

# Tutorial on Storage Technology: RAID and Beyond

Garth A. Gibson

School of Computer Science, Carnegie Mellon University, Pittsburgh, PA 15213-3891

Garth.Gibson@cs.cmu.edu, <http://www.cs.cmu.edu:8001/Web/Groups/PDL/>

## Abstract

*In stark contrast to the 25% per year increase in areal density delivered by the magnetic disk industry during the 1970s and 1980s, yearly increases today are 60%, on par with DRAM density increases. Moreover, the storage industry is also delivering substantially higher data rates, smart disk-embedded readahead and writebehind, and a new generation of high-speed serial interconnects. This industry has also embraced Redundant Arrays of Inexpensive (or Independent) Disks (RAID) technology - 1997's RAID market is expected to be 13 billion dollars. With this rapidly evolving market and technology base, parallel storage systems must evolve beyond RAID levels 1 through 5. This talk is intended for researchers and practitioners interested in current trends in storage systems. It will highlight storage technology trends, RAID technology trends, and trends toward RAID-style support for network-based parallel storage systems.*

## Outline

Technology trends in secondary storage

- areal density, disk diameter, data rates
- embedded controller caching and reordering
- peripheral interface evolution
- drive support for RAID

RAID technology trends

- improving update performance
- transparent on-line recovery
- caching for RAID
- multiple failure tolerance for very high reliability
- rapid prototyping for RAID

Network RAID

- reliable striping over multiple servers
- support for multimedia service
- application managed availability
- parallel file systems for parallel programs

Permission to copy without fee all or part of this material is granted provided that the copies are not made or distributed for direct commercial advantage, the ACM copyright notice and the title of the publication and its date appear, and notice is given that copying is by permission of the Association of Computing Machinery. To copy otherwise, or to republish, requires a fee and/or specific permission.  
SIGMOD '95, San Jose, CA USA  
© 1995 ACM 0-89791-731-6/95/0005..\$3.50

## Biography

Garth Gibson is an assistant professor in the School of Computer Science and the Department of Electrical and Computer Engineering at Carnegie Mellon University. He leads the Parallel Data Laboratory's research into high-performance, highly-available, networked, file and storage system architectures. This lab participates in the Nectar Gigabit Networking testbed and the HPCC Multicomputing Systems Software project, and constitutes the Storage and Computer Systems Integration area of the Data Storage Systems Center. Companies collaborating with this lab include AT&T-GIS (NCR), Data General, IBM, DEC, HP, Seagate, and Storage Technology. Gibson received a Ph.D. degree from the University of California at Berkeley in 1991 where he co-founded Berkeley's Redundant Arrays of Inexpensive Disks (RAID) project with Prof. D. A. Patterson and Prof. R. H. Katz. His dissertation, "Redundant Disk Arrays: Reliable, Parallel Secondary Storage," tied for second place in the 1991 ACM Doctoral Dissertation Award and was published by the MIT Press in 1992.

## Bibliography

- S. Berson, et. al., "Staggered Striping in Multimedia Information Systems," *Proc. ACM SIGMOD Conf.*, 1994.
- P. Chen, E. Lee, G. Gibson, R. Katz, D. Patterson, "RAID: High-Performance, Reliable Secondary Storage," *ACM Computing Surveys*, 26(2):145-185, 1994.
- P. Corbett, D. Feitelson, "Design and Implementation of the VESTA Parallel File System," *Proc. Scalable High-Performance Computing Conf.*, 1994.
- G. Gibson, *Redundant Disk Arrays: Reliable, Parallel Secondary Storage*, Ph.D. Dissertation, MIT Press, 1992
- G. Gibson, et. al., "The Scotch Parallel Storage Systems," IEEE COMPCON, 1995.
- J. Hartman, J. Ousterhout, "The Zebra Striped Network File System," *Proc. 14th ACM SOSP Conf.*, 1994.
- H. Kung, et. al., "Network-based Multicomputers: an Emerging Parallel Architecture," *Supercomputing'91*, 1991.
- M. Kryder, "Data Storage in 2000 - Trends in Data Storage Technologies," *IEEE Trans. Magnetics*, 25(6), 1989.
- D. Patterson, G. Gibson, R. Katz, "A Case for Redundant Arrays of Inexpensive Disks (RAID)," *Proc. ACM SIGMOD Conf.*, 1988.