NETWORK FILE SYSTEMS AND FILE SERVERS

Network file systems are the natural extension of local file systems. End users and applications can access directories and files that are physically located on a different computer – the file server – over a network file system (Section 4.2.1). File servers are so important in modern IT environments that preconfigured file servers, called Network Attached Storage (NAS), have emerged as a separate product category (Section 4.2.2). We highlight the performance bottlenecks of file servers (Section 4.2.3) and discuss the possibilities for the acceleration of network file systems. Finally, we introduce the Direct

Access File System (DAFS), a new network file system that relies upon RDMA and VI instead of TCP/IP.

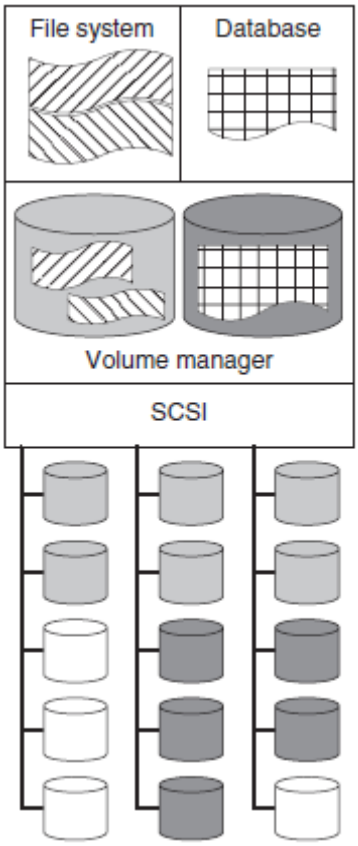

Figure 5.3 The volume manager aggregates physical hard disks into virtual hard disks, which it can break back down into smaller virtual hard disks. In the illustration one virtual hard disk is used directly from a database, the others are shared between two file systems.

5.2.1 Basic principle

The metaphor of directories and files for the management of data is so easy to understand that it was for a long time the prevailing model for the access of data over networks. So-called network file systems give end users and applications access to data stored on a different computer (Figure 4.5). The first widespread network file system was the Network File System (NFS) developed by Sun Microsystems, which is now *the* standard network file system on all Unix systems. Microsoft developed its own network file system – the Common Internet File System (CIFS) – for its Windows operating system and this is incompatible with NFS.

Today, various software solutions exist that permit the exchange of data between Unix and Windows over a network file system. With the aid of network file systems, end users and applications can work on a common data set from various computers. In order to do this on Unix computers the system administrator must link a file system exported from an NFS server into the local directory structure using the mount command. On Windows computers, any end user can do this himself using the Map Network Drive command. Then, both in Unix and in Windows, the fact that data is being accessed from a network file system, rather than a local file system, is completely hidden apart from performance differences.

Long before the World Wide Web (WWW), the File Transfer Protocol (FTP) provided a mechanism by means of which users could exchange files over the Internet. Even today, FTP servers remain an important means of distributing freely available software and freely available documents. Unlike network file systems, access to FTP servers is clearly visible to the end user. Users require a special FTP client with which they can copy back and forwards between the FTP server and their local computer.

The Hyper Text Markup Language (HTML) and the Hyper Text Transfer Protocol (HTTP) radically changed the usage model of the Internet. In contrast to FTP, the data on the Internet is linked together by means of HTML documents. The user on the Internet no longer accesses individual files, instead he ‘surfs’ the WWW. He views HTML documents on his browser that are sometimes statically available on a HTTP server in the form of files or today are increasingly dynamically generated. Currently, graphical HTTP clients – the browsers – without exception have an integrated FTP client, with which they can easily ‘download’ files.

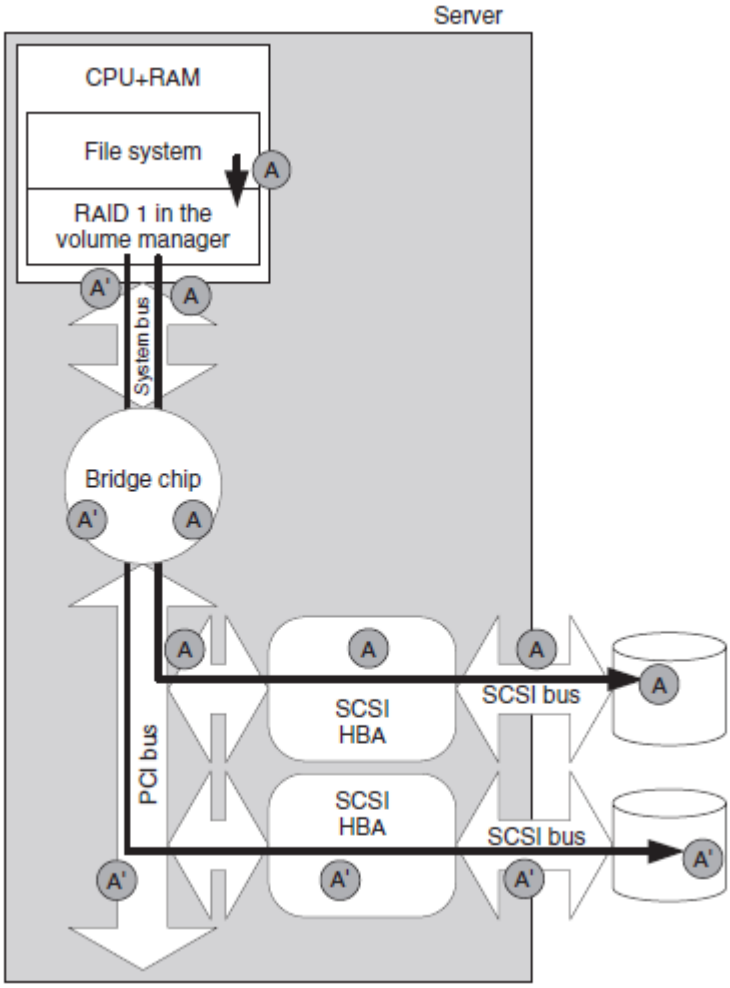

Figure 5.4 RAID in the volume manager loads the buses and CPU of the server. In RAID1, for example, each block written by the file system must be passed through all the buses twice.

Figure 5.5 Network file systems make local files and directories available over the LAN.Several end users can thus work on common files (for example, project data, source code).

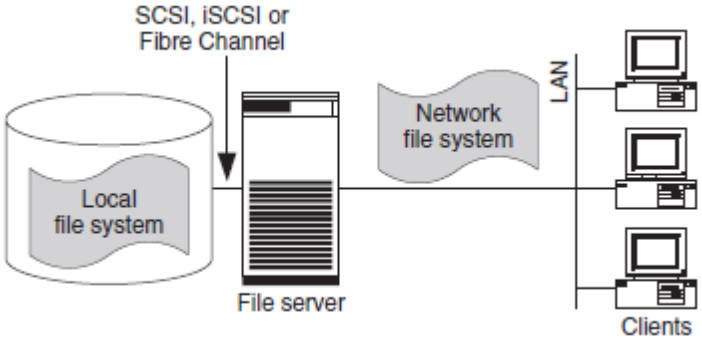

Source : http://elearningatria.files.wordpress.com/2013/10/cse-viii-storage-area-networks-06cs833-notes.pdf