

Document Image Archive Transfer from DOS to UNIX

Susan E. Hauser, Michael J. Gill, George R. Thoma
Lister Hill National Center for Biomedical Communications
National Library of Medicine
Bethesda, Maryland 20894-0001
hauser@nlm.nih.gov
Tel: 301-496-4496
Fax: 301-402-0341

Abstract

An R&D division of the National Library of Medicine has developed a prototype system for automated document image delivery as an adjunct to the labor-intensive manual interlibrary loan service of the library. The document image archive is implemented by a PC controlled bank of optical disk drives which use 12" WORM platters containing bitmapped images of over 200,000 pages of medical journals. Following three years of routine operation which resulted in serving patrons with articles both by mail and fax, an effort is underway to relocate the storage environment from the DOS-based system to a UNIX-based jukebox whose magneto-optical erasable 5 1/4" platters hold the images. This paper describes the deficiencies of the current storage system, the design issues of modifying several modules in the system, the alternatives proposed and the tradeoffs involved.

Background

The Lister Hill National Center for Biomedical Communications, an R&D division of the National Library of Medicine, has developed a prototype system for the automated retrieval and delivery of document images as an adjunct to the manual interlibrary loan service of the library. The system is integrated with the library's existing interlibrary loan system and is transparent to the requester. Since April of 1991, the system has retrieved from optical disk storage and delivered to patrons the images of over 27,000 articles by fax and mail. While the current operation has been scaled down, the system continues to deliver about 450 articles per month and about 550 page images are added to the image archive per month.

The prototype system [1] consists of several DOS-based workstations connected to a LAN and supported by a Netware 3.11 file server. The workstation functions include document capture, image quality control, document tagging, document image archive, communications gateway and document delivery. Most of the software to support these functions was developed in house. The file server serves as a temporary image store until captured images have passed quality control, and it stores the several databases that the system uses to track images and requests.

The image archive is implemented by a bank of four 12" WORM optical disk drives connected via SCSI-1 to a PC. The vendor-supplied software that mediates the operation of the drives configures the workstation as an optical disk server that communicates with other PCs on the network via the IPX protocol used by Netware. Thus, by logging into the optical disk server, other PCs on the network can write and read image files directly to and from the optical platters. All of the files on the optical disk server appear to the PC to be located at a single drive letter. The archive currently holds over 200,000 image files on 15 12" platters, for a total archive of approximately 15 Gigabytes. Because there are more active platters than there are drives, software has been written to effect a "human jukebox" for manual platter exchange.

Optical Disk Server Problems

The four WORM drives of the archive workstation range from 2 to 9 years old and all have been in continuous operation since delivery. These aging drives are no longer supported by the manufacturer. Although maintenance, troubleshooting and some repair and replacement are performed by in-house technicians, parts and high-level repair must be obtained from a third party. Compatible and reliable media are also becoming difficult to obtain. In addition, the frequent manual exchange of platters is taking its toll on both drives and media.

At the time that the optical disk server software was purchased, there were few commercial options for network access to 12" WORM drives from PCs. The optical disk server software was selected because it met our minimum requirements for remote access to optical platters and included a small set of C-callable functions that our in-house programs could use to obtain information about the status of the drives and platters. However, this DOS-based software has not proven to be robust when handling multiple requests and error recovery is generally inadequate, requiring frequent intervention by the technical staff. The original manufacturer of the optical disk server software sold their license to a company overseas with no support staff in this country. The new company has not addressed the reliability and error recovery issues, and their new version of the software cannot write to platters written to by earlier versions.

The optical disk server continues to function adequately at its current low usage level, but at the cost of several man-hours of labor per week. There is also the threat of irreparable breakdown of one or more of the aging, irreplaceable optical disk drives. For these reasons, we are exploring the transfer of the image archive to a more reliable, flexible optical disk server employing current technology.

Rationale

The degree to which images in the archive are accessed is a function of their age, the probability of more recently published documents being accessed being higher than for older documents. One approach to solving the archive problem is to permanently retire disks containing older documents, and have only three or four platters permanently placed in the drives. These would then contain those documents that have the highest probability of being requested, thus reducing wear on drives and media from manual platter exchange. This approach might extend the life of the system for a short time but is not likely to significantly reduce the amount of staff labor needed to maintain the system.

There are good reasons to preserve the entire image archive. These images represent a large investment in equipment and labor. Although the development and operation of the prototype system largely answered the original research questions regarding cost, performance and image quality, the database of document images has potential value for future research. The archive could be used in projects addressing document image processing, image compression, file format conversion, image transfer, image access, or mass storage. It could also prove useful in testing components of improved document image delivery systems.

For these reasons, an effort has begun to relocate the entire image database from the DOS-based system to another system of optical media in which media are automatically exchanged when necessary and multiple network communications are handled reliably.

New Image Storage Requirements

In the new image store, all active images should be accessible from the current document delivery system without manual intervention. To be available to the widest number of future projects, the image database should be accessible from UNIX platforms, which normally communicate via TCP/IP, as well as from the many Netware-based PCs in the division. Internet access to the database would make it available to collaborators at other sites as well. These requirements are met by the division's HP 100 optical disk jukebox [2] in conjunction with the Netware NFS Gateway software [3]. The four-drive jukebox has a current near-line capacity of 93 Gigabytes, expandable to 186 Gigabytes. It is connected to a Sun 670MP and controlled by software from Alphatronix. Each platter side appears as a UNIX file system and is directly available to any computer to which it is exported. Netware NFS Gateway software supports NFS mounting of UNIX file systems to Netware servers, where the file system is available to Netware users as a Netware volume.

The jukebox also supports other projects. Should insufficient space be available for the image database, other commercial solutions are appearing. It is expected [4] that

expandable network storage products will soon emerge that will connect directly to the network and will offer storage that is independent of the operating system. There is one optical disk jukebox system available now that connects directly to the network and supports both TCP/IP and IPX/SPX communications [5].

Software Requirements

In an ideal world, the image database could be moved to a new image store with no effect on the operation of the current document delivery system. However, because much of the in-house-developed software is tightly integrated with the current optical disk server software and the operation of the "human jukebox", no simple substitution is possible. Any change in image store will require modifications to several of the modules that comprise the system. Software modification is not a casual matter. Several of these modules are written for a C compiler that is no longer supported, while others are written for an older version of Microsoft's C compiler. All these modules use a no-longer-supported library of routines to interact with the databases that resolve the location of image files corresponding to journal articles.

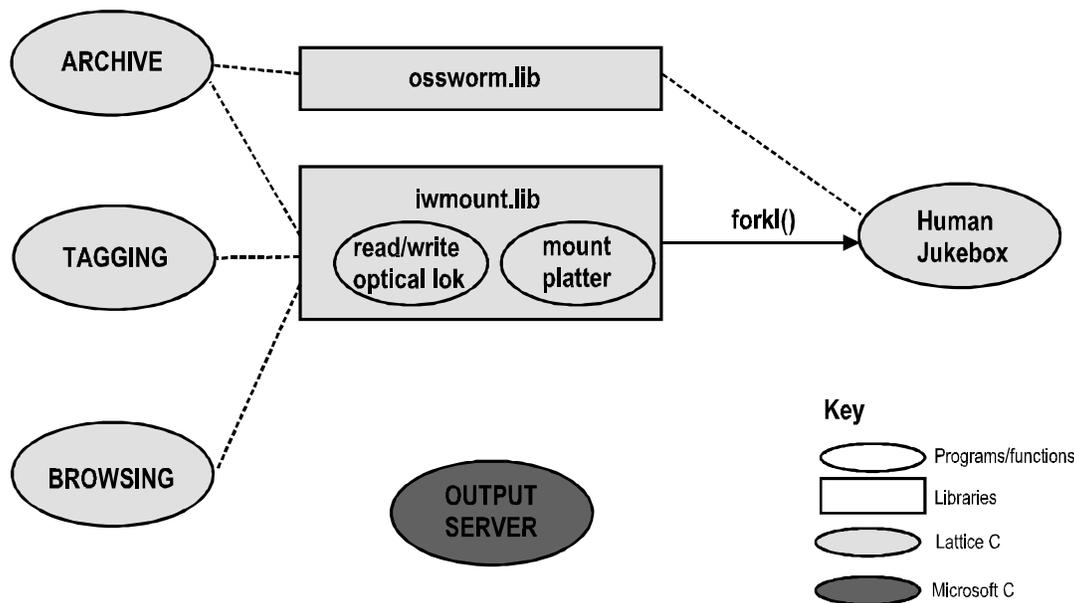


Figure 1. Modules of the current system that access the image store.

Figure 1 illustrates the software modules of the current system that interact with the image store and the libraries that are used to facilitate use of the optical disk server. The archive module moves the page images of a journal issue from the temporary store on the Netware server to permanent store on a WORM platter. For each issue, the

tagging module adds operator-supplied data that identifies the page images that correspond to individual articles in the issue. An operator can use the browsing module to match articles with requests containing ambiguous or insufficient information for the system to automatically select the article. The output server copies page image files corresponding to an article from an optical platter and either faxes the images to the requester or prints the article for delivery by mail. All but the output server interact with an operator.

Ultimately all reads and writes to the optical disk server are straightforward, but modules must first determine if the required platter is in a drive. If it is not, human intervention must be invoked through the module labeled "human jukebox" in the figure. In addition, before archiving a journal issue, the archive module must determine the remaining space on a platter to be certain that there is sufficient space for all page images of the issue. Since a file/platter locking feature is not part of the commercial optical disk server software, all modules use the special optical.lock file to prevent one module from requesting the operator to remove a platter that another module is using. Although three of the modules share a few library functions, as shown in Figure 1, in general each module is responsible for how it accesses files on the optical disk server.

Minimum modifications to the software to accommodate a new image store implemented by an optical disk jukebox will have to remove references to operator intervention and to the functions that obtain information about drive and platter status.

Other issues

File format and image organization: The page images in the current system are compressed using the CCITT Group IV algorithm. Each page is stored as a separate file with no header. All of the page images from one journal issue are stored in one subdirectory. The metadata that describes which page images are associated with each article in the issue are stored in one file in the subdirectory with the images. The subdirectory name is a number, assigned consecutively at the time the issue is archived. Thus, any module using images as they are currently stored must obtain information from the system database files to find the path to a given issue, must be able to interpret the metadata file to find pages for a given article and must have *a priori* knowledge of the image file format. To make the image database not only available to a wider audience, but also self-explanatory, changes in file format and organization will be considered.

Access time: Very fast image retrieval is not critical to the system supporting interlibrary loan since the recipients of the articles are not on line waiting for delivery. Earlier studies of jukebox performance [6] found that the time to retrieve one article is about two seconds when the platter on which the images resides is in a drive. When the platter is not in a drive, the retrieval time becomes a function of the number of

other requests waiting for service from the jukebox. In general, retrieval times from the jukebox are sufficiently fast to support the interlibrary loan prototype system. If the image database is used for some other project for which inherent retrieval times from the jukebox are too slow, apparent speed can be improved by designing a prestaging algorithm specifically for the application.

Backup: Once each 12" WORM platter of the prototype system is filled to 95% of its total capacity, a duplicate platter is made using in-house software, and a new platter is formatted for succeeding documents. Should a platter fail, which has happened, the backup can be used in its place. To date, the magneto-optical (MO) media in the jukebox have proven to be reliable. Since it is unlikely that an entire MO platter will fail, it may be sufficient to back up the document files to tape in case individual files should become corrupted. The important issue of effective backup procedures has yet to be fully addressed.

Platter spanning: The software controlling the jukebox supports platter spanning [7]. With spanning, up to 16 platter sides can be merged to become one filesystem of about 4.5 Gigabytes. The filesystem can be exported to the Netware server and made available to PC users as a single Netware volume. The current image database would require more than three such volumes.

Proposed Solutions

In addition to the hardware and software requirements discussed earlier, the design of a new image store should include as goals: a) minimum investment of labor and equipment, and b) maximum flexibility to allow future changes to the image store and future use of the image database. Meeting these goals involves tradeoffs. Minimum investment in labor and equipment implies minimum software modifications to the current document delivery prototype system and the use of current storage devices, namely the HP jukebox. Maximum flexibility to position the database for continued and future use may require new hardware procurements for the image store and extensive changes to the current software. These solutions are discussed below.

Solution 1: For Minimum Cost

To minimize software modifications, the new image database would be organized exactly like the current database, with the images of each issue residing in an arbitrarily named subdirectory, accompanied by a cryptic file containing data used to connect individual pages to the respective articles in the issue, and using a headerless file format for the images themselves. To minimize modifications, the current conglomerate of outdated compilers, databases and user interfaces would be preserved. The result may support the document delivery system for several years, but would

provide other applications only awkward access to the images. Furthermore, should the image database require another physical move, to be distributed among several servers, for example, the software would likely have to be modified once again.

Figure 2 illustrates how modules of the document delivery system that access the image store would be organized in a system designed to minimize labor and equipment costs. In this scheme, a selected subset of the images are moved to platters in the HP optical disk jukebox connected to the Sun host. Sixteen platter sides in the jukebox are spanned to create one 4.5 Gigabyte filesystem that is exported to the Netware server. Only the more heavily requested issues would be copied to the new image store, with one Gigabyte reserved for about 2 years worth of additional documents. The remaining images are permanently retired. Files in the new filesystem are organized exactly as in the current document delivery system. The entire filesystem appears to the document delivery workstations as one Netware volume which is mapped to a single drive letter, just as the current optical disk server is accessed though a single drive letter. Functions in the `iwmount.lib` library, which were previously used to operate the "human jukebox", are replaced by functions bearing the same name whose only purpose is to return a good status. In this way, the tagging and browsing modules need not be rewritten, but only relinked to the new library. Because the archive module is so tightly integrated with the current optical disk server and includes functions, such as determining available space on a platter, that are not in the `iwmount.lib` library, it must be rewritten to support the same functionality with respect to the jukebox. The required information can be obtained through functions in the Netware Software Development Kit (SDK). Since the SDK supports Microsoft, but not Lattice compilers, the new archive module is written for Microsoft. The output server module may not need to be rewritten or relinked. Because it is intended to operate automatically, even during periods when an operator is not available, it does not directly access the "human jukebox" functions.

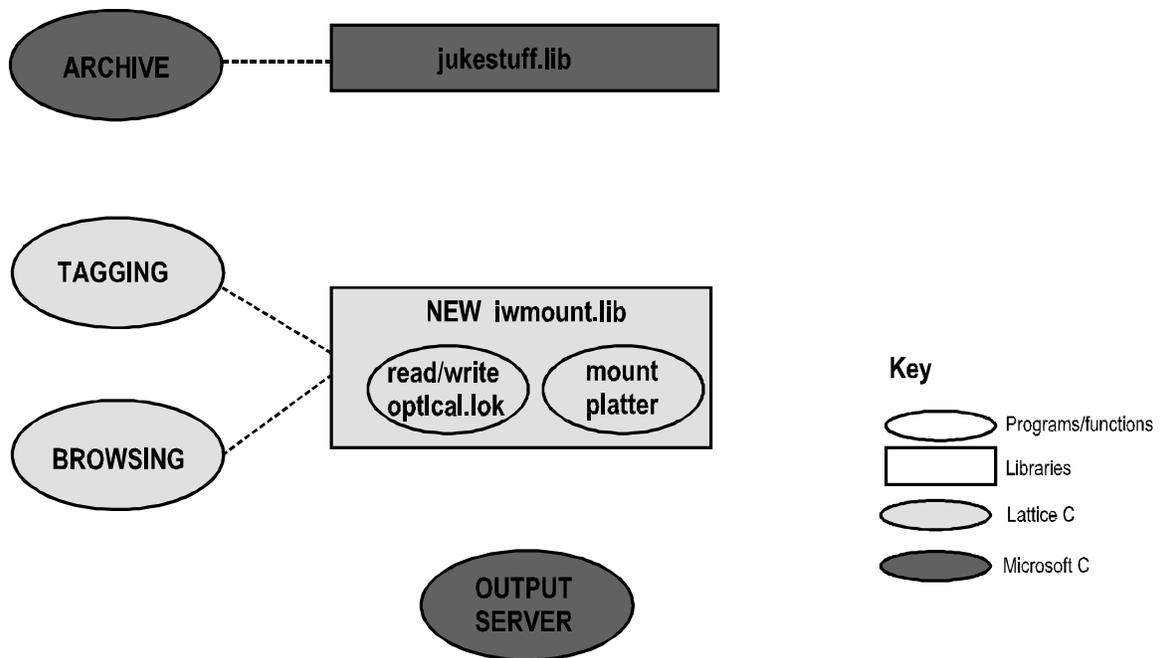


Figure 2. Proposed modules for minimum cost solution.

Solution 2: For Maximum Flexibility

To maximize flexibility and access, the image database would be reorganized for easy management and for intuitive navigation to image files that have a standard file format and header. The path to the subdirectory containing the images for one issue includes a three letter code identifying the journal title and other characters to indicate volume and issue. With each issue there is an easily interpreted text or database file containing data linking page images to articles. All images reside in an optical disk jukebox with sufficient capacity to store the existing database plus at least five years expansion. The jukebox connects directly to the network with software that supports access via both TCP/IP and IPX/SPX. Each platter side appears as one volume to Netware clients and as one filesystem to UNIX clients. All modules of the document delivery system that access the image database would be rewritten to reflect the new organization and location of the image archive. The result would permit easy access to the database by both Netware and UNIX applications. However, the high cost of the new, sophisticated image store and the many person-months of programming effort may need to be justified on programmatic grounds.

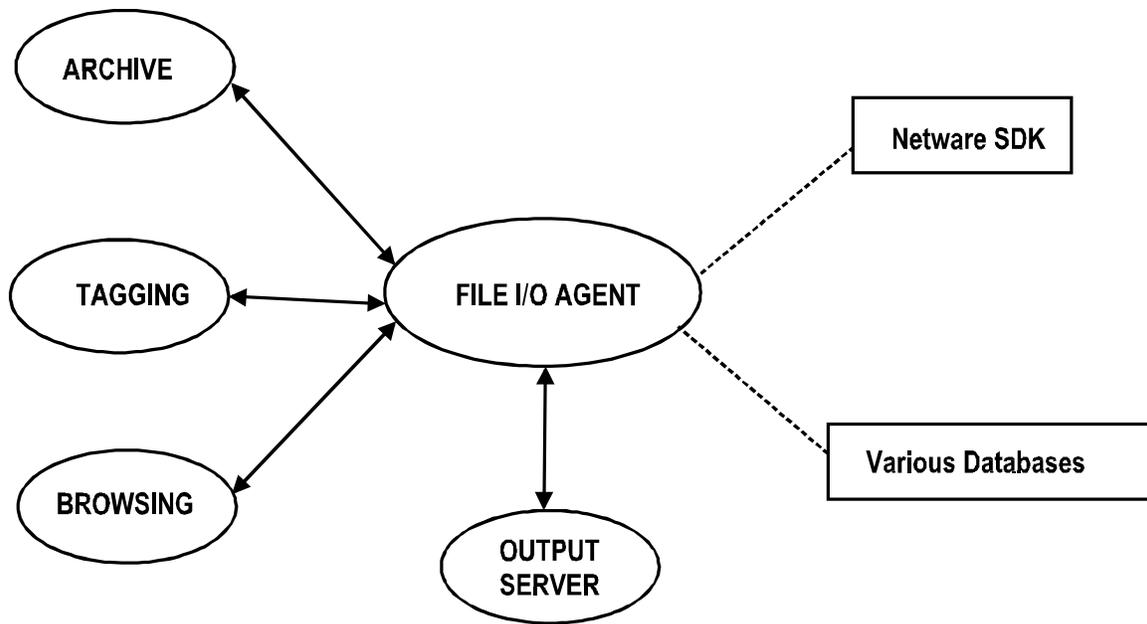


Figure 3. Proposed modules for maximum flexibility solution.

Figure 3 illustrates one concept of how modules of the document delivery system that access the image store would be organized in a system designed to maximize flexibility and access. All modules are rewritten to reflect the new system and file organization. They are no longer individually responsible for understanding image database organization or file location, but invoke a new module for all image file access. The new file I/O "agent" is responsible for specifying or discovering the location of any file and mediating all reads and writes to the image store. It creates and uses database information to determine the path to a given file and employs the functions in the Netware SDK to obtain information about the volumes containing the files. Should the image database be relocated or distributed among several sites, only the databases used by the file I/O agent would be changed.

Conclusions

Most hardware and software requirements for a new image store are satisfied by the division's HP 100 optical disk jukebox connected to a Sun UNIX platform and accessible from PC Netware clients via the Netware NFS Gateway. Moving the image database to a UNIX platform also immediately increases its exposure to a new set of clients and potential applications. But moving the image database to any new location demands changes in the software of the application for which it was originally created. The difficulty in determining the design of the new image store lies with the conflicting goals of minimizing cost and maximizing flexibility and access. The final

decision on the design approach will depend upon the importance to the organization of being able to use the image database in the future.

References

1. "System for Automated Interlibrary Loan (SAIL): System and Operations Description," Internal technical report, Communications Engineering Branch, Lister Hill National Center for Biomedical Communications, National Library of Medicine, Bethesda, MD, November 1992.
2. Hewlett-Packard. Optical Disk Library Systems: Product Family Brief, April 1991.
3. Novell, Inc. NetWare NFS Gateway Supervisor's Guide, 1993.
4. Root R. Data storage in the networks of the '90s. *Inform*, July 1994.
5. Symmetrical Technologies. SPANStor-Opti data sheet, 1994.
6. Hauser S E, Roy G, Thoma G R. Optical disk jukebox performance in multi-user applications. *Proceedings of the 1994 Optical Data Storage Topical Meeting*; Vol. 10: pp.53-55.
7. Alpatronix Inc. Inspire Jukebox: User's Manual for Sun Workstations, June 1993.