Simulation of Disaster Recovery of a Picture Archiving and Communications System Using Off-Site Hierarchal Storage Management

David E. Avrin, Katherine P. Andriole, Lloyd Yin, Robert Gould, and Ronald L. Arenson

The purpose of this communication is to report on the testing of the disaster recovery capability of our hierarchical storage management (HSM) system. Disaster recovery implementation is a requirement of every mission-critical information technology project. Picture archiving and communications systems (PACS) certainly falls into this category, even though the counterpart, conventional film archive, has no protection against fire, for example. We have implemented a method for hierarchical storage with wavelet technology that maximizes on-site case storage (using lossy compression), retains bit-preserved image data for legal purposes, provides an off-site backup (lossless bit-preserving wavelet transform), and provides for disaster recovery. Recovery from a natural (earthquake and subsequent fire) or technical (system crash and data loss) disaster was simulated by attempting to restore from the off-site image and database backup to clean core PACS components. The only existing loaded software was the operating system. The database application was reloaded locally, and then the database contents and image store were loaded from the off-site component of the HSM system. The following measurements were analyzed: (1) the ability to recover all data; (2) the integrity of the recovered database and image data; (3) the time to recover the database relative to the number of studies and age of the archive, as well as bandwidth between the local and remote site; and (4) the time to recover image data relative to compression ratio, number of studies, number of images, and time depth of the archive. This HSM system, which maximizes on-site storage, maintains a legal record, and provides off-site backup, also facilitates disaster recovery for a PACS.

PICTURE ARCHIVING and communications system (PACS) is certainly a mission-critical component of health care informatics. Good practice and common sense require disaster recovery implementation for every mission-critical information technology project, including both planning and testing. We have previously described and demonstrated a hierarchical storage system design for PACS that includes an off-site component.[1] The vast majority of commercial PACS deployed today are designed with separate database and image archive components. That is, the data base does not contain actual image data, but rather pointers to the location or filename (unique identifier [uid]) of the image data.[2] This design allows for flexible implementation of image data compression schemes. In addition to the image archive, we have included off-site backup of the actual PACS database. While there are distinctions between off-site backup procedures and disaster recovery, there can be significant overlap. It is interesting to note that there is no backup or disaster recovery for a conventional film archive.

Full disaster recovery in PACS requires availability of computer platforms and network connectivity (wide area network [WAN] and local area network [LAN]), rapid ability to reload system software, database software, custom applications, and database and image data to replacement core components to reconstruct a functioning system with clinical access to demographic and image data. The planning for and testing of disaster recovery using the University of California, San Francisco (UCSF) hierarchal storage management (HSM) system design is described below.

From the Department of Radiology, Laboratory for Radiological Informatics, University of California, San Francisco, CA.

Address reprint requests to David E. Avrin, MD, PhD, UCSF Department of Radiology, 505 Parnassus Ave, Campus Box 0628, San Francisco, CA 94143-0628. E-mail: david.avrin@radiology.ucsf.edu.

0897-1889/00/1302-1040$10.00/0
doi.10.1053/jdim.2000.6864

MATERIALS AND METHODS

The UCSF HSM scheme has been presented and described[1] elsewhere. Briefly, image and demographic data are acquired into a conventional Digital Imaging and Communications in Medicine (DICOM) PACS (Sun [Palo Alto, CA] core components, Sybase [Berkeley, CA] database), with current image data stored on a redundant array of inexpensive disks (RAID). A database (Sybase) is maintained with patient, study, series, and image level tables to support the PACS. Upon image data arrival, it is queued for wavelet encoding (Mr Sid wavelet encoding; LizardTech, Seattle, WA) at a target compression ratio of 2.5, which yields lossless encoding and compression. This encoding process is completed within hours of image data arrival.

The on-site PACS and off-site HSM/Disaster components are connected through the Internet and encrypted with effective bidirectional transfer rates equivalent to 10 Mbit/s Ethernet rates. The full encoded set is sent encrypted over the Internet to an off-site tape jukebox storage system, which maintains its own independent database. A truncated (fraction dependent on modality) portion of the encoded file is kept locally and stored in an magneto-optical disk (MOD) jukebox, for use as the long-term reference image.

The original DICOM data set is used for primary interpretation. The losslessly encoded off-site image data are available for recall in the unlikely event that the radiologist or legal personnel determines that it is required. As has been demonstrated[1] the quality of the lossy encoded image data is extremely high, and difficult to distinguish from the original, even by a trained observer.

Finally, a procedure is executed at 24-hour intervals to back up the actual database to the off-site archive. Total backup is performed once per week, with incremental backups performed daily.

Disaster (earthquake with subsequent fire is our favorite in San Francisco) was simulated for the PACS core components by restoring to clean core components (running with Sun OS only). The database application was reinstalled locally, and then the database contents were loaded from the off-site system. Encoded image data sets were then reloaded, most recent first, from the off-site system.

We measured or assessed: (1) absolute ability to recover; (2) time to recover the database relative to the number of studies and time length of the archive (age of oldest studies); (3) integrity of the recovered database; and (4) time to recover image data relative to compression ratio (related to modality mix), number of studies, number of images, and time depth of archive.

RESULTS

There are six phases for PACS recovery from a disaster: (1) availability of computer platform; (2) if necessary, reloading of system, database, and application software; (3) re-establishment of network connectivity to the remote site (WAN); (4) recovery of the database; (5) recovery of the image data set; and (6) re-establishment of on-site LAN/WAN connectivity.

Physical planning and arrangements are necessary for phase 1. There are cost-benefit tradeoffs that need to be assessed for this item. Phase 1 is not the focus of this report. Similarly, phase 2 requires physical planning, and care to maintain securely available the proper versions of all applicable software.

Phase 3 is problematic, even with the best planning. An off-site repository almost certainly must be accessed by commercial internet service provider (ISP) services, and this infrastructure may suffer in a natural disaster such as an earthquake, but not necessarily in a fire limited to the PACS site.

Recoverability

We were able to recover the system from simulated disaster, with the following exceptions: some image data waiting in the queue to be sent to the off-site component were lost; and the most recent database entries since the last incremental backup and off-site transfer of that backup were lost.

Time to Recover Database

The test database consisted of 62,042 patients who had 198,215 studies (approximately three studies per patient). These examinations contained 3,912,496 images (approximately 20 images per study). The database recovery through fast 100 Mbyte/s Ethernet connection took 2 hours 30 minutes (Table 1).

Integrity of Recovered Database

Some current data were lost, as described above, in the interval between the last incremental backup to off-site facility and the "disaster" event.

Time to Recover Image Data

Test image data set recovery was achieved with losslessly encoded image data at 2.5 compression ratio relative to the raw DICOM data. Table 2 summarizes results of this test scenario. It took 2 hours 27 minutes to recover approximately 9,000 images of 260 studies. This study volume required 2.18 Gbyte of storage for the losslessly (bit-preserving) compressed data, and represented approximately 1 day's volume for our PACS. Therefore, recovery of the actual image data requires about 2.5 hours per day of image data. Restoration occurs backwards from current time, restoring the most recent image data first, since that is of the highest clinical interest, and is subject to the highest access or recall rate by radiologists and clinicians.

CONCLUSION

Disaster recovery planning and testing are essential to PACS implementations of any scale. Assum-

Table 1. Database Recovery

No. of patients: 62,042
No. of studies: 198,215
Average no. of studies per patient: 3.1
No. of images: 3,912,496
Average no. of images per study: 19.7
Time to recover: 2 hours 30 minutes

Table 2. Image Data Recovery

Total studies: 260
Total no. of images: 9,146
Raw file size: 5.31 Gbyte
Compressed (2.5×): 2.18 Gbyte
Recovery time: 2 hours 27 minutes
Average images/study: 35
Average Mbyte/study: raw, 20.42; compressed, 8.38
Average image size: raw, 0.58 Mbyte; compressed, 0.24 Mbyte
Image data movement rate: compressed, 240 Kbyte/s

ing that a suitable hardware platform for recovery is available (through planning), the following issues must be addressed: (1) appropriate version of the operating system available to load; (2) establish remote or (WAN) network connectivity; (3) appropriate version of database system available to load; (4) recover database from off-site backup; and (5) proceed to recover image data (compressed).

Recovery probably be significantly reduced by approximately a factor of four by restoring conservatively "safe" lossy compressed data. Table 3 lists generally accepted safe modality-dependent wavelet encoding compression ratios. Our method will enable this by allowing us to restore truncated files, with an average compression ratio over the modalities of about 10:1. This would enable restoration of about 40 days of image data in a 24-hour period.

Table 3. Proposed "Safe" Lossy Encoded Compression Ratios

Truncated wavelet encoded data, with the truncation ratio modality dependent:	
Computed tomography	10:1
Magnetic resonance	5:1
Computed radiography	3:1

By simulation, we have demonstrated that it is possible to use the off-site lossless wavelet encoded component of the HSM to recover lossy encoded image data of diagnostic quality, at a rate approximating 1.5 days of image data per hour, using secure Internet techniques.

REFERENCES

1. Avrin DE, Andriole KP, Yin L, et al: A hierarchial storage management (HSM) solution for cost-effective on-line archival and disaster recovery using onsite lossy compression and offsite tape backup. Presented at 85th Scientific Assembly and Annual Meeting, Radiological Society of North America, Chicago, IL, November 1999

2. Andriole KP, Avrin DE, Yin L, et al: PACS databases and enrichment of the folder manager concept. J Digit Imaging 13:1-11, 2000