First Operation of the D0 Run II Reconstruction Farm

M. Diesburg[2], B. Alcorn[2], J. Bakken[2], T. Dawson[2], D. Fagan[2], J. Fromm[2], K. Genser[2], L. Giacchetti[2], D. Holmgren[2], T. Jones[2], T. Levshina[2], L. Lueking[2], L. Loebel-Carpenter[2], I. Mandrichenko[2], C. Moore[2], S. Naymola[2], A. Moibenko[2], D. Petravick[2], M. Przybycien[1], H. Schellman[1], K. Shepelak[2], I. Terekhov[2], S. Timm[2], J. Trumbo[2], S. Veseli[2], M. Vranicar[2], R. Wellner[2], S. White[2], V. White[2]

(1) *Dept. of Physics and Astronomy, Northwestern University, Evanston IL, USA*
(2) *Computing Division, Fermi National Accelerator Center, Batavia IL, USA*

Abstract

We report on the first operations of the reconstruction farms for the D0 experiment. Data were read from a tape robot to 50 PC's running Linux, processed, spooled to a central disk buffer for merging and then written back to the tape robot. The farms are being used successfully to reconstruct the data as it comes in. Transfer rates well over the 12.5 MB/sec needed for full data rates have been achieved.

Keywords:Linux,Farms,PC,commodity

1 Introduction

The D0 experiment began taking data for its second run (Run II) in March 2001. The detector is still undergoing commissioning; the data rates into reconstruction farm expected in early 2002 are $\sim$ 50 Hz peak with an average of 20Hz. Event sizes are expected to be 0.25 MB and CPU time/event is expected to be in the range 5-10 sec on a 750 MHz Pentium III processor. These rates and event sizes imply a peak input bandwidth of 12.5MB/sec to the reconstruction farms and a need for 250-500 processors in the final system to handle the peak rates. The output bandwidth is lower due to a smaller record size for reconstructed data but will still be significant. During the current commissioning phase the data rate from the detector is of order 8 Hz, the event size averages 350KB and the CPU time per event is 20 secs. The output records are currently larger than the input records, due to retention of intermediate cluster data for detector debugging.

2 Hardware

The D0 farm consists of 90 PC's, the first 40 have 2 500 MHz Pentium III processors, a 6 GB system disk, 2 16GB IDE disks for data, 512MB of memory, floppy and CD drives and a 10/100 MB ethernet card. The SPECINT95 for each of these processors is $\sim$21. An additional 50 machines with 750 MHz PIII processors, and 1GB of memory were added in early 2001, these new machines have a SPECINT95 rating of 36/processor. The 40 slower nodes are now used for testing while the fast 50 are used for production. The operating system is a 'farm' version of RedHat Linux 6.1 with non-essential features removed. The PC's are connected via a Cisco 6509 switch to the SAM and Enstore systems and to the D0 farm control system, a 4 processor SGI Origin 2000 with 1GB of memory and 500GB of disk (2-way striped). The Origin serves the home areas for all 90 PC's, serves as a disk buffer for output files and is the master for the FBS batch system. It has two Gigabit Ethernet cards configured to talk to the farm worker nodes and the Enstore tape system respectively.

The hardware has been found to be quite reliable, once a rigorous 1 month acceptance test has been completed. Uptimes on individual nodes are limited by maintenance days, which occur every 1-2 months. The rate for hardware problems which remove a node from the system is around 1/month for 90 machines.

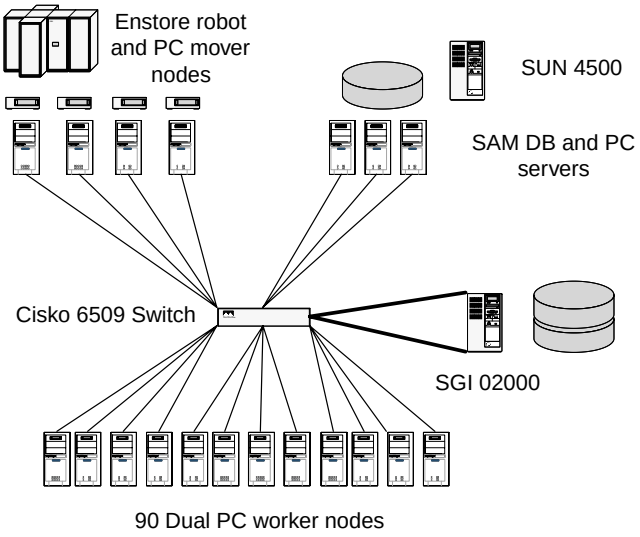

Figure 1: The 180 CPU D0 reconstruction farm. Thin lines indicate 100MB ethernet connections, thick lines are Gigabit ethernet.

3 Software and Control Systems

The Run II farm design separates data storage and book-keeping from farm operations by use of the Fermilab Sam[1] and Enstore[2] data handling systems. The Fermilab Farm Batch System[3] is used for job control. They are described in contributions to the CHEP 2000 conference. [4, 5, 6] As a result of the use of these external products the farm specific coding has been reduced to a small number of python scripts and and was done by three physicists working part time.

3.1 Data Delivery

The SAM/Enstore data delivery system consists of an ADIC tape robot with 10 Mammoth II tape drives connected to PC's running Linux. The PC's read the tape and pass the data over the network via the D0 switch to the farm worker nodes. The Enstore system handles tape mounts and data delivery via the pnfs file system while the SAM system provides a full file catalog and processing database. Once a connection to the SAM system has been established, any process on the farm can request 'the next file' in a list and it will be delivered. The SAM system does not count the file as successfully processed until the farm process 'releases' it with a status code. The SAM system can thus provide a summary for any farm job of file deliveries and the processing status of any file.

3.2 Batch System

The Farm Batch System allows a user to submit jobs with a simple control structure and dependencies. It also provides very useful system status displays. FBS is used by several D0 institutions and many experiments at Fermilab. D0 farm jobs use only the minimal FBS functionality and consist of:

Start Section In the start section, which runs on a single node, the SAM data handling system is informed that a job requiring a given list of files is starting. Directories for output files are set up on the control node. The start section stays alive for the full length

of the job and also starts a 'store' process which stores output file after they have been processed on the worker nodes and copied to the output directories on the control node. This allows merging and derandomization of the files arriving from the worker nodes.

Worker Section Once the start section has begun, worker processes are started on N worker nodes. One process is run per processor due to the large size of the D0 executable - the batch system is configured to impose this constraint. When a process starts on a worker node, the environment needed to run the reconstruction, including executables and control files is copied across the network and built from scratch. The worker process then requests data files from the SAM system, processes them and copies the output and log files to the output directory on the control node. The SAM system delivers the 'next' file in response to a request, thus if a worker machine crashes, processing of 1 (or 2 in the case of dual processors) files is lost but the system as a whole can proceed to completion. The SAM database will not receive the 'release' signal and knows that the file had not been successfully processed. When a worker section can receive no more files, because all have been delivered, it terminates and another job can start on that processor.

The worker section takes between 2 and 40 minutes to start up, depending on the number of other nodes accessing the same data tape. Prestaging of additional input files then continues in the background while the reconstruction algorithms are run. The reconstruction algorithms run very efficiently (99% CPU utilization) as long as no memory swapping occurs. Use of swap space degrades the performance to $\sim$50% and needs to be avoided.

End Section When all worker sections have terminated, the End Section is run. It informs the SAM system that no more data are needed. It then signals the 'start' section that processing is done. The start section continues until all output files have been stored back into the robot.

3.3 File Handling

All of the file transfers are done via either pnfs or the Fnal product fcp[7], a throttled rcp which allows the system to specify a maximum number of allowed transfers (currently 6) and maintains a queue of pending transfers. Using fcp rather than direct nfs mounted disks makes the system much more robust at startup, when up to 180 processes may be requesting files at once.

4 Description of current operations

The current system is closely matched to the raw data rates from the detector and is generally running at high CPU utilization. We have observed very few failures due to the farm system itself. The major failure modes are memory leaks in the executables and problems reading tapes. During normal operations we have verified that the data tranfer rates into the farms are already well above the 12.5 MB/sec expected at 50 Hz. Figure 2 illustrates data transfers into the system after submission of 3 jobs, each using 30 nodes.

5 The future

Our first operational experience has reassured us that the basic hardware and software configuration is easy to install and maintain and that it can handle data rates and CPU loads expected at full rate. We will upgrade the farm with another 130 nodes (260 CPU's) over the next six months, as the D0 data rate increases.

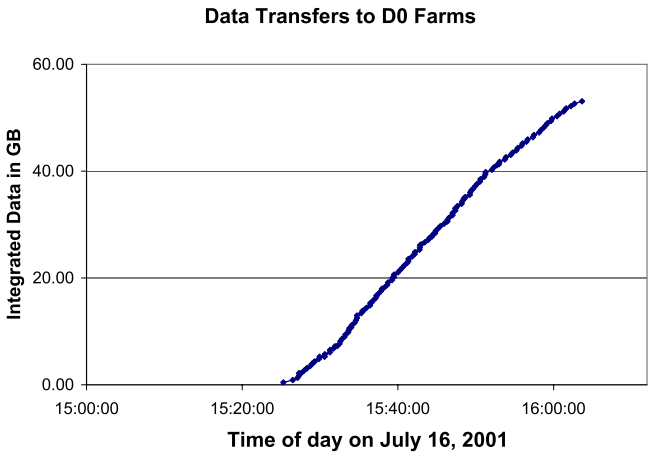

Figure 2: Integrated transfers to the D0 farms for 3 30 CPU jobs. 141 files of size 376 MB were transferred over 44 minutes from tape to the 90 worker nodes. The average data rate once all 3 jobs were running was 24.5 MB/sec. Individual transfers were 10.2 MB/sec, consistent with the expected speed for a Mammoth II tape drive.

Acknowledgements

We would like to thank the Fermilab Computing Division for their invaluable assistance and advice in contructing this system. We would also like to thank the CDF farms and data handling groups for many useful discussions.

References

[1] `http://d0db.fnal.gov/sam`
[2] `http://d0ensrv2.fnal.gov/enstore/`
[3] `http://www-isd.fnal.gov/fbsng/`
[4] L. Lueking *et al.*, The Data Access Layer for D0 Run II: Design and Features of SAM, CHEP 2000 proceedings, p462-466.
[5] D. Petravick *et al.*, Fermilab Computing Division Systems for Scientific Data Storage and Movement, CHEP 2000 proceedings, p630-633.
[6] I. Mandrichenko *et al.*, Farms Batch System and Fermi Inter-Process Communication toolkit, CHEP 2000 proceedings, FERMILAB-CONF-00-083.
[7] `http://www-isd.fnal.gov/fcp/`