

Redundant Array of Independent Disks: Practical On-Line Archiving of Nuclear Medicine Image Data

James L. Lear, Jonathan P. Pratt, and Nelson Trujillo

While various methods for long-term archiving of nuclear medicine image data exist, none support rapid on-line search and retrieval of information. We assembled a 90-Gbyte redundant array of independent disks (RAID) system using 10-, 9-Gbyte disk drives. The system was connected to a personal computer and software was used to partition the array into 4-Gbyte sections. All studies (50,000) acquired over a 7-year period were archived in the system. Based on patient name/number and study date, information could be located within 20 seconds and retrieved for display and analysis in less than 5 seconds. RAID offers a practical, redundant method for long-term archiving of nuclear medicine studies that supports rapid on-line retrieval.

Copyright © 1996 by W.B. Saunders Company

KEY WORDS: nuclear medicine computers, personal computers, image archiving, picture archiving and communication system (PACS).

PERSONAL computer-based networks for acquisition, distribution, and archiving of nuclear medicine image data have been available for several years. The use of such systems over standard film-based systems offers several advantages. Because the use of such systems offers advantages compared with use of film-based imaging, they are becoming widely accepted.¹

A problem with the use of all digital data acquisition is how to store the data. A busy nuclear medicine department may perform 5,000 to 10,000 cases per year with an average requirement of approximately 1 Mbyte of storage per study. Therefore, over a period of 10 years, 100 Gbyte of data may be generated.

Such quantities of data can be stored on magnetic tape, optical disks, or CD ROM disks. However, these methods have limitations. Magnetic tape is inexpensive, but data access times are very long (hours). Optical and CD ROM disks are faster, but present disks are limited to approximately 1 Gbyte each. Although jukeboxes help reduce this limitation, they are fairly expensive and mechanical problems can occur.

Redundant arrays of independent disks (RAID) have been developed to facilitate data storage on magnetic drives. RAID uses multiple

disks to store data, so potential storage space is greater than with single-disk systems.

For several years, RAID systems have been available in two configurations known as levels 0 and 1. With level 0, information is split between two disks. This approach allows data to be stored or retrieved at a rate approximately twice that of a single-disk system. However, a disadvantage is that if one disk of the pair fails, all data are lost. In other words, system failure will occur at twice the frequency of a single-disk system. Also, as the data transfer speeds of single-disk drives are increasing continually, this year's level 0 speed is usually achieved by next year's single disks.

With level 1, information is mirrored onto two disks. Thus, data are stored in a redundant fashion. If one disk fails, the other preserves the data. However, this approach has limited speed and requires twice the disk space of a single-disk system to store data. Also, the added expense of the RAID control system is difficult to justify in comparison to inexpensive software that can write information onto two standard single disks, albeit more slowly.

RAID levels 3 and 5 have been recently introduced to solve the problems associated with levels 0 and 1. Although levels 3 and 5 differ slightly between each other, they provide similar functionality. Data are stored in a redundant manner, but not mirrored as with level 1. Rather, data and parity information are spread across several disks. If a single disk fails, data can be reconstructed from the information on the other disks. Therefore, this approach supports both rapid data transfer and redundancy,

From the Department of Radiology University of Colorado Health Sciences Center, Denver.

Supported in part by National Institutes of Health Grant No. NS-26657 and by a Scholar Grant from the Radiological Society of North America.

Address reprint requests to James L. Lear, MD, Nuclear Medicine (A034), University of Colorado Health Sciences Center, 4200 E Ninth Ave, Denver, CO 80262.

*Copyright © 1996 by W.B. Saunders Company
0897-1889/96/0901-0005\$3.00/0*

while requiring only an additional 25% of disk space compared with single-drive systems.

One potential disadvantage with most RAID 3 and 5 systems, with respect to their use with personal computers, is that they require the use of one computer address per disk. Personal computers have only a limited number of addresses available, usually 5 or 6. Thus, the total number of disks that can be used in standard RAID systems is limited. Also, personal computers generally cannot directly access data stored in drives with greater than 4 Gbyte of space.

We developed a modified approach to RAID level 5 data storage for use with nuclear medicine image archiving. The approach was intended to allow the use of vast, redundant disk storage space on a standard personal computer, the POWER Macintosh 8100/80 (Apple Computer, Cupertino, CA).

MATERIALS AND METHODS

Ten high-speed disk drives, each with 9 Gbyte storage capacity, were assembled into two systems of five disks each.

Proprietary controllers (MegaDrive Inc, Beverly Hills, CA) that permitted the use of multiple disks per small computer serial interface (SCSI) address were used as opposed to standard controllers. This approach required the use of two SCSI addresses. Importantly, the controllers also supported "hot-swapability"; if a disk failed, another could be inserted in its place while the system was running.

Software was used to segment the 72 available Gbyte (90 Gbyte \times 80%) into 3-Gbyte partitions. The software caused each 3-Gbyte partition to appear to the computer as a single-disk drive. Each "pseudo-disk" appeared on the computer desktop so that all 72 Gbyte of space could be accessed.

After the system was assembled, it was tested by performing continual data storage and retrieval for 10 days. Simulated failures (removal of disk drives) were performed twice daily.

After testing, the system was optimized for our caseload. Eight years of digital data that we had previously stored on optical disks were loaded into the system. Data were stored in the 4-Gbyte partitions based on the dates the studies were acquired. Within each partition, the studies were stored in 1-month folders. Within each 1-month folder, the data were arranged according to study type, eg, thallium scans, bone scans, etc. Individual studies were placed within the appropriate study-type folder and labeled by patient name, patient number, study type, and date using a maximum of 31 characters (the Macintosh limit).

RESULTS

The system passed all initial tests without problems. After simulated failures, data could

be reconstructed in less than 45 minutes (while the system remained in use).

By compressing only cardiac first-pass data, we could store our 7,000 annual cases using approximately 6 Gbytes of disk space. Therefore, data were stored in 6-month blocks using one 3-Gbyte partition per block, eg, 1/88-6/88, 7/88-12/88.

The time to locate a file and retrieve data was very rapid. A study could be located within any partition in approximately 20 seconds using the standard Macintosh Finder. Retrieval of the study and display on a monitor required less than 5 seconds.

No failure occurred during 6 months of clinical use.

DISCUSSION

RAID technology has reached a level that can support long-term archiving and rapid retrieval of nuclear medicine image data. The cost of a practical system is reasonable. For example, the complete system (computer, controllers, and drives) used in this study was assembled at a cost of approximately \$50,000. Importantly, this cost will decrease by greater than 50% by late 1995 with the commercial release of 18-Gbyte drives.

Of course, no system can guarantee perfect safety of data storage. For example, the RAID system could be damaged by fire or water. Therefore, we back up all data on standard 1.3-Gbyte optical disks that are stored in a fireproof box.

In addition to its use in nuclear medicine, RAID level 5 will likely have important use in storage of digital information in a general radiology department. A system with 250 Gbyte of redundant data storage capacity can be assembled using 20, 18-Gbyte drives at a cost of approximately \$100,000. Assuming an average file size of 10 Mbyte, more than 25,000 studies can be stored in such a system. This number of cases represents a month of all studies in a busy department. Thus, RAID can be used as a buffer between immediate and long-term storage in a complete picture archiving and communication system network.

REFERENCE

1. Lear JL, Pratt JP, Roberts DR, et al: Gamma camera image acquisition, display, and processing with the personal microcomputer. *Radiology* 175:241-245, 1990