Appears in *Proc. of the 30th Hawaii International Conf. on Systems Sciences*, Jan. 1997

Task Force on Network Storage Architecture: Task Force Agenda

Garth A. Gibson

School of Computer Science, Carnegie Mellon University, Pittsburgh, PA 15213-3891
garth.gibson@cs.cmu.edu http://www.cs.cmu.edu/Web/Groups/NASD/HICSS.html

Task Force Abstract

Storage systems represent a vital market that is growing faster than the personal computer market. Its primary constituents are magnetic and optical disk drives, magnetic tapes, and large-capacity (robotic) assemblies of drives and cartridges. Storage hardware sales in 1995 topped $40 billion, including more than 60,000 terabytes of hard disk storage. In recent years, the amount of storage sold has been almost doubling each year; in the near future it is expected to sustain an annual growth of about 60 percent. This enormous growth rate has been accompanied by a 35-50 percent per year decrease in the cost per byte of storage. Consequentially, insuring the continued vitality of storage's architecture in future computing systems is essential.

While many of these storage products are being directly attached to personal and home computers, enabling local copies of substantial code and information bases to improve privacy, speed and availability, 65% of the disk array products are already being used in LAN file servers and this fraction is expected to rise to 75% over the next few years. With today's distributed file system technology, all storage bytes are copied through file server machines, introducing a potentially serious bottleneck for high-bandwidth, low-latency, and scalability. Moreover, the sustained bandwidth of storage devices is rapidly outstripping current interconnection technologies and rendering inexpensive store-and-forward server solutions impractical. For example, the rapid technological improvements in linear bit density and magnetic disk rotation rate is driving data rate up at 40% per year, insuring 25-40 MB/s sustained disk bandwidth by the end of the decade. Consequentially, the quality of storage service is increasingly interdependent with local area networking.

This work was supported in part by DARPA contract N00174-96-0002, by NSF grant ECD-8907068, and by Data General, Symbios Logic, Hewlett-Packard, and Compaq. The US government has certain rights in this material. The views contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of any funding agency.

The goal of this task force is to chart out this interaction between local area network and storage architecture. Our primary tasks will be to:

- document the trends and forces that are driving this convergence;
- define networked storage and distinguish its flavors;
- identify critical technologies and major technological obstacles;
- plot avenues for community-wide adoption and deployment; and
- identify viable alternatives and chart litmus tests between alternatives.

Seeding this task force discussion are the positions of four communities concentrating on networked storage architectures. Dave Anderson reports that Seagate's customers are driving network-attached devices because of the improved scalability, uniform fault-tolerance model, conformance to an increasingly open systems architecture, and especially demanding applications such as video. Rod Van Meter, Steve Hotz and Greg Finn, of USC/ISI, up the ante by calling for network-attached storage to speak internet protocols, especially IP, because of its advantages in wide-area connectivity, physical media independence, mature and rich development community, and the recent techniques for eliminating bottlenecks in per-byte costs. John Wilkes, Hewlett-Packard Labs, sees an evolution of processor-centric computing to storage-centric computing in no small part because of the cost of storage management (capacity planning, fault-tolerance, heterogeneity, load balancing, backup). In this storage-centric architecture, Wilkes offers a quality of service abstraction called attribute-managed storage, to allow storage subsystems to better manage themselves. Finally, Carnegie Mellon's network-attached secure disks project argues that raising the level of abstraction of the storage interface enabled SCSI to hide the rapid changes in storage technology over the last decade and it is time to raise it again. At this higher level of abstraction storage would "understand" data groupings that correspond to client files, enabling the much deeper prefetching, dynamic allocation, compression/search, efficient security enforcement and aggressive self-management.