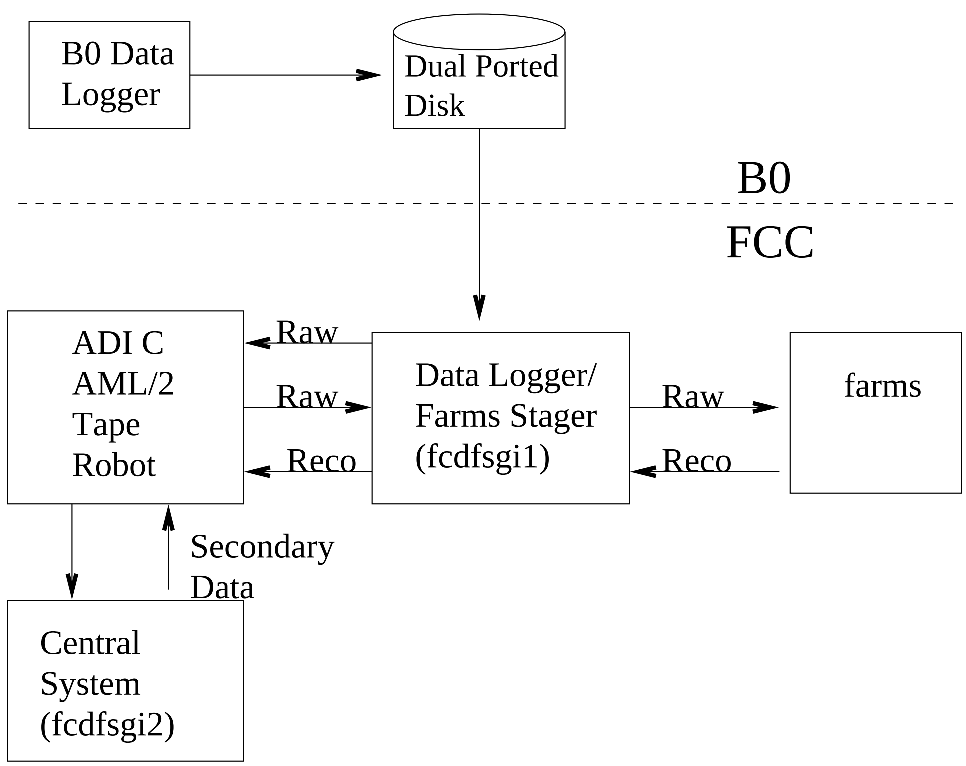

Figure 1: Dataflow in the existing data handling system (see text)

2.2.1 Dataflow

Dataflow in the existing data handling system is pictured in Fig. 1. Raw data from collisions is logged to a dual ported disk by the B0 data logger, and is written from the dual ported disk to the ADIC tape robot. The writing is performed using AIT-2 drives directly attached to fcdfsgi1, which shares the dual ported disk with b0dap31 via the CXFS file-system. The raw data is staged back from the tape robot to fcdfsgi1 to be reconstructed by the production farms, which write the reconstructed data back to the tape robot using fcdfsgi1 again. The raw and reconstructed data are read from the tape robot to fcdfsgi2 where physics group produce secondary datasets and write them back to tape.

2.2.2 Tape Handling Software

The legacy tape handling software at CDF is run by mt_tools, a tape writing package that supports the AIT-2 tape technology, ANSI tapes, tape partitioning and handles errors. The mt_tools package makes calls to OCS and FTT for tape allocating, reading and writing primitive operations. OCS and FTT are CD supported packages which are no longer being developed at the lab. The software that controls the mounts in the ADIC robot is OCS_robot, written on top of OCS.

2.2.3 Higher Data Handling Software

The highest level of DH software is DHMods, the AC++ input/output module for data handling. DHMods is the user interface to access data on disk or tape. DHMods communicates with the Data File Catalog (DFC) via the DataFileDB package. The DFC contains the metadata of files, filesets, tapes and datasets. DHMods knows the status of the data from the DFC, and from the Disk Inventory Manager (DIM). The DIM dynamically caches data on disk that is copied off of tape, allowing users to access data on disk repeatedly without having to go back to tape every time. The data is read from tape into the DIM by the stager, which in turn calls the tape handling software.

Writing to tape is done by the Fileset/Tape Daemon (FSTD) on fcdfsgi1. This software collects data into filesets before writing complete tapes consisting of roughly 5 filesets. DFCTestTapeWrite allows physics groups and users to cluster files into filesets, and is currently used by physics groups on fcdfsgi2 to write secondary datasets to tape, performing a similar function on fcdfsgi2 that FSTD plays on fcdfsgi1.

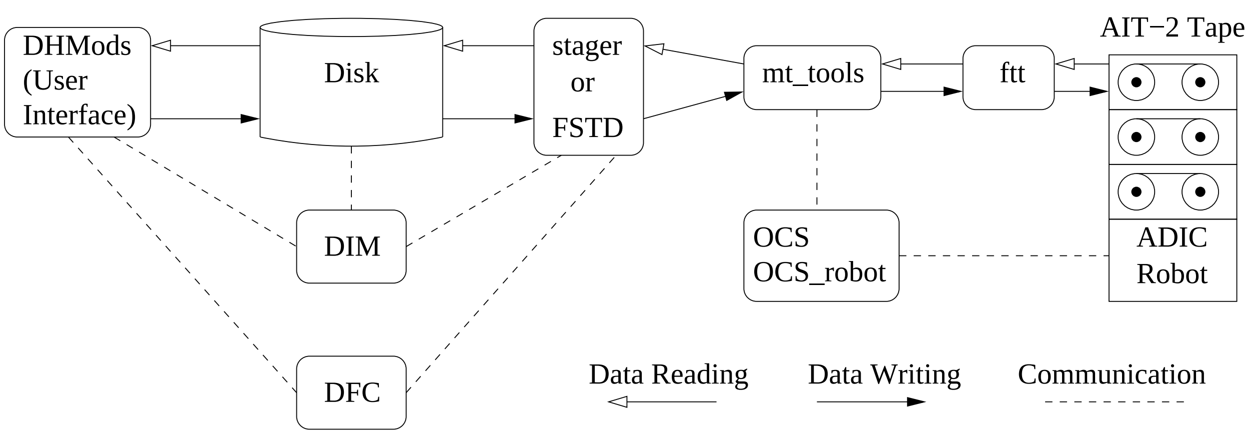

Figure 2: The data flow, higher level software (DHMods, DFC, DIM, Stager, FSTD) and tape handling software (mt_tools, ftt, OCS, OCS_robot) in the existing data handling system. Adapted from figures in [2].

2.3 Problems with Legacy Tape Handling

2.3.1 Errors during Operations

Errors during tape operations were expending roughly two FTEs of effort in November 2001, and inoperability of the system was interfering with both the reconstruction of data on the farms and user analysis on fcdfsgi2.

There were tape reading problems, including short reads of tapes resulting incomplete data for farms and user analysis. The majority of tape reading problems was due to confusion in the system as to the status of tape drive allocations and mounts. Sometimes the drive would get deallocated but the tape would remain mounted in the drive, unbeknownst to the system, causing problems when the next job would go to mount a tape in the drive. Other times both the tape would remain mounted and the drive would remain allocated, although the job would complete successfully, leaving the drives unavailable to other jobs. A recurring problem for months was when LSF jobs were killed by the user, leaving the drives allocated and the tapes mounted, unavailable to other jobs. All these problems required manual intervention to cleanup the tape queues so that drives could be freed. These tape reading problems were primarily due to the tape handling software. In late November 2001 manual intervention was required every few hours.

In addition to reading problems, more than 1% of all tape writes failed, requiring a half day of manual intervention to cleanup and resubmit the tape write manually. Though seldom seen in November, these write failures should occur every 2 days when we reach raw data taking and production rates anticipated in only a few months. The write failures were due to problems with the media, and were therefore irreducible in nature.

In early December we discovered that the ADIC robot had damaged 38 AIT-2 tape cassettes. Although all data was recovered on the data handling test stand, the potential for data loss in the future was undeniable.

2.3.2 Expertise & Support

In addition to operational problem, there was a shortage of expertise and support for the legacy tape handling system. The designer and primary expert for both the hardware and software of the tape handling system resigned on October 12. The choices that CDF made in tape handling are not supported elsewhere in the laboratory. The tape handling package mt_tools and the mount package OCS_robot are CDF in-house software packages, and were simply not being supported at the required level. The underlying tape handling software, OCS and FTT, were no long being developed at the lab and their support was severely limited. The AIT-2 commodity tape drives and media were rejected by the lab due to their unreliability and fragility. The legacy system was simply not supportable over the long run.

2.3.3 Strategy to Address Problems

We presented the situation described above to the CDF Data Handling Review on November 19 [1], and recommended the following two front strategy. We would patch the existing problems with the tape handling system sufficiently to allow us to operate it for a few months until a replacement was available. The effort to patch the legacy system, described in Appendix A, would proceed in parallel with the construction of a replacement system, described by the remainder of this document.

2.4 Conceptual Description of Tape Handling Replacement

We have begun the following tape handling replacement at CDF. The Sony AIT-2 drives, firmware and media in the ADIC AML/2 robot were supported by the CDF department in the computing division (CD/CDF). This hardware is replaced with STK 9940 drives and media in an STK powerderhorn robot supported by the ISD department of the computing division (CD/ISD). These are additional resources being devoted to CDF by the computing division. The legacy tape handling software described in section 2.2.2 were supported by CD/CDF. This tape handling software is replaced by Enstore, supported by CD/ISD. The higher level software described in section 2.2.3 remains primarily the same and remains supported by the CDF data handling group. Small modifications were necessary to the stager, the FSTD, and DFCTestTapeWrite to use Enstore instead of the legacy tape handling software.

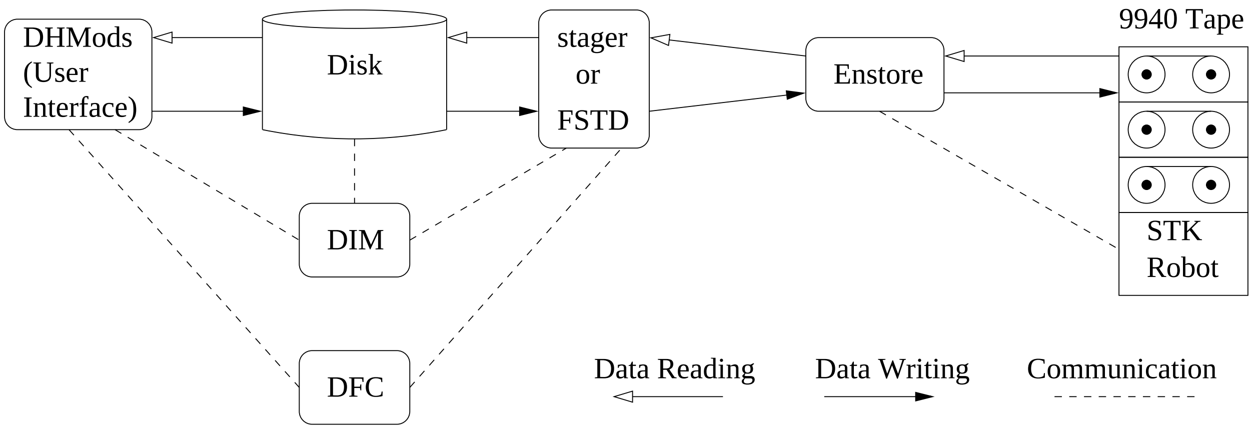

Figure 3: The data flow, higher level software, and tape handling software in the replacement data handling system. Same as figure 2 but the tape handling software has been replaced with Enstore and the robot, tapes and drives have been replaced with an STK system. Adapted from figures in [2].

2.5 Benefits of Tape Handling Replacement

2.5.1 Support

We are migrating to a tape handling system that is fully supported by the computing division. An example is the public STKEN robot that contains data for the following groups: Auger, Boone, BTEV, CKM, CMS, D0,E791, E815, E831,E872, KTeV, Miniboone, Minos, SDSS, Selex and Theory. The STKEN robot now also contains test data and Run 1 data for CDF. The STKEN robot has been in use for about a year and a half. The tape handling system will be supported by computing professionals in CD/ISD. They will support the drives in the robot and the robot operations as they do now for the rest of the lab, which avoids CD/CDF wasting effort on a fragile ADIC + AIT-2 system. CD/ISD will support the Enstore software as they do now for the rest of the lab, this avoids CD/CDF wasting effort on tape handling software used only by CDF. Details of support are discussed in Appendix B.

2.5.2 Robustness

The Enstore tape handling software has been in use for over a year and a half servicing multiple experiments, and the operational expereience gained has resulted in a robust tape handling system. Further, it is designed in a client-server architecture that is not prone to the same kind of errors that occur frequently with the legacy tape handling software. For example, the problem that occurs in our legacy system when an LSF job is killed by the user, and the tape remains mounted in the drive preventing the drive from being used by others, does not occur in the Enstore software. After a number of tries for Enstore to communicate with the killed job, the software would simply time out and dismount the tape. In short, the Enstore software handles error conditions significantly better than our legacy tape handling software does.

2.5.3 Hardware Reliability

The STK 9940 drives and media are data center quality, in contrast to the AIT-2 drives and 8 mm cassettes. The STK drives are in service at FNAL, BNL, CERN and DESY (9840 variant). At Fermilab CD/ISD has encountered less than 1 write failure in every 1000 tapes written (under 0.1%) in contrast to the AIT-2 media which has more than 1 write failure in every 100 tapes (more than 1%). The relative quality of the media is obvious on inspection. The AIT-2 8 mm media looks exactly the same as a cassette tape you would commonly use in your video camcorder: a small and light plastic cassette. The STK 9940 media is intended only for data center drives and robotic use: a large and heavy cartridge with components that have been engineered for repeated use.

2.5.4 Network Access

Enstore requires that the tape drives be network accessible, which potentially solves a number of problems with data distribution. For example, the new central analysis facility [4] will be a user PC farm that will be able to get it's data over the network, directly connected to the same switch as the drives in the STK robot. In contrast, the software of the legacy data handling system required that the node accessing the data be directly attached to the tape drives, which would be problematic for the new system for a number of reasons. First it would require that the data handling software be ported and maintained under Linux as well as IRIX. Second it would require a large number of additional tape drives be installed and maintained for the new CAF. Third it would allow only a few dedicated nodes to have access to the tape robot, while network access potentially allows any of the nodes to have access if necessary.

Network access to the drives is a step on a logical path to making the data on tape available to the CDF trailers and the universities. This could potentially be done using either the DH software on the desktop or directly using Enstore encp commands to read files. Priorities would be granted to raw data archiving and data reconstruction, but there would still be bandwidth available for other users. To be able to read using Enstore without caching to disk the network to the trailers and universities would need to be improved, however, caching is available in Enstore via dCache, a joint FNAL-DESY project to allow slower network connections to read data and not monopolize tape drives for long periods of time. The ability to read data over the network does not preclude the possibility of copying tapes and distributing them to universities. Indeed, the performance of AIT-2 drives and media is adequate for occasional use at a university, although not for repeated and heavy production use at Fermilab.

2.5.5 Miscellaneous Data Archive

The Enstore system can also be used to archive data of various sorts, without the use of heavy bookkeeping machinery like the data file catalog. Since Enstore has it's own flexible metadata database, pnfs, user's and physics groups could save miscellaneous data to tape from their static data disks and be able to search for them again in pnfs. The ability to store all kinds of data, such as ntuples, or large files of various formats, may reduce the needed size of static disks which otherwise grow constantly without any ability to empty them. Users are still encouraged to group files for tape writing, as this results in more efficient reading later on.

3 Software

The Enstore software will replace the existing tape handling software, as described in section 2.4. A review of this replacement concluded that it was both feasible and beneficial [2], and a detailed discussion of adapting CDF software to use Enstore can be found there. Here we briefly discuss Enstore and the process of mapping the concepts of the legacy system onto our use of Enstore.

3.1 Enstore

3.1.1 Overview

Enstore [3] is distributed software to move files to and from tape over the network. There are two levels of access. The command encp can be used within FNAL site to copy files with deterministic tape-rate performance. Buffered access is available through dCache for off-site and low performance on-site reading. The files are viewable through an NFS exported file catalog called pnfs, which has the look and feel of a UNIX file system, but only the metadata is exported. Thus users can do an "ls" of the contents of Enstore, and get the path that files are stored in, but only metadata is displayed over NFS and the user cannot directly copy the files over NFS. The actual data files are accessed using encp, modeled after the UNIX "cp" command.

3.1.2 Fault Handling

The Enstore system detects broken components. Tape drives having many error configure themselves out of the system. Tapes having errors on many drives are flagged for investigation. Tape drive cleaning is automated.

There is internal monitoring in Enstore. Plots showing system peformance as a function of time are regularly accumulated and displayed, as shown in
http://www-stken.fnal.gov/enstore/cron_pics.html. There are alarms and lists of work for normal administration. For example, http://www-stken.fnal.gov/enstore/enstore_alarms.html. Detected faults are integrated into an automated 24x7 pager system. See for example, http://www-stken.fnal.gov/enstore/enstore_saag.html.

3.2 Mapping the Concepts

In order to utilize Enstore and replace the legacy tape system quickly, we plan to map the CDF data handling system to Enstore without fundamentally changing either. Enstore writes files to tapes, but the user does not control completely which tape the file goes on, which is different from the legacy system. Enstore does not support filesets, collections of files mapped to a physical partition on the tape in the legacy system. The developers of Enstore are willing to consider the concept of filesets in the future, but physical partitioning of tapes is not necessary and will not be part of the tape handling system in the future.

We plan to logically map the physical concepts of the legacy data handling system onto Enstore. We'll create logical filesets, not corresponding to partitions on a tape or even necessarily to a single tape always, and logical tapes, simply collections of logical filesets but not necessarily corresponding to one physical tape. Users request datasets, logical tapes, logical filesets, runs or files, as in the legacy data handling system. The files requested may be spread over more than one physical tape. As in the legacy system, the smallest unit of data read from a tape to the DIM is a fileset. A fileset will most likely be on a single tape because a fileset is roughly 20% of a tape and the filesets are written sequentially to tape. About 20% of the time a fileset will be spread over two tapes.

3.3 Proof of Principle

The first test of this mapping was completed in early November 2001. Using the public STKEN robot, Paul Hubbard and Dmitri Litvintsev demonstrated with only a few hours of effort that Enstore could be used by our legacy data handling system to both write and read data on tape. Only the stager and DFCTestTapeWrite were modified: the DIM, DFC and DHMods were not changed. This showed that Enstore could be used by our data handling system without substantially modifying either.

4 Hardware

4.1 Capacities and Costs

4.1.1 Legacy System

The existing ADIC AML/2 robot has 20,000 slots for tapes. Each AIT-2 8 mm cassette tape has a capacity of 50 GB, and CDF obtains an average usage of 45 GB. At $60 per tape, the media cost to CDF is $1.3/GB. AIT-2 drives read at 6 MB/s. For CDF the tapes are partitioned, and the partition marks must be written to the tape before the data can be written, so the head must pass over the entire tape twice to write the data once, yielding a write rate of only 3 MB/s. The AIT-2 tape drives come in special canisters for installation in the robot, and consequently cost roughly $5,000 each for the robot, but only roughly $3,000 each for stand-alone models employed by CDF universities. CDF currently has 32 AIT-2 drives: 12 on fcdfsgi2, 8 on fcdfsgi1, 2 on fcdfsun1, and 10 not yet used. The 8 drives on fcdfsgi1 would be capable of sustaining the rough DC rates 9 MB/s raw data logging, 9 MB/s production farms reading, and 9 MB/s production farms writing, if the operational efficiency of the DH system were close to 100% which has not been the case in the recent past. Of the 12 drives on fcdfsgi2, 11 are allocated to users, and could sustain 66 MB/s at 100% operational efficiency.

4.1.2 Replacement System

The STK powderhorn silo has 5,500 slots for tapes. They usually cost $200K each but have been on sale for a few months at $85K. We already have one STK silo, which we have been using for Run 1 data, and we have emptied it to use for run 2 data. Each STK cartridge has 60 GB capacity, and Enstore gets an average usage of 60 GB since it fills up tapes with files. The tapes cost $78 each, which is a media cost of $1.3/GB, the same as CDF is now getting with AIT-2. Each STK 9940 drive has a read and write rate of 10 MB/s, the read and write rates are the same because we will not partition the tapes unlike the legacy system. The STK 9940 drives cost $27K each retail. We have already purchased 10 drives for our existing silo and the order cost CDF $250K. STK is planning for a new drive, the STK 9940B, for release to the market in summer 2002. The STK 9940B is intended to sustain a bandwidth of 30 MB/s, will write 200 GB on the same cartridges as the STK 9940, and will be able to read data written by the older STK 9940 drives.

4.2 STK Silo Logistics

Space in FCC is an issue for new STK Silos. D0 has purchased two Silos, and ISD has purchased two silos, both of which are going into the Mezzanine between the 1st and 2nd floors. The Mezzanine is the only place with space for new Silos at the moment, and construction has just been completed and the other Silos are being installed with expected completion sometime in February, so CDF could not quickly obtain and install a new Silo. There are currently two other STK silos at the lab. The public STKEN robot with 2700 free slots on the 2nd floor of FCC is currently being used to log D0 collision data as well as supporting 15 other groups at the lab and CDF use in roughly 100 slots. The CDF Run 1 STK silo, on the second floor of FCC right next to the STKEN robot, was the natural choice for immediate CDF use.

4.3 Staged Proposal

On November 19 we made the following proposal to the CDF internal data handling review and updated it on December 19.

4.3.1 Prototype

The prototype stage shown in figure 4 consists of 2 STK 9940 drives in the run 1 STK Silo, renamed the CDFEN robot. Each drive is connected to a mover PC which in turn uses 100 Mb/s Fast Ethernet to connect to the CDF offline switch. The total bandwidth capability is 20 MB/s, and the ultimate capacity of the robot is about 330 TB. To make the prototype available as soon as possible, the prototype shares five Enstore administrative servers with the public STKEN robot. The prototype also shares with the STKEN robot the robot control computer which runs the ACSLS robot control software. The prototype will primarily be used by fcdfsgi2, and to handle the rates a Gigabit NIC will be attached between fcdfsgi2 and the CDF offline switch. A single processor on fcdfsgi2 will be dedicated to network traffic, allowing a sustained rate of up to 30 MB/s for fcdfsgi2. The prototype was made available to offline operations for data migration and testing on January 7.

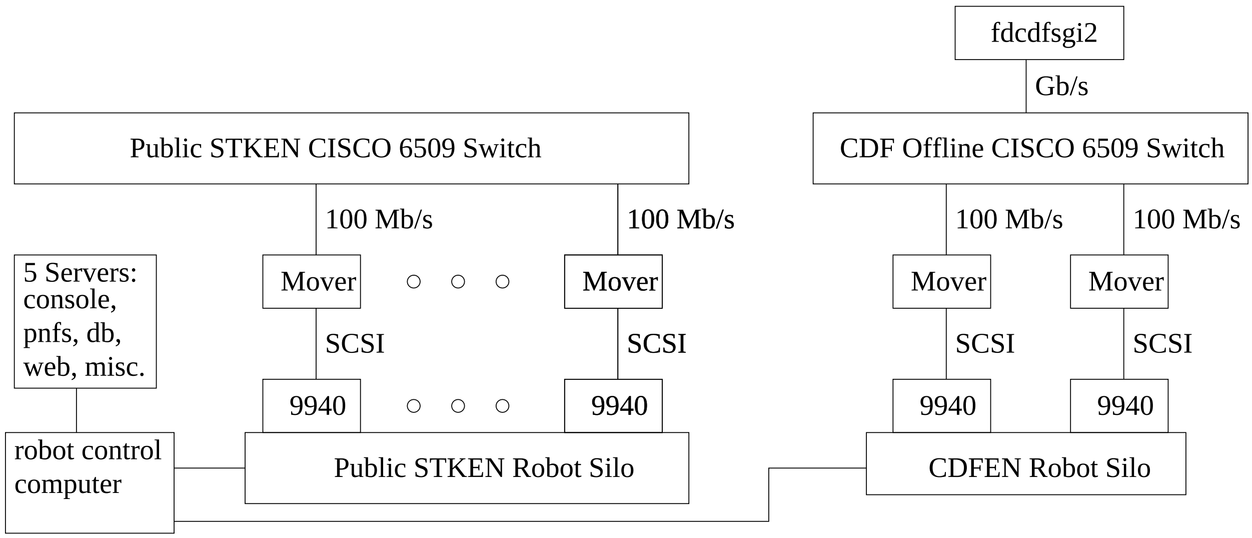

Figure 4: The prototype stage of CDF tape handling system. The prototype consists of 2 9940 drives and mover nodes in the CDFEN robot shown at right, sharing the servers and robot controller with the STKEN robot shown at left. There is a Gb/s connection from the offline switch to fcdfsgi2 with a dedicated CPU for routing data from CDFEN to fcdfsgi2.

4.3.2 Stage 1

Stage 1 showin in figure 5 consists of 10 STK 9940 drives in the CDFEN robot with it's own Enstore administrative servers. The drives were installed on December 31. The total bandwidth capability of 100 MB/s should be adequate for raw data logging, production farms reading and writing, and users for up to roughly the fall of 2002. The capacity of roughly 330 TB should be sufficient until fall of 2002 as well. The system should be available to offline operations in late February of 2002.

4.3.3 Stage 2

Stage 2 will likely exploit the STK 9940B drives that we hope will be available in Summer 2002, and will expand the data capacity needed after fall of 2002. If the combined accelerator and detector data taking efficiency is 30%, then we log roughly 1 TB of raw and produced data a day, and the CDFEN robot fills up around September. We roughly plan to purchase an additional STK silo and equip it with 7 STK 9940B drives for a total data rate of 210 MB/s in a new silo with 1.1 PB of capacity, expanding the total capacity to 1.4 PB. The existing silo can be upgraded with 5 STK 9940B drives to sustain analysis data rates of 150 MB/s for the run 2 data existing at that time.

If the STK 9940B drive is not available, we could buy a single silo and equip it with 21 STK 9940 drives, and add 5 more STK 9940 drives to the existing robot. This will yield the same rate capability of 210 MB/s in the new robot and 150 MB/s in the old robot, but only 0.66 PB of total capacity. We would then have to buy one additional robot in 2003 to expand to at least 1 PB of capacity.

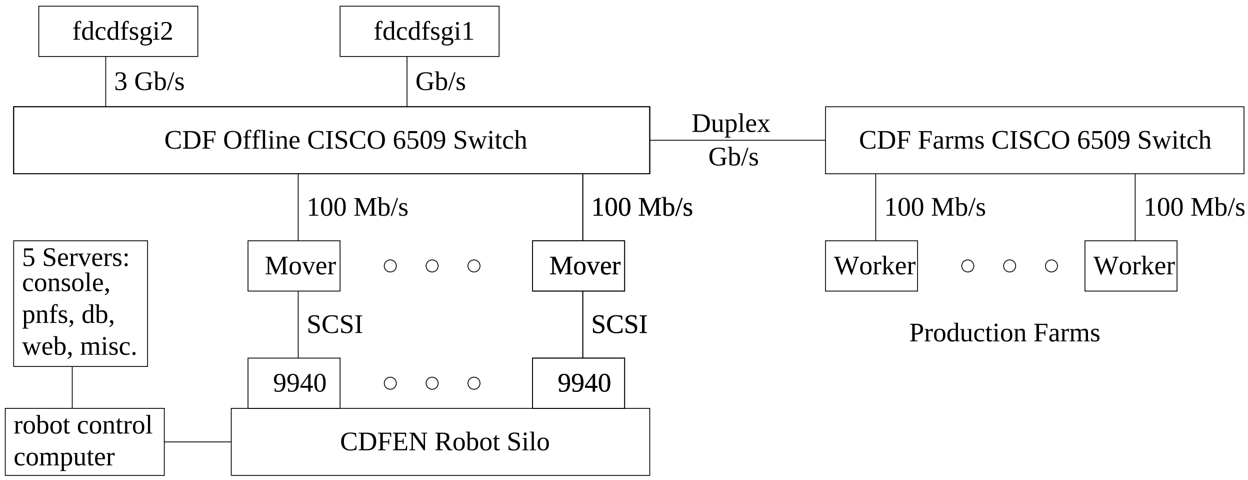

Figure 5: Stage 1 of CDF tape handling system consists of 10 STK 9940 drives and mover nodes in the CDFEN robot with it's own servers and robot control computer. There will be three Gb/s connections from the offline switch to fcdfsgi2, and a duplex Gb/s link between the offline and farms switch to allow the farms worker nodes to directly stream data to and from the Enstore robot. The online, not shown, should eventually also stream directly to Enstore. On the way to this configuration there will be an intermediate step for both the farms and the online in which they will employ a Gb/s connection from fcdfsgi1 to write data to Enstore.

4.4 Networking and Data Flow

The network backbone of the system is the CDF offline Cisco 6509 switch with 32 Gb/s capacity. The CDF Enstore robot is attached to the offline switch, as is fcdfsgi1 and fcdfsgi2. For each 30 MB/s a Gigabit NIC needs to be attached between the CPU and the offline switch, and one of the SGI processors needs to be dedicated to network traffic.

4.4.1 Production Farms

The production farms currently write data to the ADIC AML/2 robot using fcdfsgi1. To transition to use Enstore as quickly as possible, the farms will initially write to Enstore using fcdfsgi1 as well, hopefully in late January. A Gigabit NIC on fcdfsgi1 will be temporarily needed for this. The final plan for the farms is to read and write data all directly from the worker nodes to Enstore, and this is scheduled for late February. A Gigabit link between the farms Cisco 6509 switch and the offline switch is being installed to accommodate the traffic.

4.4.2 Online

The online currently writes data to the ADIC AML/2 robot using fcdfsgi1. To ensure stability in online data logging, the online will initially write data to Enstore via fcdfsgi1. This should begin in March, and a Gigabit NIC on fcdfsgi1 will be temporarily needed for this. Once stability has been demonstrated, we hope the online machines will write data directly to Enstore to eliminate fcdfsgi1 as a possible point of failure.

5 Tasks and Schedule

5.1 Prototype

5.1.1 Run 1 Silo Migration and Conversion

The run 1 STK silo contained 1.3 TB of secondary and Monte Carlo datasets. This was roughly 2000 cartridges of 200 MB and 800 MB capacity. To accommodate the data 1.5 TB of RAID disk were added to the silo staging pool on cdfsga. A CDF team lead by Eric Wicklund and including Andy Beretvas, Bill Robertson, and Glenn Cooper, copied all the data from the Run 1 Silo to the staging disk on cdfsga. The copying began around December 1 and completed on December 21, 2001. The run 1 data is also being copied to the public STKEN robot for protection. Between December 22 and 24 the CSD department of the computing division removed all the data tapes from the Run 1 silo, boxed them, and put them in storage. By Christmas CDF had turned over to ISD the run 1 STK silo for their support in Run 2. Between December 27 and 31, STK removed all the old drives and cabinets from the silo and installed 10 STK 9940 drives. The new CDFEN robot was turned over to ISD to complete the remaining installation and future operations.

5.1.2 Mover Nodes

Each STK 9940 drive requires a mover node to move data between the drive and the network. The node is a dual PC with a Lancewood motherboard running Linux. The two nodes for the prototype were donated by the CDF farms group, setup and connected to the network in mid-December and connected to the drives via SCSI at the end of December.

5.1.3 Network Interface Card for fcdfsgi2

The network interface card for fcdfsgi2 arrived at CDF on January 3. It is scheduled to be installed on January 11. D0 has experience dedicating a single processor to route traffic over the NIC, ncecssary to achieve 30 MB/s per NIC, and will help us do that configuration on January 17.

5.1.4 Tasks by Department

The list of tasks by department of CD or STK is roughly [5]

1. CDF, get two 9940 drives. Status: 10 drives purchased and installed.
2. CDF, Copy data from STK Silo to cdfsga. Status: completed.
3. CSD, Empty tapes from STK Silo. Status: completed.
4. CDF, order 9940 media. Status: 30 tapes here, 60 soon, 910 on order.
5. CSD, install 9940 media. Status: complete for 30 tapes.
6. STK, install drive bay, 9940 drive installation. Status: complete.
7. STK, transfer control of silo to STKEN Sun acsls computer. Status: complete.
8. ISD, configure 2 mover computers. Status: complete.

9. ISD, create CDF pnfs database. Status: complete.
10. ISD, allocate tapes. Status: done.
11. ISD, test silo + drives + movers. Status: completed on Jan. 7.

5.2 Stage 1

The main difference between the prototype and Stage 1 is eight additional mover nodes and five administrative server nodes. The setup of the mover nodes is relatively straightforward and has already begun. The server nodes will take roughly 7 weeks to setup and the work for that began in mid-December and should be complete by late February. The detailed project plan is presented in Appendix C

5.2.1 Tasks by Department

A breakdown of the tasks by department of CD are [5]

1. CDF, get remaining 8 STK 9940 drives from STK. Status: done in prototype phase.
2. STK, transfer control of silo to new Sun ACSLS computer.
3. FSS, provide placement and power for mover and server nodes.
4. DCD, Networking including
   - 15 fast Ethernet connections in CDF offline switch.
   - 2-6 fast Ethernet back-end connections for robot control.
   - new subnet for Enstore in CDF offline switch.
   - laying fast Ethernet trunk cables to Enstore.
   - Install new Gigabit connection between CDF offline and farm switch.
5. CDF, buy Gigabit NICs
6. D0 and CDF, dedicate a single SGI processor to each NIC.
7. DCD, configure offline switch for each Gigabit interface.
8. CDF, deliver farm Lancewoods for Enstore movers + servers. Status: done.
9. ISD, configure movers for Enstore.
10. ISD, configure servers for Enstore.
11. ISD, configure new pnfs databases.
12. ISD, configure new file and volume databases.
13. ISD, configure backup of nodes and metadata.

5.2.2 Some ISD Tasks in Detail

Some of the steps that ISD needs to complete, in a little more detail, are listed below [6]

1. Construct the servers, and place all materials. The one exception is that the ACSLS computer will not control the CDF silo, which allows the prototype phase to operate continuously during the stage 1 construction period.

2. Configure all the Enstore software on the CDFEN system.

3. Test, allowing the prototype phase to continue to operate continuously from the STKEN system. The test is roughly

   - Take the drives not in the prototype system and some media.
   - Configure in a ”media changer” to point to FNTT, the STKEN acsls computer, which is controlling the CDF library during this phase.
   - Beat upon the system to make sure the servers all work and reduce the technical risk before change over.
   - When satisfied tear down the temporary databases and scratch any media used in the test

4. Have a shutdown. After the shutdown all CDF data will be in the CDFEN system. Some down-time here for both the prototype and Stage 1 systems.

   - Move all CDF tapes in the STKEN silo to the CDFEN silo.
   - Use the CDF ACSLS computer to control the new silo.
   - Move the pnfs meta data to the CDFEN pnfs computer.
   - Move the volume and file data bases to the new computer.

5. Commence operations in the new library.

5.3 Stage 2

We will need one additional Silo for Stage 2 in the fall, and we will probably order it soon to take advantage of the $85K price being recently offered by STK. We will then wait until summer to order additional drives, hoping to get the STK 9940B drives. The drive testing that will take place this winter and spring 2002 will tell us whether the STK9940B drives are acceptable, or whether we must order a larger number of STK 9940 drives.

5.4 Use and Migration

The legacy and CDFEN tape handling system will be supported simultaneously for roughly a 3 month transition period. The option for use and migration to the CDFEN robot that presents the least technical risk is the one that involves logging the raw data to the CDFEN robot as the last step in a multi-step migration process, outlined below.

1. Begin copying data from ADIC to CDFEN robot in January. Takes a couple of months to copy all data. We plan to start with data from the end of 2001 and work backwards, coming back for the most recent data later.

2. Test users read and write secondary datasets to CDFEN in January. Data that is in both robots will be read from CDFEN by default, allowing CDFEN to be used almost immediately.

3. Farms begin read/write testing to CDFEN in late January. This first step will use fcdfsgi1 for I/O to CDFEN as in the legacy system.

4. Farms begin writing produced data to CDFEN only in February. By late February the worker nodes should be reading and writing directly from CDFEN, eliminating the need for the farms to use fcdfsgi1.

5. Raw data logged to CDFEN from fcdfsgi1 by March. This step is taken after most problems have been ironed out by user and farms testing.

6. Cease use of ADIC robot completely by April 1. We should have copied all data from the ADIC robot by this date, and will no longer need to operate it.

Appendix

A Patching the Legacy Tape Handling System

A.1 Team of CD Tape Handling Experts

As discussed in section 2.3.1 by November the level of tape handling errors was severe, requiring human intervention every few hours. The computing division pledged CDF support, and a team of tape handling experts was assembled. From the ODS department Margaret Votava lead the team, along with Luciano Piccolo from ODS. Marc Mengel from OSS provided valuable expertise. These outside experts assisted Mark Leininger, Paul Hubbard, Richard Glosson and Dmitri Litvintsev from the CDF department of the computing division and Tony Vaiciulis from the CDF collaboration in analyzing the problems.

A.2 Votava Summary and Recommendation

On November 29 Margaret Votava sent around an e-mail message detailing the nature of the problems and making a recommendation. Here are the first two paragraphs summarizing the situation and recommendation

> The system overall does not handler error conditions well. Fixing this properly would lead to hundreds of code changes which would require a large development and testing effort. Fortunately, un-handled errors seem to leave the system tape drives in one of a small number (currently believed to be 3) "confused" states. The end resulting stage can currently be recovered from manually deallocating drives and putting tapes back in the shelves and in some cases stopping and restarting servers. The latter is extremely disruptive to the community at large.
>
> Setting up a test stand, getting latest versions of the code in hand, making changes to make the code easier to debug, and writing/testing error recovery code to undo these "confused" states, seem to us to be the most effective way to solve these problems.

Below is the teams summary of tape handling problems from the report:

> The kind of confusion which occur can be grouped into 3 "confused" states
>
> 1. Tape present in drive, drive marked not allocated in OCS – this makes new load attempts fail, as the code did not handle this case.
> 2. No tape present in drive, but drive marked allocated in OCS – the drive does not get any further use due to its still being allocated.
> 3. OCS devd process hung, unkillable, while talking to some drive.

The report concludes with a recommendation that the team work for a few days with CDF to address these problems, and division management agreed.

A.3 Patches and Closeout

In late November and early December the team put in place their recommendations. They assembled the data handling test stand on b0sig02, and installed all legacy tape handling software. They established the correct versions of all the software and made sure they were tagged in CVS. The most significant problems were isolated and repeated on the test stand. Patches were installed in the code that checked whether a drive was empty and deallocated before attempting to allocate the drive or mount another tape, and diagnostic messages were added to the software to allow the system to be debugged. All patches were tested on the test stand. On December 4 new versions of the software that solved the majority of the operational problems were installed on fcdfsgi1, fcdfsun1 and fcdfsgi2. This resulted in a significant reduction in operational problems, from requiring human intervention every few hours to only every few days. The team of CD experts gave a closeout on December 5 and remain available for future consultation if necessary.

Patching work on the legacy system is ongoing, but we believe that the work that took place up to December 4 gave CDF enough breathing room to allow us to patch the system our self as needed as we operate it. We owe a great debt of thanks to Margaret, Marc and Luciano for their assistance.

B Tape Handling MOU between CD and CDF

The computing division has and will continue to provide and support the Central Tape Handling systems and services for the CDF experiment. Currently this consists of a legacy system utilizing CDF software to operate directly attached AIT-2 drives and media in an AML/2 robot, supported primarily by the CDF department of the computing divison (CD/CDF) with assistance as necessary from other departments. We are in the process of replacing this system with one utilizing Enstore to operate network attached STK 9940 drives and media in an STK robot, supported primarily by the Integrated Systems Development department of the computing division (CD/ISD) with assistance as necessary by other departments. The current status of planning is through a prototype and stage 1 phase. This MOU is intended to outline the responsibilities of the CDF experiment and the various departments of the computing division for the new system.

B.1 Hardware and Media

CD/CDF coordinates with the CDF experiment and CD/ISD for the design and planning for the proposed system and all system upgrades. After consulting with the CDF experiment and CD/ISD, CD/CDF will allocate the budget and procure hardware and tape media for the system which will be delivered to CD/ISD. CD/ISD will install and support the majority of the tape handling hardware. They are currently providing a stand-alone Enstore system to CDF built around the powderhon STK silo. The system is built with normal hardware processes known to CD/ISD (mover computers), funded by CD/CDF budget, and which conforms to CD/ISD normal practice. CD/CDF will coordinate the installation of the tape media, with assistance from the Computing Services Department of the computing division (CD/CSD). CD/ISD will coordinate the support of the STK Silo, the tape media, the STK 9940 drives, the mover nodes, the server nodes, and all other hardware between the tapes and the network local to the Enstore system.

B.2 Network

CD/CDF will coordinate with PPD/CDF, the CDF experiment, CD/ISD and the Data Communications Department of the computing division (CD/DCD) to insure that sufficient networking for data transfer is planned and purchased for the CDF central computing, the CDF trailers and B0. The networking is supported by CD/DCD in coordination with CD/ISD and CDF/CDF.

B.3 Software

Enstore, the tape handling software at Fermilab, is supported by CD/ISD in coordination with user requests, including those of CD/CDF and the CDF experiment. The interface to the user is currently the pnfs file system hierarchy, the encp copy command, and miscellaneous administration commands. CD/ISD will furnish these to CDF as it currently furnishes them to the rest of the Fermilab user community and will entertain requests for modification in the same fashion. Higher level software that determines how the tapes are used by CDF are the responsibility of the CDF experiment and CD/CDF. For example, the data handling input/output module in AC++, the schema of the data file catalog, the disk inventory manager, stager, and fileset tape daemon all remain the responsibility of the CDF experiment and CD/CDF. Enstore replaces the low level tape handling software products FTT, OCS, OCS_robot and mt_tools within the legacy data handling system. CD/ISD will provide consultation and documentation to facilitate the CDF experiment and CD/CDF in merging Enstore into the CDF data handling system.

B.4 Operations and Administration

CD/ISD will provide the operational and administrative effort to run the system. The operations of tape handling will be provided by CD/ISD. Tape writing is handled automatically by the robot, and supported 24/7. There is a pager rotation of tape handling experts available 24/7 in case of tape writing problems, that will respond within an hour to pages, adequate to support the needs of raw data logging and production farms data writing at CDF. The functioning of the hardware is monitored on a 24/7 basis CD/CSD and supported on a 24/7 basis by CD/ISD. Tape reading is handled automatically by the robot. Tape reading is supported on a 8/5 basis by CD/ISD, providing response within one business day for tape reading problems. CD/CDF and the CDF experiment will monitor tape usage with basic tools provided by CD/ISD and will procure the necessary tapes and provide the tape label numbers. CD/CSD will label and install tapes, monitor for problems in the system, and page CD/ISD experts as necessary when a problem is detected. CD/ISD provides CD/ISD with the web page of contacts for paging.

B.5 Upgrades

CDF/CDF will coordinate with the CDF experiment and CD/ISD to determine the requirements and understand the capacity of existing and emerging technologies. CD/ISD will collaborate with CD/CDF on the testing of new technologies that are needed to meet the requirements of the CDF experiment, and will collaborate with CD/CDF on the design and commissioning of the new systems. For example, there will soon be an STK 9940B drive that should triple the bandwidth and storage capacity of the STK 9940 drive,and CD/ISD intends to purchase a test unit.

C Detailed Stage 1 Project Plan

Tables 1 and 2 are a transcription of a Microsoft Project document written by D. Petravick, J. Bakken and B. Alcorn of ISD on January 4 that details the stage 1 plan.

ID	Status	Task	Duration	Start	Dept	Req.
1		**Networking**		12/28/01	DCD	
2		Establish Private LAN	1 Day	1/9/02	DCD	
3	Done	Establish Data LAN	1 Day	12/28/01	DCD	
4		Cable private LAN Switch to patch panel 6x100 mbps	1 Day	1/10/02	DCD	2
5	Done	Cable public switch LAN to patch panel 17x100 mbps	1 Day	12/28/01	DCD	3
6	Done	**Power and Data Center Infrastructure**		12/28/01	FSS	
7	Done	Siting and specifying for all equipment	1 Day	12/28/01	FSS	
8	Done	Power	1 Day	1/3/02	FSS	7
9	Done	208V for the Drive cabinet	1 Day	1/3/02	FSS	7
10	Done	Permanent power (120 V) for the servers	Movers	1 Day	FSS	7
11		**CDFEN Silo**		12/28/01		
12	Done	STK install 8 additional 9940 drives	1 Day	12/28/01	STK	
13		Deinstall STKEN ACSLS	2 hrs	2/15/02	STK	
14		STK intall CDF ACSLS up to the point of commissioning	2 hrs	1/4/02	STK	
15		Tape handling test	2 days	2/15/02	STK	13
16		Insert production tapes	2 days	2/19/02	ISD,CSD	15
17		Provide backup device for fntt2	2 hrs	1/4/02	STK	
18		**CDFEN computers - Physical Installation**		12/27/01		
19		**Racking**		12/27/01		
20	Done	Verify/extend thepatch panel will support the 6 cables for the private lan	1 hr	1/3/02	ISD/IA	
21		Equip and furnish rck for ACSLS computer	2 hrs	12/27/01	FSS	
22		Strain relief for 10 SCSI cables	1 day	2/11/02	ISD/IA	
23		Brush-by protection for the rear of th 19" relay rack	2 days	2/7/02	ISD/IA	
24		**SRV0, SRV1 and SRV spare – Enstore Databases**	10.25 days	1/2/02	ISD/IA	
25		Build one prototype	5 days	1/2/02	ISD/IA	
26		Replicate twice	4 days	1/9/02	ISD/IA	25
27		Private LAN 100 mbs ethernet card (3 ea Intel eepro 100)	1 hr	1/16/02	ISD/IA	25
28		Standardize BIOS, set BIOS, ditto BMC SDR	1 day	1/15/02	ISD/IA	25
29		**SRV2 (mover case)**	1.25 days	1/18/02	ISD/IA	
30		Augment disks - add 1 ea IDE disk	1 day	1/18/02	ISD/IA	
31		Private LAN 100 mbs ethernet card (1 as Intel eepro 100)	1 hr	1/22/02	ISD/IA	
32		Standardize BIOS, set BIOS, ditto BMC SDR	1 hr	1/23/02	ISD/IA	25
33		**SRV3 (mover case) console server**	2.75 days	1/22/02	ISD/IA	
34		PCI Cyclades card, cable to pods, configure	1 day	1/22/02	ISD/IA	
35		Monitor, keyboard, mouse	2 hrs	1/24/02	ISD/IA	
36		Augment disks, add two disks	3 hrs	1/23/02	ISD/IA	
37		Private LAN 100 mbs ethernet cards	1 hr	1/24/02	ISD/IA	
38		Standardize BIOS, set BIOS, ditto BMC SDR	2 hrs	1/24/02	ISD/IA	
39		**SRV4 (mover case)**	0.38 days	1/28/02	ISD/IA	
40		Private LAN 100 mbs ethernet card (1 as Intel eepro 100)	1 hr	1/28/02	ISD/IA	
41		Standardize BIOS, set BIOS, ditto BMC SDR	2 hrs	1/28/02	ISD/IA	
42		**8 additional movers**	2.5 days	1/30/02	ISD/IA	
43		Identify quantity needed and order HVD cards + spares	4 hrs	1/30/02	ISD/IA	
44		AHA2944 UW cards installed	1 day	1/30/02	ISD/IA	43
45		Standardize BIOS, set BIOS, ditto BMC SDR	1 day	1/31/02	ISD/IA	
46		**1 spare mover**	0.25 days	2/13/02	ISD/IA	
47		AHA2944 UW cards installed	1 hr	2/13/02	ISD/IA	43
48		Standardize BIOS, set BIOS, ditto BMC SDR	1 day	2/14/02	ISD/IA	
49		**APC power controllers (3)**	5.63 days	2/6/02	ISD/IA	
50		Procure "horizontal" 20 amp APC units	2 hrs	2/6/02	ISD/IA	50
51		Install into relay rack, and route power cords to machines	1 day	2/6/02	ISD/IA	50
52		Serial cables to SRV3	2 hrs	2/14/02	ISD/IA	
53		**CDFEN computers - Software Installation**		1/16/02		
54		**Operating system + support + ENstore Distribution**		1/16/02		
55		Get Kerberos principals for 15 new systems	1 day	2/12/02	ISD/IA	
56		Account for changes in images, merge into standard system images	1 day	2/18/02	ISD	57,60,etc.
57		**Movers + Mover Spare**	3 days	2/1/02	ISD/IA	
58		Install, test, debug system image, clean up cron state	2 days	2/1/02	ISD/IA	42
59		Install, test, debug, kerberos	1 day	2/5/02	ISD/IA	
60		**SRV2**	0.63 days	1/23/02	ISD/IA	29
61		Install, test, debug system image, clean up cron state	2 hrs	1/23/02	ISD/IA	
62		Install, test, debug, kerberos	2 hrs	1/23/02	ISD/IA	
63		Make private LAN work	1 hr	1/23/02	ISD/IA	
64		**SRV1**	0.63 days	1/16/02	ISD/IA	
65		Install, test, debug system image, clean up cron state	2 hrs	1/16/02	ISD/IA	24
66		Install, test, debug, kerberos	2 hrs	1/16/02	ISD/IA	
67		Make private LAN work	1 hr	1/16/02	ISD/IA	
68		**SRV0**	0.63 days	1/16/02	ISD/IA	
69		Install, test, debug system image, clean up cron state	2 hrs	1/16/02	ISD/IA	24
70		Install, test, debug, kerberos	2 hrs	1/17/02	ISD/IA	
71		Make private LAN work	1 hr	1/17/02	ISD/IA	
72		**SRV3**	0.63 days	1/24/02	ISD/IA	29
73		Install, test, debug system image, clean up cron state	2 hrs	1/24/02	ISD/IA	
74		Install, test, debug, kerberos	2 hrs	1/25/02	ISD/IA	
75		Make private LAN work	1 hr	1/25/02	ISD/IA	
76		**SRV4**	0.63 days	1/28/02	ISD/IA	29
77		Install, test, debug system image, clean up cron state	2 hrs	1/28/02	ISD/IA	
78		Install, test, debug, kerberos	2 hrs	1/28/02	ISD/IA	
79		Make private LAN work	1 hr	1/29/02	ISD/IA	
80		**Spare SRV[01]**	0.38 days	1/18/02	ISD/IA	24
81		Install, test, debug system image, clean up cron state	2 hrs	1/18/02	ISD/IA	
82		Install, test, debug, kerberos	1 hr	1/18/02	ISD/IA	

Table 1: Stage 1 Project Plan. The task ID, status, name, duration, start date,responsible department, and required task to be completed first are listed.

ID	Status	Task	Duration	Start	Dept	Req.
83		**Enstore Configuration**		12/27/01	ISD	
84	Done	DNS ailias for web server	1 hr	12/27/01	ISD	
85		**Database SRV0**	1.75 days	1/17/02	ISD	68
86		Make directories with proper protections for databases	2 hrs	1/22/02	ISD	
87		get database backups working	3 hrs	1/17/02	ISD/IA	
88		get system image backups working	3 hrs	1/17/02	ISD/IA	
89		**NAMESPACE SRV1**	4 days	1/16/02	ISD/CDF	24
95		**Web Server (server 2)**	3.25 days	1/23/02	ISD	60
96		**Web**	2.5 days	1/23/02	ISD	
97		configure apache	5 hrs	1/24/02	ISD	
98		introduce CDFEN into the pages	8 hrs	1/24/02	ISD	
99		clobber URLs to point to cdfen "fix_url"	3 hrs	1/25/02	ISD	
100		**config server**	0.5 days	1/25/02	ISD	
101		make a configuration file for CDFEN	4 hrs	1/25/02	ISD	
102		**log server**	0.13 days	1/28/02	ISD	
103		make directories with proper protections	1 hr	1/28/02	ISD	
104		**Tape inventory file area**	0.13 days	1/28/02	ISD	
105		make directories with proper protections	1 hr	1/28/02	ISD	
106		**Rate leeper**	0.13 days	1/28/02	ISD	
107		make directories with proper protections	1 hr	1/28/02	ISD	
108		**Inquisitor**	0.13 days	1/28/02	ISD	
109		make directories with proper protections	1 hr	1/28/02	ISD	
110		**9 movers**	0.25 days	2/6/02	ISD	57
111		Create directories from areas not saved via system backups	1 hr	2/6/02	ISD	
112		get system image backups working	1 hr	2/6/02	ISD	
113		**SRV 3 – console server**	1.13 days	1/25/02	ISD	72
114		create, protect directory structure	2 hrs	1/28/02	ISD	
115		Make PNFS meta-data backup to tape work	2 hrs	1/28/02	ISD	
116		Make vile, volume cleark backup to tape work	3 hrs	1/28/02	ISD	
117		Make sure conserve software is configured and works for every port	4 hrs	1/25/02	ISD/IA	
118		Make conserve work for APC power controllers	2 hrs	1/25/02	ISD/IA	
119		Make data base integrity check for file, volume clerk work	2 hrs	1/28/02	ISD	
120		Make the volume and file inventory work	3 hrs	1/28/02	ISD	
121		get system image backups working	1 hr	1/28/02	ISD/IA	
122		**SRV 4**	0.13 days	1/29/02	ISD	76
123		get system image backups working	1 hr	1/29/02	ISD/IA	
124		**Spare SRV[01]**	0.25 days	1/18/02	ISD/IA	80
125		look at it, make sure we are happy the hardware works	2 hrs	1/18/02	ISD/IA	
126		**FNTT2**	0.88 days	1/29/02	ISD/IA	
127		make it work w/ private LAN	2 hrs	1/29/02	ISD/IA	4,76
128		Rhost files	2 hrs	1/29/02	ISD/IA	
129		**STKENSRV5**	0.5 days	1/29/02	ISD/IA	
130		Provide extra disks for CDFEN backup staging	4 hrs	1/29/02	ISD/IA	
131		**Software development**		1/4/02	ISD	
132		Cut version of ENCP having UPS table support for CDFEN	1 day	1/22/02	ISD	
133		Support CDFEN symmetrically to te others in the core	4 days	1/11/02	ISD	
134		Develop PNFS data demangler	5 days	1/4/02	ISD	
135		Make Enstore work with ACSLS V6	2 days	2/11/02	ISD	126
136	Done	**Scripts and Administration**	5 days	12/24/01	ISD	
137	Done	Modify administration scripts to add CDFEN case	5 days	12/24/01	ISD	
138		Burn in of CDFEN servers		2/6/02	ISD	83
139		Configure in media changer pointing to FNTT2	1 hr	2/13/02	ISD	
140		Allocate and burn in drives, tapes	2 hrs	2/7/02	ISD/IA	
141		Exercise and resolve problems	5 days	2/6/02	ISD	
142		Scratch burn in tapes, pruge system of meta-data, etc	4 hrs	2/7/02	ISD/IA	
143		Add operational support, operators/helpdesk monitoring, etc.	1 day	2/6/02	ISD/IA	
144		**Transition to stand alone operation**	2.25 days	2/14/02	ISD	138
145		**Connect FNTT2 to control CDF Silo**	1.38 days	2/14/02	ISD	
146		Physically connect	1 hr	2/15/02	ISD/IA	
147		Inventory	7 hrs	2/14/02	ISD/IA	
148		**Transfer meta-data of CDFEN silo tapes to CDFEN Enstore system**	2.25 days	2/14/02	ISD	
149		Stop writes	1 hr	2/18/02	ISD/IA	
150		**Extract meta-data**	2.25 days	2/14/02	ISD	
151		PNFS	2 hrs	2/14/02	ISD	
152		File and volume clerk	2 hrs	2/18/02	ISD	
153		**Insert meta-data**	0.25 days	2/18/02	ISD	
154		PNFS	2 hrs	2/18/02	ISD	
155		File and volume clerk	2 hrs	2/18/02	ISD	
156		**Purge meta-data**	0.25 days	2/18/02	ISD	
157		PNFS	2 hrs	2/18/02	ISD	
158		File and volume clerk	2 hrs	2/18/02	ISD	
159		**Transfer data on STKEN silo tapes to CDFEN system**	1.75 days	2/14/02	ISD	
160		Stop CDF writes into STKEN	1 hr	2/14/02	ISD/IA	
161		**Export the meta-data from STKEN**	0.63 days	2/14/02	ISD	160
162		Export the PNFS data	2 hrs	2/14/02	ISD	
163		Export the file/vol clerk data	2 hrs	2/14/02	ISD	
164		Actually move the tapes	4 hrs	2/15/02	ISD/IA	
165		**Import the meta-data to CDFEN**	0.5 days	2/14/02	ISD/IA	161
166		Import the file/vol clerk	2 hrs	2/14/02	ISD	
167		Import PNFS data	2 hrs	2/15/02	ISD	
168		**Purge state from STKEN**	0.5 days	2/15/02	ISD	165
169		Purge the file/vol clerk	2 hrs	2/15/02	ISD	
170		Purge PNFS	2 hrs	2/15/02	ISD	
171		**CDF Side Work**		1/9/02	CDF	
172		Order, install and configure two additional gigabit ethernets for fcdfsgi2		1/9/02	CDF	
173		Order, install and configure one gigabit ethernet for fcdfsgi1		1/9/02	CDF	
174		Check and burn from other nodes	4 days	2/18/02	CDF	144
175		Consider with respect to critical systems plans, etc.	1 day	1/9/02	CDF	

Table 2: Stage 1 Project Plan (continued). The task ID, status, name, duration, start date,responsible department, and required task to be completed first are listed.

References

[1] The presentations shown to the CDF internal DH review of November 19, and the reviews recommendations, may be found at http://www-cdf.fnal.gov/upgrades/cdfdh/doc/dhreview/dhreview_sept2001/index.html

[2] R. Kennedy, Review of a Model for Adapting CDF Tape Handling Software to use Enstore, CDF Note 5772, November 19, 2001.

[3] Enstore presentations and related projects may be found at http://hppc.fnal.gov/enstore.

[4] B. Ashmanskas et al., Final Report CDF CAF Review Fall 2001, CDF Note 5802, December 7, 2001.

[5] Updated from J. Bakken, private communication, December 2001.

[6] D. Petravick, private communication, December 19, 2001.