The Significance of Storage in the 'Cost of Risk' of Digital Preservation

Richard Wright*, Matthew Addis, Ant Miller***

* BBC Research and Innovation
Broadcast Centre
201 Wood Lane
London W12 7TP, UK
richard.wright@bbc.co.uk;
ant.miller@bbc.co.uk

**IT Innovation Centre
University of Southampton
2 Venture Road
Southampton SO16 7NP, UK
mja@it-innovation.soton.ac.uk

Abstract

As storage costs drop, storage is becoming the lowest cost in a digital repository – and the biggest risk. We examine current modelling of costs and risks in digital preservation, concentrating on the Total Cost of Risk when using digital storage systems for preserving audiovisual material. We present a managed approach to preservation, and the vital role of storage and show how planning for long-term preservation of data should consider the risks involved in using digital storage technology. Gaps in information necessary for accurate modeling – and planning – are presented. We call for new functionality to support recovery of files with errors, to eliminate the all-or-nothing approach of current IT systems, reduce the impact of failures of digital storage technology and *mitigate against loss* of digital data.

Significance of Storage

As storage costs continue to drop by roughly 50% every 18 months, there are two effects:

- *Storage looks free (but isn't)*: the cost of storage devices becomes negligible, but power, space, cooling, management and replaced costs remain significant.
- *Storage is abundant*: much more storage is used

The following figure shows how hard drive storage has increased over the last 25 years (Hankwang 2008).

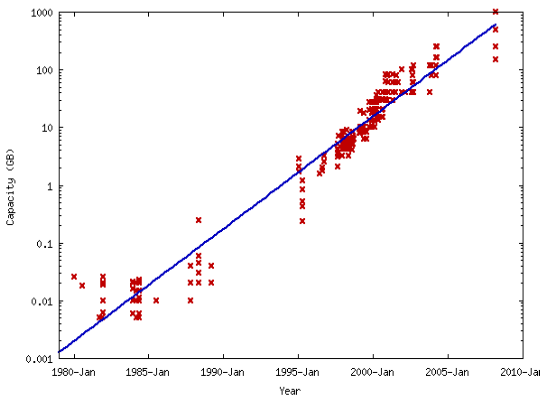

The largest available size (for a desktop computer) has increased from 5 MB to one terabyte – a factor of 200 000 (which is about 18 doublings in about 25 years, so very close to doubling every 18 months).

The 'growth of risk' is of course much larger: a factor of 200 000 in disc size, times the increase in the usage of discs (about 10 000 over the same period; Computer World, 2008).

This "growth of storage" also divides into two effects:

- the number of storage units (globally, and used by any given institution) increases
- the amount of data stored on each unit also increases

The increase in storage units (devices) means that statistics on failure rates that were once seen as 'safe' are now appreciable risks. An advertised Mean Time Between Failure of 1000 years looks very safe to a person buying a new hard drive (though is will be obsolete in 5 years). Schroeder and Gibson (2007) give results on a survey of major datacentres holding 100 000 discs, and found annual failure rates ranging from one to 13 %, averaging around 3% - far higher than an MTBF of 1000 years.

Comment [mja1]: 'Storage is free' is a dangerous statement to make – if storage is free then keeping multiple copies is free and hence there is no cost in reducing risk – you can do LOCKSS for free.

This failure rate means that owners of 1000 of those same hard drives will need systems (eg big RAID6 arrays) and processes (eg continual hot-swapping and rebuilding) to ensure these failures are managed..

The increase in storage units results in more and more users being responsible for, or dependent upon, storage systems that have thousands of individual storage devices (hard drives, optical discs, data tapes). The increase in the amount of data stored on each device makes the failure of each device more significant in terms of the volume of data potentially lost. A 3.5" floppy disc with 1.4 megabytes (MB) of data represented a few dozen files. A 650 MB CD could hold 500 times more data: thousands of files, or one hour of audio. A USB-attached terabyte hard drive is 700 000 times bigger than a floppy, and 1400 times bigger than a CD. It could, for example, hold the entire contents of an institution's audio collection (such as several years' work by many people, collecting oral history recordings).

Cost Modelling

We will present an approach to risk that combines the dimensions of cost, risk (uncertainty) and value (benefits). This model builds upon and extends work on cost modelling by both the digital library and audiovisual communities. Early on in the development of digital libraries there was the fundamental work on preservation strategies by Beagrie and Greenstein (1998), Hendley (1998), Granger, Russell and Weinberger (2000) – and eventually something about the audiovisual sector from EU PRESTO project (Wright, 2002). The state of the art was brought together, and specifically labelled 'life cycle', in the important paper of Shenton (2003).

Since then, there have been entire projects and conferences devoted to *life-cycle models and costs*. At a conference organised by the Digital Preservation Coalition and the Digital Curation Centre (DPC/DCC 2005) there were reports from the LIFE and eSPIDA projects, both specifically about costs, though the eSPIDA work was more generally concerned with a formal method for including intangible benefits (value) in business cases. More pertinent to the present paper, it also specifically introduced the issue of uncertainty into the modelling process.

Specific digital library and digital preservation cost models reported at the 2005 DPC/DCC conference included work from Cornell University, TNA in the UK, and the Koninklijke Bibliotheek in the Netherlands as well as two papers arising from PrestoSpace. In all these models and studies, and for digital library technology in general, little is said about storage (except in the PrestoSpace work). Digital libraries assume that storage will be there (somewhere), and will work and continue to work. In estimating Total Cost of Ownership (TCO), the complexity of the models just mentioned is devoted to digital library processes, not storage devices (or their management). In digital library/repository TCO models, storage cost is generally modelled as a single number per year, and the model simply 'adds up' those numbers.

Cost-of-Risk Modelling

Estimation of cost involves uncertainties. Some uncertainties can be represented as variances in cost estimates (uncertainty about how much costs may vary from the predicted value), but a whole range of uncertainties are related to things that may or may not happen, and should be formally identified as **risks**.

A risk is the likelihood of an incident along with the business consequences (positive or negative) (Addis, 2008a).

Examples of possible incidents include:

- Technical obsolescence, e.g. formats and players
- Hardware failures, e.g. digital storage systems
- Loss of staff, e.g. skilled transfer operators
- Insufficient budget, e.g. digitisation too expensive
- Accidental loss, e.g. human error during QC
- Stakeholder changes, e.g. preservation no longer a priority
- Underestimation of resources or effort
- Fire, flood, meteors …

Traditional risk modeling (and its use in project management) looks at lists of such incidents, and their attendant likelihoods (assessing likelihood may have the largest uncertainty of the whole process!) as contained in a risk register, and then proceeds to predict the consequences – the impact – of each item.

Possible consequences for preservation from the above list of incidents would include:

- Corruption or loss of audiovisual content
- Interruption to services
- Inefficiencies and increased costs
- Corner cutting and increased risks
- Failure to meet legal obligations
- Loss of reputation or loss of customers

A more comprehensive approach to the whole issue of uncertainty in preservation is to include the concept of value (benefit). The work of eSPIDA has already been mentioned.

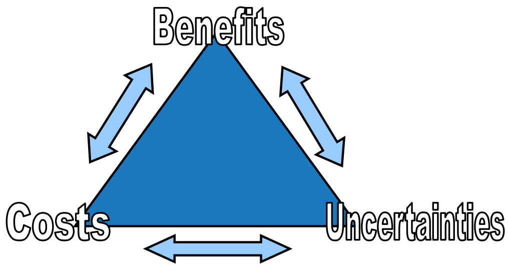

The combination of uncertainty, cost and benefits forms a three-way interaction, as shown in the above diagram. The key point about this approach is that it is as applicable to the whole issue of business-case planning, not just to the more narrow issues of risk analysis and cost modeling.

A typical preservation scenario, which can be optimized by use of the cost-of-risk approach, is given in the following diagramme:

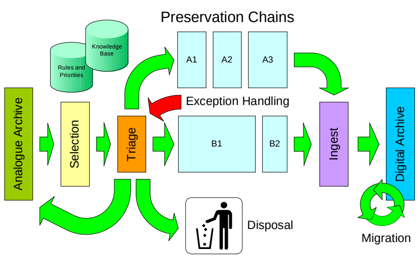

This integrated approach to cost, risk and value allows all the factors affecting preservation planning, funding and management to be considered in one set of interactions, rather than being taken separately.

For quantitative modeling, all three factors need to be converted to a common unit unit of measurement. As cost and benefits are already commonly thought of in financial terms, the task is then to also express the uncertainties in monetary units: the cost-of-risk.

Full details require a much longer presentation. There has already been a great deal of detailed work, specifically relevant to preservation, in the DRAMBORA project (DRAMBORA 2008), and much more detail is in Addis (2008a).

The following diagram shows the consideration of risk as the central metaphor in strategic planning.

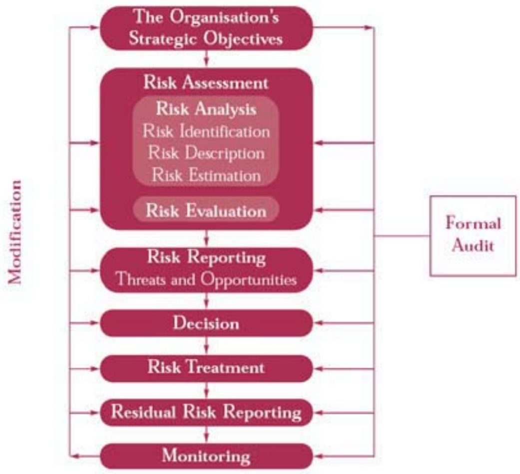

Minimisation of Risk and Cost of Risk – and Mitigation of Loss

The effort within the digital library community to define and construct trusted digital repositories pays little attention to storage. The *trust* issue is defined and examined mainly at the institutional level, not at the level of IT systems and certainly not at the level of individual device or file failures. Yet the *only* physical reality of the content of a trusted digital repository lies in its files, sitting on its storage. The 'atomic level' of success or failure of a repository is the success or failure of an attempt to read individual files. Such success or failure is clearly fundamental to the concept of trust for the whole repository.

Effort of the storage area of the IT industry is entirely focused on reducing the likelihood of read errors (device failure or file read error). There is no concept, within standard IT systems, of a partially-recoverable file. If the inevitable low-level errors cannot be corrected by the built-in error detection and correction technology, the read fails and the file fails to open. There is nothing that the ordinary user can do at this point, and even the all-powerful system manager can only look at backups to see if there is another copy of exactly the same file. There is technology to attempt to read corrupted files or failed hard drives, but such technology falls in the category of *heroic measures*: sending the file or drive to an external company that will attempt a recovery using proprietary technology, at a substantial price (see reference: Recovery Tool Box).

Physically, a file with a read error is not an all-or-nothing situation. There will still be a stream of data (somewhere in the *stack* of operations between the user and the hardware) which is likely to be mainly correct, and is also likely to even have indications of which bytes are incorrect (because of lateral parity errors). For simple error detection and correction schemes, a common situation underlying an inability to read a file is a single block of data that has two or more such errors, so that the longitudinal parity check is ambiguous. At that point, a whole file of many blocks of data is called unreadable, because two bytes – at known locations – fail their parity check and so are known to be erroneous.

Returning to the definition of risk as having two factors: *probability* and *impact*: the ability to read *most* of the data in a corrupted file would, in certain cases, greatly reduce the *impact* of the error. This is the area of risk reduction that is being examined by the UK project AVATAR (Addis et al, 2008b; AVATAR is also looking at the whole issue of optimization and management of storage, from the perspective of archiving and long-term preservation).

Reducing the impact of a storage failure is a method for *mitigation of loss* (Knight, 2007). The issue of loss and recovery from loss has been identified as a neglected area in digital preservation thinking, but its importance has been highlighted by the growing awareness of the phenomenon of bit rot (see reference).

Despite the best efforts of the IT industry, despite mean time between failure of hard drives exceeding one million hours, and despite tests of storage functionality yielding read-error estimations of one failure in 10^{17} read attempts – errors do occur. The author has, in 2008, been personally experiencing one file read failure per month – and in each case these are total failures, with no possibility of mitigation (beyond the commercial route of heroic measures).

Redundancy and Risk

Standard practice for reducing risk of loss is to have another copy. The use of second (or higher) copies is a method of reducing impact: a file read error or a device failure has much less impact if recourse can be made to a backup copy or system.

At a more sophisticated level, arrays of hard drives are used to gain the benefits of redundancy at lower cost. RAID (see reference) technology achieves protection for the loss of one of N drives in a set of N+1 – so the net cost is N+1 drives, rather than the 2N that would be required by simple redundancy.

RAID has now advanced (e.g. RAID6) to the point where multiple disks can fail without data loss, which means data can still be accessed safely whilst individual disks are being replaced and live rebuilding takes place. This allows disk systems to be built that are resilient to hardware failures, human errors and data read errors. For large data centres, the problem is shifted from risk of loss from device failure to having the right support processes to 'feed' large systems with a constant supply of new drives and have the people in place to do so.

At the same time as redundancy is added to storage systems to reduce risk, redundancy is being taken out of the files stored on those systems, as a way to save space. Compression, lossless or lossy, is based on the innate redundancy (entropy) of the original data. When the redundancy is removed from a file, a complex transformation has to be applied to the resulting data in order to transform it back to the original (or close to the original, in the case of lossy compression).

To Encode or Not to Encode

The process of compressing (encoding) a file has profound consequences for attempts to mitigate against loss. A consequence of removal of redundancy is that the remaining data is all very significant – because a compression process is entirely an attempt to eliminate insignificant data. If one byte of the resultant file is then damaged, that byte is then very likely to be used involved in computations (the decoding or decompressing process) that will affect many other bytes. Encoding a file severely affects the ability to use corrupted data as a method of reducing the impact of error.

As an example: an uncompressed audio .WAV file is simply a header followed by a sequence of numbers – one number per sample of the desired audio waveform. If the audio is sampled at 44.1 kHz (the rate used on CDs), each sample represents about 23 micro-seconds of data. Losing one byte of data results in one bad sample, but there is no spread to any of the rest of the data.

Hence an uncompressed audio file can be perfectly usable despite loss of one byte. Indeed, experiments have shown that a .WAV file with 0.4% errors is almost undistinguishable from the original, whereas an MP3 file with the same level of errors either will not open at all, or will have errors affecting most of the audio, and rendering it unusable.

The same logic applies to video, images – and even to text if represented as a sequence of characters (with embedded mark-up, as in the old days of 'printer control characters' as escape sequences within a text 'stream').

An extensive study of the consequences of byte-level errors on different file types, compressed and uncompressed, was recently presented by Heydegger (2008). His results include the following data for image files; in each case exactly one byte had been changed:

- a 10 MB TIFF = .000 01% errors (meaning just that one byte affected)
- a lossless JP2 had 17% errors for a saving of 27% in storage
- a lossy JPEG had 2.1% for a saving of 62% in storage

Comment [mja2]: Is this really 17% for a saving of 27% storage? - that's a pretty crap trade off.

As an example of the affect of data loss on imager files, here are two examples: a BMP (uncompressed) and a GIF (compressed). Each had one byte in 4k changed – meaning 3 bytes total for the GIF, and 12 for the BMP

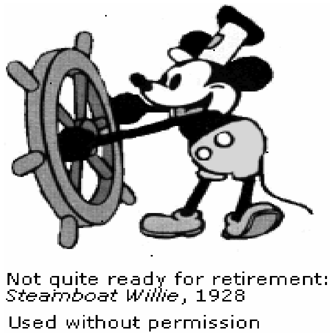

BMP with one error every 4K bytes

GIF file with one error every 4K bytes.

From the above results, it is evident that removing redundancy increases impact, the "cost of error". The compression increases the proportional damage caused by an unrecoverable read error. However if there is no mechanism for using files despite read errors, then it is of no practical significance whether a one-byte error causes major damage, or only very local and very minor damage. If the file can't be read in either case, the error-magnification factor caused by compression is hidden.

If less-than-perfect files can be passed back to the user, or to a file restoration application, then the increase in 'cost of error' caused by compression can be legitimately compared with the decrease in cost of storage.

An unsolved issue in preservation strategy is whether it is better (lower 'cost of risk' for the same or less total risk) to use lossless compression and then make multiple copies (externalized redundancy) as a way to reduce the impact of storage errors – or to avoid compression and exploit the internal redundancy of the files. The problem at present is that there is little or no technology (within conventional storage systems, or conventional digital repositories) to support the second option.

The question of which strategy to take depends on more than just the ability of file systems to return files with partial errors. A holistic approach to risk management means dealing with disaster recovery (fire, flood, theft etc.), human error (accidental corruption, deletion, miscataloguing etc.), and technology obsolescence (formats, software, devices etc.). All present powerful drivers for multiple copies in multiple places using multiple technical solutions. If an offsite copy of uncompressed video is created to address DR, then lossless compression may allow two offsite copies for the same cost. Three copies in three places may well be enough to reduce the risk of loss due to individual storage failures to a level where no further measures are needed beyond those of conventional storage systems, e.g. RAID.

However, until file reading systems are willing and able to return files despite errors, and include media-specific reconstruction techniques to 'fill in' where errors are known to exist, there will be no effective way to exploit file-error recovery as a method to mitigate against loss. This prevents a whole class of 'cost of risk' strategies from being used to complement conventional techniques.

The frustration for audiovisual archivists is that digital technology has taken us one step forward, and now is taking us two steps back. The ability of analogue videotape recorders to cope with loss of data (dropout) was limited, and black lines would appear in the resultant images. Digital tape recorders had much better built-in compensation: the *concealment* option would allow a missing line to be replaces by a neighbouring line, and expensive machines could even replace entire frames with an adjacent (in time) frame. Now file-based digital technology has no ability to cope with loss, beyond the 'external redundancy' option of multiple copies.

One could accept that files are, and will remain, 'all or nothing' entities – you either get everything in them or you lose the lot. The strategy then becomes one of splitting assets, e.g. a video sequence, into multiple files and then implementing safety measures at the 'application' level. For example, an audiovisual program could be split into separate files for shots, scenes, frames, regions of interest, audio, video or many other ways. The most important parts would then be assigned to one or more storage systems with appropriate levels of reliability – avoiding the 'all eggs in one basket' problem. The advantage here is that how to 'split' an asset into pieces can be done based on an understanding of what the asset is – something that a file system or storage device will never have. The downside is increased technology and management costs – a violation of the 'simplest is best' principle.

Comment [mja3]: I deleted this paragraph as I don't think it's true. Device costs (disks or tapes) may become negligible, but TCO for storage (inc. maintainance, upgrade, power, space, people etc.) certainly won't – see intro section.

We hope that current work in preservation theory and methodology, with use of file description metadata , will support and encourage the ability of storage systems to return less-than-perfect files in a usable fashion.

Examples of work with relevance to file description include Planets (file characterization) and Shaman:

- MPEG-21 DIDL = Digital Item Declaration Language (see File Description reference)
- XCEL, XCDL = eXtensible Characterisation Languages (Becker, 2008; Thaller, 2008)
- Shaman = multivalent approach (Watry, 2007)

Conclusions

Comprehensive and integrated planning for preservation can be accomplished through use of a three-factor model, based on costs, benefits and uncertainties. The cost-of-risk concept allows all three factors to be quantified on a common, monetary scale.

Reduction of the cost-of-risk, and the best chance for mitigation of loss, is by **always taking the simplest option** – beginning with not compressing the data.

Storing only uncompressed data would appear to add cost rather than reduce it – but storage costs are typically a small part of a preservation project or strategy (labour is always the dominant cost), and storage cost is dropping by 50% every 18 months.

The full benefit of uncompressed files (in terms of mitigation of loss and consequent reduction of impact) will remain irrelevant unless and until the storage industry and digital repository architects produce systems that allow access to less-than-perfect files.

References

Addis, M (2008a) Cost Models and Business Cases (for Audiovisual Curation and Preservation) presentation at HATII-TAPE Audiovisual Preservation Course, Univ of Edinburgh, May 2008. www.hatii.arts.gla.ac.uk/news/tape.html

Addis, M et al, (2008b) Sustainable Archiving and Storage Management of Audiovisual Digital Assets; to be presented at IBC 2008, September, Amsterdam.

AVATAR-m: http://www.it-innovation.soton.ac.uk/projects/avatar-m/

Beagrie, N & Greenstein, D (1998) A strategic policy framework for creating and preserving digital

collections. British Library Research and Innovation Report 107. London: British Library.
www.ukoln.ac.uk/services/elib/papers/supporting/pdf/framework.pdf

Becker, C et al (2008) A generic XML language for characterising objects to support digital preservation. Proceedings of the 2008 ACM symposium on Applied computing, Fortaleaza, Brazil. pp 402-406
http://portal.acm.org/citation.cfm?id=1363786&jmp=indexterms&coll=GUIDE&dl=GUIDE

Bit Rot: http://en.wikipedia.org/wiki/Bit_rot

Chapman, S (2003) Counting the Costs of Digital Preservation: Is Repository Storage Affordable? Journal of Digital Information, Volume 4 Issue 2 Article No. 178, 2003-05-07

Computer World (2008) Seagate ships one-billionth hard drive.
http://www.computerworld.com/action/article.do?command=viewArticleBasic&taxonomyName=storage&articleId=9079718&taxonomyId=19&intsrc=kc_top

DPC/DCC (2005) Workshop on Cost Models for preserving digital assets
www.dpconline.org/graphics/events/050726workshop.html

DRAMBORA (2008) www.repositoryaudit.eu/

eSPIDA: (costs, benefits and uncertainties in project proposals) www.gla.ac.uk/espida/

File Description methodologies:
DIDL: http://xml.coverpages.org/**mpeg21-didl**.html
XCEL, XCDL: see Becker, 2008 and Thaller, 2008

Granger, Russell and Weinberger (2000) Cost models and cost elements for digital preservation CEDARS project
www.leeds.ac.uk/cedars/colman/costElementsOfDP.doc

Hankwang (2008)
http://en.wikipedia.org/wiki/Image:Hard_drive_capacity_over_time.png

Hendley, T (1998) Comparison of methods and costs of digital preservation. British Library. Research and Innovation Report, 106
www.ukoln.ac.uk/services/elib/papers/tavistock/hendley/hendley.html

Heydegger, V (2008) Analysing the Impact of File Formats on Data Integrity. Proceedings of Archiving 2008, Bern, Switzerland, June 24-27; pp 50-55.

Knight, S (2007) Manager Innovation Centre and Programme Architect National Digital Heritage Archive, NLNZ. Remarks 'from the floor' on the significance of efforts to mitigate against loss, at the SUN PASIG meeting, Nov 2007, Paris.
http://sun-pasig.org/nov07_presentations.html

LIFE: Life Cycle Information for E-Literature
www.life.ac.uk/

Moore's Law: http://en.wikipedia.org/wiki/Moore's_law

Palm, J (2006) The Digital Black Hole:
www.tape-online.net/docs/Palm_Black_Hole.pdf

PrestoSpace: www.prestospace.eu
http://digitalpreservation.ssl.co.uk/index.html

RAID: Redundant Array of Inexpensive Discs – an efficient method of achieving device-level redundancy.
en.wikipedia.org/wiki/Redundant_array_of_independent_disks

Recovery Tool Box http://www.recoverytoolbox.com/
This company is just one of many offering tools that *may* be able to repair a corrupted file.

Risk definition and risk management:
JISC: www.jiscinfonet.ac.uk/InfoKits/risk-management/
PRINCE2: www.ogc.gov.uk/methods_prince_2.asp

Shenton, H (2003) Life Cycle Collection Management LIBER Quarterly, 13(3/4)

Schroeder, Bi and Gibson, G A (2007) Disk failures in the real world: What does an MTTF of 1,000,000 hours mean to you? 5th USENIX Conference on File and Storage Technologies, Feb 2007
www.usenix.org/event/fast07/tech/schroeder/schroeder_html/

Thaller, M (2008) Characterisation (Planets presentation)
http://www.planets-project.eu/docs/presentations/manfred_thaller.pdf

Watry, P (2007) Digital Preservation Theory and Application: Transcontinental Persistent Archives Testbed Activity. International Journal of Digital Curation, 2-2. www.ijdc.net/ijdc/article/viewArticle/43/0

Wright, R. (2002) Broadcast archives: preserving the future. PRESTO project
http://presto.joanneum.ac.at/Public/ICHIM%20PRESTO%2028_05_01.pdf

Wright, R (2007) Annual Report on Preservation Issues
www.prestospace.org/project/deliverables/D22-8.pdf