

NETWORK FILE MIGRATION SYSTEM

BACKGROUND OF THE INVENTION

This invention relates generally to digital data processing, and, more particularly, relates to systems for efficient writing, storage, and retrieval of data in local area networks (LANs) utilizing file servers.

The use of storage-intensive computer applications such as high-performance, high-resolution graphics has grown significantly in recent years, with indications that it will continue to grow through the next decade. Fueling user demand has been the introduction of lower cost 32-bit workstations and an increase in the base of applications software available for those systems. Because of their computational and graphics power, these workstations are employed in data-intensive applications such as electronic publishing, computer-aided design (CAD) and scientific research.

Paralleling these developments has been the emergence of industry standard communication protocols which permit users to operate in a multi-vendor environment. Each protocol defines the format of messages exchanged between devices in a network, such that the devices cooperatively execute selected operations and perform given tasks. In particular, file access protocols permit at least two machines to cooperate with a file server. The file server stores files and selectively enables remote client devices to read and write these files.

One such protocol is the Network File System (NFS) protocol, developed by Sun Microsystems, which allows users to share files across a network configuration such as Ethernet. It is most frequently used on UNIX systems, but implementations of NFS are utilized on a wide range of other systems. The NFS protocol can be described as a request-response protocol. More particularly, it is structured as a set of interactions, each of which consists of a request sent by the client to the server, and a response sent by the server back to the client. Generally, the response indicates any errors resulting from processing the request, provides data sought by the request, or indicates that the request has been completed successfully. Requests are reissued by the client until a response is received. A response which indicates that a request has been performed is referred to as an acknowledgement.

All file servers and most workstations on modern networks are configured with local disk storage, and consequently all file servers and most workstations experience a variety of storage management problems. Probably the most pervasive of storage management problems is lack of sufficient free space to accommodate users' working sets of active files.

The term "working set" was first used in the context of virtual memory, where it refers to the amount of physical memory required by an application to execute during a limited period of time. If a computer does not have enough physical memory to accommodate an application's working set, performance degrades because the system incurs excessive overhead while swapping virtual pages of memory in and out of physical memory. In contrast, if the system has an adequate amount of physical memory, then performance is good because each page of virtual memory is more likely to be loaded in physical memory when it is needed.

The same concept applies to files. If a user's local disk has sufficient storage capacity to hold those files that the user is likely to need over the course of, say, one

month (that is, his working set of active files), then the user's productivity will rarely be affected by having to wait for a file to be made accessible. But, if the local disk does not have enough room for the active files, then files must in effect be swapped to and from archival storage to make more room. Typically, users' disks are occupied by inactive files—those not part of the working set—so that insufficient space is available for the local set.

More particularly, up to 80% of the network's disks are occupied by inactive files that have not been referenced in over 30 days. This happens in spite of the fact that the size of each user's working set of active files remains fairly stable over time, because inactive files steadily accumulate and there are no tools to efficiently manage that accumulation of inactive data.

Several serious effects are caused by the chronic shortage of storage for active files. In particular, significant time is spent archiving files. Space is made available in small quantities in filesystems scattered around the network. Consequently, new files tend to be scattered randomly around the network instead of being located where they can best serve their users. This severely degrades network performance. Files that get archived are often lost because traditional archiving methods are primitive and lack the ability to easily locate and restore old files. New disks and file servers are constantly being added to the network. In addition to the cost of acquiring the additional equipment, each new piece of equipment must be managed. In particular, the task of backing up the network's servers and workstations becomes very difficult. Each new piece of equipment adds complexity and reduces the overall reliability of the network.

To be a viable solution, a storage system must satisfy several requirements. For example, the system must be automatic. The cost of manually managing distributed storage quickly exceeds the cost of acquiring the storage. This, coupled with the reality that most network management staffs are severely short-handed, dictates that the degree of manual intervention must be minimized. Also, the mechanism must reduce the network backup problem.

Moreover, the system must be transparent to users and to applications programs. Users must not be exposed to the fact that some of their files have been physically migrated to archival storage. From the user's perspective, all files must remain logically in place at all times. From the perspective of applications programs, changes to the programs must not be required to accommodate the storage system.

The system must be efficient, in that it does not degrade network performance, and does not degrade the performance of the servers and workstations being managed.

The system must also be accurate. A critical requirement is that a system's scarce, expensive, high-performance local storage be available for the data most suited to it—i.e., the active data. Only demonstrably inactive files should be moved to archival storage.

Moreover, the system must be scalable, in order to provide a solution for the networks most in need of a solution: the large networks with many gigabytes of data distributed across many clients. The system should also be portable to a heterogeneous set of platforms; and reliable, in that the data that are moved to on-line archi-