Task Force on Network Storage Architecture: Management of future storage

John Wilkes
Hewlett-Packard Laboratories, 1501 Page Mill Road, Palo Alto, CA 94304-1126
wilkes@hpl.hp.com http://www.hpl.hp.com/personal/John_Wilkes/

Position Statement

Information storage systems are the bedrock on which a modern company rests: data has to be available to whoever needs it, whenever they need it, from wherever they may be. As a result, emphasis is beginning to shift from a processor-centric view of the world towards an storage-centric one: it is the information that holds the real value - the processing merely exposes it. Realizing this model requires a storage system that can provide continuous, guaranteed accessibility to potentially vast quantities of shared information, with protection against both malicious attack and accidental failures, and scalability across a wide range of capacity and performance needs. Network-based storage, combined with smart storage management, can provide it.

Large storage systems are exceedingly complex: a listing of the storage units and their redundant I/O paths for a one terabyte sized system can easily run to tens of pages. Planning for their installation can take 2-3 months. Even so, configuration choices for such systems are usually done by simple rules of thumb. Even once a suitable storage system can be constructed, keeping it in good shape is a constant struggle against all sorts of pressures:

- the sheer number and complexity of devices,
- a dynamically-changing workload,
- a dynamically changing device pool (faults + growth + reconfiguration),
- fault recovery (eg backup/restore),
- enforcement of quality of service guarantees, and
- enforcing security.

The result is sub-optimal system configurations that suffer from lackluster performance, high cost, insufficient fault-tolerance, and unacceptably high administrative burdens. Many large installations will do anything to avoid expanding their storage subsystems. According to some sources, it already costs 7-12 times as much to support one gigabyte of storage for a year as it does to purchase the disk it lives on. Although storage hardware is a big business (about $60 billion in 1995), storage management is bigger (about $120 billion in 1995)[1].

It is the thesis of this position statement that these storage management problems will form an barrier to the widespread development and deployment of large-scale, robust storage systems unless they can be eliminated.

A way to achieve this is attribute-managed storage, which uses quality-of-service (QoS) attributes such as capacity, performance, availability, and cost to describe the goals and needs of parts of the storage space to a central storage utility. The storage utility is then responsible for mapping these demands onto the collection of storage devices at its disposal. What makes this possible is the development of high-speed back-end networks such as FibreChannel, that remove the effects of logical and physical distance between clients and Network-Attached Storage Devices (NASD). For efficiency, clients can be provided direct, protected access to the relevant portions of these NASD that pertain to their needs.

To achieve these goals requires research along a number of fronts:

- attribute specifications and QoS measures,
- device specifications,
- assignment algorithms to partition work between devices,
- runtime monitoring and enforcement techniques,
- secure, network attached storage peripherals to meet QoS goals, and
- distributed, scalable, fault-tolerant algorithms for the above.

1 Network Storage Conference ‘96, Michael Peterson, Strategic Research Co.