Digital disaster recovery for audiovisual collections: testing the theory

Mick Newnham, Trevor Carter, Greg Moss: National Film & Sound Archive of Australia.

Background

Audiovisual media have recorded the 20th century in way no other era has been recorded previously. Film, audio and video have enabled significant people and events to be witnessed by millions of people and it is hard to imagine a world where moving image and recorded sound did not exist. The problem of preserving this amount of information in the original analogue formats has been monumental and despite the best efforts only a fraction of the original recordings made survive worldwide. The skills required to adequately preserve and make accessible the remaining records have been honed for only the past two decades. And now the world has moved into the digital realm.

The quantity of data created by even a small digital audiovisual collection is massive by any measure. Individual files may be in excess of 10 terabytes[1]! This has created a new set of problems and demanded audiovisual archivists acquire a new set of skills while still requiring the original skills to manage the legacy collections. The costs required to digitise a legacy collection are largely beyond reach of all but the best resourced archive, and yet this is required if a collection is to be preserved and accessible. Consequently hard financial decisions about the way a collection is to be managed into the future need to be made by those responsible. Risk management is a crucial part of the decision making process.

Digital preservation risk management practices make use of a variety of strategies including the obvious legacy solution carried across from analogue collection management such as multiple copies and geographically separated storage. For many small collecting organisations risk management means a second data tape stored in a different part of the building as this is all that may have been budgeted for; and it appears to satisfy the requirements of managing the risk.

However good risk management of a digital collection requires a great deal more than a second copy.

A disaster should be thought of as any incident or event that may potentially prevent permanent access to the record. As such an effective disaster plan must encompass a full regime of assessment, checks and testing. This requires the full support of the entire organisation from the top levels down and can be best managed when disaster mitigation is thought of as equally important to a collecting organisation as access.

Digital collection disaster management is a continuum starting with ensuring the original file is intact to begin with by a thorough quality checking procedure, the plan follows through with strategies for managing the changing environment of files types and hardware evolution and minimising the potential for loss by negligence or malicious attack and onto reducing the impact of catastrophic disasters on a regional scale.

Risks

To determine the best way to approach a disaster recovery is to act proactively and mitigate the impact on the collection as much as is possible.

The first step is a thorough analysis of the risks that may affect the collection, the probability of each and the impact of each type of disaster.

The following information has been collated from published sources on the internet[2][3] and most accurately refers to the experience of the USA and European researchers for a typical business IT situation.

[1] Motion picture films digitised at 4K

[2] Lainbart J., Robinson S., van Zanelhoff M.; "Managing Threats in the Digital Age"; IBM Institute for Business Value

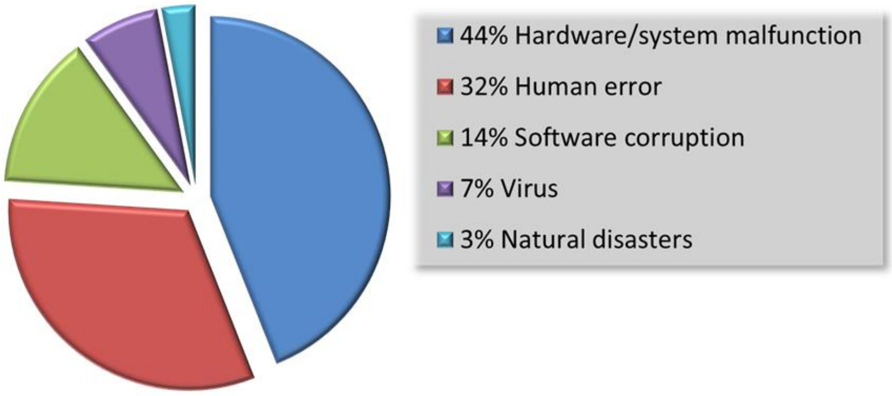

These may be considered as *generic* factors. Audiovisual archives, where much of the digital collection will be as a result of the digitisation of analogue materials to less ubiquitous formats, can add several other risks to be managed.

The original digitisation of the audio, video or film requires close scrutiny to ensure that there has been no corruption of the content introduced when compared to the original source. It is important to distinguish between blemishes/artefacts that may have existed in the original and any that appear on the digital duplicate. Apart from the obvious issues of compression artefacts due to poor or compromised file format choice, degrading noise or dropouts may be introduced by some minor issue within the network or transcoding errors from the raw data to the final file format. Therefore a crucial step in considering a recovery plan is to ensure that any recovery is not attempting to repair existing issues.

The large quantity of information contained in an audio visual collection requires any digital archive to be a comparatively large digital storage system. Systems in the order of 5-10 terabytes would struggle to hold even a moderately sized collection of standard definition video, even using losslessly compressed video files.

Example 1

A 10TB system could store approximately 420 hours of PAL standard definition video losslessly compressed in Motion JPEG2000
Assumptions:

Bit depth	8 bit
Average data rate	55Mb/sec.

Given the rate of change in storage technology a close watch on how new generations and shifts in technology is required to ensure that support can be maintained. The frequency that would be required to maintain technological currency is not fixed but will depend on the global conditions and market. Failure to maintain the collection with formats and equipment that still has technological support raises the serious and highly probable risk that even a slight system malfunction has the potential to cause significant or even total loss of the collection.

The comparatively high costs of migrating a digital collection to a new format/hardware system means that for many collecting organisations there will be the temptation to push the frequency of the migrations to the limit of support. As the end-of-life of the format/hardware approaches there is an increased risk of

[3] “Causes and Cost of Data Loss” www.anysoftwaretools.com, accessed 21 February 2012

being caught out with a sudden change in support if the global market shifts faster than predicted.

Reliability of the power supply is a major factor. If there are interruptions to the supply during crucial file writing stages data will be lost or corrupted. The level of file corruption may be slight, but if it is not detected before the checksums are added further routine checking will not report an error.

Mitigation of risks

The basic principles of organisational disaster planning apply to all collecting institutions regardless of whether the collection is analogue/object based or digital. The overall risks such as fire, flood and security all need to be addressed at the organisational level.

The potential disasters require more than identification. Most disaster guides use a matrix of potential and impact to rank the risk presented by the disaster to create a priority listing of disaster mitigation and recovery strategies. A highly likely factor, for example water entering the building from poor guttering, may only have a very minor impact on the collection, whereas a factor with a low probability, for example a major fire, would have a major impact.

Additional digital collection risks require additional mitigation strategies; however this should be considered part of the overall preservation strategy and incorporated in the workflow.

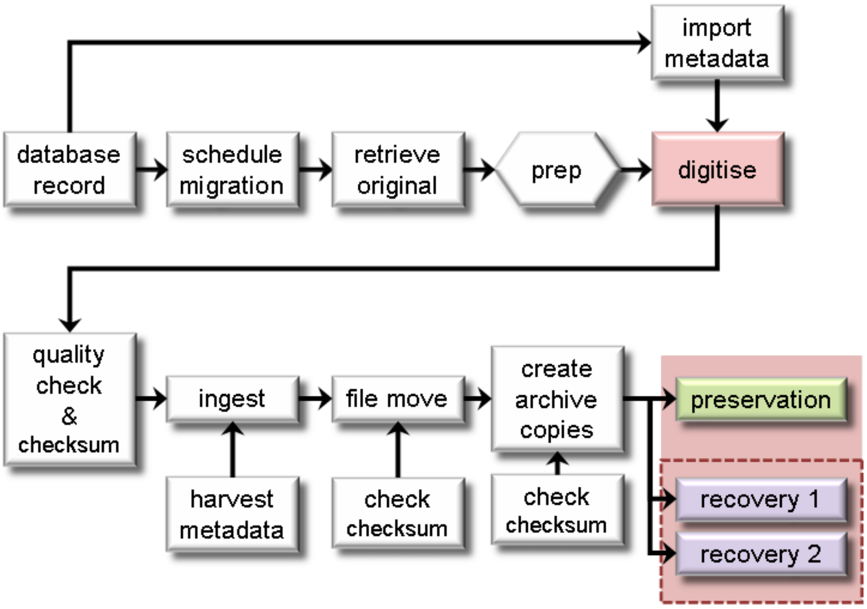

In the figure above, following the workflow from the initial digitisation step; after the digitised content passes the quality check step a *checksum*, also known as a *hash sum*, is added to the file. A checksum is a numerical value calculated in a specific way from the bits in the file. At any point when the accuracy of the file needs to be checked, for example in a migration from one server to another, the checksum value is recalculated and compared with the original. If the figure match then the probability of the file be uncorrupted is very high. The actual value is unique to the file and will change if the file is edited or migrated to a new format.

There are many different algorithms or methods of creating a checksum. The simplest form is a longitudinal parity check where the bits are broken into “words” with a fixed number of bits. The words then are used to calculated an *exclusive or* function for the file.

Case Study:
Using checksums to locate corrupted files

A series of born digital video files from a popular television show totalling 3.2TB arrived on removable media. The connection speed of USB2 meant that the transfer took several days to complete. The process appeared to have proceeded smoothly as no errors were reported by the Cycle Redundancy Check (CRC), a standard method used when copying material on hard drives and networks to insure error free copying. A comparison of the properties as reported by Windows XP showed the exact number of files and reported the same disk space was taken up by both the original and new copy of each file. As material had up until this time arrived in such small numbers no automated quality checking (QC) functions were in use. Having someone watch of each episode would have taken around 100 hours to complete and, as there was no suggestion of problems, was felt to be unnecessary.

To make sure the files codec used in the QuickTime wrapper were supported by the system 5 episodes were picked at random and played back in real time. While watching the 4th of these files an odd artifact was noted toward the end of the file.

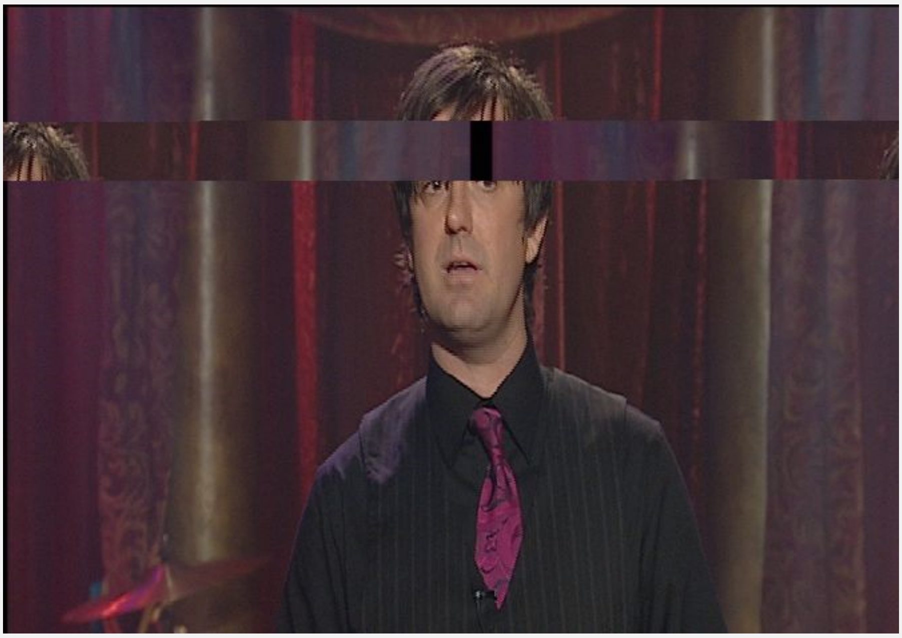

This artifact was just on the one frame and was thought to be a transcode error made at the time of producing the backups. To let the producer's know about this issue the original file was played to compare with the copy. In this process it was identified that the copy had the artifact but the original was fine. The original file was copied again and the new copy was fine. This appeared then to be a random error, even so the likelihood of this occurring again was high. At this point it was decided to use a checksum to confirm that the information was correct.

Once the checksum was applied 7 files were identified as having problems. When these seven files were copied across again 5 were still found to have errors, however the errors were occurring in different location!

In order to achieve no errors the process of copying and validating had to be done a further three times to insure all files had been copied without errors.

Testing showed that it was a faulty USB cable that was causing the problem but it clearly illustrated the point that a checksum was the only way to insure data validity.

An important component of any preservation strategy is to spread the risk across several copies of the content. In the analogue environment this may have meant creating several duplicates of the original tape, disk or film. Each analogue generation would be slightly degraded in technical quality due to inherent losses in the transfer, however this was unavoidable and was accepted. However in the digital domain additional copies may be regarded as identical. The crucial consideration for the additional copies is that each is checked for errors prior to committing the file to the Media Asset Management (MAM) system.

The number of digital copies is not standardised however as a rule of thumb most major organisations create three full resolution copies. In Example 1 given above, while it is feasible to store approximately 420 hours of standard definition PAL video on the 10TB server, this would be as:

- 420 hours of a single copy, with no backup or recovery copies
- 210 hours for the prime and recovery copy, or
- 140 hours for three copies

However, the last two configurations would provide no physical separation between the copies.

At this stage data tape is the most cost effective storage solution for large collections of digital files. While magnetic tape is a well proven technology the formulation may vary from manufacturer to manufacturer. While the frequency of migration means that "life" is not required to be measured in decades as was the case with analogue media[4] it may still be considered a risk to rely on a single manufacturer. Prior to accepting a batch of data tape for use samples should be tested to simulate the conditions under which the tape will be used and stored. While this carries a cost it is not as expensive as trying to recover lost data in the event of a problem.

Each of the three copies needs to be stored discretely. It makes little sense to have all three copies stored on a single tape or a single server. To spread the risk further each copy should be stored in as separate location, the further the geographic separation the better. Ideally each location should be subject to a different set of environmental or physical risks.

Example

Storing copies in different parts of the same building is more effective in reducing risk than storing both copies in a single room, however this practice is less effective in minimising the risk than separately storing the two copies in different buildings. The physical risk is further reduced if the two buildings have a great deal of geographical separation.

same room > different rooms > different building > different city

Reduction of environmental risks →

The decision on the file format chosen to store the data should involve consideration of the support for the format and especially if it is an open source or proprietary format. Proprietary formats may have features that are desirable

[4] The practice of relying on media life has to issues of format obsolescence that have proven to be very difficult to manage, for example 2" video tape is very difficult to duplicate as playback equipment and skilled operators are very few, even globally.

however the life of the format and the support is at the discretion of the format owner. Open source formats are frequently supported by a standard formulated by a community of experts. This standard describes the way the file is played and since the source is open there are more options in recovering the data should a problem occur. An open source format has a community supporting the format's development. Change is fed into the development from a wide range of perspectives and changes voted on by the community. This open communication offers a greater certainty and information on the timetable for migration.

Archives are about information, not only the information contained in the records but also about the records. This later information may be regarding the provenance of the records or technical details about the characteristics of the record itself. The expression *metadata*, or information about information, is used to describe these characteristics. In a digital collection metadata is crucial in managing the collection at every stage. In the past the technical information regarding the physical object was recorded in a database or appended to a catalogue. This practice could be continued for the files in a digital collection, however it would be far more difficult to manage the collection without having the metadata linked directly to the file. The importance of metadata is so great to a digital collection that should the metadata be lost or corrupted managing the collection of any size would be virtually impossible. For this reason the there should be sufficient metadata embedded within the file structure so if the external copies of the metadata are corrupted or lost the file can still be identified and played back. There are two levels at which the metadata may be embedded:

- at the file level, many file formats have sections designed to store metadata. TIFF for still images, BWF for audio and AVI or MOV for video are examples of file formats that have been designed to carry metadata as well as content. MXF is a format designed to wrap coded audio and video along with extensive metadata to create a single file that has the potential to be exchanged between users.
- the second level is by "wrapping" groups of related files together in a single file package, such as TAR or Zip.

As has been mentioned a digital collection requires an active maintenance program. Some the maintenance may be automated within the server system, for example continuously, on-demand or "on-use" verifying of checksums to ensure that no corruption of the file has occurred. A major maintenance event is the migration of the data to new current hardware and formats on a regular or as needed basis to maintain the data in a readily accessible and supported format. The frequency of the migration is open to the changes in technological development, even open formats change. For example the Linear Tape Open (LTO), first released in 2000, is currently on the 5th generation with the 6th generation expected to be released later this year.

Any computer system, either storage or digitisation, will lose data if there are no measures in place to secure a controlled shutdown in the event of loss of power supply. If you are making the one possible replay pass from a degraded carrier and your system goes down before the file is fully written, then you run the risk of losing the original and the duplicate file. Files must be properly written to be viable.

Instantly activated backup supplies or at least an Uninterruptable Power Supply (UPS) system to enable a controlled shut down and prevent loss of data are crucial steps to mitigate this risk. Additionally data storage systems should always operate in a copy - then verify - then delete method, so that a file is never in limbo if a system fails at a crucial point. SAN controllers often have internal battery supplies so that the controller will cache data in the event of a mains supply failure.

Recovery after a disaster
The nature of the disaster will influence the approach taken to the recovery operation. If there has been hardware failure recovery may be as simple as recovering from an error on a RAID system[5], or recovery may involve complex conservation treatments to tapes or hardware and computer forensics. It is down to the successful development and implementation of the preservation strategy as to how complex, and expensive, the disaster recovery operation will be.

There is no "one size fits all" disaster recovery plan as each disaster will have differing direct effects and scale. There is a wealth of information on developing organisational disaster recovery plans available. In general terms the recovery plan should incorporate several distinct stages:

- *Evaluation of the impact of the disaster*
  An information gathering step where the scope and scale of the disaster is assessed and the information feeds into the development of a plan of how to best to proceed and permits a draft budget to be formulated.
- *Salvage*
  More relevant for physical objects that have been affected by a major disaster where the storage environment has been damaged to some extent by water, fire or collapse. For a digital collection this may involve salvaging data tapes or server equipment or locating files resident on a hard drive where the operating system has become corrupted.
- *Stabilisation*
  Again this step is more directly relevant to physical objects. After a disaster chemical and biological factors will affect the stability of the object. Stabilisation in this instance could apply to data tapes or to preventing corrosion affecting hardware.
- *Specific recovery actions*
  These are the specialist skills needed to make the carrier (tape or hard disk) able to be played and recovering the data. The recovered data is then checked and stored on new media or system.

Recovering data from a severely damaged hard disk is a highly specialised skill and falls into the area of forensics. If the drive is still able to be run there is software readily available that may be able to recover files from an otherwise inaccessible disk. However the recovery is a very slow process and many products are designed for consumer level operating systems and drives, not the larger enterprise level installations.

[5] Redundant Array of Independent Disks (RAID) systems use a variety of strategies including total redundancy to store the data across several layers of disks that operate as a single disk. However if a problem should occur with one of the disks in the array the information on that disk can be recovered and the faulty disk replaced.

Digital disaster recovery for audiovisual collections: testing the theory

Mick Newnham, Trevor Carter, Greg Moss: National Film & Sound Archive of Australia.

Case Study:
Full recovery from the failure of a RAID system

Recently a 48 Terabyte fibre channel connected hard disk storage system was installed in a video production and archiving environment. The video work area has three edit suites equipped with four workstations, evenly divided between Apple OS-X and Microsoft Windows XP environments. The disk system is not used as a permanent collection archive but hosts production work in progress and reusable stock footage. It is also used a drop box for work to and from the institutions MAM system

The 24 disk storage system is equipped with enterprise class 2 Terabyte drives and is protected to the RAID 5 level. (that is, one disk failure can be automatically corrected). In the event that a disk fails, the system administrator is alerted and a system rebuild automatically begins to recreate the array without the failed disk. 2 spare disk drives are held locally and may be hot swapped into the system by non-skilled staff.

A single disk drive in the array failed and the faulty disk was replaced with a known good spare. The array rebuild began and was scheduled to take more than 15 hours. *In some instances, rebuilding a disk array may take many days or even weeks depending on the size of the hard disks.* While the rebuild was underway, a second disk drive began to exhibit data errors. As a result, the parity system of the RAID 5 system was unable to simply recreate all the data without errors.

All the media files stored in the storage system appeared to be intact after the initial rebuild, but had random vision and sound errors which were only detectable by real-time replay. *A checksum would have proven this in a very simple manner, but production networks do not usually use checksums because of the overhead of generation and verification*. The system had alerted the administrators to the corruption of the volumes, but could not list the corruptions to file level.

The disk storage is fully supported by its manufacturer, and has a very sophisticated file system, designed to protect data even when multiple drives fail. Over a period of 3 – 4 days, the manufacturer's engineers were able to remotely access the system and to reconstruct 100% of the data without error by using their special and proprietary support tools. New replacement disk drives were provided under the existing warranty agreement and are available again for immediate use if required. The failed disks were shipped back to the manufacturer so that diagnosis of the failures could be undertaken.

Conclusions:
If you are using a disk based storage system for archival data, RAID redundancy measures are essential and can protect against data loss.

The bigger the disk drives in a system, the greater the rebuild time will be, and consequently the higher the probability that another disk drive will fail completely or partially during the rebuild period.

A RAID disk system should have the level of redundancy appropriate to the risk of loss of data for your application. RAID 5 has 1 redundant disk, RAID 6 has 2 redundant disks, etc.

Maintain support contracts with experts in your disk system. They can help when you have failures.

Backup your data if it is valuable and irreplaceable.

In a well planned and executed preservation strategy that has included contingencies for disaster recovery in the event of a disaster, the recovery would hopefully only require the selection of the best surviving copy of the files from the various recovery tapes. In practice this will involve checking each of the files checksums to ensure no corruption has occurred selecting the intact files from the source and reconnecting the damaged or lost media with the database links and metadata to which they belong. In addition to restoring any lost digital media, the collection metadata should also be updated to reflect that the media file had been restored from a new source, in the event that it later turns out that the media is incorrectly identified or corrupted. At worst this may mean ingesting these files onto a rebuilt server system. This action is equivalent in time and effort to a full migration

As the most cost effective storage media, data tape is the most probable media storage for a digital collection to be affected by a disaster. Most data tapes use a traditional binder style tape with magnetic particles (MP tape). Fortunately there has been a body of knowledge developed around the recovery of MP tapes. While it is strongly recommended to gain specific conservation advice before attempting any recovery treatment the basic procedure is to:

1. remove any dirt from the outside of the tape cassette shell
2. check for physical damage to the cassette shell
3. open the cassette shell and clean any loose dirt lodged inside. Check the condition of the tape and the tape path.
4. it may be necessary to use a low relative humidity treatment[6] to dehydrate the binder if the tape has been subjected to water or high relative humidity.
5. reassemble the cassette
6. clean the tape surface, this will require cleaning equipment specific to the format. If the tape is stretched or distorted during cleaning is highly probable that the data will not be able to be recovered.
7. if the tape appears to be in good condition and is thoroughly cleaned then it may be cautiously played on correctly adjusted equipment and the data transferred. There may only be one opportunity to do this depending on the condition of the tape.

Even after the best effort has been made to recover the data the disaster may have caused too much damage for simple playback. At this point it may still be feasible to recover some of the data by reconstructing the data bit by bit.

Conclusions

Prevention is the best option. Manage the risks with a well considered preservation strategy that incorporates effective disaster risk mitigation, rather than risk the costs or total loss that can so easily occur as a result of a disaster.

[6] This is commonly referred to as "baking", however this term is misleading and has led to the total loss of the tape/data due to the temperature being too high. As a rule of thumb the temperature should not exceed 50°C. Some references do suggest temperatures higher than 50°C to "speed" the process up, this is not recommended.