Task Force on Network Storage Architecture: Network attached storage is inevitable

Dave Anderson

Seagate Technology Incorporated, 8053 East Bloomington Freeway, Bloomington, MN 55420

Position Statement

We see networked attached storage as being a common means of storage connection and access by the year 2000. This will not, at first, be manifested in direct communication between any node on a network and an individual disc drive. Rather, the likely first step is that storage will be linked with clusters of file servers as part of a server-wide network. Later, broader sharing of storage access may appear.

The reason for taking this position is quite simple: our customers are telling us that they want to and will do this. The acceptance and adoption of Fibre Channel as both an intersystem and I/O channel interconnect is a reflection of and a stimulus for this trend.

I believe there are four reasons why systems suppliers - and their customers - find this direction in storage attachment to be right for large segments of the computer systems market.

First, networked attached storage lends itself to better scalability than the traditional approach of storage attached via a local channel. One of the most significant obstacles that must be overcome to provide increasing transaction performance, which includes most file server activity, is a CPUs difficulty scaling the number of I/O interrupts it can service. Storage hidden behind one or two servers, limits access to that storage according to the number of I/O interrupts those two processors can service. Although CPU cycle time and more sophisticated ALU architectures have produced dramatic improvements in the quantity of instructions microprocessors can execute, there has not been a corresponding advance in the ability of processors to scale I/O interrupt processing.

Second, network attached storage supports superior fault tolerant models by making the availability of each element independent of the availability of any other. It is no coincidence that Tandems model for a fault tolerant system was a network of independent components rather than storage tightly coupled to CPUs. Having a cluster of processors with multiple paths to common storage, eliminates the problem of some storage becoming unavailable due to the CPU behind which it sits failing. Since magnetic discs are at most dual ported, they can only be attached directly to two processors, if one should fail the processing power supporting that segment of the storage is cut in half. This problem is avoided by having a cluster of processors access common storage via redundant network links. It could be argued that one link failing cuts in half the access bandwidth to the storage, but I suspect that the probability of failure in a cable is less than a processor, especially when the software reliability is factored in.

Third, network attached storage continues the industry direction toward increasingly open system architectures. It used to be that a computer manufacturer would design the entire system. Gradually more and more of the architecture has been standardized - operating system, system bus, peripheral interface, etc. Multi-CPU architectures have remained proprietary for the most part. I see the movement to networked attached storage part of a trend toward - and supporting - the use of clusters of standard systems as the building blocks in multiprocessing configurations.

Fourth, there are specific applications that make networked attached storage compelling. The most dramatic instance is that of video delivery. While to date this applications has generated considerably more press that profits, it is certainly one that many of our customers are actively interested in. A server really need do no more than be a traffic cop, directing streams of data from the storage to the video interface, such as an ATM switch or cable head. If passing the data through the server can be avoided, and the server only be responsible for initiating streams and managing errors, there is a significant reduction in the processing power - and cost - needed in the server.

The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of Seagate Technologies Incorporated.