THE USE OF WRITE-ONCE READ-MANY OPTICAL DISKS FOR TEMPORARY AND ARCHIVAL STORAGE

By Brenda L. Groskinsky

U.S. GEOLOGICAL SURVEY

Open-File Report 92-36

Portland, Oregon
1992

U. S. DEPARTMENT OF THE INTERIOR

MANUEL LUJAN, JR., Secretary

U.S. GEOLOGICAL SURVEY

Dallas L. Peck, Director

For additional information
write to:

District Chief
U.S. Geological Survey, WRD
10615 S.E. Cherry Blossom Drive
Portland, Oregon 97216

Copies of this report can
be purchased from:

U.S. Geological Survey
Books and Open-File Reports Section
Federal Center, Box 25425
Denver, Colorado 80225

CONTENTS

THE USE OF WRITE-ONCE, READ-MANY OPTICAL DISKS FOR TEMPORARY AND ARCHIVAL STORAGE

By

Brenda L. Groskinsky

ABSTRACT

A study was conducted by the U.S. Geological Survey to evaluate a prototype write-once, read-many disk drive for the purpose of temporary storage and archiving of the U.S. Geological Survey's records and documents. Storage of documents and data is a requirement for all activities in the U.S. Geological Survey. Two types of storage are (1) routine backup of digital data and (2) archival storage.

A 2-gigabyte, write-once, read-many drive was connected to a 32-bit workstation. An optical image scanner was connected to the same workstation so that images on paper could be converted to digital form for storage on write-once read-many optical disks.

Write-once, read-many, optical disks may offer an improvement over magnetic tape storage for backup procedures because of their great storage capacity. The amount of data routinely stored for temporary backup of the workstation during 1 month was copied to a write-once, read-many, optical disk. Time requirements of the magnetic and write-once, read-many, optical-disk systems used for temporary backup were compared; functionality and any differences in software were considered. Write-once, read-many, disk-backup procedures may be appropriate for those sites that use magnetic-tape procedures for backup and lack the personnel required to change the tapes.

Write-once, read-many, optical-disk technology may offer an improvement over printed paper or microfiche for archival storage, because the disks require minimal storage space and do not require humidity- or temperature-controlled environments. More than 300 well logs selected from the Portland, Oregon, metropolitan area were scanned, indexed by range and township numbers, county, and ownership into a directory system maintained by the write-once, read-many, drive software, and stored on write-once, read-many disks for retrieval by monitor or laser printer. Image data were easily stored and retrieved to magnetic disk from the write-once, read-many, optical disk.

INTRODUCTION

Storage of documents and data is a requirement for all activities in the U. S. Geological Survey (USGS). Two types of storage are (1) routine backup of digital and (2) archival storage. Routine backup of digital data is temporary storage of digital computer data. Backup of digital computer files or data is done as a precautionary measure against digital data-loss caused by accidental file deletion, computer disk failure, or by some catastrophic event. Routine backup schedules for temporary data storage differ from site to site in the USGS, but generally are performed daily by downloading data from magnetic disks to magnetic tapes. Archival storage of data in the USGS is storage of records and image data for periods longer than 10 years for legal posterity reasons. In most cases, image data are not in digital form. Currently, archival data are stored on paper or microfiche.

USGS offices are accumulating large amounts of data that require permanent and temporary storage. The network of USGS minicomputers has distributed tremendous computing power to field offices, as well as the ability to store massive quantities of current and historic data. The new computation power and ease of data access have redistributed the responsibility for maintaining and preserving USGS data to field office managers.

Temporary routine backup of on-line data is labor intensive and costly. Backup routines differ among USGS offices, but full backups are customarily performed weekly, with backups of new and modified files performed daily. The storage capacity of hard-disk volumes has steadily outpaced the growth in magnetic-tape technology. A reel of magnetic tape can store 60 to 180megabytes of data, whereas 600-megabyte disk drives are commonplace and disk drives nearly twice this capacity are available for new computer systems. Although most of the backup procedure can be automated, manual intervention is still required to monitor and change magnetic-tape reels. For large offices, backup routines commonly take 6 or more hours for completion. This is an inefficient use of daytime computer resources.

Although much of the USGS total archival data is inactive and stored off-site, some noncomputerized file systems are actively used on-site. Retrieval of archival information stored on paper or microfiche is a cumbersome and costly manual process.

Optical storage devices may offer improvement over current methods to meet increasing demands for archival and routine backup storage of USGS's image and digital data. One optical-disk-storage system, write-once-read-many (WORM), has undergone considerable technological advancement in recent years. Data are stored in binary form as tiny depressions burned by a laser into the surface of a plastic or glass disk. Once formed, these depressions are suitable for long-term (archival) data storage. A WORM disk drive is capable of storing 100 megabytes to 2.4 gigabytes of data, depending on the disk size. WORM-drive systems are available for personal computers, microcomputers, minicomputers, and mainframes. WORM archives do require computer software and magnetic-drive hardware for retrieval, and the viability of a WORM archive is dependent on the existence of these two computer resources. Depending on the supporting software, WORM drives appear to the host computer system as either an additional hard disk (random access) or tape drive (sequential access). It is anticipated that WORM technology will greatly increase usability and will decrease costs as compared to conventional manual methods of archival data storage. The contents of the paper or microfiche can be stored on a WORM disk as scanned images or as digital-text data. WORM disks have the potential of providing longer storage life than either paper or microfiche.

WORM technology many offer an improvement for backup procedures because of their great storage capacity. A backup system that uses a WORM drive does not require manual intervention, except for the initial disk load onto the drive, and could automatically perform routine-backup operation during times of low demand for computer resources, such as at night and over the weekend.

Although printed paper and microfiche can remain unaltered for time periods longer than 10 years, they require temperature- and humidity-controlled environments. These conditions are not necessary for archival data stored on WORM disks. Because a controlled environment is not necessary and storage for WORM disks is very minimal compared to file cabinets (one 12-inch by 12-inch by 2-inch disk equals two full file cabinets), the archive disks could be stored locally on-site, with backup copies off-site.

Current Federal government users of WORM-drive technology for data storage and retrieval include the Library of Congress, U.S. Patent Office, and National Aeronautics and Space Administration (NASA). NASA did a pilot project in February 1986 to evaluate a WORM-drive system for document storage. They concluded that the interface between mainframe data bases and WORM storage systems is effective and reliable (Manns and Wilder, 1987, p. 11). The Patent Office has 180 gigabytes of on-line data, 75 percent of which is now stored on WORM disk drives. In addition to their own use of WORM drives for storage, NASA is considering the use of WORM disks containing archival space data as a means of data distribution. Research by NASA has shown that an automated library mechanism using nonerasable media would best support their archival and distribution requirements (Kempster and Martin, 1987, p. 30).

PURPOSE AND SCOPE

The purpose of this paper is to describe the results obtained from the prototype testing of a WORM storage system on a 32-bit Unix workstation. These results include evaluation of the WORM and scanner driver software that was installed on the workstation, time assessments of the archive and backup processes, and an assessment of the transfer process of textual and image data between USGS computer environments to the WORM drive. Specifically, this study evaluated (1) the ability to transfer digital and image data between USGS computer environments, (2) optical-disk temporary storage as a more efficient method of routine backup within USGS, and (3) optical-disk procedures for the archival and retrieval of USGS digital and image data.

APPROACH

To meet the objectives, the Oregon USGS Office purchased a 2-gigabyte (1-gigabyte per side) WORM disk drive and connected it to a SUN 3/50 diskless workstation. An additional purchase requirement was an image scanner, to convert images to digital form. The scanner was connected on the Small Computer System Interface (SCSI) bus of the workstation along with the WORM drive. The scanner was purchased with driver software that was installed on the workstation. The driver software was necessary for scanner usage and also for retrieval of scanned, image documents. The scanner and the WORM driver software required kernel modification.

Experiments in accessing the optical disk from USGS computers connected by local and wide-area networks were conducted. The amount of data routinely stored for temporary backup of the SUN system during 1 month (approximately 2 gigabytes) was copied to optical disk. The amount of time magnetic and optical systems used for temporary backup was compared, taking into account functionality and any differences in software.

More than 300 well logs (8 x 11 inch carbon copy of well information reported by a well driller) selected from four townships in the Portland, Oregon, metropolitan area were scanned, indexed, and stored. The WORM drive appears as a tape drive to the workstation operating system except that the driver software provides a directory system which is stored on the WORM disk. The directory system allows a user to index files to be written to the WORM disk. The well-log index is generated by the directory system of the WORM driver software. The directory system allows for three input keys. First and secondary file names are used along with a third category which allows for 205 American Standard Code for Industry Interchange (ASCII) characters. Two commands were used to store well logs on the optical disk. The first command made an entry into the index. The second command wrote the well log to the optical disk. The well-log index is modified every time a new well log is stored on the WORM disk.

Well logs are retrieved by searching the WORM directory system for the log needed and then copying the well-log file to the magnetic disk of the workstation. The directory system has a standard Unix wild-card searching option for first and secondary file names. If the ASCII character category needs to be searched, the WORM directory can be copied to the magnetic disk of the workstation as a file. This file can be searched for key information found in the ASCII character category. The retrieved well log can be viewed on the monitor or printed. Retrieved well logs are deleted from the magnetic disk once their use is complete.

RESULTS AND DISCUSSION

Connecting the WORM drive and the scanner on a diskless rather than a diskfull workstation caused some initial problems accessing the added peripheral equipment due to the differences in the operation-system kernel between diskless and diskfull workstations. It was determined by the WORM drive vendor that a special user identification was needed to allow the WORM drive to write properly to the directory located on the WORM disk. Because the WORM drive and the scanner were purchased from different vendors, the installation of both devices on the same workstation through a single SCSI bus required additional assistance from the vendors to complete the software and hardware installation. Several files. which are used by the operating system to describe devices to the kernel, needed to be modified. After the modification of these files a new kernel was generated to incorporate both the WORM and the scanner device drivers.

Once the installation of the equipment was complete the ability of the devices to perform between USGS computing devices was tested. The WORM drive simulates a magnetic tape drive. The standard UNIX utilities that were tested using the WORM drive are *dd, tar, dump* and *restore*. These utilities performed satisfactorily for both reading and writing of files to and from the WORM disk over the Local Area Network (LAN) of the 32-microcomputer system. LANs are commonly used in workstation environments as a means to share resources. When using the *dump* utility to write to the WORM disk from a remote workstation, it was necessary to access the WORM disk through a remote shell procedure allocating the appropriate partition on the WORM disk to the user-id of the remote system. After this procedure was complete, the WORM disk could be read or written to by specifying the workstation name associated with the WORM drive as an argument to the *dump* command. An example of a procedure using *dump* from a remote workstation follows:

rsh osiris *odt* all

rsh osiris *odt* newf dumptest -s date_of_dump

dump 0ubsf 126 50000 osiris:/dev/rsod0a /dev/sd0g

The first command allocates the default partition (/dev/rsod0a) of the WORM drive to the remote user using the WORM driver command *odt*. The name of the workstation where the WORM resides is osiris. The second command creates a directory entry on the WORM disk. The third command executes the utility *dump* which copies the contents of the remote magnetic disk partition /dev/sd0g to the WORM drive on the workstation osiris. The block size is specified as 126 which is the maximum data which can be transferred for the workstation during a single read or write (63 Kilobytes = 126 x 1024 byte blocks). The size option of *dump* is set to 50,000 specifying that one side of a WORM disk is equivalent to a 50,000 foot magnetic tape.

A successful *tar* command can be performed over the network during a remote login. Typically a *tar* command is used to archive files to magnetic tape. In comparison, the *dump* command is used to backup entire file systems.

Even though the WORM drive simulates a magnetic tape drive, it differs from a tape drive in that it allows multiple users. As many as six additional users can read data or search through directories on the WORM disk while one user is writing to the WORM disk. This feature was tested while one remote user was writing data from a remote magnetic disk during a normal backup procedure using the Unix *dump* command to the WORM disk and another user was reading files to magnetic disk from the WORM disk using the Unix *tar* command.

Data from a 327-megabyte disk were copied remotely to the WORM drive and to a 9-track tape system to compare time requirements for backup operations. Because the WORM and tape drives are located on remote workstations, the additional time factor for networking was common to both devices. The backup of a 3-partitioned magnetic disk to the WORM disk took nearly twice as long as backup using a 9-track tape. However, the WORM disk has has the capacity to hold up to 2-gigabytes of data, as compared with a 9-track tape which, at a density of 6,250 bytes per inch (a 2,400 foot reel), will hold 180 megabytes of data. A full-backup of most computer facilities cannot be contained on one 9-track, 2,400 foot reel tape. The WORM drive allows the backup routine to run overnight during nonpeak computering time. The 9-track backup routine requires user attention to load and unload new tape reels.

More than 300 well logs selected from the Portland, Oregon, metropolitan area were scanned, indexed, and stored for retrieval by monitor or laser printer. Data-base management software was considered as a means to index the archival image data; however, this system would require index storage on the magnetic disk of the workstation which is a media that does not meet the requirements of archival storage. Therefore, the index system for the archive is stored on the WORM disks themselves using the WORM driver software. This methodology eliminates storage of the index on the magnetic disk of the computer workstation.

The archive of the well logs is created with three distinct processes. The first process is to electronically scan the logs. This step is the most time consuming. Images are scanned and stored on magnetic disk at a rate of 3 to 5 minutes per well log. The start-up process of the scanner software, which is menu driven, requires the user to set page sizes specific to the well logs. The image scanner used for this prototype has a software memory allocation error which causes the software to fail after approximately seven image scans. The software can then be restarted, but it is necessary to establish the page-size setting again. Initially, the entire 8x11 inch log sheet is not shown in its entirety. Once the well log is scanned, the user is required to zoom by a factor of 75 percent on the image to show the entire 300 dots per inch (DPI) document. The operator has the ability to alter the contrast for every image that is scanned. The contrast setting can in some cases improve the quality of the image. The user, however, is required to set the contrast setting prior to the scan. A typical well log image has a file size ranging from 300,000 to 400,000 bytes, when stored as an encoded 300 DPI raster file. Encoding the image file allows more efficient use of the optical-disk space. The range in file size is caused by the contrast setting of the scanner and the amount of information that is written on the well log.

The scanned documents are then entered into the directory or index system of the WORM disk. The WORM driver software allows the user to write two file names and up to 205 ASCII characters for each individual file. The directory entry made with each well log includes a primary file name (which is unique as are all file names on most computing environment) which has the well-log number, the township, the range and the section number. The county where the well is located is used as the secondary file name to the index. The owner of the well is entered into the ASCII character section of the directory entry. A typical directory entry is made with the following command:

odt newf log105.01N-02S-14ab -s Washington -l owner name.

In the above example, the primary file name is log105.01N-02S-14ab and the secondary file name is the county name Washington.

The last step of the archive process is to write the scanned file to the WORM disk. Once the directory entry for the scanned image file is written to the WORM disk, the positioning of the WORM disk where the file is to be written is complete and the file can be directly written to the WORM disk with the Unix *tar* utility:

tar cvf /dev/rsod0a log105.01N-02S-14ab.

Collectively, the last two processes of making the directory entry and writing the file to the WORM disk takes approximately 2 to 3 minutes for each well log. The WORM drive has a raw-data transfer speed of 479 Kilobytes/second and can transfer data on the SCSI bus at 1.5 Megabytes/second. On the workstation used, (Sun 3/50), the WORM drive has been benchmarked at sustained reads of 220 Kilobytes/second and sustained writes of 150 Kilobytes/second. The WORM drive checks each block to make sure it is blank before writing which causes a slower write time (Delta Microsystems, 1988). The capacity of one side of a 2-gigabyte WORM disk is approximately 3,000 scanned pages of well logs.

The retrieval process of the archival image data involves searching the directory index system of each WORM disk in which the archive is located for the appropriate log and extracting the log file (which is stored as an encoded raster file) from the WORM disk. The user can display the directory entries on one side of the optical disk at a time by typing:

odt ls -l.

The user can use the standard wild-card options common with the Unix directory listing command, *ls*, on the WORM disk directory entries. This option permits the user to search for specific townships, ranges, sections or counties. Once the directory listing is displayed on the screen or redirected to a file, the user can search for the owner name. When the desired log is found, the user positions the WORM disk to the file using the command:

odt pos log104.01N-025-14ab.

The command in the above example positions the WORM disk according to the unique primary file name.

The well-log file is written onto the workstation's magnetic disk using the Unix *tar* utility:

tar xvf -f /dev/rsd0a.

After the log is written back to the magnetic disk of the workstation, it is necessary to decode the raster file with the use of the menu-driven scanner software. The page-size and zoom options must be adjusted to examine the entire well log. To print the well-log file on a postscript laser printer, the user is required to save the file using the scanner software as an encapsulated postscript file to the magnetic disk of the workstation. For each well log file, this process takes several minutes and creates a disk file up to 1 million bytes in size. The scanning software allows the user to send the encapsulated postscript file to the printer directly with several remote printing exceptions. Some large files cannot be printed remotely if the file exceeds the maximum allowable size for the Unix queuing facility to transfer the file to a remote workstation connected to the selected printer. It was discovered in this prototype that the printing of a 1 million byte well log (encapsulated postscript file) through the serial port on the workstation took as long as an hour to print.

SUMMARY AND CONCLUSION

A prototype WORM optical disk was evaluated for the purpose of temporarily storing and archiving of U.S. Geological Survey's important records and documents. The time assessment of the backup process showed the WORM disk to be half the speed of 9-track tape device under similar conditions. The speed at which the WORM disk wrote and extracted files, including making a WORM disk directory entry, took less than 2 minutes, in most cases, for the well-log application. The time involved in printing a large retrieved image well-log file (as long as 1 hour for an 8x11 inch image page) was long. The cost factor (personnel time compared with equipment) of doing the daily or incremental backups on a WORM disk needs to be accessed by individual sites. Because of the size capability, a WORM disk backup procedure may be appropriate for those sites that lack the personnel required to change tapes during a magnetic disk backup.

The assessment of the transfer process of textual and image data from a USGS computer environment to the WORM disk was favorable. Once the data are on the magnetic disk of the computer platform, the data can be easily written to the WORM disk. Once the image data are stored on the WORM disk, the data can be easily retrieved to magnetic disk. Problems retrieving image data from magnetic disk to a printed file occur because of the large file sizes involved. Software is required to transfer image data to a printable form. Advantages using the WORM driver software for the archival application were that the driver software allows the user to put a directory system of files on the WORM disk and allows simultaneous access by multiple users. The scanner software used in this prototype has a memory allocation error which caused the software to fail. The restarting process added appreciably to the average scan time. There was no a software link between the WORM disk and scanner drivers. A software link between devices would have allowed for more efficient use of the equipment because the transfer of image data between the scanner and the magnetic disk would have been eliminated.

SELECTED REFERENCES

Antonuccio, Anthony, 1986, Tape backup systems: Byte, v. 11, no. 5, p. 227-232.

Delta Microsystems, Inc., September 3, 1988, SS-2000WA optical disk drive users manual, part number 801-SS-200WA, Revision 1.0, p. 32.

Kempster, L. S., and Martin, J. B., 1987, In search of: NASA space data storage solutions: 1987 Electrical and Electronics Engineers Symposium on Mass Storage Systems, Institute of Electrical and Electronics Engineers Computer Society Press, p. 27-31.

Manns, Basil and Wilder, Dean, 1987, Interfacing a mainframe database retrieval system with an optical disk image storage system: 1987 Electrical and Electronics Engineers Symposium on Mass Storage Systems, Electrical and Electronics Engineers Computer Society Press, p. 7-11.

Malloy, Rich, 1986, A roundup of optical disk drives: Byte, v. 11, no. 5, p. 215-224.

McLeod, Jonah, 1988, Optical storage may fulfill its promise - after years of trying: Electronics, v. 61, no. 10, p. 75.

-----, 1987, The optical-disk industry gropes for standards: Electronics, v. 60, no. 6, p. 82-85.

Microtek Lab Inc., 1989, ScanVu user's guide.

Rosch, Winn L., 1987, Backup choices from tapes to disks to WORMs: PC Magazine, v. 6, no. 12, p. 102-108.

-----, 1987, WORMs for mass storage: PC Magazine, v. 6, no. 12, p. 135-148.

-----, 1987, LaserBank 400: best of WORM, hard-disk worlds: PC Week, v. 4, no. 16, p. 21.

Stephens, Mark, 1987, WORM disc drive features 800 magabyte storage capacity: Infoworld, v. 9, no. 126, p. 19.

Teja, Ed, 1985, Optical disk drives expand workstations capacities: Electrical Design News, v. 30, no. 12, p. 34-41.