

Redundant Arrays of IDE Drives

David Sanders, Lucien M. Cremaldi, Vance Eschenburg, Christopher N.
Lawrence, Chris Riley, and Donald J. Summers

University of Mississippi

Donald L. Petravick

Fermilab




Introduction

- \$4000 per Terabyte Storage is Available
- Scalable for use at both Small and Large Institutions — From 1 TB to 250 TB, the same as a \$ Million tape silo.
- Fast Access to Data
- Redundant — RAID5
- Commodity Hardware



Data Storage Cake




Definitions

- RAID — Redundant Array of Inexpensive Disks
- RAID level 0 — Concatination
- RAID level 1 — Mirroring
- RAID level 4 — Parity
- RAID level 5 — Striped-Parity
- EIDE — Enhanced Integrated Drive Electronics



Why Use Commodity Hardware?

© 1998 Randy Glasbergen. E-mail: randyg@norwich.net www.norwich.net/~randyg/toon.html



**“Frankly sir, we’re tired of being
on the cutting edge of technology.”**



Hardware

- System Disk — 100 GB Maxtor
- Eight 100 GB Maxtor Disks
- 2 Promise Ultra100 PCI cards
- 24” EIDE Cables
- CPU — 1.4 GHz AMD Athlon
- Motherboard — Asus A7A266
- 512 MB DDR memory
- Second Power Supply (15A at 12V)



RAID5 Box for BABAR



Disks

Disk	RPM	\$/GB	GB/platter	Amps@12V
80 GB Maxtor	5400	1.88	20	2.000
100 GB Maxtor	5400	2.20	33	0.640
160 GB Maxtor	5400	2.05	40	1.800
75 GB IBM	7200	3.00	15	2.000



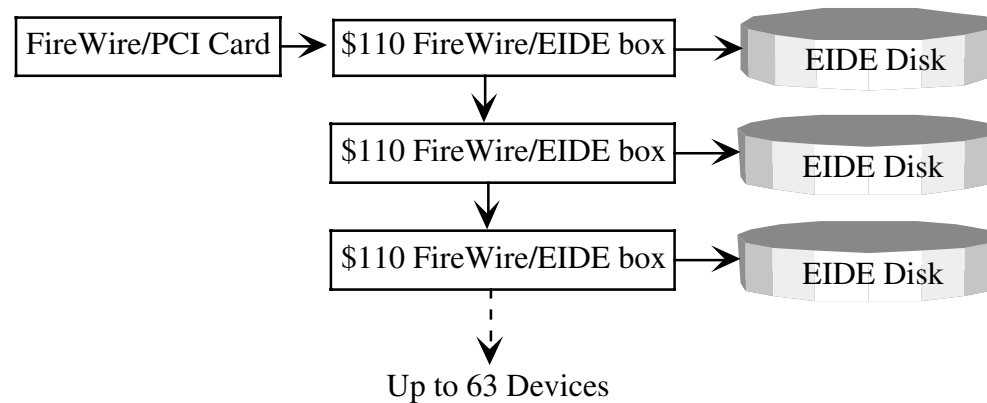
Software

- Linux 2.4.13 Kernel (2.4.5 Tested)
- raidtools available with most distributions
- Journaling File systems (ReiserFS and ext3)
- NFS to mount on other computers
(Linux, Sun Solaris, DEC Ultrix, Mac OSX)
- HDPRM speed test (~28 MB/s)
- Simple write test (22 MB/s)



FireWire™

- 63 Disks per chain
- 10 Terabytes per chain using 160 GB disks
- Hot swappable (Using Firewire™ Hubs)
- We tested RAID5 in Linux
(Oxford 911 chipset)

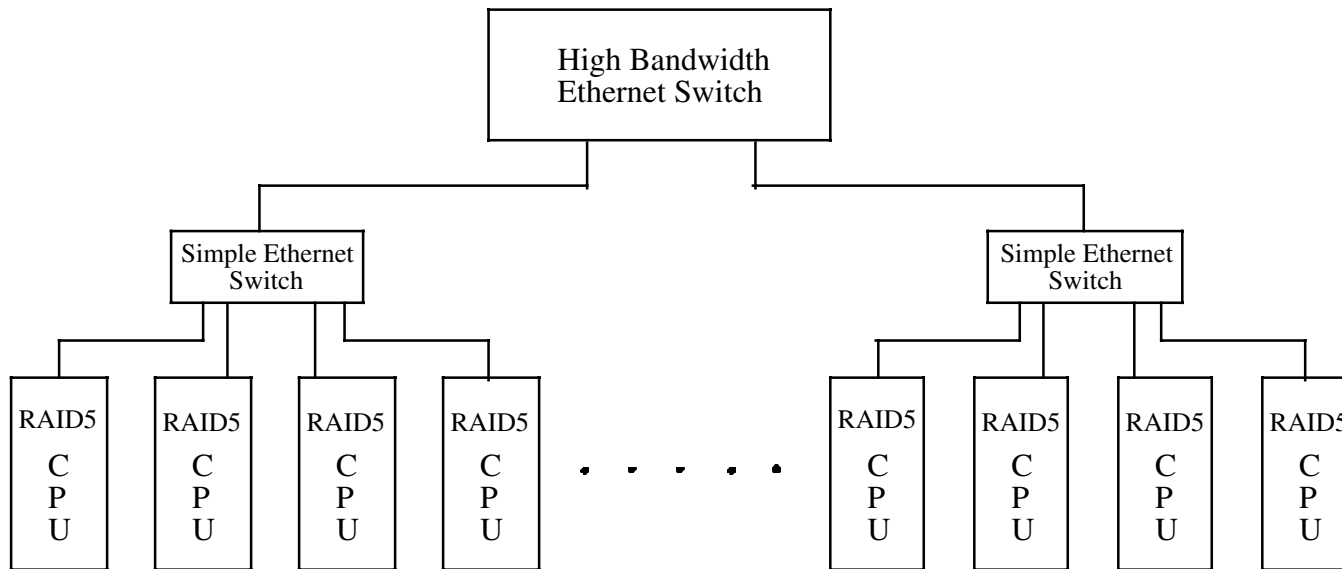


High Energy Physics Data Analysis Strategy

- Use Parallel Processing
- Split data and store on many RAID5 PCs
- Analysis for a subset of data takes place locally on the PC where the data resides
- Network is only used to combine results
- Or use NFS to mount RAID5 array on many PCs (Less efficient due to network overhead)



High Energy Physics Cluster



Moving Data to Institutions

- Internet transfers
- Hot pluggable EIDE disks in Firewire cases
- DVD-R disks — 4.7 GB DVD-R disks are now \$6, the DVD writer is only \$500, and the DVD-ROM reader is \$60.



Summary

- \$4000 per Terabyte RAID5 arrays of EIDE Drives tested, without tape backup.
- They are Scalable —Same cost/TB as a tape silo, but scalable down to 1 TB.
- Uses Commodity Hardware.
- FireWire has been tested in Linux

Supported by the U.S. Department of Energy under DE-FG05-91ER40622 and DE-AC02-76CH03000.

