



Deliverable ID3.2.1

Threats to data integrity from use of large-scale data management environments



Matthew Addis (IT Innovation Centre)

22 February 2010

Document administrative table

Document Identifier	PP_WP3_ID3.2.1_ThreatsMassStorage_R0	Release	0.10
Filename	PP_WP3_ID3.2.1_ThreatsMassStorage_R0_v1.00.doc		
Workpackage and Task(s)	WP3 Data Management WP3T2 – Content Quality Appraisal and Risk Management		
Authors (company)	Matthew Addis (T Innovation Centre)		
Contributors (company)	Christoph Bauer (ORF), Richard Wright (BBC), Jean-Hugues Chenot (INA), Marcel Mattheijer (B&G).		
Internal Reviewers (company)	Richard Wright (BBC), Nir Kashir (ExLibris), Laurent Boch (RAI), Roberto Borgotallo (RAI)		
Date	22/02/2010		
Status	Release		
Type	Part of Deliverable		
Deliverable Nature	Report		
Dissemination Level	Public		
Planned Deliv. Date	31 December 2009		
Actual Deliv. Date	22 February 2010		

This IsPartOf Part of D3.1

This HasPart

Abstract

Maintaining data integrity when using IT infrastructure for the long-term storage of audiovisual files is a major challenge. This report investigates the threats to files from the use of mass storage technologies (e.g. hard drives in servers and data tapes in robots); how can file corruption can be identified; and how the risk of loss be can assessed.

DOCUMENT HISTORY

Release	Date	Reason of change	Status	Distribution
0.1	16 Apr 2009	Outline of file corruption types and detection by ORF	Outline	Confidential
0.2	18 Nov 2009	First Draft of deliverable including structure of sections for risk assessment and problems of using mass storage	Working Draft	Confidential
0.6	18 Dec 2009	Near final draft. Only section 4 and 5 need some more work	Near Complete Draft	Confidential
0.8	7 Jan 2010	Finished section 4,5. Added content to section on HDD. Finished introductions	Complete Draft	Confidential
0.9	26 Jan 2010	Further edits in response to the four internal reviews	Final Draft	Confidential
0.10	10 Feb 2010	Final version with typo corrections and other format fixes	Final	Confidential
1.00	22 Feb 2010	Finalised - delivered	Release	Public

Table of contents

Scope.....	5
Executive summary.....	6
1. Assessing the risks to data integrity when preserving and accessing AV files	11
1.1. Information Security.....	11
Security standards relevant to data integrity.....	11
1.2. Digital Preservation.....	13
1.3. Preservation, access and 'active archives'.....	13
1.4. Risk Management.....	14
1.5. DRAMBORA.....	16
1.6. OCTAVE Allegro.....	21
1.7. Combined use of DRAMBORA and OCTAVE Allegro.....	23
2. Example Risk Analysis focusing on Data Integrity.....	27
2.1. Risks of loss of data authenticity and integrity.....	28
2.2. Risks of data destruction or degradation.....	31
2.3. Risks to data through loss of services.....	32
2.4. Risks to loss of data integrity through mismatch of expectations.....	35
2.5. Summary.....	37
3. Failure modes and data corruption from mass storage technology.....	37
3.1. Overview.....	37
3.2. Latent Faults and Visible Faults.....	38
3.3. Mean Time Between Failures (MTBF).....	40
3.4. Storage capacity: trends and the implications.....	42
Hard Disk Drives (HDD).....	44
HDD throughput and error rates: trends and the implications.....	48
3.5. Data Tape.....	50
3.6. Use of AV compression.....	52
3.7. IT methods for data integrity checking.....	55
4. Errors for video, audio and images and their detection.....	58
4.1. Examples of data-integrity violation.....	59
Example 1: Artefacts in Video-Files.....	59
Example 2: Artefacts in Audio-Files (ORF).....	62
Example 2: Artefacts in Picture-Files (ORF).....	72
4.2 Summary.....	80
5. File Quality-control in current use.....	82
Annex 1: OCTAVE Allegro worksheets.....	84
Annex 2: OCTAVE Allegro Risk Analysis Example.....	91
Annex 3: Quality Control case studies.....	97
BBC	97
Final Quality Checks (Technical Review) at BBC-Scotland.....	97
Spot Checking of Audio-Productions.....	98
Identification of critical areas by using D3-replay logs.....	98
ORF	98
Specifications for QC in P-CMS and on Storage-entry.....	98
B&G.....	101
INA.....	101

Scope

The European Commission supported PrestoPRIME project (www.PrestoPRIME.org) is researching and developing practical solutions for the long-term preservation of digital media objects, programmes and collections, and finding ways to increase access by integrating the media archives with European on-line digital libraries in a digital preservation framework. This result will be a range of tools and services, delivered through a networked Competence Centre.

Maintaining data integrity when using IT infrastructure for the long-term storage of audiovisual files is a major challenge. This report investigates the threats to files from the use of mass storage technologies (e.g. hard drives in servers and data tapes in robots); how can file corruption can be identified; and how the risk of loss be can assessed.

The report addresses the following questions:

- How can risk assessment methodologies be applied to the risk of data loss, in particular when using mass storage for digital preservation? What are the range of risks that exist and how can they be categorised? How can existing efforts from the digital preservation community (e.g. DRAMBORA and TRAC) and the information security community (e.g. OCTAVE from CERT) be combined and used together?
- What evidence is there that mass storage technology is not a 'safe' solution for digital preservation of audiovisual content, in particular at the scale of Europe's AV archives? What types of corruption of failures occurs, how likely are they, and what measures exist to reduce them?
- How can file corruption be identified? What techniques are currently used by AV archives? What products are available in the market place? What are the necessary Quality Control (QC) processes and at what points does quality control need to be applied? How should QC be done at the syntactical level (wrapper, codec, video, audio), semantic level (video, audio), display of results, performance, and the like.

The conclusions of this report are clear:

- Mass storage technology from the IT Industry simply doesn't have the levels of reliability needed for long-term preservation of large audiovisual data files. Ways in which loss can occur are manifold, hard to predict, and most worrying can take place silently, even in storage systems explicitly designed to prevent data loss.
- A meaningful strategy for assessing the threats to data preservation from the use of IT storage technology has to consider the risk of loss, the cost of mitigating this risk, and the benefits of doing so. We call this a 'cost of risk of loss' approach.
- Maintaining integrity of digital audiovisual assets is a proactive activity and has to be supported by appropriate corruption detection tools, a quality control process, and a knowledge base of what can go wrong, how likely this is, and what to do about it.

Executive summary

There are a wide range of threats to long term data preservation when using IT systems, in particular mass storage technology. Risk assessment methodologies from both digital preservation and information security communities provide a structured approach to identifying and assessing these risks. In this report, we focus on AV essence and the long term safety of this essence by considering risks to essence in four main areas.

- **Risks of loss of data authenticity and integrity.** These risks are mostly concerned with the loss of ability to track and record the origins of data and then everything that is done to data during digital preservation. Without this provenance trail, there is the risk that changes to integrity or authenticity happen but go unnoticed.
- **Risks of data destruction or degradation.** These risks are concerned with the loss or corruption of data, for example from imperfect storage technology, deliberate or accidental damage, or loss of access to data due to technical obsolescence.
- **Risks to data through loss of services.** If there is a loss or interruption to the services and processes that are involved in preservation or access to digital content, then this has the potential to put the content itself at risk of loss. For example, this might be the loss of a service that routinely checks and maintains data integrity in a storage system.
- **Risks to loss of data integrity through mismatch of expectations.** If preservation is provided as a service, e.g. within an organisation or by a third-party then there the potential for a mismatch in expectations or understanding between the providers of the service and the community for which the services are being provided. If the changes in expectations are too rapid, or not communicated properly, then data can be put at risk. For example, the required level of data integrity might not be properly defined, or the sudden need for higher levels of integrity might be beyond the capabilities of current systems

Data corruption and failure modes exist for all common mass storage technologies, e.g. hard drives and data tape. There is currently a lack of awareness in the AV archive community that commodity IT technology is not a safe option for long-term storage – at least not without care and proactive data integrity management.

- When faults develop in mass storage systems they are not always immediately detected by these systems, let alone automatically corrected. These so called ‘latent’ faults happen for a variety of reasons and are not uncommon, particularly in hard disk based storage systems. Without careful consideration they can easily lead to silent and irrecoverable data loss – this is commonly called ‘bit rot’
- There are many failure modes in all levels of storage systems (hardware, firmware, software) that can result in data loss. The reliability of the underlying storage media (disk drive or data tape) is not necessarily the dominant factor.

- Proper analysis and planning of data safety in digital storage systems, be it hard disk servers or data tapes on shelves, has to consider not only the types and frequencies of faults, but most importantly how fast those faults can be detected and then repaired. These rates dictate the overall likelihood of data loss.
- Anyone seeking to achieve long term preservation of data using mass storage technology needs to be aware of latent faults, e.g. bit rot, their consequences and to counter them by proactive data integrity monitoring and management.
- The archive community needs to collect and share information on their experience with media and system reliability, especially in the first few years of use. This is the only way to generate meaningful statistics on the failure modes and their frequency.
- The trend towards ever higher capacity of media (tapes, hard disks) for the same cost is very attractive to archives for obvious reasons. However, increase in reliability is not keeping pace with this increase in capacity. Furthermore, data access rates are also not keeping pace with capacity, which impacts the rate at which data corruption can be detected and repaired. The result is that the cost of storage is going down, but the risk of loss when using that storage is going up.
- Strategies for data distribution across storage media or systems need to evolve to ensure that the probability of data loss is kept within acceptable limits. Simple strategies that work today, e.g. direct data tape replication to create two copies may not be so applicable in the next decade or beyond.

There is particular interest in both Hard Disk Drive (HDD) and data tape as preservation storage technologies. Our findings and recommendations here are clear:

- Modern data tape, e.g. LTO, is relatively reliable with problems tending to come from drives rather than tapes. Two or more tape copies are still needed for safety.
 - Regular migration is essential. Whilst media lifetime of data tape can exceed 15-30 years in good conditions, tape drives become obsolete in 5-7 years and new drives have limited backwards compatibility with older data tape generations.
 - Full details of tape reliability are still unclear. There is a need for the AV archive community to share statistics on reliability 'in the field' of data tape, including during migration as well as in operational systems. In particular, information on latent faults is lacking.
 - Lifetime (head life, media read/write cycles etc.) is a more important as a metric than MTBF. Lifetimes should be respected and a conservative approach can pay dividends, especially when a tape is used frequently.
- Hard disk drives as a preservation storage media should not be used without great care and extra systems to ensure long-term content safety and integrity.

- Annual Failure Rates (AFR) between 1 and 10% are common in the first few years of life. HDD should only be used within mass storage systems that manage data integrity and component failures.
- Data corruption also occurs in hard disk based storage systems, including those explicitly designed to reduce data loss. Most worryingly this corruption can be completely silent and hence go undetected for long periods of time.
- Long-term data integrity requires an ongoing and proactive programme of data integrity checking and repair at an end-to-end systems level. No component of the system (networks, storage, memory, processing) should be assumed to be somehow 'safe', i.e. immune from failures and data corruption problems.

Data corruption does happen when using IT storage with catastrophic effects on AV files.

- Data corruption causes major problems for AV content in file format. This applies to video, images and audio alike. In the worst case, which is not infrequent, files simply can't be opened or played by their respective applications.
- Considerable further work is needed to investigate how data corruption of preservation video formats translates into visible artefacts.
- The use of encoding, in particular compression, can massively amplify even low levels of data corruption, or result in whole files becoming useless.
 - For single images (and intra-frame encoded video) a single byte of corruption to a compressed image can render the whole image completely useless. The sensitivity to data corruption is not correlated to the level of compression, e.g. lossless JPEG2000 is just as sensitive to data corruption as lossy compression.
 - For corruption of intra-frame encoded video there is at least a possibility to use concealment techniques, e.g. interpolation between adjacent frames, to correct the effects of data corruption – provided that the number of frames affected in a sequence is low. The same is unlikely to be true for inter-frame encoded video due to the temporal propagation of errors between frames.
 - For audio, major audible artefacts can be generated and persist well beyond the temporal location where data loss first takes place. Files compressed with a variable bit-rate are most vulnerable. Files compressed by constant bitrates are more stable and robust to corruption. PCM-files are extremely vulnerable to data loss, resulting in unusable content in almost all cases.
 - More investigation is needed for both audio and video formats, in particular video that uses inter-frame encoding. The expectation is that all compressed formats in common use are not likely to be at all robust to data corruption, even at low levels of data corruption.

- Compressed formats are in general much more sensitive to data corruption than uncompressed formats. Due to the 'amplification' effect that compression has on data corruption, the percentage saving in storage space is often much less than the percentage increase in the amount of information that is affected by data corruption.

Given the threat to AV content when using mass storage technology, proactive data integrity management is an essential activity.

- Checksums provide a fast and simple way to monitor data integrity in mass storage systems. Simple algorithms such as Aldler32 and CRC32 are sufficient in scenarios where accidental corruption needs to be detected. These can be computed at very high speed (hundreds of MB per sec) using modest commodity PC hardware and hence have little overhead.
- The serious consequences of data corruption to AV content, especially when compressed means that regular integrity checking should be considered the norm. The demonstrable existence of silent data corruption in mass storage systems (bit rot), especially those systems based on hard drives, and including systems explicitly designed to prevent loss, means that no IT system should ever be considered 100% safe and independent regular integrity checking is necessary.

It would be natural to expect proactive data integrity checking and repair to be an integral and well developed part of archive operations. However, when looking at and comparing the content QC processes of the broadcast and archive partners in the PrestoPRIME consortium, there is not a common set of tools, processes or techniques in place.

- Each archive surveyed has its own approach to QC, with some being much more developed and automated than others.
- Many of the QC processes that are in place focus on ensuring the quality and standards of content admitted into the archive, for example identification and checking of file formats (wrappers, metadata, video, audio) against standards at the syntactic level and also checking content (often manually) for visual or audible quality problems (e.g. during digitisation, transfer or format migration).
- Less attention is paid to proactively monitoring data integrity after data has entered the archive. This is partly because many archives still operate an 'items on shelves' model and the bulk of their content is not yet in digital file form.

This is changing and several of the archives surveyed recognise the need to review and further develop their QC processes in the area of proactive data integrity management.

However, given that the national broadcasters and archives in PrestoPRIME are at the vanguard of new techniques and best practice for digital preservation of AV, yet appear not to have well developed data integrity management processes in place, then this all suggests that there is a general lack of awareness of the problems that archives face in this area. This leads to our final recommendation.

- There is a lack of awareness in the AV archive community of the threats to data integrity that come from the use of mass storage technology. PrestoPRIME needs

to take action to promote awareness of this issue and follow this up with practical guidelines, tools and solutions.

1. Assessing the risks to data integrity when preserving and accessing AV files

In this section of the report we:

- (a) Examine existing risk assessment techniques from both the digital preservation and information security communities
- (b) Show how they can be combined and used for a methodical and holistic approach to assessing all the ways in which data integrity might be lost
- (c) Demonstrate this approach by looking in more detail the risk of integrity loss from mass storage systems

1.1. Information Security

A useful definition of the term information security is provided in the United States Code ¹. Information security is protecting information and information systems from unauthorized access, use, disclosure, disruption, modification, or destruction in order to provide:

- **integrity**, which means guarding against improper information modification or destruction, and includes ensuring information non-repudiation and authenticity
- **confidentiality**, which means preserving authorized restrictions on access and disclosure, including means for protecting personal privacy and proprietary information
- **availability**, which means ensuring timely and reliable access to and use of information.

Clearly all of these aspects of information security are applicable to archiving of audiovisual assets.

Security standards relevant to data integrity

Security standards offer guidelines and general principles for information security management within an organisation. We present here two ISO standards that show the different phases for the security policy definition process. Information Security policy, in the broad sense of the term, refers to “the set of laws, rules and practices that regulate how an organisation manages, protects, and distributes sensitive information”².

¹http://www.law.cornell.edu/uscode/html/uscode44/usc_sec_44_00003542----000-.html

²

ITSEC, "Information Technology Security Evaluation Criteria (ITSEC): Provisional Harmonised Criteria, Document COM(90) 314, June 1991 http://www.ssi.gouv.fr/site_documents/ITSEC/ITSEC-uk.pdf

The ISO 27001³ standard, titled "Information Security Management - Specification With Guidance for Use", is considered as the specification of best practice for an Information Security Management System (ISMS).

The objective of the standard itself is to provide a model for establishing, implementing, operating, monitoring, reviewing, maintaining, and improving an Information Security Management System. This model ensures that the design and implementation of an organisation's ISMS is influenced by their needs and objectives, security requirements, the process employed and the size and structure of the organisation.

The standard defines its 'process approach' as "The application of a system of processes within an organisation, together with the identification and interactions of these processes, and their management". It employs the "Plan-Do-Check-Act (PDCA)" model to structure the processes.

The ISO 27002⁴ standard, formerly known as ISO 17799, is a code of practice for information security. It outlines hundreds of potential controls and control mechanisms, which may be implemented, in theory, subject to the guidance provided within ISO 27001.

This standard established guidelines and general principles for initiating, implementing, maintaining, and improving information security management within an organisation. The actual controls listed in the standard are intended to address the specific requirements identified via a formal risk assessment. The standard is also intended to provide a guide for the development of organisational security standards and effective security management practices and to help build confidence in inter-organisational activities.

The different chapters of the standard cover:

- Security Policy: addresses the guidelines and procedures that reflect the ongoing commitment of the organisation concerning information security.
- Organisation of Information Security: addresses the need for a management framework that creates, sustains, and manages the security infrastructure. It also covers external party use of local information.
- Asset Management: inventory and classification of information assets.
- Human Resources Security: security aspects for employees joining, moving and leaving an organisation.
- Physical Security protection of the computer facilities.
- Communications and Ops Management: management of technical security controls in systems and networks.
- Access Control: restriction of access rights to networks, systems, applications, functions and data.

³

ISO/IEC 27001:2005, Information technology - Security techniques - Information security management systems – Requirements. http://www.iso.org/iso/catalogue_detail?csnumber=42103

⁴

ISO/IEC 27002:2005, Information technology -- Security techniques -- Code of practice for information security management. http://www.iso.org/iso/iso_catalogue/catalogue_tc/catalogue_detail.htm?csnumber=50297

- Information Systems Acquisition, Development, Maintenance: building security into applications.
- Information Security Incident management: anticipating and responding appropriately to information security breaches.
- Business Continuity: protecting, maintaining and recovering business-critical processes and systems.
- Compliance: ensuring conformance with information security policies, standards, laws and regulations.

It should be clear from the above list that Information Security standards should be an integral part of assessing the degree to which an archive is able to ensure data integrity when using file-based systems to store and access AV assets.

1.2. Digital Preservation

Preservation of audiovisual material is defined by The Coordinating Council of Audiovisual Archive Associations (CCAAA) as the totality of the steps necessary to ensure the permanent accessibility – forever - of an audiovisual document with the maximum integrity⁵. There is a clear overlap between the objectives of information security and the objectives of preservation, particularly in the area of integrity and accessibility of content within an archive.

For many audiovisual archives the ability to access and use content is fundamental to the purpose of the archive (as opposed to archives that are more oriented to keeping content for compliance reasons with no expectation of needing to reuse the content in the archive).

The ability to access content, i.e. having confidence in being able to maintain availability of an archive, means having a wide range of processes and systems in place. These include dealing with disasters (fire, flood, theft etc.), human error (accidental corruption, deletion, mis-cataloguing etc.), and technology obsolescence (formats, software, devices etc.). These areas are all part of a digital preservation strategy and hence the digital preservation community in addressing these challenges has a lot to offer AV archives when seeking to maintain data integrity.

1.3. Preservation, access and ‘active archives’

Archiving of audiovisual content is rapidly becoming an integral part of content production, distribution and consumption processes. Archiving no longer sits at the tail end of the content lifecycle as a place where content ‘ends-up’ for ‘safe keeping’. There is a business need for continual and ‘online’ access to archive contents, both for reuse within an organisation, e.g. a national broadcaster, and also for public access or commercial use by people outside of the organisation.

The increased level of integration of archive systems with wider content production and distribution systems necessitates accompanying security integration to ensure secure and seamless exchange of content between these systems. This has an important role to play

⁵

http://www.ccaaa.org/ccaaa_heritage.pdf

in data integrity as it safeguards against unauthorised access and hence possible modification of AV content.

This business need for access to archive content that also brings with it a potential conflict with an archive's mission to maintain the integrity of that content to the highest degree possible. This extends to the use of service oriented models, including the delivery of archive hosting through third-party services.

When resources are limited, as they almost invariably are in archiving, and need to be shared between preservation actions and the services that provide access, there is the opportunity for one of these functions to suffer at the expense of the other.

Both digital preservation and Information security are about protecting assets, and these assets include systems as well as data. Techniques from both communities clearly have a role to play in assessing the threats that access to content place on the ability to maintain data integrity.

1.4. Risk Management

Risk management takes a structured approach to the identification, analysis and management of risks to protect an organisation's critical assets and processes in a way that is commensurate with the organisation's risk appetite and risk tolerances.

Before going further it is worth defining some terminology:

- Risk is the probability and impact of something happening. Risks can be positive or negative – risk isn't just about bad things happening.
- Risk tolerance is the extent to which an organisation is prepared to accept or even seek out risk.
- Risk management is process of assessing and dealing with risk and typically involves the selection and application of one or more treatments for risk.
- Risk treatments can generally be put into one of four classes: accept the risk, remove the risk, transfer the risk to someone else, or reduce the risk.
- Assets are the subject of risk management and are anything of value to the organisation. Assets include digital content (obviously!), but also services (e.g. archive services provided by a third party) and less tangible things such as reputation (e.g. recognised as providing safe keeping of national heritage).
- Impact (outcomes) of risk can generally be put into the following classes: destruction (direct loss of the asset); disclosure (loss of confidentiality); modification (loss of integrity); or interruption (loss of availability).
- Risks come from various sources (threats) and these include actors both internal and external (e.g. content owners, archive providers, storage providers, staff, competitors, hackers) and systems (e.g. bugs in hardware or software) as well as other sources that may be completely outside of an organisations control (e.g. fire, flood or earthquakes).

Risk management is a cyclic activity⁶ that involves identifying an organisation’s objectives, assessing the risks, and then treating and monitoring these risks. It involves understanding what needs to be protected, why it needs to be protected, what happens if it is not protected, what potential consequences need to be prevented and at what cost.

The various steps in the risk management process are shown in Figure 1. Information security and long-term⁷ digital preservation are both (related) risk management problems. Risk management techniques are used in both domains as a framework for analysing risk and making informed choices on the balance between cost, safety and security.



Figure 1 Risk management process (reproduced from The Risk Management Standard from The Institute of Risk Management (IRM),The Association of Insurance and Risk Managers (AIRMIC) and ALARM The National Forum for Risk Management in the Public Sector.

Many things threaten the security or integrity of digital AV assets over long periods of time. There are not just technological risks, but risks that arise from human, financial or legal factors, e.g. the loss of staff or skills, insufficient resources, changing needs of stakeholders, accidental loss, deliberate circumvention of protection measures and many more. More than the digital assets are at risk: the reputation of individuals or organisations, the ability to maintain and deliver services, and regulatory compliance are all potentially at stake.

⁶

For example, see This Risk Management Standard, which is the result of work by a team drawn from the major risk management organisations in the UK - The Institute of Risk Management (IRM),The Association of Insurance and Risk Managers (AIRMIC) and ALARM The National Forum for Risk Management in the Public Sector. http://www.theirm.org/publications/documents/Risk_Management_Standard_030820.pdf

⁷

OAIS defines long term as: Long Term: A period of time long enough for there to be concern about the impacts of changing technologies, including support for new media and data formats, and of a changing user community, on the information being held in a repository. This period extends into the indefinite future.

Risk management has been long been used in 'mission critical' domains, and information security, e.g. as advocated by CERT⁸, and is now emerging in the digital preservation domain. DRAMBORA (Digital Repository Audit Method Based on Risk Assessment) provides an approach to auditing digital repositories⁹ and has been developed through a partnership between the UK Digital Curation Centre¹⁰ and the EU DPE¹¹ project.

The TRAC¹² (Trustworthy Repositories Audit & Certification) Criteria and Checklist has been developed by the Research Libraries Group (RLG)¹³ and the National Archives and Records Administration (NARA)¹⁴ through a joint task force to specifically address digital repository certification and sets criteria to identify digital repositories capable of reliably storing, migrating, and providing access to digital collections. This work is currently being used by NARA and the EU SHAMAN projects as part of the development of preservation testbeds¹⁵. The EC Digital Preservation Europe (DPE) project has very recently released its PLATTER¹⁶ tool for planning trusted digital repositories that builds on both TRAC and DRAMBORA.

Lead by the CCSDS (the same body that gave us OAIS) another working group is currently advancing the establishment of an ISO standard on which a full audit and certification of digital repositories can be based. The venture aims to combine the efforts of TRAC, DRAMBORA, Nestor¹⁷ and ISO/IEC 27001:2005¹⁸ and to standardise the results in the same way as the OAIS Reference Model (ISO 14721)¹⁹. The working group is expected to publish two documents, one containing metrics to audit and certificate digital repositories and one specifying the requirements for the bodies that actually provide these audit and certification. Consequently the consistency of the audit and certification process is ensured additionally by investigating the expertise and qualification of the auditors as well.

1.5. DRAMBORA

DRAMBORA provides a self audit toolkit to facilitate the auditor of a digital repository in:

- Defining the mandate and scope of functions of the repository;
- Identifying the activities and assets of the repository;

8

CERT. http://www.cert.org/work/organizational_security.html

⁹ Drambora interactive : <http://www.repositoryaudit.eu/>

¹⁰ UK Digital Curation Centre (DCC) The DCC provides a national focus for research and development into curation issues and to promote expertise and good practice, both national and international, for the management of all research outputs in digital format. <http://www.dcc.ac.uk/>

¹¹ <http://www.digitalpreservationeurope.eu/>

¹² <http://www.crl.edu/PDF/trac.pdf>

¹³ The RLG is now part of the OCLC see: <http://www.oclc.org/programs/default.htm>

¹⁴ <http://www.archives.gov/>

¹⁵ The International Journal of Digital Curation. Issue 2, Volume 2 | 2007. Digital Preservation Theory and Application: Transcontinental Persistent Archives Testbed Activity Paul Watry, University of Liverpool November 2007. <http://www.ijdc.net/ijdc/article/view/43/50>

¹⁶ DPE deliverable D3.2 Repository Planning Checklist and Guidance. http://www.digitalpreservationeurope.eu/publications/reports/Repository_Planning_Checklist_and_Guidance.pdf

¹⁷ Nestor, <http://edoc.hu-berlin.de/series/nestor-materialien/8en/PDF/8en.pdf>

¹⁸ International Standards Organization, <http://www.iso.org>

¹⁹ OAIS Blue Book, CCSDS 650.0-B-1, <http://public.ccsds.org/publications/archive/650x0b1.pdf>

- Identifying the risks and vulnerabilities associated with the mandate, activities and assets;
- Assessing and calculating the risks;
- Defining risk management measures;
- Reporting on the self audit.

The self audit toolkit is designed to help and guide the auditor along a similar route of analysis to that which an external auditor would use to examine and analyse the work of the repository. The process followed is shown in Figure 2.

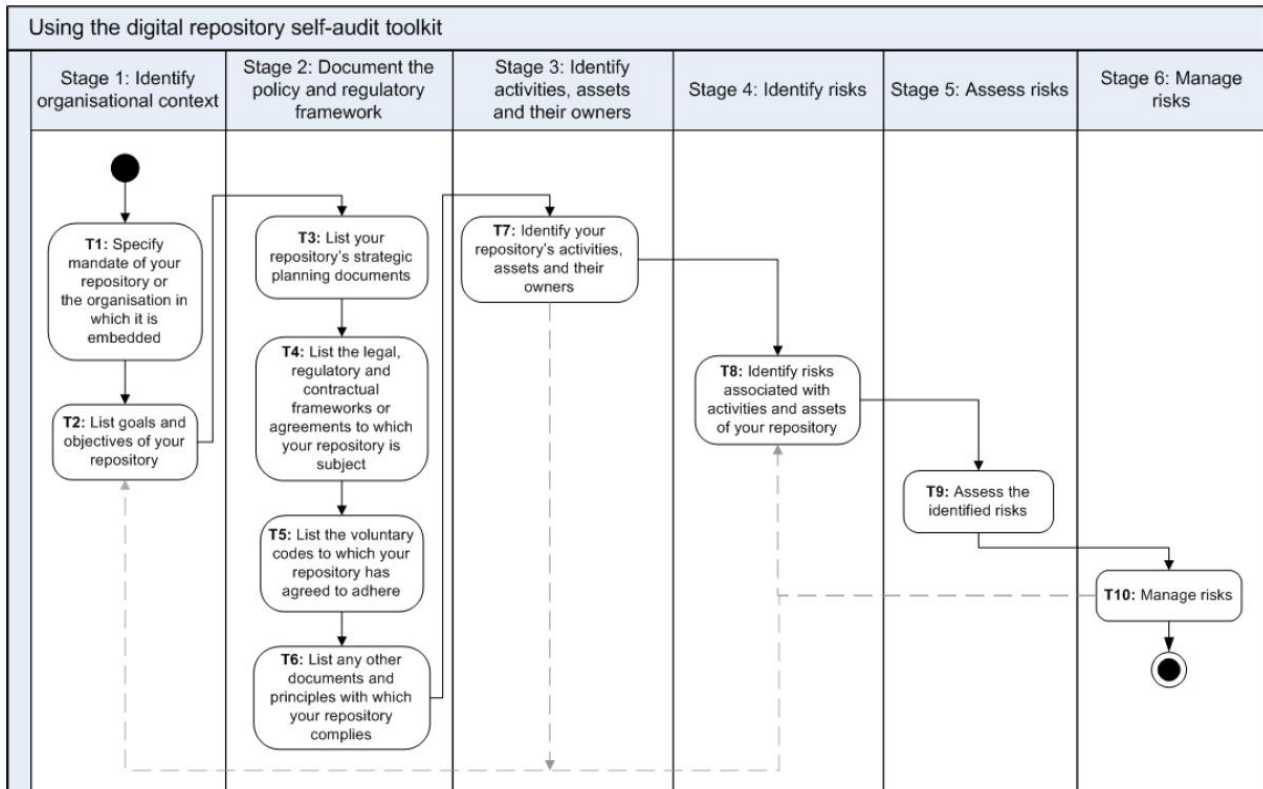


Figure 2 DRAMBORA audit process (reproduced from DRAMBORA v1.0)

The bulk of this process involves building up a register of risks, which includes who within the organisation is responsible for dealing with them and what mitigation approaches can be taken. An example is shown in Figure 5. Each risk has a numerical score for probability and impact to allow risks to be quantified and prioritised as shown in Figure 3 and Figure 4.

<table border="1"> <thead> <tr> <th>Risk Probability Score</th> <th>Interpretation</th> </tr> </thead> <tbody> <tr> <td>1</td> <td>Minimal probability, occurs once every 100 years or more</td> </tr> <tr> <td>2</td> <td>Very low probability, occurs once every 10 years</td> </tr> <tr> <td>3</td> <td>Low probability, occurs once every 5 years</td> </tr> <tr> <td>4</td> <td>Medium probability, occurs once every year</td> </tr> <tr> <td>5</td> <td>High probability, occurs once every month</td> </tr> <tr> <td>6</td> <td>Very high probability, occurs more than once every month</td> </tr> </tbody> </table>	Risk Probability Score	Interpretation	1	Minimal probability, occurs once every 100 years or more	2	Very low probability, occurs once every 10 years	3	Low probability, occurs once every 5 years	4	Medium probability, occurs once every year	5	High probability, occurs once every month	6	Very high probability, occurs more than once every month	<table border="1"> <thead> <tr> <th>Risk Impact Score</th> <th>Interpretation</th> </tr> </thead> <tbody> <tr> <td>1</td> <td>Zero impact, results in zero loss of digital object authenticity and understandability</td> </tr> <tr> <td>2</td> <td>Negligible impact, results in isolated but fully recoverable loss of digital object authenticity and understandability</td> </tr> <tr> <td>3</td> <td>Superficial impact, results in widespread but fully recoverable loss of digital object authenticity and understandability</td> </tr> <tr> <td>4</td> <td>Medium impact, results in total but fully recoverable loss of digital object authenticity and understandability</td> </tr> <tr> <td>5</td> <td>High impact, results in isolated loss, including unrecoverable loss of digital object authenticity and understandability</td> </tr> <tr> <td>6</td> <td>Considerable impact, results in widespread loss, including unrecoverable loss or loss that is recoverable only by third party of digital object authenticity and understandability</td> </tr> <tr> <td>7</td> <td>Cataclysmic impact, results in total and unrecoverable loss of digital object authenticity and understandability</td> </tr> </tbody> </table>	Risk Impact Score	Interpretation	1	Zero impact, results in zero loss of digital object authenticity and understandability	2	Negligible impact, results in isolated but fully recoverable loss of digital object authenticity and understandability	3	Superficial impact, results in widespread but fully recoverable loss of digital object authenticity and understandability	4	Medium impact, results in total but fully recoverable loss of digital object authenticity and understandability	5	High impact, results in isolated loss, including unrecoverable loss of digital object authenticity and understandability	6	Considerable impact, results in widespread loss, including unrecoverable loss or loss that is recoverable only by third party of digital object authenticity and understandability	7	Cataclysmic impact, results in total and unrecoverable loss of digital object authenticity and understandability
Risk Probability Score	Interpretation																														
1	Minimal probability, occurs once every 100 years or more																														
2	Very low probability, occurs once every 10 years																														
3	Low probability, occurs once every 5 years																														
4	Medium probability, occurs once every year																														
5	High probability, occurs once every month																														
6	Very high probability, occurs more than once every month																														
Risk Impact Score	Interpretation																														
1	Zero impact, results in zero loss of digital object authenticity and understandability																														
2	Negligible impact, results in isolated but fully recoverable loss of digital object authenticity and understandability																														
3	Superficial impact, results in widespread but fully recoverable loss of digital object authenticity and understandability																														
4	Medium impact, results in total but fully recoverable loss of digital object authenticity and understandability																														
5	High impact, results in isolated loss, including unrecoverable loss of digital object authenticity and understandability																														
6	Considerable impact, results in widespread loss, including unrecoverable loss or loss that is recoverable only by third party of digital object authenticity and understandability																														
7	Cataclysmic impact, results in total and unrecoverable loss of digital object authenticity and understandability																														
<p>Figure 3 DRAMBORA risk probability scale</p>	<p>Figure 4 DRAMBORA risk impact scale</p>																														

DRAMBORA not only provides a well structured way of doing risk assessment in the context of digital repositories and hence is well suited to our needs in PrestoPRIME, but it comes complete with a set of candidate risks and example mitigation approaches. Many of these are relevant to maintaining data integrity. The complete list is given below and the most relevant ones are highlighted in yellow.

Organisation Management

- R01 Management failure
- R02 Loss of trust
- R03 Activity is overlooked or allocated insufficient resources
- R04 Business objectives not met
- R05 Repository loses mandate
- R06 Community requirements change substantially
- R07 Community requirements misunderstood or ineffectively communicated
- R08 Enforced cessation of repository operations
- R09 Community feedback not received
- R10 Community feedback not acted upon
- R11 Business fails to preserve essential characteristics of digital information
- R12 Business policies and procedures are unknown
- R13 Business policies and procedures are inefficient
- R14 Business policies and procedures are inconsistent or contradictory
- R15 Legal liability for IPR infringement
- R16 Legal liability for breach of contractual responsibilities
- R17 Legal liability for breach of legislative requirements
- R18 Liability for regulatory non-compliance
- R19 Inability to evaluate repository’s successfulness
- R20 False perception of the extent of repository’s success

Staffing

- R21 Loss of key member(s) of staff
- R22 Staff suffer skill loss
- R23 Staff skills become obsolete
- R24 Inability to evaluate staff effectiveness or suitability

Financial Management

- R25 Finances insufficient to meet repository commitments
- R26 Misallocation of finances

- R27 Liability for non-adherence to financial law or regulations
- R28 Financial shortfalls or income restrictions
- R29 Budgetary reduction

Technical Infrastructure and Security

- R30 Hardware failure or incompatibility
- R31 Software failure or incompatibility
- R32 Hardware or software incapable of supporting emerging repository aims
- R33 Obsolescence of hardware or software
- R34 Media degradation or obsolescence
- R35 Exploitation of security vulnerability
- R36 Unidentified security compromise, vulnerability or information degradation
- R37 Physical intrusion of hardware storage space
- R38 Remote or local software intrusion
- R39 Local destructive or disruptive environmental phenomenon
- R40 Accidental system disruption
- R41 Deliberate system sabotage
- R42 Destruction or non-availability of repository site
- R43 Non availability of core utilities (e.g. electricity, gas, network bandwidth, water)
- R44 Loss of other third-party services
- R45 Change of terms within third-party service contracts
- R46 Destruction of primary documentation
- R47 Inability to evaluate effectiveness of technical infrastructure and security

Acquisition and Ingest

- R48 Structural non-validity or malformation of received packages
- R49 Incompleteness of submitted packages
- R50 Externally motivated changes or maintenance to information during ingest
- R51 Archival information cannot be traced to a received package

Preservation and Storage

- R52 Loss of confidentiality of information
- R53 Loss of availability of information and service
- R54 Loss of authenticity of information
- R55 Loss of integrity of information
- R56 Unidentified information change
- R57 Loss of non-repudiation of commitments
- R58 Loss of information reliability
- R59 Loss of information provenance
- R60 Loss or non-suitability of backups
- R61 Inconsistency between redundant copies
- R62 Extent of what is within the archival object is unclear
- R63 Inability to validate effectiveness of ingest process
- R64 Identifier to information referential integrity is compromised
- R65 Preservation plans cannot be implemented
- R66 Preservation strategies result in information loss
- R67 Inability to validate effectiveness of preservation
- R68 Non-traceability of received, archived or disseminated package

Metadata Management

- R69 Metadata to information referential integrity is compromised
- R70 Documented change history incomplete or incorrect

- R71 Non-discoverability of information objects
- R72 Ambiguity of understandability definition
- R73 Shortcomings in semantic or technical understandability of information

Access and Dissemination

- R74 Non-availability of information delivery services
- R75 Authentication subsystem fails
- R76 Authorisation subsystem fails
- R77 Inability to validate effectiveness of dissemination mechanism
- R78 Loss of performance or service level

Full details of each of these are provided in the DRAMBORA documentation. An example is provided below in Figure 5.

Risk Identifier:	R55	
Risk Name:	Loss of integrity of information	
Risk Description:	Repository is incapable of demonstrating that the integrity of information has been maintained since its receipt, and that what is stored corresponds exactly with what was originally received.	
Is this Risk Relevant?:	<ul style="list-style-type: none"> • Does repository commit to preservation of information integrity? 	
Example Risk Manifestation(s):	<ul style="list-style-type: none"> • Records documenting government expenditure have been subjected to unauthorised or unanticipated changes, rendering them no longer representative of originally deposited content 	
Nature of Risk:	Physical environment	
	Personnel, management and administration procedures	
	Operations and service delivery	X
	Hardware, software or communications equipment and facilities	
Owner:	Preservation	
Escalation Owner:	Preservation	
Stakeholders:	Management; financiers; staff; depositors; users; producers	
Mitigation strategy(ies):	Avoidance strategies: <ul style="list-style-type: none"> • Ensure policies and procedures are conceived with due consideration of integrity requirements • Maintain and review policies and procedures to ensure adequate recording and comparison of checksums to demonstrate that archived information has suffered no loss of integrity since its deposit or receipt • Ensure software and hardware systems and preservation strategies are capable of preserving information integrity In the event of risk's execution: <ul style="list-style-type: none"> • Invoke treatment strategies to alleviate loss of reputation or trust 	
Risk Relationships:	→R01 [contagious] →R02 [contagious]	
Risk Probability:	4	
Risk Potential Impact:	3	
Risk Severity:	12	

Figure 5 Example of DRAMBORA risk assessment sheet (reproduced from DRAMBORA v1.0)

Many of these are directly applicable to assessing the risks that surround the long-term storage of digital content, with some examples shown below of how these might be interpreted for audiovisual preservation.

DRAMBOR A Risk ID	Title ²⁰	Example
R30	Hardware Failure	A storage system corrupts files (bit rot) or loses data due to component failures (e.g. hard drives).
R31	Software Failure	A software upgrade to the system loses or corrupts the index used to locate files.
R32	Systems fail to meet archive needs	The system can't cope with the data volumes and the backups fail.
R33	Obsolescence of hardware or software	A manufacturer stops support for a tape drive and there is insufficient head life left in existing drives owned by the archive to allow migration
R34	Media degradation or obsolescence	The BluRay optical discs used to store XDCAM files develop data loss.
R35-R38	Security	Insufficient security measures allow unauthorised access that results undetected modification of files.
R39	Disasters	All content is in a small space through use of high density storage systems (e.g. tape robot) which makes the archive vulnerable to large-scale loss in a fire or flood.
R40	Accidental System Disruption	An operator accidentally deletes one or more files.
R55, 56, 59	Loss of integrity or authenticity	There is no audit trail for the changes made to content, which mean preservation actions are not taken or are inappropriate.
R60	Unsuitable backups	The backup tapes can't be read.
R61	Inconsistent copies	There are two copies of the content but they are different due to corruption of one of them, but which one is correct can't be identified.
R64, R69	Content Identifiers	The identifier used to locate a particular file in the system is lost or corrupted.

1.6. OCTAVE Allegro

EBIOS²¹, MEHARI²² and OCTAVE²³ are all examples of risk management methods for information security.

²⁰ In some cases the title has been shortened or paraphrased to make it easier to understand.

²¹ Expression of Needs and Identification of Security Objectives (EBIOS), <http://www.ssi.gouv.fr/en/confidence/ebiospresentation.html>

²² Méthode Harmonisée d'Analyse de Risques, MEHARI 2007 Concepts and Mechanisms, CLUSIF, April 2007 https://www.clusif.asso.fr/fr/production/ouvrages/pdf/MEHARI-2007-concepts_principles_2007.pdf

²³ Operationally Critical Threat, Asset, and Vulnerability Evaluation (OCTAVE), <http://www.cert.org/octave/>

OCTAVE (Operationally Critical Threat, Asset, and Vulnerability Evaluation) is a risk-based strategic assessment and planning technique for security proposed by CERT (Computer Emergency Response Team) Coordination Centre ²⁴. It is self-directed, that is, a small team of people from the operational (or business) units and the IT department work together to address the security needs of the organisation. The team draws on the knowledge of many employees to define the current state of security, identify risks to critical assets, and set a security strategy.

OCTAVE is an asset-driven evaluation approach. The analysis team:

- Identifies information-related assets (e.g., information and systems) that are important to the organisation.
- Focuses risk analysis activities on those assets judged to be most critical to the organisation.
- Considers the relationships among critical assets, the threats to those assets, and vulnerabilities to those threats (both organisational and technological) in an operational context - how they are used to conduct an organisation's business and how those assets are at risk due to the security threats.
- Creates a practice-based protection strategy for organisational improvement as well as risk mitigation plans to reduce the risk to the organisation's critical assets.

OCTAVE-S²⁵ is more structured where security concepts are embedded in OCTAVE-S worksheets, allowing less experienced practitioners to use them. However the analysis team should have an extensive knowledge of the organisation's business and security processes.

OCTAVE Allegro²⁶ is a third OCTAVE variant with the goal of producing more robust results without the need for extensive risk assessment knowledge. This approach focuses on information assets in the context of how they are used, where they are stored, transported and processed and how they are exposed to threats, vulnerabilities and disruptions as a result.

OCTAVE Allegro is also well suited for use by individuals who want to perform risk assessment without involving the whole organisation, or security specialists. This makes it ideally suited to PrestoPRIME where don't have direct access to all the possible users of the technology we plan to develop and hence need to do our own analysis based on our internal belief of what they might want.

The OCTAVE Allegro method consists of eight steps organized into four phases:

- Phase 1: Participants develop risk measurement criteria consistent with organisational drivers: the organisation's mission, goal objectives, and critical success factors

²⁴ Computer Emergency Response Team (CERT), <http://www.cert.org>

²⁵

OCTAVE-S: <http://www.cert.org/octave/osig.html>

²⁶ Richard A. Caralli, James F. Stevens, Lisa R. Young, William R. Wilson. Introducing OCTAVE Allegro: Improving the Information Security Risk Assessment Process, Technical report, May 2007, <http://www.cert.org/archive/pdf/07tr012.pdf>

- Phase 2: Participants create a profile of each critical information asset that establishes clear boundaries for it, identifies its security requirements, and identifies all of its containers.
- Phase 3: Participants identify threats to each information asset in the context of its containers.
- Phase 4: Risks to information assets are identified and analyzed and the development of mitigation approaches is begun.

1.7. Combined use of DRAMBORA and OCTAVE Allegro

The stages in Octave Allegro are very similar to those used in DRAMBORA and both techniques use a top down approach based on business objectives and asset-based risk assessment. This in turn means that we can, by and large, use the two approaches interchangeably.

The approach we take in this report is to use example risks and mitigation approaches from DRAMBORA and combine them with further analysis using OCTAVE Allegro in order to create an overall set of security requirements for PrestoPRIME that focus on the risks to AV essence.

OCTAVE Allegro is well suited to users who are not security experts, so works well in the project for extracting data integrity requirements from the PrestoPRIME motivating scenarios. A potential PrestoPRIME adopter may choose to use other methods (e.g. EBIOS, MEHARI or some other method) depending on their needs and access to security expertise. In which case, this document provides a starting point for their analysis (much in the same way that we have used DRAMBORA as our starting point).

OCTAVE Allegro supplies standard worksheet templates to be used in risk analysis (A worked example of these templates is in Annex 1: OCTAVE Allegro worksheets).

Worksheets 1 to 6 support defining the risk measurement criteria. Risk measurement criteria are a set of qualitative measures against which the effects of a realised risk can be evaluated. In addition to evaluating the extent of an impact in a specific area, an organisation must recognize which impact areas are the most significant to its mission and business objectives.

Worksheet 7 supports identifying and classifying these impact areas.

Worksheet 8, titled "Critical Information Asset Profile", contains the description of a critical information asset and its security requirements. It indicates what security requirement(s) are the most important and thus need to be considered.

9a, 9b and 9c Worksheets support identifying the boundaries of the threat environment and the scope of the risk assessment by introducing the container notion. These worksheets identify the containers of the information asset. OCTAVE Allegro considers a container to be a location/system in which the asset is stored, transported or processed,

and which protects the asset from possible threats. The containers may be technical (9a), physical containers (9b) or people (9c). Moreover, these containers may be internal to the organisation that owns the asset, or external at another organisation.

The “information asset worksheet”, Worksheet 10, identifies the threat and the consequence for each asset. Obviously, we may need multiple sheets for each asset. A threat scenario questionnaire) is supplied to help seed the identification of possible threats to the information asset.

OCTAVE Allegro comes with some example threats that could be relevant. Of these, the ones at the technical container level are very relevant to threats that apply to data integrity in the context of PrestoPRIME (modification in the terminology of Allegro). For example, a technical container might be the mass storage system used to store and manage AV data.

Threat Scenario Questionnaire 1		Technical Containers	
<p>This worksheet will help you to think about scenarios that could affect your information asset on the technical containers where it resides. These scenarios may pose risks that you will need to address. Consider each scenario and circle an appropriate response. If your answer is “yes” consider whether the scenario could occur accidentally or intentionally or both.</p>			
<p>Scenario 1: Think about the people who work in your organization. Is there a situation in which an employee could access one or more technical containers, <i>accidentally</i> or <i>intentionally</i>, causing your information asset to be:</p>			
Disclosed to unauthorized individuals?	No	Yes (accidentally)	Yes (intentionally)
Modified so that it is not usable for intended purposes?	No	Yes (accidentally)	Yes (intentionally)
Interrupted so that it cannot be accessed for intended purposes?	No	Yes (accidentally)	Yes (intentionally)
Permanently destroyed or temporarily lost so that it cannot be used for intended purposes?	No	Yes (accidentally)	Yes (intentionally)
<p>Scenario 2: Think about the people who are external to your organization. This could include people who may have a legitimate business relationship with your organization or not. Is there a situation where an outsider could access one or more technical containers, <i>accidentally</i> or <i>intentionally</i>, causing your information asset to be:</p>			
Disclosed to unauthorized individuals?	No	Yes (accidentally)	Yes (intentionally)
Modified so that it is not usable for intended purposes?	No	Yes (accidentally)	Yes (intentionally)
Interrupted so that it cannot be accessed for intended purposes?	No	Yes (accidentally)	Yes (intentionally)
Permanently destroyed or temporarily lost so that it cannot be used for intended purposes?	No	Yes (accidentally)	Yes (intentionally)

Figure 6 Example threats from OCTAVE Allegro at the technical container level

Threat Scenario Questionnaire – 1 (cont)		Technical Containers			
<p>Scenario 3:</p> <p>In this scenario, consider situations that could affect your information asset on any technical containers you identified. Determine whether any of the following could occur, and if yes, determine whether these situations would cause one or more of the following outcomes:</p> <ul style="list-style-type: none"> • Unintended disclosure of your information asset • Unintended modification of your information asset • Unintended interruption of the availability of your information asset • Unintended permanent destruction or temporary loss of your information asset 					
A software defect occurs	No	Yes (disclosure)	Yes (modification)	Yes (interruption)	Yes (loss)
A system crash of known or unknown origin occurs	No	Yes (disclosure)	Yes (modification)	Yes (interruption)	Yes (loss)
A hardware defect occurs	No	Yes (disclosure)	Yes (modification)	Yes (interruption)	Yes (loss)
Malicious code (such as a virus, worm, Trojan horse, or back door) is executed	No	Yes (disclosure)	Yes (modification)	Yes (interruption)	Yes (loss)
Power supply to technical containers is interrupted	No	Yes (disclosure)	Yes (modification)	Yes (interruption)	Yes (loss)
Problems with telecommunications occur	No	Yes (disclosure)	Yes (modification)	Yes (interruption)	Yes (loss)
Other third-party problems or systems	No	Yes (disclosure)	Yes (modification)	Yes (interruption)	Yes (loss)
Natural or man-made disasters (flood, fire, tornado, explosion, or hurricane) occur	No	Yes (disclosure)	Yes (modification)	Yes (interruption)	Yes (loss)

Figure 7 Example threats from OCTAVE Allegro at the technical container level

Threat Scenario Questionnaire – 1 (cont)		Technical Containers			
<p>Scenario 3:</p> <p>In this scenario, consider situations that could affect your information asset on any technical containers you identified. Determine whether any of the following could occur, and if yes, determine whether these situations would cause one or more of the following outcomes:</p> <ul style="list-style-type: none"> • Unintended disclosure of your information asset • Unintended modification of your information asset • Unintended interruption of the availability of your information asset • Unintended permanent destruction or temporary loss of your information asset 					
A software defect occurs	No	Yes (disclosure)	Yes (modification)	Yes (interruption)	Yes (loss)
A system crash of known or unknown origin occurs	No	Yes (disclosure)	Yes (modification)	Yes (interruption)	Yes (loss)
A hardware defect occurs	No	Yes (disclosure)	Yes (modification)	Yes (interruption)	Yes (loss)
Malicious code (such as a virus, worm, Trojan horse, or back door) is executed	No	Yes (disclosure)	Yes (modification)	Yes (interruption)	Yes (loss)
Power supply to technical containers is interrupted	No	Yes (disclosure)	Yes (modification)	Yes (interruption)	Yes (loss)
Problems with telecommunications occur	No	Yes (disclosure)	Yes (modification)	Yes (interruption)	Yes (loss)
Other third-party problems or systems	No	Yes (disclosure)	Yes (modification)	Yes (interruption)	Yes (loss)
Natural or man-made disasters (flood, fire, tornado, explosion, or hurricane) occur	No	Yes (disclosure)	Yes (modification)	Yes (interruption)	Yes (loss)

Figure 8 Example threats from OCTAVE Allegro at the technical container level

For each risk, severity parameters are calculated in each impact area. Severity (high, medium, low) semantics are identified in worksheets 1 to 5. The estimated severity value (high=3, medium=2, low=1) is multiplied by the organisation impact area priority (worksheet 7) to obtain the score. This relative risk score enables to classify risks and prioritize them in the context of organisation’s mission and business objective. For example, if reputation is most important to an organisation, risks that have an impact on the organisation reputation will generate higher scores than risks with equivalent impacts in another area. The second part of worksheet 10 indicates the mitigation approach. This includes identifying administrative, technical, and physical controls and measures to be applied on certain containers in order to mitigate risks.

A worked example of using the OCTAVE Allegro sheets for a hypothetical AV storage scenario is provided in Annex 2: OCTAVE Allegro Risk Analysis Example.

2. Example Risk Analysis focusing on Data Integrity

A meaningful risk assessment can only be done in the context of a specific organisation, with a specific set of assets and specific business objectives. Therefore, the example analysis in this section is somewhat illustrative. The analysis also short-cuts the early stages of the process for sake of brevity.

In a risk assessment context, in order to cover a wide range of possible risks, the example analysis covers a set of scenarios rather than just one. A thorough risk assessment would treat each scenario individually. We include the following scenarios:

1. **Archive as a Service within the enterprise.** In this scenario, we consider the case where there is an archive function within a business that is tasked with providing archive services to that business according to an OAIS model (e.g. ingesting content into the archive, providing access to content in the archive, ensuring the content remains useable over time). The archive is responsible for the long-term integrity of the content it holds and for maintaining a guaranteed level of service to its customers. This includes supporting WORM (write once, read many) so that when data is in the archive then there is a guarantee that it can't be changed. Imagine the resources used within the archive to ensure that content is safe (e.g. mass storage systems and a programme of data integrity checking and technology migration) are the same as those used to deliver archive services to the customers.

2. **Distributed preservation copies.** In this scenario, in order to achieve required levels of data safety an archive decides to keep multiple copies of content in multiple locations in order to protect against failures and disasters at all levels. To do this the archive chooses to use both an in-house storage solution and one or more storage service providers to store the content offsite.

3. **Preservation Service Provider.** In this scenario, imagine an archive service provider that hosts and operates one or more archives on behalf of its customers. Services provided to the customers include migration of content to address technical obsolescence and the delivery of content in appropriate formats for use. The service and the SLA outlive the technologies used to deliver the service (media formats, ingest and access mechanisms, service delivery platforms etc.)

4. **Preservation supply chain.** Imagine one or more content owners that use a third party archive service provider for the hosting of their archive content. The service provided is in terms that archives understand (assets, ingest, access, retention, safety, security) and supports the OAIS model (SIPs for ingest, DIPs for access, AIPs for internal preservation etc.). The archive service provider is expert in preservation, media formats, archive business models, integration of archives into cross-organisation media processes (production, post-production, distribution etc.). However, the archive service provider uses storage sourced from one or more storage service providers. These storage service providers (e.g. Amazon S3, eVault, EMC Mozy etc.) specialise in the technical aspects of data storage and work in terms of data volumes, bandwidth, latency, availability etc

2.1. Risks of loss of data authenticity and integrity

These risks are concerned with the loss of ability to track and record the origins of data and then everything that is done to data during digital preservation. Without this provenance trail, there is the risk that changes to integrity or authenticity happen but go unnoticed.

Loss of authenticity and integrity of content								
Presto Prime RiskID	Risk Title	Risk Description (what the threat is, i.e. what might happen)	DRAMB ' RiskID	Risk origin (where the risk comes from - containers and people)	Information Assets at risk	Impact area (Reputation, productivity, financial, legal, health and safety etc.)	Container	Mitigation
1	Loss of authenticity of information	Repository is incapable of demonstrating that information objects are what they purport to be and hence what their original integrity was	R54	Lack of, or failure to follow, proper process				Enforce authentication and access control so only trusted individuals have ability to manipulate assets (both within and external to the organisation) Record all actions to content that take place (who did what and when) to create a complete audit trail
2	Loss of integrity of information	Repository is incapable of demonstrating that the integrity of information has been maintained since its receipt, and that what is stored corresponds exactly with what was originally received.	R55					
3	Unidentified information change	Repository is incapable of tracking or monitoring where one or more changes to archived information has taken place, so integrity is lost	R56					
4	Loss of non-repudiation of commitments	Repository is incapable of ensuring that commitments cannot later be denied by either of the parties involved, which could include integrity of	R57					
				Failure to record all actions performed within the archive		Loss of reputation (archive service provider)	Archive Service Platform	Use technical measures such as digital signatures (e.g. hashing) and integrity monitoring to detect changes in digital content, both within storage systems and in transit over networks
				Failure of archive storage systems or processing of content	Audiovisual content	Legal penalties imposed due to breach of contract	Archive Storage Server	Log any attempted breaches, deliberate or accidental, and whether they were successful or not to allow security effectiveness to be

		data supplied or stored						measured.
5	Loss of information reliability	Repository is incapable of demonstrating the reliability of its information holdings	R58	Failure to record attempts (deliberate or otherwise) to breach systems	Descriptive Metadata	Extra time and resources needed to repair loss of integrity or reputation	Unsecured Networks	Regular security audits of technology, processes, staff skills etc.
6	Loss of information provenance	Repository is incapable of demonstrating the provenance of its information holdings, and their traceability from receipt and through each interaction that they have been subject to.	R59	Failures at remote storage service providers	Contracts	Loss of ability to use content (customer)	Third-party storage services	Evaluate and take into account any increased risk from using data encryption in storage systems as a potential degradation amplifier.
7	Inability to evaluate effectiveness of technical infrastructure and security	Repository is incapable of effectively determining the extent to which its technical infrastructure and security provisions are capable of facilitating business objectives (inc. Integrity guarantees)	R47	Deliberate attack by disgruntled employees		Failure to record details of transactions with consequent denial by customer or service provider that they have agreed obligations		Use appropriate integrity assurance processes that match the frequency, timescales and severity of the ways in which integrity could be lost
8	Identifier to information referential integrity is compromised	Where identifiers are applied to information, the repository is incapable of locating the archival package that corresponds to a given ID	R64	Deliberate attack by hackers or other third-parties				Ensure integrity records (e.g. checksums or signatures) are kept safe and are themselves subject to integrity control
9	Inability to evaluate repository's successfulness	Repository is incapable of effectively determining the extent to which it has successfully achieved its business objectives.	R19	Failure of preservation systems to correctly apply preservation actions				Ensure integrity control is comprehensive and consistent, i.e. applied to all forms of data (metadata, identifiers, checksums, logs, credentials, audiovisual

								content)
--	--	--	--	--	--	--	--	----------

2.2. Risks of data destruction or degradation

If we broaden the notion of risks to data beyond the strict digital preservation notion of 'integrity' and include the more general risks of loss or corruption of data, which are equally important for archives, then a further set of risks arise

Degradation or destruction of content								
Presto Prime RiskID	Risk Title	Risk Description (what the threat is, i.e. what might happen)	DRAMB ' RiskID	Risk origin (where the risk comes from - containers and people)	Information Assets at risk	Impact area (Reputation, productivity, financial, legal, health and safety etc.)	Container	Mitigation
10	Exploitation of security vulnerability	Shortcoming in repository's security provisions is identified and used to gain unauthorised access.	R35	Failure of preservation systems to correctly apply preservation actions Failure of archive storage systems or processing of content Customer supplies content that with mistaken levels of quality or	Audiovisual content	Loss of reputation (archive service provider)	Archive Service Platform	Enforce authentication and access control so only trusted individuals have ability to manipulate assets (both within and external to the organisation). Limit access to systems are only strictly necessary. Record all actions to content that take place (who did what and when) to create a complete audit trail and check regularly for inappropriate behaviour. React quickly to breaches.
11	Unidentified security compromise, vulnerability or information degradation	Security exploitation or vulnerability occurs and is not monitored or identified by repository staff.	R36					
12	Physical intrusion of hardware storage space	Intruder gains access to area within which repository technical hardware is physically located	R37					
13	Remote or local software intrusion	Repository suffers software intrusion conducted either from onsite or from a remote location, by bypassing network security	R38					
					Descriptive Metadata	Legal penalties imposed due to breach of contract	Archive Storage Server	Evaluate and take into account any increased risk from using data encryption in storage systems as a potential degradation amplifier.

		provisions.		completeness				
14	Accidental system disruption	Business activities are adversely affected by non-deliberate intervention, or intervention that was not intended to result in these outcomes.	R40	Customer supplies content that isn't encrypted properly	Security credentials	Extra time and resources needed to recover content or reputation	Unsecured Networks	Use appropriate processes that match the frequency, timescales and severity of the ways in which corruption or loss could occur.
15	Deliberate system sabotage	Business activities are adversely affected by measures intended to have these effects.	R41	Deliberate attack by disgruntled employees		Loss of ability to use content (customer)		Ensure measures are comprehensive and consistent, i.e. applied to all forms of data (metadata, identifiers, checksums, logs, credentials, audiovisual content)
16	Structural non-validity or malformedness of received packages	Received packages fail to correspond to what repository expects or is capable of preserving.	R48	Deliberate attack by hackers or other third-parties				Check all content being ingest into the archive for completeness and correctness and agree this with content submittor.
17	Obsolescence of hardware or software	Core technology is no longer current or is incongruent with that of most comparable organisations	R33					Ensure systems are resilient to attack by internal staff, recognised customers and service providers, and third-parties.
18	Media degradation or obsolescence	Storage media deteriorates, limiting the extent to which it can be written to and read from.	R34					Determine, monitor and manage both media degradation and technical obsolescence timescales

2.3. Risks to data through loss of services

If there is a loss or interruption to the services and processes that are involved in preservation and access to digital content then this has the potential to put the content itself at risk of loss. This is particularly important if the service being provided is one of maintaining data integrity.

Degradation or loss of service(s)									
Presto Prime RiskID	Risk Title	Risk Description (what the threat is, i.e. what might happen)	DRAMBORA RiskID	Risk origin (where the risk comes from - containers and people)	Information Assets at risk	Impact area (Reputation, productivity, financial, legal, health and safety etc.)	Container	Mitigation	
24	Loss of performance or service level	Repository is incapable of meeting service level goals in accordance with its business objectives.	R78	Archive service provider or storage service provider goes bust	Services providing access	Loss of reputation (archive service provider)	Archive Service Platform	Secure and allocate resources based on business priorities. Establish mechanisms to regularly review and, if necessary, adjust policies and procedures in order to ensure business objectives are realised	
25	Business objectives not met	One or more integral business outcomes are not achieved, or are achieved inadequately.	R04	Bugs or failure of authentication or access control systems				Establish mechanisms to review and adjust resource allocations and to monitor and control workloads placed on resources	
26	Non-availability of information delivery services	Repository is unable to provide access to information packages.	R74	Bugs or failure of archive storage systems or processing of content				Legal penalties imposed due to breach of contract	Maintain 'spare capacity' to facilitate subsequent resourcing of originally overlooked or overloaded activity
27	Activity is overlooked or allocated insufficient resources	An integral business activity is mismanaged leading to its noncompletion.	R03	Deliberate attacks or accidental damage to systems				Audiovisual content	Extra time and resources needed to restore services or rebuild reputation

28	Enforced cessation of repository operations	Repository is forced to cease its business activities.	R08	Legal action taken that causes services to be suspended or permanently taken down.	Descriptive Metadata	Loss of ability to access content at all or in a timely way (customer)	Archive Storage Server	Evaluate, monitor and manage use of resources in inter-enterprise scenarios, e.g. whether storage as a service from Amazon S3 is sufficient to support ingest SLA given to a customer of an archive service provider Evaluate effects of system changes prior to their implementation
29	Loss of availability of information and/or service	Repository is unable to provide a comprehensive range of services or access to all of its information holdings for which access ought to be available.	R53	Insufficient resources to meet service level agreements		Increase in risk that content is lost or is corrupted		Allocate a proportion of staff time to monitoring the ongoing suitability of hardware and software resources and assessing the potential value of emerging technologies. Pre-empt hardware failure with anticipatory investment.
30	Hardware failure or incompatibility	System hardware is rendered incapable of facilitating current business objectives.	R30	Incorrect prioritisation or allocation of resources		Access to assets by those not originally intended, e.g. third-parties if a company is liquidated.		
31	Software failure or incompatibility	System software is rendered incapable of facilitating current business objectives.	R31	Inefficient use of resources				

2.4. Risks to loss of data integrity through mismatch of expectations

If preservation is provided as a service, e.g. within an organisation or by a third-party then there the potential for a mismatch in expectations or understanding between the providers of the service and the community for which the services are being provided. If the changes in expectations are too rapid, not communicated properly, then data can be put at risk. For example, if the required level of data integrity is not properly defined, or if the sudden need for higher levels of integrity are beyond the capabilities of current systems.

Mismatch of expectations								
Presto Prime RiskID	Risk Title	Risk Description (what the threat is, i.e. what might happen)	DRAMB' RiskID	Risk origin (where the risk comes from - containers and people)	Information Assets at risk	Impact area (Reputation, productivity, financial, legal, health and safety etc.)	Container	Mitigation
32	Community requirements change substantially	Community expectations or requirements are substantially altered, and no longer correspond to business activities.	R06	Failure to understand security needs, e.g. confidentiality, or to communicate service offered		Loss of ability to use content		Include security as part of contract with service provider including limits on how it can be changed or clauses that allow it to be changed on request
33	Community requirements misunderstood or miscommunicated	Repository is incapable of determining the expectations of its stakeholder communities and therefore unable to tailor business activities appropriately	R07					Failure to anticipate or react to changes in services used from third parties
34	Business policies and procedures are inefficient	Rationale and/or practical approach adopted for business fail to demonstrate optimal efficiency.	R13	User community adopts new software systems which provide no support for legacy encrypted data formats that were previously dominant.	Audiovisual content	Inappropriate security used that results in loss of integrity, confidentiality or availability	Archive Service Platform	Maintain flexible approach to operational objectives to react to emerging community requirements

35	Business policies and procedures are unknown	Fundamentals of why and how repository's business activities are conducted are undocumented and unknown, or known only by specific individuals.	R12	Community becomes increasingly unfamiliar with the semantics of a previously well-known and widely employed security techniques	Descriptive Metadata	Service provider is inefficient or not competitive through not understanding user needs	Archive Storage Server	Continuously monitor user needs to detect mismatch in services or expectations.
36	Change of terms within third-party service contracts	Conditions with which third-party services are delivered change substantially.	R45	Users want to change security of their archived content to better fit with their operational systems (e.g. new access protocols or authorisation systems)		Incorrect or inappropriate hardware/software/people used with consequent increase in costs		Establish long-term agreements with service providers that are flexible enough to evolve with changing needs
37	Hardware or software incapable of supporting emerging repository aims	Technical infrastructure, while adequate for meeting current aims, is incapable of meeting new requirements resulting from organisation's natural evolution.	R32	User demands on service are not what the service provider expects, e.g. increase in ingest over time				Review use of internal resources and external services on a regular basis (including options for changing providers) and optimise combination.

2.5. Summary

There are a wide range of threats to long term data preservation when using IT systems, in particular mass storage technology.

Risk assessment methodologies from both digital preservation and information security communities provide a structured approach to identifying and assessing these risks.

We have divided risks into four main areas.

- **Risks of loss of data authenticity and integrity.** These risks are mostly concerned with the loss of ability to track and record the origins of data and then everything that is done to data during digital preservation. Without this provenance trail, there is the risk that changes to integrity or authenticity happen but go unnoticed.
- **Risks of data destruction or degradation.** These risks are concerned with the loss or corruption of data, for example from imperfect storage technology, deliberate or accidental damage, or loss of access to data due to technical obsolescence or which are equally important for archives, then a further set of risks arise
- **Risks to data through loss of services.** If there is a loss or interruption to the services and processes that are involved in preservation or access to digital content, then this has the potential to put the content itself at risk of loss. For example, this might be the loss of a service that routinely checks and maintains data integrity in a storage system.
- **Risks to loss of data integrity through mismatch of expectations.** If preservation is provided as a service, e.g. within an organisation or by a third-party then there the potential for a mismatch in expectations or understanding between the providers of the service and the community for which the services are being provided. If the changes in expectations are too rapid, or not communicated properly, then data can be put at risk. For example, the required level of data integrity might not be properly defined, or the sudden need for higher levels of integrity might be beyond the capabilities of current systems

3. Failure modes and data corruption from mass storage technology

3.1. Overview

Currently there are two main options for storage of audiovisual files that are viable for long-term, cost-effective audiovisual archiving: data tape and hard disk drives (HDD). There are other technologies, e.g. optical disks, solid state memory or even recording digital information onto film, and whilst each has its niche application, none satisfy the general

requirements for cost, accessibility, safety and scalability. The various options, including their use within file-based preservation strategies and most importantly the Total Cost of Ownership (TCO) over time are all explained in some detail in PrestoPRIME D2.1.1 “Preservation Strategies”.

This report concentrates on data tape and HDD, including the use of these media within automated systems (disk servers and tape robots), and concentrates on a single and very important issue:

How reliable are these technologies, in what ways do they fail, what are the consequences of these failures on the content stored within them, and what can be done to mitigate these failure modes?

Before going into the details, the main conclusions are that:

- Hard disk drives (HDD) are complex electrical and mechanical devices that are prone to failure and cannot be considered as reliable on their own. HDD on shelves is not a viable archiving approach from a safety perspective. For this reason (and others), a range of techniques have been developed by the IT Industry to help manage the failure modes of disk based storage, both within the drives themselves and by systems, e.g. RAID arrays, in which they sit. However, these systems, despite best efforts, are not perfect and sometimes introduce their own errors and data corruption modes.
- Data tape on the other hand is relatively reliable. Tapes are simple and robust from a mechanical and electrical point of view (tape cartridges do have some electronics in them). Tapes can be stored for relatively long periods without intervention, e.g. on shelves. Periodic migration is still required due to obsolescence of tape drives, with robots providing a way to automate this process as well as to provide online access. The main threat to content on data tape is not the media itself but the drives used to access it, which can damage tapes or mark them as unusable.

We have a simple message to archives: current IT storage technology cannot be considered 100% safe. There is a very clear need to have at least two copies of content for safety reasons, ideally on two different technologies, in two different locations and operated by different teams of people. If this is done effectively, along with a properly managed programme of regular migration, then digital archiving can achieve very high levels of data integrity and longevity. But this takes planning and proactive management.

3.2. Latent Faults and Visible Faults

It would be nice to think that as soon as an error occurs in storage then it is detected and reported, i.e. it becomes *visible* and hence can be *repaired*. However, many faults in IT storage systems are *latent* i.e. the error goes undetected for some period of time. This is particularly dangerous since without detecting the error there is no hope of correcting it, yet there is always the chance of more errors occurring until the point of no return – unrecoverable data loss.

Latent faults can occur for a wide range of reasons.

- Errors may be detected at and reported at some low level in the storage stack, e.g. in a disk controller (e.g. SMART²⁷) and then subsequently ignored in higher levels of the stack, e.g. operating system.
- Errors may occur in storage media that isn't being actively used and these naturally remain latent until the next time of access, which could be several years (deterioration of CD-R is a well known example here).
- A storage system may think it has correctly written or read data, but in fact some form of error has occurred (e.g. misdirected write) that isn't picked up.
- Bugs in software, firmware or hardware may cause errors in the data where the system thinks that everything has happened correctly.

A very simple example is shown in Figure 9. This shows what might happen with two copies of a file on two different storage systems or media (e.g. two USB external hard disk drives) Starting on the left, there both copies are in known good condition. However, one of the copies might become corrupted (orange state) for some reason (e.g. disk develops some bad sectors). This is a latent (undetected) fault until some attempt is made to access the data, at which point the fault becomes visible (yellow state). Only then can an attempt be made to repair the corrupted data (e.g. by copying from the other good disk). However, in between the time that the first fault develops and its detection and repair (or replacement), there is the chance that the second copy also develops a fault. If this happens before repair of the first copy is finished then both copies are corrupted and there is no longer any way to repair the data, i.e. it is permanently lost.

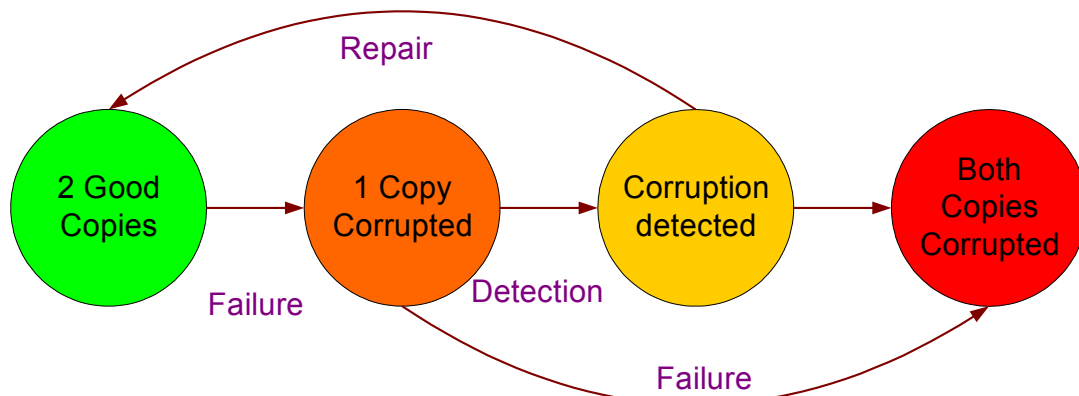


Figure 9 Data corruption and repair

This example illustrates the difference between latent and detected faults, and most importantly shows that it is important to consider the detection and repair processes in any system that aims to maintain data integrity when using unreliable storage.

In the world of HDD based storage, this approach is recognised and well developed, e.g. through the use of RAID²⁸ arrays and scheduled 'data scrubbing' operations by RAID systems (proactive integrity checking and repair)²⁹. This is not to say that these systems

²⁷ <http://en.wikipedia.org/wiki/S.M.A.R.T.>

²⁸ <http://en.wikipedia.org/wiki/RAID>

²⁹ Parity Lost and Parity Regained, Andrew Krioukov et al

http://www.usenix.org/events/fast08/tech/full_papers/krioukov/krioukov_html/index.html

are perfect, which they are not - see the section on hard disks below – only that proactive integrity management is a necessity in hard drive based systems and hence techniques have been put in place that attempt to do this in these systems for some time. Some specific details of latent faults in these systems, in particular the problem of 'bit rot', are provided below for HDD storage. Very little information is available on latent faults developing in data tape, although evidence suggests this is not a common phenomenon.

In conclusion:

- When faults develop in mass storage systems they are not always immediately detected let alone corrected. These so called 'latent' faults happen for a variety of reasons and are not uncommon, particularly in hard disk based storage systems. Without careful consideration they can lead to irrecoverable data loss.
- Proper analysis and planning of data safety in digital storage systems, be it HDD in servers or data tapes on shelves, has to consider not only the types and frequencies of faults, but most importantly how fast those faults can be detected and then repaired. This dictates the overall probability of data loss.
- Anyone seeking to achieve long term preservation of data using mass storage technology needs to be aware of latent faults, their consequences, and to employ a proactive approach to data integrity monitoring and management.

Examples of more sophisticated modelling of multi-copy storage systems is provided in "The Modelling System Reliability For Digital Preservation: Model Modification and Four-Copy Model Study"³⁰ by Yan Han and Chi Pak Chan.

An extensive and excellent discussion of visible vs. latent faults is provided by Mary Baker et al in "A Fresh Look at the Reliability of Longterm Digital Storage"³¹

3.3. Mean Time Between Failures (MTBF)

The storage industry makes heavy use of the term Mean Time Between Failure (MTBF) when advertising the reliability of both storage media (e.g. disks) and storage systems (e.g. servers, robots).

MTBF for hard disks is typically 1,000,000hrs. A LTO tape drive might have a MTBF of 200,000 hrs at 100% duty cycle.

But what do these mean?

- A simple interpretation is that there is approximately 10,000 hrs in a year and hence a MTBF of 1 million years means that a hard drive will typically *last* 100 years. This is wrong.
- A slightly more sophisticated interpretation would be that the 1 in 100 hard drives will fail each year, i.e. the chances of a hard drive failing is 1% *each* year. This is also wrong.

³⁰ http://www.bl.uk/ipres2008/presentations_day2/44_Han.pdf

³¹ <http://www.lockss.org/locksswiki/files/Eurosyst2006.pdf>

Neither interpretation considers what MTBF really means and most importantly what it *doesn't* tell you.

MTBF is a statistic on the *average* time between failures for a population of things, e.g. drives or tapes. So, for the hard drive MTBF example above, if you have 100 hard drives, each with a MTBF of 1 million hours, then on average a drive will fail in this set every 10,000 hrs, i.e. one drive a year.

It is 'on average' that is the key here.

Imagine there are three hypothetical hard drive manufacturers.

1. Manufacturer 1 does a test on 10 of their drives, 9 of which fail after 1 year and one of which that lasts for 100 years. They claim a MTBF of approximately 10 years
2. Manufacturer 2 does a test on 10 of their drives, 5 of which fail after 1 year and 5 of which fail after 20 years. They claim a MTBF of approx 10 years.
3. Manufacturer 3 does a test on 10 of their drives, 1 of which fails after 1 year and all the rest fail after 10 years. They claim a MTBF of approx 10 years.

Now suppose you are interested in using hard drives for storing your content for 5 years (after which you have to migrate to a new server so you only need the drives to get you this far). You look at the MTBF for each manufacturer and they all say 10 years. However, choosing drives from Manufacturer 1 is clearly a lot more likely to cause you major data loss in the first than if you choose drives from Manufacturer 3. But how do you know which is best?

In reality, there is little to choose between HDD of a given type (e.g. SATA, SAS) from different manufacturers – they all have roughly the same reliability. However, this reliability as seen 'in the field' can be a lot lower than implied by MTBF figures from the manufacturers. This is examined in more detail in the section on disk storage below.

The other thing to note in the example above is the idea of a manufacturer doing a test and observing that their drives last 10, 20 or even 100 years (approx. 1 million hours). They of course can't physically observe drives for this long. They instead build simulations and models to predict lifetime or do stress testing to 'accelerate' aging so they can draw some conclusions in a much more limited time (a few months). Therefore, a MTBF of 1 million hours for a HDD is of course an estimate rather than a hard and fast number.

Finally, MTBF should not be confused with useful life of a component or system. So, for example, the MTBF for car tyres is very high, e.g. the time on average between a car having a tyre 'blow out' when you drive on it is probably 100s of years or longer (otherwise there would be a lot of accidents on the roads). However, the useful life of a car tyre is a lot shorter, say 50,000km which can be less than a few years of driving. It is clearly important not to use the MTBF number to imply how long something will last in normal service.

The same applies to storage media and storage systems. For example, the MTBF of a tape drive might be 250,000 hrs but the life of the head in the drive might only be 16,000

hrs. Likewise, the MTBF for tape cartridges (e.g. the average time between a tape snapping) can be very high (typically not even quoted), but the useful life of a tape is typically only a few thousand hours of play/record time.

In conclusion:

- MTBF is of limited value when it comes to assessing reliability of storage media or systems. What would be much more useful is failure rates for the first 5 years of life. This information is typically only available from field studies of media or systems used 'in the wild' and not 'in the manufacturer's lab'. It is also very important to distinguish between MTBF and the normal service life of components, which can be considerably shorter.
- What the archive community really needs is to do is to collect and share information on their experience with media and system reliability, especially in the first few years of use. This is the only way to generate meaningful statistics on the failure modes and their frequency for digital mass storage technology.

For a more extensive discussion on MTBF, see "Disk failures in the real world: What does an MTTF of 1,000,000 hours mean to you?"³² by Bianca Schroeder and Garth Gibson, and "Bit Preservation: A Solved Problem"³³ by David Rosenthal, which considers in detail how meaningful the concept of Mean Time To Data Loss (MTTDL) is for long-term digital preservation.

3.4. Storage capacity: trends and the implications

The storage capacity of media, e.g. HDDs, is increasing by 100 times every decade on average (a trend that has existed for the last 30 years – see D2.1.1), but reliability of media items is not increasing at the same rate (if at all). This has some very serious consequences for archives.

The number of hours of programme material that will fit on a single item of media is going up. For example, a single 1TB HDD can hold 1000hrs of CD quality audio or 40 hours of 50MBit/sec MPEG2 compressed SD resolution video. This contrasts with the analogue world of video or audio tapes where it is typical to have a 'one item per tape' approach. If you lose one tape, then you lose one item. In the digital domain, if you lose a data tape or a hard drive you could easily lose 10, 100 or even 1000 items. In a decade this could be up to a million items. This makes multiple copies ever more important.

For example, PrestoSpace found the average duration of an item in AV archives was approximately 20 minutes. 20 minutes of uncompressed SD video requires approx 30GB of storage. Therefore, you can get at approx 30 items on a 1TB hard drive or LTO4 data tape. Within a decade, at current trends, it will be possible to fit several thousand hours of uncompressed SD resolution content on a single item of media. If this drive or tape fails then a whole collection can be lost.

³² http://db.usenix.org/events/fast07/tech/schroeder/schroeder_html/index.html

³³ http://www.bl.uk/ipres2008/presentations_day2/43_Rosenthal.pdf

This is of course addressed to some extent by having multiple copies. As number of hours of programme material per HDD or tape goes up, then so too does the number of copies to achieve a given level of reliability. Thankfully, the falling costs of storage make this viable.

Consider the hypothetical scenario below.

1. Imagine you have a collection of 1000 items of video material, which are each one hour in duration. Suppose that each item requires approx 400GB of storage (e.g. 880MBit/s HD in ProRes format). If each item goes on an LTO3 tape (up to 400GB uncompressed per tape), then there will be 1000 tapes. Suppose you have two copies of each tape. Now, imagine that you are migrating the tapes in a few years time. Suppose for purpose of this example, the rate at which tapes fail to play back during migration (due to problems with either the tapes or the drives) is 1%. The chances of both the primary and backup copy of a tape independently having problems is small (1 in 10,000) and hence across all 1000 tapes there is a 90% chance that no data at all will be lost during the migration. The chance of one item being lost out of the 1000 is 10% and the chance of two or more items of video being lost is only 1%. The chance of all the video being lost is vanishingly small.
2. Imagine the same scenario in 10 years time. Following the LTO roadmap, which has held for the last 10 years, then capacity of LTO tapes will have increased 30 times³⁴. It will now possible to get at least 30 items of video on each tape. So, the 1000 hours of video are now on approx 30 tapes, not 1000. If the same rate of problems are seen when trying to play the tapes back, then there is a 0.3% chance of both copies being lost of at least one out the 30 tapes. The chances of a loss of video item taking place is lower than when the video was on 1000 tapes, but when loss does takes place then a lot more items are affected.
3. Imagine another 10 years down the line. Now all the video will fit on a single tape. If there are two copies of this tape and the failure rates are still the same, then there is a 1 in 10,000 chance of losing all the video. This is an all or nothing scenario.

So, over the space of 20 years, the archive has gone from a relatively high chance of losing only a small amount of video in the collection during a migration, through to a relatively small, but still significant, chance of losing absolutely everything! Most archives would opt for the first option where they might expect to lose the occasional file but their overall collections are very safe. Maintaining this mode of loss during successive migrations requires proactive measures – it is not enough to simply follow a 2 copies strategy in perpetuity.

There are several approaches to this problem.

- Firstly, make more tape copies. In 10 or 20 years time when storage costs have fallen then it is more affordable to have 3, 4 or more copies. This makes the probability of major loss very small.

³⁴ LTO tape roadmap has seen a new LTO version every 2 years, each of which doubles capacity. Therefore, in 10 years there will be 5 new versions which equates to 32 times more capacity. Having said that, LTO5 which is the next generation due for release is behind schedule so there is some evidence that the LTO roadmap is slowing down.

- Secondly, don't make exact tape for tape copies. Rearrange the files on the backup copies. So for example if there are 10 files on the master tape, then spread these 10 files across 10 backup tapes. Do this likewise for the other master tapes. This means that if a master tape is damaged then there are 10 tapes to recover the files from. The chances of all these being problematic is very small, so the majority of the files can always be expected to be recovered. It does however mean more effort in doing backups/restores.
- Thirdly, use HDD storage instead of data tape. This allows a wide range of existing data distribution strategies to be employed (RAID, RAIN etc.) to increase data safety for minimal cost. This is currently more expensive than tape options, but the gap is likely to close over time.

The second strategy can potentially be included as a natural part of archive growth. With increases in data volumes, the majority of archives are expanding. Even though data density for media is exponentially increasing, this is to a large extent being counteracted by files that are much larger e.g. due to higher resolution, frame rate, colour depth etc. The result is that most archives of any significant size will always have hundreds or thousands of data tapes for a considerable time to come. This gives the opportunity of spreading the older and smaller files across the tapes that contain the larger and newer files – gap filling in effect. This makes efficient use of storage and at the same time lowers the risk of catastrophic loss for older collections of material as the copies are distributed across multiple items of media.

In conclusion:

- The trend towards ever higher capacity of media (tapes, hard disks) for the same cost is very attractive to archives for obvious reasons. However, increase in reliability is not keeping pace with this increase in capacity, which has the result of making mass storage cheaper, but also more likely to cause large scale data loss.
- Strategies for data distribution across storage media or systems need to evolve to ensure that the ways in which loss might occur are kept within acceptable limits. Simple strategies that work today, e.g. direct data tape replication to create two copies may not be so applicable in the next decade or beyond.

Hard Disk Drives (HDD)

There is already a large body of literature about the problems of long-term data storage using HDD based systems. David Rosenthal regularly provides a useful round-up of the literature in his blog³⁵, for example on the topic of storage reliability and silent data corruption³⁶. A useful guide to hard disk technology in general can be found on StorageReview.com³⁷. It is not the purpose of this report to review or reproduce the results of this large body of work in detail. Instead we make the following main points. Reliability of hard drives as seen 'in the wild' is considerably lower than might be inferred from manufacturers MTBF figures. Annual Failure Rates (AFR) over the first five years of life (a reasonable service lifetime for a drive) have been observed to vary from <1% to

³⁵ <http://blog.dshr.org/>

³⁶ <http://blog.dshr.org/2008/03/more-bad-news-on-storage-reliability.html>

³⁷ <http://www.storagereview.com/guide/index.html>

over 13%. Seminal work in this area includes studies by NetApp and University of Wisconsin³⁸ that looked at over 1.5 million drives in a 41 month period and Google³⁹ that looked at hundreds of thousands of drives over a 5 year period. There are many interesting findings of these reports, e.g. that temperature of operation has little bearing on longevity and neither does whether the drive is heavily used or not or whether the drive is 'enterprise class' (e.g. SAS drives) or 'consumer grade' e.g. SATA drives.

Interesting as these findings are, they are only of limited use. Disk drives are used within systems and it is the overall reliability of the systems that counts. Disks will of course fail in these systems, but if the right protection mechanisms are in place then this need not mean data loss. Therefore, the failure rate 'headlines' of the NetApp and Google study do not translate into rates of data loss (unless you are crazy enough to only have a single disk copy of data, e.g. a 'HDD on shelves'; preservation strategy'). The rates are of course useful in knowing the costs of keeping the system in good shape, e.g. by replacing failed drives and what labour this will entail.

So, when people have asked 'What does a MTBF of 1,000,000 hours mean to you?' and then come back with details of observed AFR that are 10 times higher than expected, it is also worth asking 'What does an AFR of 10% mean to you?'

If you have two HDD copies of data and decide to only replace failed drives at the end of each year, then there is a 1% chance of data loss for each pair of disks. On the other hand, if you check data integrity every week⁴⁰ and replace failed or problematic drives as soon as problems are detected, then the probability of loss falls to 0.02% per annum – clearly a lot better.

What is important is the levels of data safety that can be achieved in practice. Here other studies are very relevant. For example, CERN investigated the rate of data corruption in their systems⁴¹ by checking data integrity before/after writing/reading from disk. This revealed endemic silent data corruption in hard drive storage, including in 'enterprise class' systems that are explicitly designed to prevent data loss. As many as 1 bit in every 10⁹ was on average irreversibly corrupted. At the file level, one file in 1500 was affected. Errors occurred in the very systems, e.g. RAID controllers, that are designed to mitigate against failures lower down in the stack, and to protect against bit or sector level errors on hard drives (some reasons why data loss occurs with hard drives is explained well by Josh Eddy in his whitepaper on silent data corruption⁴²).

The practical evidence of CERN and others that many factors contribute to data corruption is backed up by an analysis by Jiang et al in 2008⁴³ again based on NetApp statistics, but this time for 1.8 million drives in 155,000 systems. The report includes a host of interesting findings. For example:

Finding (1): In addition to disk failures (20-55%), physical interconnect failures make up a significant part (27-68%) of storage subsystem failures. Protocol failures and performance failures both make up noticeable fractions.

³⁸ http://www.usenix.org/event/fast08/tech/full_papers/bairavasundaram/bairavasundaram_html/

³⁹ http://labs.google.com/papers/disk_failures.pdf

⁴⁰ Sustained data read rates from modern hard drives can easily exceed 50Mbyte/sec. Doing a simple checksum on files on the drive can be done at this speed or faster using standard servers. So, it is possible to do a complete checksum test of all the files on a 1TB drive in just over 5 hours.

⁴¹ <http://indico.cern.ch/getFile.py/access?contribId=3&sessionId=0&resId=1&materialId=paper&confId=13797>

⁴² <http://raidinc.com/pdf/Silent%20Data%20Corruption%20Whitepaper.pdf>

⁴³ <http://www.usenix.org/events/fast08/tech/jiang.html>

Implications: Disk failures are not always a dominant factor of storage subsystem failures, and a reliability study for storage subsystems cannot only focus on disk failures. Resilient mechanisms should target all failure types.

This shows that many of the failures in storage systems are not down to the disks within them, indeed Jiang also concludes that disk failure rates should not be used to predict system failure rates.

Surveys and investigations have been done into what it takes to achieve very high levels of data integrity in disk based storage systems, e.g. as reported by Krioukov et al in 2008⁴⁴ and shown in Figure 10.

RAID	Scrub	Sector Checksum	Block Checksum	Parent Checksum	Write-Verify	Physical ID	Logical ID	Version Mirror	Chance of Data Loss
✓									0.602%
✓	✓								0.602%
✓	✓	✓							0.322%
✓	✓		✓						0.041%
✓	✓			✓					*0.486%
✓				✓					*0.153%
✓	✓		✓		✓				0.002%
✓	✓		✓			✓			0.038%
✓	✓		✓				✓		*0.033%
✓			✓				✓		*0.010%
✓	✓		✓			✓	✓		*0.031%
✓			✓			✓	✓		*0.010%
✓	✓		✓				✓	✓	*0.004%
✓			✓				✓	✓	*0.002%
✓	✓		✓			✓	✓	✓	0.000%

Table 3: Probability of Loss or Corruption. *The table provides an approximate probability of at least 1 data loss event and of corrupt data being returned to the user at least once, when each of the protection schemes is used for storing data. It is assumed that the storage system uses 4 data disks, and 1 parity disk. A (*) indicates that the data loss is detectable given the particular scheme (and hence can be turned into unavailability, depending on system implementation).*

Figure 10 Extract from paper by Krioukov et al on the measures needed to reduce the chance of data loss in disk based storage systems.

To reduce the probability of data loss down to low levels requires a range of different measures. These also need to be implemented in a way that is bug free (which we know from the CERN study is something that manufactures find hard to achieve – and not surprisingly either given the complexity of the systems concerned).

So, the rate at which data loss will actually take place in a given storage system is highly dependent on the exact configuration of that system as well as its design and manufacture. This makes it impossible to assign a single number to the reliability of hard disk based storage. The headline grabbing statistics on disk failure rates do not help.

⁴⁴ http://www.usenix.org/events/fast08/tech/full_papers/krioukov/krioukov.pdf

It is also briefly revisiting the 1 bit in 10^9 data corruption statistic from CERN. A 1TB data file contains approximately 10^{13} bits. Therefore, if the CERN observed data corruption was at the individual bit level and randomly spread, then we would expect 10^4 corruptions in the data file. Looking ahead to the section later on the sensitivity of compressed audio and video to bit level data corruption, 10^4 bits could be expected to render the video completely useless. However, there is plenty of experience in using HDD storage for large video files, at least in the production and post-production process if not archiving. People here are not regularly complaining of large scale video corruption – so what is happening? The answer is that many of the corruptions are not at the bit level but are instead grouped together in blocks, e.g. 64kB blocks lost due to bugs in a RAID controller. Corruption at the 64k block level with an average of 1 bit in 10^9 would affect less than 1 in 100 files at the TB size. This is more consistent with practical experience as well as the file level corruption rates seen by CERN. Other studies confirm that corruption is typically at the block level rather than bit level, and it tends to be spatially correlated, e.g. successive blocks on a disk are more likely to be corrupted than blocks at random, and it affects media in batches (e.g. a bad batch of hard drives from a particular manufacturer). This is why corruption of large files exists, but is not endemic. Further studies are needed on how this pattern of corruption files translates to loss of content.

So, from all of this we draw the following conclusions:

- Hard disks as storage media should not be considered reliable. Annual Failure Rates (AFR) between 1 and 10% are common in the first few years of life. This means that a strategy of unpowered HDD on shelves is almost certain to fail and instead HDD should always be used within mass storage systems that manage data integrity and component failures.
- Data corruption also occurs in hard disk based storage systems, including those explicitly designed to reduce data loss. Most worryingly this corruption can be completely silent and hence go undetected for long periods of time.
- Even if data corruption could be eradicated from disk based storage systems, then there are wider aspects to consider too, e.g. data transmission over networks, temporary data storage in memory, errors in data management systems or even human errors when operating these systems.
- If you go to such extreme measures that data corruption is eradicated from disk-based systems then this utopia is likely to cost so much that it is either unsustainable or doesn't scale to the data volumes found in typical audiovisual archives.
- Any organisation using HDD based mass storage systems with a requirement for long-term data integrity should employ an ongoing and proactive programme of data integrity checking and repair at an end-to-end systems level. It should not be assumed that any component of the system (networks, storage, memory, processing) is somehow 'safe', i.e. immune from failures and data corruption problems.

That said, we conclude by remarking that HDD are staggering pieces of engineering and it is remarkable the reliability they achieve given their complexity and the need for very precise engineering.

The comments about reliability of HDD in this section are in no way intended to detract from the achievements of HDD manufacturers in these feats of engineering.

HDD throughput and error rates: trends and the implications

Whilst the storage capacity of HDD shows a clear trend of doubling every 18 month, the improvement in the rate at which this data can be retrieved from HDD is not keeping pace.

The throughput of HDD, i.e. the rate at which data can be accessed, determines the rate at which that data can be checked for data corruption. It also determines that rate at which repair can take place if corruption is identified. The throughput of HDD is also not keeping pace with the increases in capacity (see Elerath's paper for more information). If the rate of detection and repair falls in proportion to the volume of data being stored then there are two problems:

- Corruption checking and repair takes up proportionally more 'system time' and can impinge on the operational use of the system, for example serving data to archive users.
- The longer it takes to check and repair the data then the higher the chance of further failures occurring during the checking or repair process that then cause permanent data loss.

Furthermore, improvements in the error rates when reading data from media is also not keeping pace with the rate of increase in storage capacity of this media. The read error rates reported by Elerath from NetApp in 2007⁴⁵, e.g. 1 error in 8×10^{14} Bytes read being considered 'medium', are little improved over error rates reported by Chen et al in 1994⁴⁶ of 1 error in 10^{14} bits read, which was over a decade earlier. In the same period disk capacity has increased 100 fold. This lack of improvement in error rates has serious implications on the design and implementation of systems that ensure data safety.

For example, a Bit Error Rate of 1 bit in 10^{14} is typical in practice for a HDD⁴⁷. There are approx 10^{13} bits of data on a 1TB drive. Therefore, the probability of a read error occurring when reading all the data from the drive is approx 10%. Unless dramatic improvements are made to the level of read errors (which are already amazingly low considering), then we will soon reach the point where it is more likely than not to encounter a read error when reading all the data from a hard drive.

The increased checking and repair time along with probability of read errors are particularly relevant to HDD based systems. For example, RAID5⁴⁸ has for a long time provided a way to protect against HDD failures. If a HDD fails in a RAID5 set of disks then

⁴⁵http://entertainmentstorage.org/articles/Hard%20Disk%20Drives_%20The%20Good,%20The%20Bad%20and%20The%20Ugly.pdf

⁴⁶ <http://www.pdl.cmu.edu/PDL-FTP/RAID/computin.pdf>

⁴⁷ This is distinct from the higher error rates reported in the CERN study mentioned earlier in the report. The BER for hard drives is the probability of an error occurring when reading the data from the drive. The error rate seen by CERN includes all parts of their system (network, memory, RAID, disk etc.) and is considerably higher as a result.

⁴⁸ <http://en.wikipedia.org/wiki/Raid5>

the data on the remaining disks can be used to rebuild the full dataset. This repair operation requires reading of all of the remaining disks to recreate the data that was lost on the failed disk. A modern SATA 1TB HDD can typically sustain a read rate of 100MB/sec⁴⁹, which means reading the full disk will take approx 2.5hours. If this is done at the same time as the disk is being used to serve users then the elapsed time will be considerably longer. Having to read data from a set of these disks in a RAID5 configuration and then write to a replacement disk for the one that failed could easily result in a 24 hour rebuild time.

During this time, if a further disk fails then all the data is lost. If we assume an AFR for a HDD of 10% (see below) then for an array of 4 remaining disks, then there is approx a 0.1% chance of failure during a 24hr rebuild. Furthermore, there is a 40% chance that there needs to be at least one of these rebuilds each year due to the rate at which HDD will fail in the array. The chances of complete data loss each year are small, but not negligible.

Also during this rebuild time, if there is a read error from one of the disks then this error will be included in the rebuilt dataset. Because of the disk that has already failed, there is no longer any way to detect or correct these errors. Using the 10^{-14} bit error rate above, then the probability of reading all 4 disks without any read errors will be 65%, in other words, there is a 35% change of some form of data corruption occurring during the rebuild. This is clearly not acceptable for preservation scenarios.

The main threat to data in the system is not hard disk failures but data corruption or data read errors. In addition to the read error rate, latent faults also have to be considered, e.g. as discussed above. This means regular scrubbing is important in order to minimise the chances of data errors during RAID rebuilds. However, with throughput not keeping pace with capacity improvements, RAID scrubbing is an increasingly costly and time consuming activity.

This problem has led to the development of RAID6⁵⁰ and similar models e.g. NetApp's RAID-DP⁵¹ which uses a double parity approach. These systems can accommodate two failures in the RAID set before data loss or corruption takes place, either as a result of disk failures or read errors. The chances of this happening are very small and RAID6 is now becoming a standard approach. However, for the same reasons that RAID5 is becoming obsolete and superseded by RAID6, the days of RAID6 are also numbered. Triple parity RAID is the next step, for example see the detailed explanation in the recent ACM article by Adam Leventhal "Triple parity RAID and beyond"⁵²

The point is that there is a difference in the rates at which capacity advances compared to throughput and read error rates. This means that new solutions will have to be evolved to ensure that acceptable levels of data safety are maintained. This in turn means a need to keep watch on storage technology trends and make sure appropriate solutions are adopted.

⁴⁹ <http://www.tomshardware.com/charts/2009-3.5-desktop-hard-drive-charts/h2benchw-3.12-Max-Read-Throughput,1009.html>

⁵⁰ <http://en.wikipedia.org/wiki/Raid6>

⁵¹ <http://whitepapers.techrepublic.com.com/abstract.aspx?docid=310984>

⁵² <http://queue.acm.org/detail.cfm?id=1670144>

A much simpler message is also clear – don't expect to read all the data from a large HDD without a relatively high chance (1-10%) that there will be some form of error in the data somewhere. Whilst manufacturers are improving read error rates, this improvement is not as fast as the increase in capacity of storage media. Therefore, whilst the probability of error per byte of data is going down, the probability of an error per media item, e.g. HDD, is actually going up. For HDD, RAID techniques can protect against this, but only during the time that they are fully operational and not in a rebuild cycle.

In conclusion :

The rate at which HDD throughput and read error rates are improving does not match the rate at which capacity is increasing. This has several implications for HDD based systems used for preservation.

- HDD failures still need to be protected against, e.g. by using RAID systems, but it is now read errors and latent errors that are the main threat to permanent data corruption.
- There can be a very significant chance of encountering data corruption problem unless RAID6 or similar schemes are used. HDD in JBOD or simple RAID setups (e.g. RAID5 or RAID0) are not sufficiently protected against these errors. The other alternative is to keep multiple separate copies of the data and proactively monitor data integrity, e.g. using checksums (see later in this report).
- The techniques needed to ensure adequate data safety will continue to evolve (e.g. the transition from RAID5 to RAID6 and beyond into triple parity approaches). This requires continual technology watch and trend analysis to ensure that data safety levels do not deteriorate in archive storage systems.

3.5. Data Tape

2008 saw LTO data tape pass the 100 million mark for the number of cartridges shipped. This is a huge number. Averaging over different LTO generations, this equates to some 10,000 Petabytes of capacity. Yet, and in stark contrast to HDD, there are no large-scale field-studies of tape reliability where the results are publicly accessible (well, not ones that the authors are aware of).

Firstly, such a study would be very useful to the archive community in understanding the risks of using data tape and what counter measures to employ. For this reason we make a call for more sharing of information in this area in the community.

Several archives in PrestoPRIME, and some of the other archives that they know, are major users of data tape and have done so for some time. We collected together some of their findings.

- Modern data tape, in particular LTO, appears to be a relatively reliable storage medium. There were no cases of widespread data loss due to failures or degradation of data tapes themselves. Indeed there were some cases to the contrary with large archives (100s of TB) reporting no losses of data at all⁵³ One

⁵³ The BNF in France has used LTO1,3,4 over the last 8 years and currently has 600TB of data. They have

archive that actively checks data integrity during migration confirmed they did not detect any data integrity loss⁵⁴. These archives have thousands of tapes⁵⁵ and have gone through one or more migration cycles, so these findings are significant.

- Many of the problems with data tape come from tape drives rather than the tapes themselves.
 - Several archives reported problems with drives malfunctioning that then caused damage to data tapes. However, it was also possible to manage this problem by having multiple tape copies, ensuring different drives were used for the different copies, and repairing/replacing problematic drives when problems were spotted.
 - Several archives reported problems when moving between generations of tape, in particular when a new drive was used to play an old tape from a previous generation. If the drive had problems then it marked the tape as 'unusable' which prevented any subsequent attempts to play it back, including in the original drives. The problem was with the drives and not with the tapes. The way round this problem was to involve the manufacturer and get them to resolve the problem including transfer content from the 'unusable' tapes.
 - Data tapes that are read or written frequently do 'wear out' and become more problematic as they head towards the end of their life – this can be before the manufacturers stated life. This can then result in either data on the tape not being readable or drives marking the tapes as unusable.

A combination of the effects above can occur and some archives commented that it is hard to separate out problems with tape from problems with drives. This is particularly true where tapes are used frequently (written and/or read). Care is needed to avoid unnecessary wear. A retrieval of a file from a tape can involve more than one pass. In extreme cases where the I/O capability of the tape drive isn't matched by the systems in which it is used (server, SAN, network etc), then repeated stop/reverse/restarts may occur ('shoe shining'⁵⁶) which result in tapes becoming worn quite quickly.

Overall it is fair to say that modern data tape is reliable. Practical experience of archives with data tape, including during migration projects, have proved the value of second copies tapes. However, the percentage of files that needed to be recovered from backup were typically very small (a few percent and often less for modern LTO tapes) and the number of files lost because there was a problem with the backup too was much lower again, e.g. 1 in 100,000 for one archive. These problems mostly exist for earlier generations of data tape. Reassuringly, several archives reported no problems at all in their systems based on LTO3 or LTO4 tape, including during migrations.

yet to lose any of their data, including during migrations.

⁵⁴ One archive that has used checksums to verify fixity during migration found less than 0.1% of files had problems during the first transfer attempt. All these problems were corrected during a second attempt. No data integrity was lost for over 1million files during the migration, of which over 10,000 were video.

⁵⁵ For example, the BBC has 10,000 LTO3 tapes. Only 20 of these are currently causing any problems and all similar tapes in the past have had these problems resolved (e.g. by using different drives or by getting support from the manufacturer). No content has been lost so far.

⁵⁶ http://en.wikipedia.org/wiki/Shoe_shining

The reliability of tape along with the 2:1 cost saving compared with HDD means that data tape will likely become the mainstay of large scale AV archives for some time to come, especially where there is not call to access the data frequently. The difference between tape and disk in terms of cost is closing quickly.

The conclusions are clear:

- Data tape is relatively reliable with problems tending to come from drives rather than tapes. Two or more tape copies are still needed for safety.
- Regular migration is essential. Whilst media lifetime of tape can easily exceed 15-30 years in good conditions, the drives become obsolete in 5-7 years (e.g. see LTO roadmap).
- The details of tape reliability are still unclear. There is a need for the AV archive community to share statistics on reliability 'in the field' of data tape, including during migration as well as in operational systems. In particular, information on latent faults is very lacking. Even if this shows that there are few problems then this is very reassuring to those considering use of data tape.
- Lifetime (head life, media read/write cycles etc.) is a more important as a metric than MTBF. Lifetimes should be respected and a conservative approach (e.g. stop when you half way through) can pay dividends on avoiding problems with tape, especially when used frequently.

3.6. Use of AV compression

The use of compression, e.g. video coding, is common place in the AV industry, including preservation, either because content is born this way in the production process, or because cost constraints make uncompressed storage prohibitively expensive. However, the use of compression makes the content more sensitive to corruption at the bit level.

The CERN study found that a single bit error would make a compressed data file (e.g. a zip file) unreadable, with a probability of 99.8%. As shown later in this report, just a few bytes lost or corrupted from compressed files in the AV world, e.g. a compressed video sequence, can render one or more frames completely unusable, or, in the worst case, the whole file is unusable.

As an example, consider uncompressed audio. A .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 shown⁵⁷ that a .WAV file with 0.4% errors is almost undistinguishable from the original⁵⁸, whereas an MP3 file with the same level of errors either will not open at all, or will have errors affecting most of the

⁵⁷ Informal experiments by Richard Wright, BBC R&D; unpublished

⁵⁸ Short errors, e.g. a byte or two in length, will effectively produce short spikes or dropouts in the waveform, which produce very high frequency artefacts when the audio is listened too, which will be by and large inaudible.

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”).

Not encoding, in particular not using compression, typically results in files that have minimal sensitivity to corruption. In this way, the choice not to use compression is a way to mitigate against loss. On the other hand, using compression, be it lossless or lossy, can save on storage space, and in turn allow more copies to be held for the same cost, which is an alternative approach. The use of lossy compression plus multiple copies vs. no compression and fewer copies can be considered as a trade-off between a one-off up-front loss (lossy compression) against a much lower probability that further loss will take place. Further work is needed (and will be done by PrestoPRIME) to explore these trade-offs.


Heydegger has developed a ‘robustness indicator’ on the sensitivity of image formats to bit level corruption and then investigated how compression affects robustness⁵⁹. This work is notable as it includes JPEG2000, which is emerging as a strong candidate for preservation in the AV community including digital cinema⁶⁰.

Tests by Heydegger showed that corrupting only 0.01% of the bytes in a compressed JPEG2000 file, including lossless compression, could result in at least 50% of the original information encoded in the file being affected. In some cases, corrupting just a single byte in a JPEG2000 image would cause highly visible artefacts throughout the whole of that image.

This sensitivity to corruption is traded for a saving in storage space, although the trade-off isn’t always simply one for the other. For example, Heydegger found that one byte of corruption had the following effect:

- a 10 MB uncompressed TIFF file had just .00001% errors (meaning just that one byte was affected)
- a lossless JPEG2000 file had 17% errors for a saving of 27% in storage
- a lossy JPEG2000 file had 2.1% errors for a saving of 62% in storage

As an example of the affect of data loss on image files, here are two examples: a BMP file (uncompressed) and a GIF file (compressed). The BMP file has 1400 errors, one in every 256 bytes. The GIF file has a single error.



⁵⁹ http://old.hki.uni-koeln.de/people/herrmann/forschung/heydegger_archiving2008_40.pdf

⁶⁰ Enhanced Digital Cinema project (EDCINE) <http://www.edcine.org/intro/>

Figure 11 BMP file with one error per 256 bytes (1400 errors)	Figure 12 GIF file with a single error (in 14 KB)
--	--

From the above results, it is evident that removing redundancy through compression increases impact of corruption, i.e. 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 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. As the cost of storage devices reduces, and as storage management improves in efficiency, preservation strategies based solely or largely on storage costs are less and less satisfactory.

It is also possible to encode files in a way that deliberately increases their robustness to corruption, JPEG2000 wireless (JPWL) being an example⁶¹. Redundancy and error checking are built in to improve robustness to errors introduced during transmission over wireless channels. However, whilst this approach, and source/channel coding more generally⁶² is used for robust transmission through space, i.e. from one geographical location to another⁶³, including both error recovery and error concealment, it has yet to be applied to long-term transmission through time where the channel is the storage system and noise is introduced by that channel, e.g. silent corruptions. The AV encoding schemes and the way the channel introduces errors are not necessarily the same for storage and network transmission. Further work is needed to better understand how to encode AV content, especially high bit rate video, so it is more resilient to the failure modes of IT storage technology.

The frustration at the moment 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 compensation⁶⁴: the concealment option would allow a missing line to be replaced by a neighbouring line, and expensive machines could even replace entire frames with an adjacent (in time) frame.

As shown below, corruption of a compressed video frame will, in general, render the whole frame unusable. If a video uses intra-frame compression only (e.g. JPEG2000 lossless, DV, MPEG2 D10) then corruption is contained to the specific frame where the error occurred. Various concealment techniques are then possible, e.g. replacing the corrupted frame with an interpolation of known good frames from either side (using motion compensation and other sophisticated techniques that are already well developed). However, if the video uses inter-frame compression (e.g. lossy MPEG2 or H264) then multiple successive frames can be affected and the potential for concealment is vastly reduced.

⁶¹ <http://www.jpeg.org/jpeg2000/j2kpart11.html>

⁶² http://en.wikipedia.org/wiki/Information_theory

⁶³ There is a lot of research in this area, for obvious reasons (e.g. because DVB-T, DVB-S and DVB-H all involve lossy transmission channels. One of the many examples is here

<http://ieeexplore.ieee.org/stamp/stamp.jsp?tp=&arnumber=4743912&isnumber=4743805>

⁶⁴ The example here <http://www.avpreserve.com/dvanalyzer/dv-preservation-data-or-video/> shows how well error concealment can work in practice (DV in this case).

However, despite intraframe compression being a 'safer' approach than interframe compression, generally speaking, file-based digital storage technology has little ability to cope with loss⁶⁵ (corruption; uncorrectable errors).

In conclusion:

- The use of compression for audio or video greatly increases the sensitivity of AV content to data corruption.
- The ability to recover from data corruption depends on both the spatial and temporal way in which data corruption affects the content, and here the use of inter-frame encoding will significantly reduce the potential for techniques such as concealment to be applied if corruption does occur.
- Even with lossless compression schemes with intra-frame only encoding, the loss of a single byte in a frame can render the entire frame unviewable.
- Whilst video compression is attractive in saving storage space and hence cost, this has to be balanced against the need for more frequent and proactive integrity management efforts to ensure content is not affected by data corruption. New cost models are needed that allow archives understand the Total Cost of Ownership (TCO) of maintaining a given level of data integrity.

3.7. IT methods for data integrity checking

Checksums⁶⁶, also known as hashes, are commonly used in the IT world for checking that data has not been accidentally or deliberately altered in a range of circumstances, e.g. corruption during transmission over a network or guaranteeing authenticity.

A checksum is a short and fixed size datum that is computed from the contents of a file or other block of data. The size of the checksum does not vary with the size of the file. Variations in the data cause variations in the checksum, often with an 'avalanche effect'⁶⁷ whereby even a very small change in the data, e.g. a single bit changing, will cause the checksum to change completely. This feature of checksums makes them ideal for detecting even small levels of data corruption in large files. So, for example, a 128bit MD5 hash value⁶⁸ can easily be computed for files that are TB or larger in size and the hash value will change if only one bit in that file changes. The integrity of the data can be checked at any later time by re-computing the checksum and comparing it with the stored one. If the checksums do not match, the data was almost certainly altered (either intentionally or unintentionally). Simple, non-secure hash functions are able to detect accidental changes to data, for example Cyclic Redundancy Checks⁶⁹ (CRC). A 'CRC' is calculated for each block of data and sent or stored with the data. When a block is read or received, the operation is repeated; if the new 'CRC' does not match the one calculated earlier, then the block contains a data error and corrective action taken, e.g. rereading or requesting the block to be sent again).

⁶⁵ The exception is the case where the format used for preservation and the ways in which loss occur both match error correction or error concealment techniques developed for content transmission, e.g. packet loss during network transmission. However, this is generally not the case.

⁶⁶ <http://en.wikipedia.org/wiki/Checksum>

⁶⁷ http://en.wikipedia.org/wiki/Avalanche_effect

⁶⁸ <http://en.wikipedia.org/wiki/Md5>

⁶⁹ http://en.wikipedia.org/wiki/Cyclic_redundancy_check

Beyond CRC, in particular in applications where security is required to prevent deliberate tampering of data, cryptographic hash functions⁷⁰ are more common place. These have the properties that it is infeasible to calculate what data will generate a given hash or to find two different data blocks that yield the same hash. In this way, it becomes very difficult to alter data without detection. Cryptographic hashes can of course be used for simple integrity checking to guard against accidental corruption – with the added benefit that they guard against attempts at deliberate and undetected corruption. The downside is the extra computational cost involved. Some examples of typical hash functions include MD5 (Message-Digest algorithm 5), which uses a 128-bit hash value and conforms to Internet standard (RFC 1321⁷¹). SHA (Secure Hash Algorithm) is another which has several variations (SHA-0, SHA-1, SHA-2). SHA-1 is employed in several widely used security applications and protocols including TLS⁷² and SSL⁷³, PGP⁷⁴, SSH⁷⁵, S/MIME⁷⁶, and IPsec⁷⁷. In the majority of archive scenarios, cryptographic hashing is unnecessary since the primary purpose is to monitor accidental loss of data integrity, e.g. from ‘bit rot’ in storage. Given the very large data volumes often involved in archiving, low computational overhead of generating and comparing hashes is the most important factor. Benchmarking of hash functions, e.g. as available through the Crypto++ library⁷⁸ shows that simple hash functions are capable of being executed at very high data rates using modest PC server type hardware. For example, tests by the Crypto++ library providers shows that Adler32 can be executed at over 900MB/sec on a 2GHz processor. CRC32, MD5 and SHA algorithms could all be executed at over 100MB/sec data rates. Adler32 shouldn’t be regarded as safe against deliberate attack, but is more than sufficient for basic integrity monitoring of large data files.

For the data rates seen in benchmarking, the overhead of doing checksum tests is minimal, e.g. after network data transfers or when reading data from storage. This is of course only true if the data is being read anyway and checksums are ‘piggy backed’ into the process. If the data has to be specifically and additionally read just to do a checksum then the data storage system becomes the bottleneck. For example, a system that can support data read rates of 1Gbit/sec will still take many days to deliver the 100s of TB in a medium AV archive and months for Petabyte size data sets.

Several archives are using hash based integrity checking to good effect. For example, the Austrian Mediathek used hash comparison during migration-processes where files were copied between storage systems. Less than 0.1 percent of files failed the check and all of these were transferred successfully on a second attempt. Whilst data corruption rates were low, corruption did exist, but hash checking was successful detection technique.

⁷⁰ http://en.wikipedia.org/wiki/Cryptographic_hash_function

⁷¹ <http://tools.ietf.org/html/rfc1321>

⁷² Transport Layer Security; cryptographic protocol providing security for communications over networks

⁷³ Secure Sockets Layer; cryptographic protocol providing security for communications over networks

⁷⁴ Pretty Good Privacy; computer program, often used for signing, encrypting and decrypting e-mails

⁷⁵ Secure Shell; network protocol for exchanging data using a secure channel between two networked devices

⁷⁶ Secure/Multipurpose Internet Mail Extensions; standard for public key encryption and signing of e-mail

⁷⁷ Internet Protocol Security; protocol suite for securing IP-communications

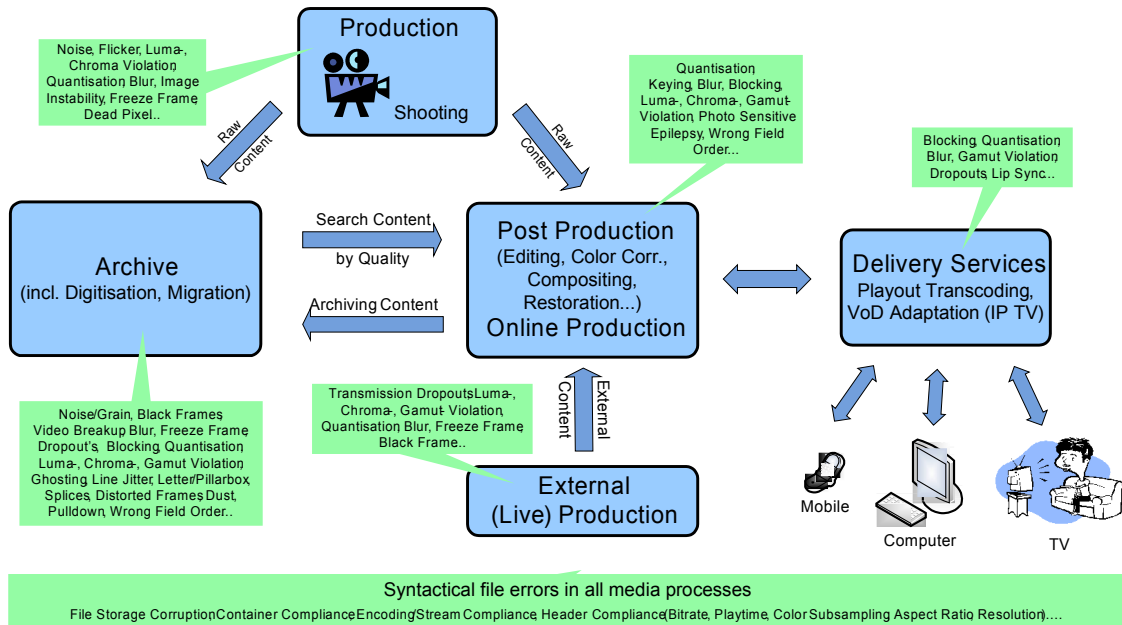
⁷⁸ <http://www.cryptopp.com/>

In conclusion:

- Checksums provide a fast and simple way to monitor data integrity in mass storage systems. Simple algorithms such as Aldler32 and CRC32 are sufficient in scenarios where accidental corruption needs to be detected and can be computed at very high speed (hundreds of MB per sec) using modest PC hardware.
- The serious consequences of data corruption to AV content, especially when compressed, combined with the existence of silent data corruption in mass storage systems (bit rot), especially those based on hard drives, mean that regular integrity checking should form a natural part of archive data management.

4. Errors for video, audio and images and their detection

Video Quality Artefacts in the Media Processes



© JOANNEUM RESEARCH2009

Figure 13 Diagram showing the many places where video quality can be compromised in the production, post-production, delivery and archiving lifecycle.

Various video artefacts and file errors can arise in the media processes. The diagram above in Figure 13 provides a simplified overview of video processes relevant for broadcasters (production, post, delivery and archiving) including information on which artefacts or errors might occur within these processes.

Syntactical file errors can occur in all the processes. They can be caused by corruption of stored files, e.g. due to bit rot by non standard compliant encoding of media containers, e.g. of MXF or MOV containers or by non standard compliant encoding of the video streams itself, e.g. of MPEG streams. All these errors have in common that they can be syntactically checked for correctness.

Another class of artefacts are those which are superimposed on the video content itself, while those video files are syntactically correct. In production sensor noise, dead pixels, luma-, chroma violations, blur, image instability and flicker can degrade the video content. In post production keying, blur, wrong field order, quantisation, blocking and luma-, chroma, gamut violation can degrade the video content. Externally produced content can suffer from luma-, chroma-, gamut -violations, quantisation, blur or blocking. Externally produced live content can suffer from transmission dropouts, freeze frame and black frame. Transcoding for delivery services can induce severe blocking, quantisation and blur. Delivered content can suffer from dropout and other transmission related artefacts like bad lip sync. Archived content suffers, dependent from the original media, of film degradations or video storage and transmission degradations. Original film content can suffer from noise/grain, dust, scratches, lost frames, splices and image instability. Original video suffers from noise, various types of dropouts, video breakup, freeze frame, luma-, chroma-, gamut violation, ghosting, line jitter, incorrect pull-down, and wrong field order.

Archived content which has been digitized or migrated can suffer additionally from blur, blocking or quantisation. Newly archived content can suffer from all degradations introduced within the production and post-production processes. For the case that delivered content is archived also those artefacts introduced within the delivery process additionally can degrade the video content.

In order to ensure proper quality checking in all the video archive related processes a holistic approach is required. It is important to ensure syntactical validity of encoded and stored video files, as well as it is important to avoid artefacts to become superimposed on the video content. Suitable syntactic and content based video checking tools shall support this goal.

4.1. Examples of data-integrity violation

Much of this report has discussed the possibilities of data corruption, e.g. through 'bit rot'. This section provides a very early stage look at the consequences of data loss in terms of how it affects the usability and quality of the audiovisual content held in files.

Analysis of the effects of data corruption on the usability of audiovisual content is actually a complex and involved task. The first step is a more detailed analysis of the failure modes of storage (e.g. silent data corruption at the bit, byte, sector, block, RAID and other levels), their frequencies, and how they get combined in end-to-end storage systems. Only then can a realistic 'corruptor' be developed with which to inject these errors into AV files so the consequent effects can be analysed. This is work that will be done in subsequent stages of the project.

In this section of the report we present some very early stage work done already by PrestoPRIME partners as well as some pointers to work done by others. The results are very much incomplete, but they do indicate that there are many problems when AV content is corrupted.

Example 1: Artefacts in Video-Files

There appears to be surprisingly little publicly available information on the effect of data corruption during storage on video content. ORF is undertaking a study in this area, although unfortunately the results are not yet available for this report. Clearly more work is required in this area and this is something that PrestoPRIME will undertake.

Of the information that is available, this can be divided into two areas.

The first area concerns approaches to transmission of digital AV content over lossy channels, e.g. digital video broadcasting (DVB⁷⁹), JPEG2000 transmission over wireless networks⁸⁰, or video delivery over IP as used in video on demand scenarios. Studies in this area tend to focus on different encoding and error correcting schemes and how they can minimise the problems.

As an example, Figure 14 shows the effect of packet loss on MPEG2 video when transmitted over an IP network (extracted from Testing Video-on-Demand Services over

⁷⁹ http://en.wikipedia.org/wiki/Digital_Video_Broadcasting

⁸⁰ <http://www.jpeg.org/jpeg2000/j2kpart11.html>

Ethernet/IP by Bruno Giguère). This example is relevant to AV archiving in several respects. Firstly, compressed MPEG2 content is a common production and distribution format and is often found in AV archives. Secondly, the loss mechanisms may not be dissimilar: losing packets on an IP network is a similar granularity of data loss to losing or corrupting sectors on a hard disk drive. It is clear that even a small level of data loss, e.g. a fraction of a percent, can have major consequences on the visual quality of the content.



Figure 14 Effects of packet loss on video quality for MPEG2 video transmitted over an IP network as part of a VoD system. Reproduced from Testing Video-on-Demand Services over Ethernet/IP by Bruno Giguère.⁸¹

The second, and much less well developed, area of where the effect of data corruption on video content has been studied is how data loss during storage can affect video files. Here work by Volker Heydegger is particularly relevant as he includes JPEG2000 as one of the file formats he has analysed, which of course is now emerging as a strong candidate for a long-term video preservation format.

In Heydegger's 2008 paper⁸² "Analysing the Impact of File Formats on Data Integrity" the severity of even very low levels of corruption on JPEG2000 images is illustrated nicely.

⁸¹ <http://www.exfo.com/en/Library/WaveReview/WRArticle.aspx?Id=116>

⁸² http://old.hki.uni-koeln.de/people/herrmann/forschung/heydegger_archiving2008_40.pdf



Figure 15 Two JPEG2000 images, both with the same degree of corruption (one single byte); the second image shows no visual difference to the rendered version of the original uncorrupted file (not illustrated) although there are actual changes in pixel data. Reproduced from “Analysing the Impact of File Formats on Data Integrity” by Volker Heydegger, 2008.

As shown in Figure 15, even a single byte of corruption, if it occurs in the wrong place, can render an entire image completely useless. Heydegger has developed a ‘robustness indicator’ which provides what is essentially a measure of how a corruption of the encoded image translates to deviations in the data when the image is decoded. Since small changes to the encoded image data can result in very large changes to the decoded image, the robustness indicator is in a sense a measure of the extent to which data corruption is ‘amplified’ or ‘magnified’ through the use of encoding. Here the results are worrying to say the least. For JPEG2000 images, which are by no means unique in their characteristics, a one byte corruption (approx 0.001% of the data) of the encoded image resulted on average in a 17% change in the data for the decoded image for lossless compression, rising to 33% for lossy compression depending on the compression ratio. How this high level of data deviation from an uncorrupted image translates into visible artefacts is very variable, e.g. as shown by Figure 15, but it is clear from Heydegger’s work that corruption levels of 0.001% or lower to compressed JPEG2000 images (lossy or lossless) can have catastrophic consequences. Sadly, as shown in the section below where ORF report on their studies of data corruption of other image formats, JPEG2000 is not unique in this characteristic – indeed, as mentioned by Heydegger, JPEG2000 is better than average in its ability to cope with low levels of data corruption.

In Heydegger’s 2009 paper⁸³ “Just One Bit in a Million: On the Effects of Data Corruption in Files”, Heydegger expands on his 2008 work and studies how a range of image formats (jpeg, gif, tiff, bmp etc.) respond to bit level corruption, i.e. ‘bit rot’. Whilst Heydegger concentrates on image formats rather than video formats, it is worth noting that the majority of ‘preservation grade’ video formats use intra-frame encoding only, i.e. they are in effect a series of single images. Examples include MPEG2 D10, the DV family, JPEG2000. Therefore, Heydegger’s work is an indicator of what can be expected for preservation video formats.

There are several important findings within Heydegger’s work.

⁸³ http://www.hki.uni-koeln.de/files/heydegger_ecdl2009_camera_final.pdf

- Compressed formats are much more sensitive to corruption than uncompressed formats (Heydegger's observation no.3). This applies to almost all formats studied and data corruption investigated. JPEG2000 with resilience features 'turned on' could only achieve a level of robustness that matched the worst case of the uncompressed formats studied.
- Uncompressed formats are also sensitive to data corruption. Although better than compressed formats, there is still an 'amplification' effect when using uncompressed formats. The percentage of pixels in the decoded image that have been affected by corruption is higher than the percentage of the data in the encoded file that has been corrupted. More investigation is needed to determine how this translates into visible artefacts.
- The level of compression is not necessarily correlated with the sensitivity to data corruption; for example lossless JPEG2000 was found to be less robust than some levels of JPEG2000 lossy compression.

From the work of Heydegger and others we draw the following conclusions

- Considerable further work is needed to investigate how data corruption of preservation video formats translates into visible artefacts.
- Encoding, in particular use of compression, is a data corruption amplifier. The corruption of just one byte to an image file can cause whole frame to be completely useless.
- Compressed video formats are likely to be much more susceptible to data corruption in storage than uncompressed formats based on the evidence for single image formats.
- For inter-frame encoded video formats the problem is likely to be a lot worse as the 'amplification' effect is not just spatial (confined to one frame) but temporal (extends across multiple frames. This is a major problem as error concealment techniques have only a limited ability to repair temporal errors.
- For intra-frame encoded video formats, provided the level of corruption is low, e.g. so that only a few frames in several thousand are corrupted, then there is a good chance that error concealment techniques (e.g. motion compensated interpolation between frames) can at least make the content usable again, if not restore the original content perfectly.

Example 2: Artefacts in Audio-Files (ORF)

To have a first view on the robustness of different Audio-codecs for both production and preservation issues and to have some first experience on the impact of data-loss on byte-level, some basic tests on the impact of brute-force corruption of different audio-file qualities has been undertaken.

The audio-files have been edited by a simple Hex-Editor (HxD V1.7.5.0); the corrupted files have been checked on playback-capabilities via different audio programmes and

players (Steinberg® Wavelab™ 5.1; Microsoft® Windows™ Media Player 10.2; foobar 0.9.6). To demonstrate the impact also via printed media, spectrum analysis (FFT) has been used in Wavelab™ to produce pictures showing the artefacts and corruptions.

Some results of the following tests will be shown here:

1. Files violated by randomly deleting x bytes 4 times (1/2/3/4 bytes)
2. Files violated by deleting a big coherent block of x bytes (16/32/64/128 bytes)

The originals have been produced by digitizing an analogue original vinyl record (Satie, Trois Gymnopedies No.2 Lent et triste) via an ESI® Juli@™ Soundcard, EMT 950 Discplayer and Steinberg® Wavelab™ 5.1

Format	Details	File-Size in KByte
BWF PCM	96 kHz, 32 bit, Stereo, uncompressed	119.631
BWF PCM	48 kHz, 24 bit, Stereo, uncompressed	33.646
BWF PCM	44.1 kHz, 16 bit, Stereo, uncompressed	27.478
MP2 (Musicam)	48 kHz, 256 kbps, Stereo, compressed	4.986
MP3 VBR	44.1 kHz, 32-192 kbps, JointStereo, VBR, compressed	2.282
WMA	44.1 kHz, 112kbps, Stereo, CBR, compressed	2.685

General results :

Byte-violation by deletion causes significant problems on both compressed AND uncompressed files; especially when the deleted word-length is NOT divisible by the word-length used during digitization. Files compressed with a variable bit-rate are most vulnerable to all sorts of byte-violation. It's interesting, that the files compressed by constant bitrates are more stable and robust to byte-violation than uncompressed PCM-files; uncompressed PCM-files are extremely vulnerable, leaving virtual "dead" and unusable files.

Experiments on data corruption (altering, but not deleting) bytes has not yet been done and is an important area for future work. Data loss is typical when transmitting data, e.g. packet loss over networks, but data corruption is more likely in storage systems.

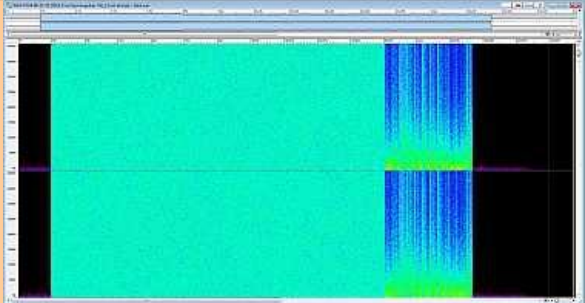
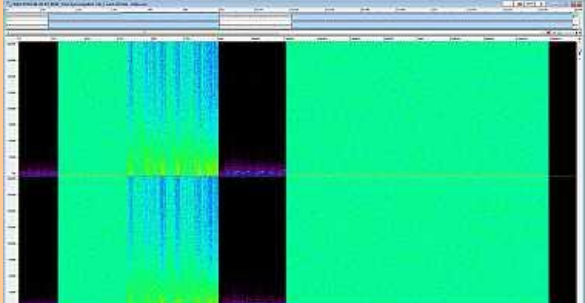
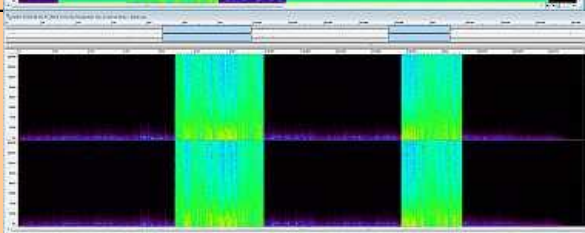
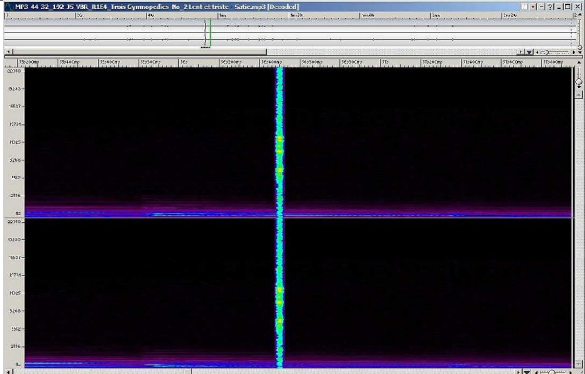
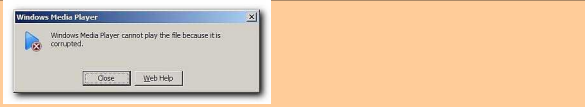
The results in detail (excerpts only, showing typical results):

No artefacts audible/visible

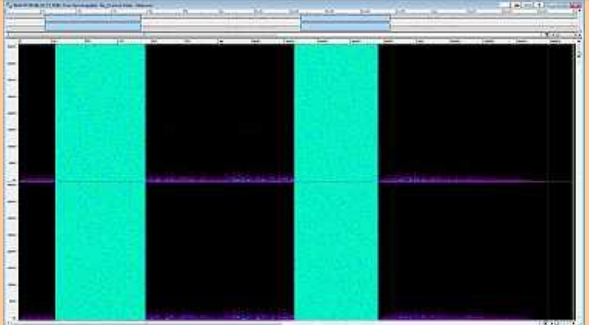
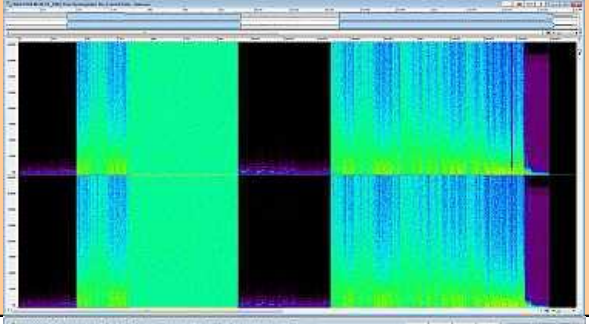
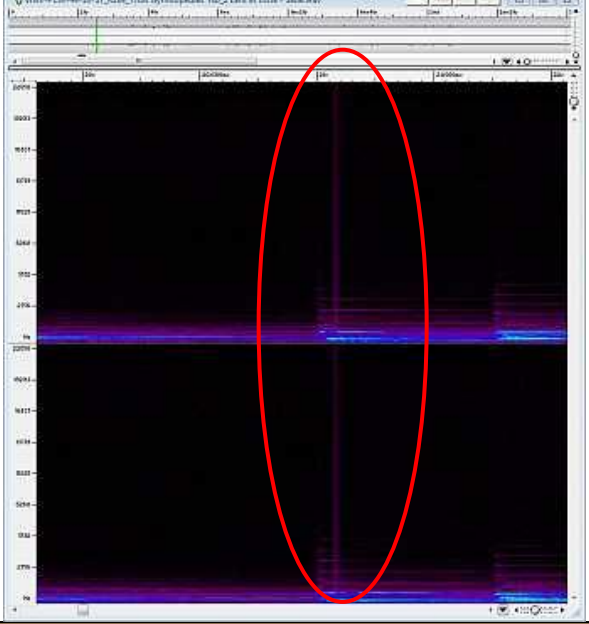
File can't be opened or is virtually destroyed

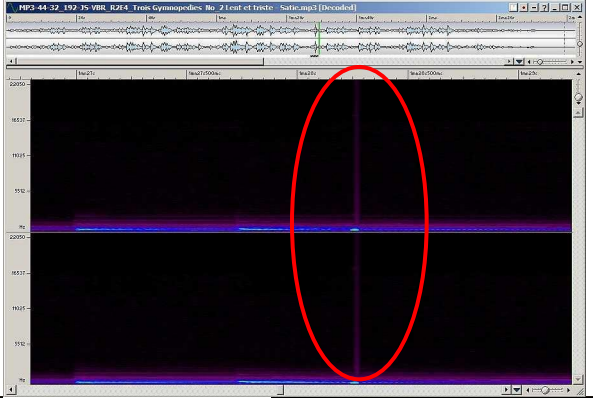
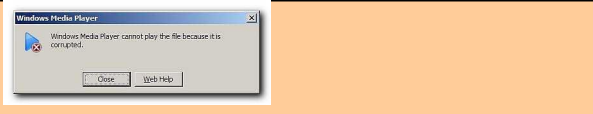
Syntax:

A) Random deletion of block (1 byte long); 4 times

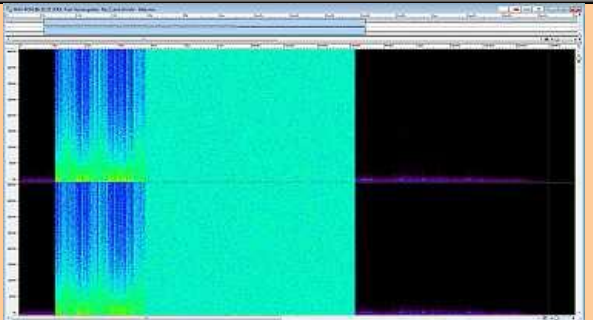
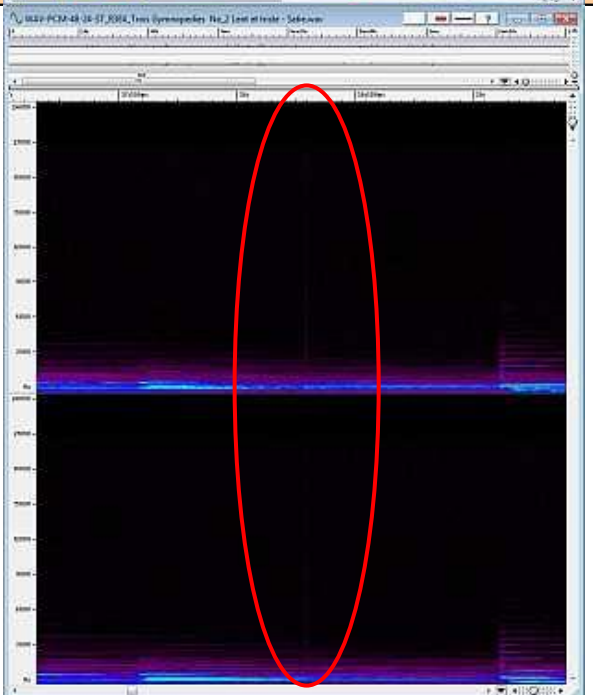
Format	Details	Random 1Byte Block-Error / Start	Result
BWF PCM	96 kHz, 32 bit, Stereo, uncompressed	007312E0	
BWF PCM	48 kHz, 24 bit, Stereo, uncompressed	0034A3B0	
BWF PCM	44.1 kHz, 16 bit, Stereo, uncompressed	007742A0	
MP2 (Musicam)	48 kHz, 256 kbps, Stereo, compressed	00071640	Can't open file / Playback stops
MP3 VBR	44.1 kHz, 32-192 kbps, JointStereo, VBR, compressed	00058980	
WMA	44.1 kHz, 112kbps, Stereo, CBR, compressed	00050C60	

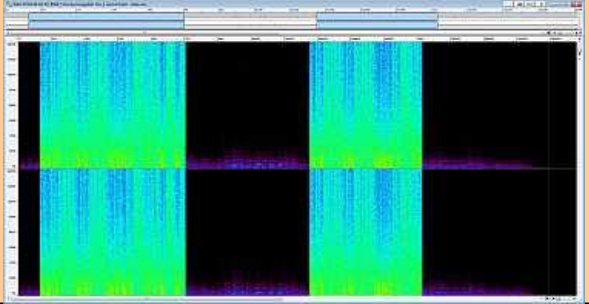
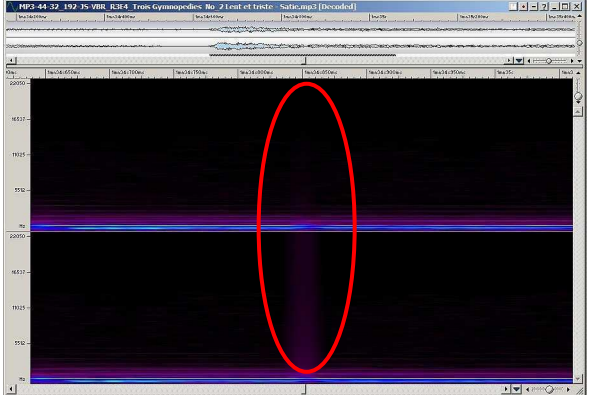

B) Random deletion of block (2 bytes long); 4 times

Format	Details	Random 1Byte Block-Error / Start	Result
BWF PCM	96 kHz, 32 bit, Stereo, uncompressed	00812510	
BWF PCM	48 kHz, 24 bit, Stereo, uncompressed	004CDAB0	
BWF PCM	44.1 kHz, 16 bit, Stereo, uncompressed	001DA800	
MP2 (Musicam)	48 kHz, 256 kbps, Stereo, compressed	00107390	Can't open file / Playback stops

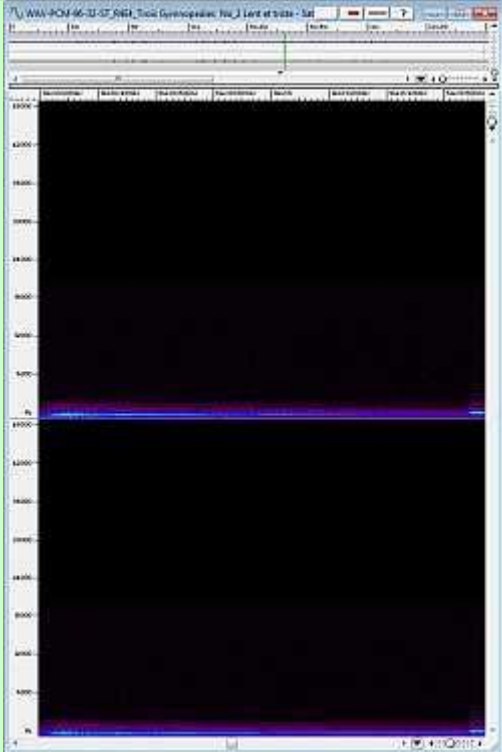
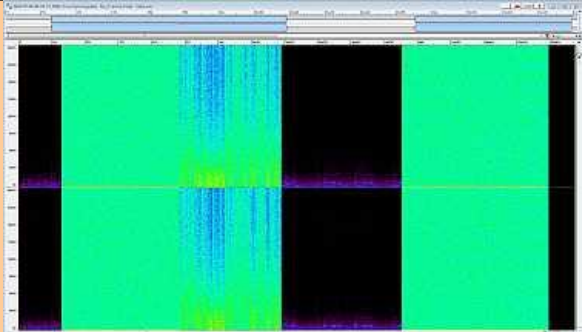
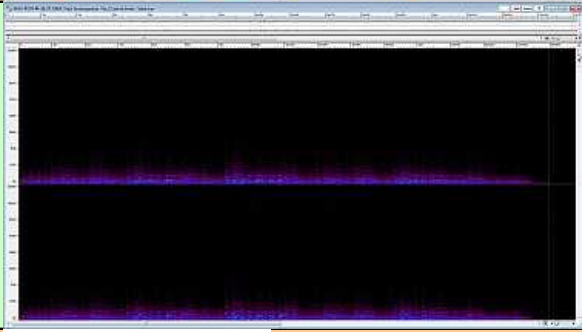

Format	Details	Random 1Byte Block-Error / Start	Result
MP3 VBR	44.1 kHz, 32-192 kbps, JointStereo, VBR, compressed	000821B0	
WMA	44.1 kHz, 112kbps, Stereo, CBR, compressed	0005FC90	

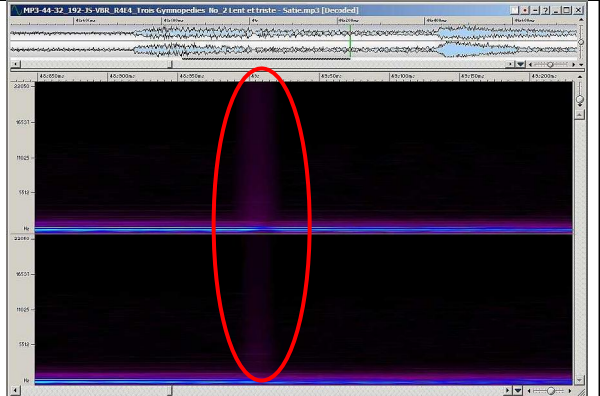
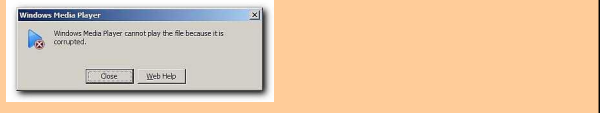
C) Random deletion of block (3 bytes long); 4 times

Format	Details	Random 1Byte Block-Error / Start	Result
BWF PCM	96 kHz, 32 bit, Stereo, uncompressed	007ER440	
BWF PCM	48 kHz, 24 bit, Stereo, uncompressed	0034A440	

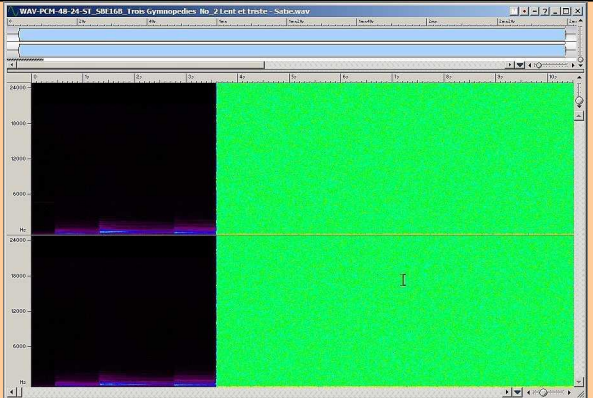
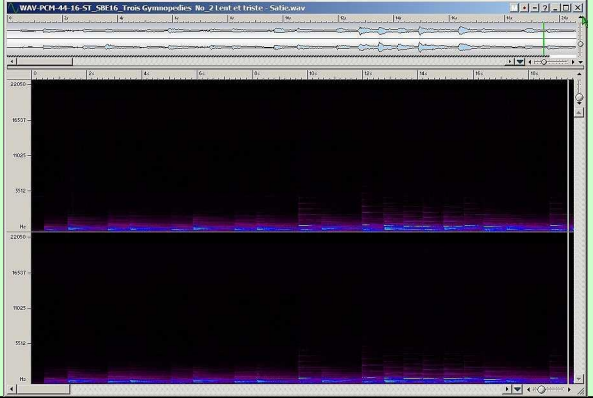
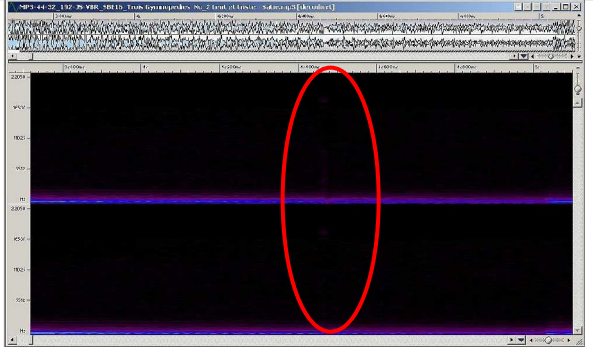

Format	Details	Random 1Byte Block-Error / Start	Result
BWF PCM	44.1 kHz, 16 bit, Stereo, uncompressed	00119910	
MP2 (Musicam)	48 kHz, 256 kbps, Stereo, compressed	000AE580	Can't open file / Playback stops
MP3 VBR	44.1 kHz, 32-192 kbps, JointStereo, VBR, compressed	0005EF80	
WMA	44.1 kHz, 112kbps, Stereo, CBR, compressed	00060BE0	

D) Random deletion of block (4 bytes long); 4 times

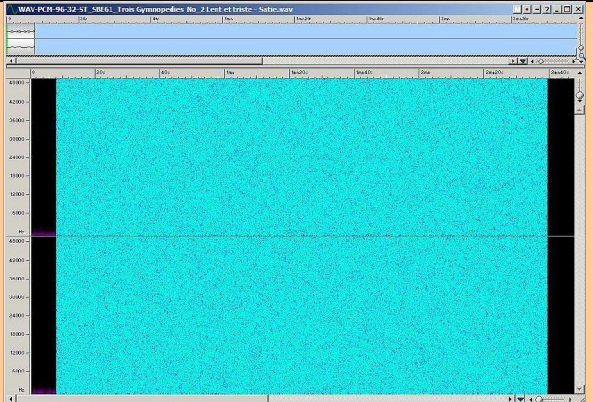
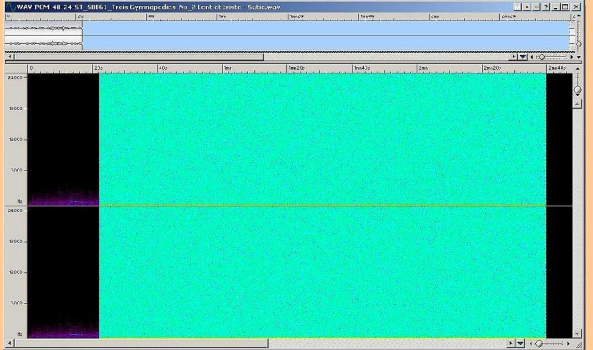
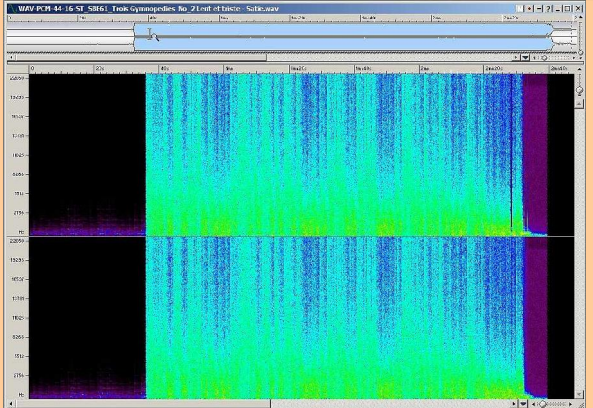
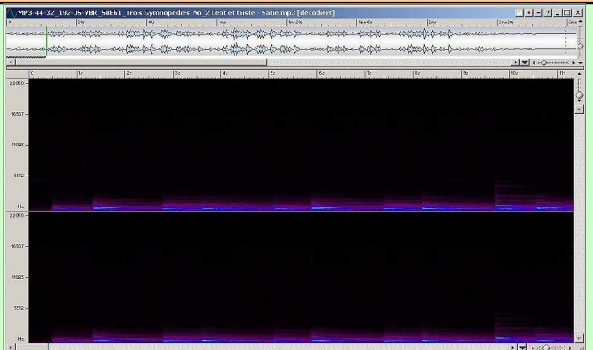

Format	Details	Random 1Byte Block-Error / Start	Result
BWF PCM	96 kHz, 32 bit, Stereo, uncompressed	00627FB0	
BWF PCM	48 kHz, 24 bit, Stereo, uncompressed	0038DA00	
BWF PCM	44.1 kHz, 16 bit, Stereo, uncompressed	00192370	
MP2 (Musicam)	48 kHz, 256 kbps, Stereo, compressed	000E7DB0	

MP3 VBR	44.1 kHz, 32-192 kbps, JointStereo, VBR, compressed	000AF970	
WMA	44.1 kHz, 112kbps, Stereo, CBR, compressed	000574F0	

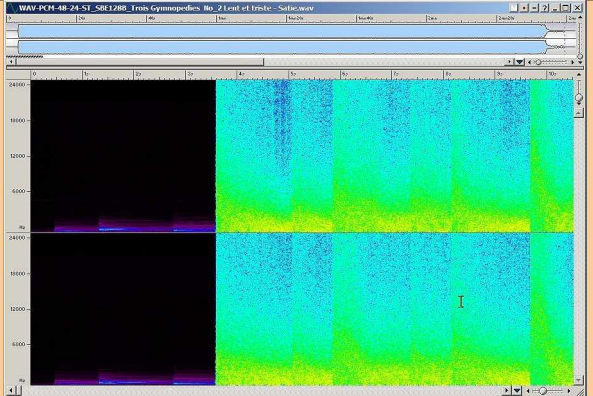
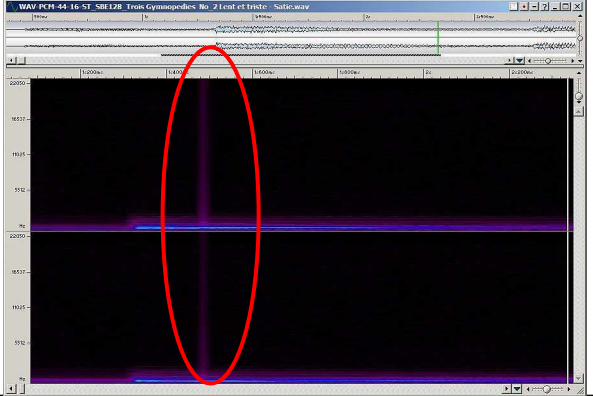
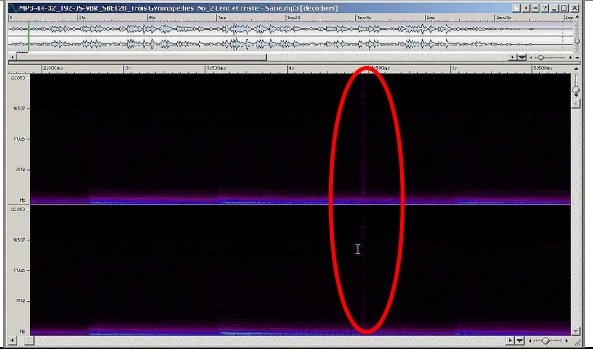
E) 1 big block (16bytes long) deleted at start position hex 00040000

Format	Details	16Byte Block-Error / Start	Result
BWF PCM	96 kHz, 32 bit, Stereo, uncompressed	00040000	Can't open file
BWF PCM	48 kHz, 24 bit, Stereo, uncompressed	00040000	
BWF PCM	44.1 kHz, 16 bit, Stereo, uncompressed	00040000	
MP2 (Musicam)	48 kHz, 256 kbps, Stereo, compressed	00040000	Can't open file / Playback stops (Depends on Playback-Software)
MP3 VBR	44.1 kHz, 32-192 kbps, JointStereo, VBR, compressed	00040000	
WMA	44.1 kHz, 112kbps, Stereo, CBR, compressed	00040000	

F) 1 big block (61bytes long) deleted at start position hex 00100000

Format	Details	16Byte Block-Error / Start	Result
BWF PCM	96 kHz, 32 bit, Stereo, uncompressed	00100000	
BWF PCM	48 kHz, 24 bit, Stereo, uncompressed	00100000	
BWF PCM	44.1 kHz, 16 bit, Stereo, uncompressed	00100000	
MP2 (Musicam)	48 kHz, 256 kbps, Stereo, compressed	00100000	Can't open file / Playback stops
MP3 VBR	44.1 kHz, 32-192 kbps, JointStereo, VBR, compressed	00100000	
WMA	44.1 kHz, 112kbps, Stereo, CBR, compressed	00100000	

G) 1 big block (128bytes long) deleted at start position hex 00610000

Format	Details	16Byte Block-Error / Start	Result
BWF PCM	96 kHz, 32 bit, Stereo, uncompressed	00610000	Can't open file
BWF PCM	48 kHz, 24 bit, Stereo, uncompressed	00610000	
BWF PCM	44.1 kHz, 16 bit, Stereo, uncompressed	00610000	
MP2 (Musicam)	48 kHz, 256 kbps, Stereo, compressed	00610000	Can't open file / Playback stops (Depends on Playback-Software)
MP3 VBR	44.1 kHz, 32-192 kbps, JointStereo, VBR, compressed	00610000	
WMA	44.1 kHz, 112kbps, Stereo, CBR, compressed	00610000	Can't open file / Playback stops (Depends on Playback-Software)

Example 2: Artefacts in Picture-Files (ORF)

During an internal decision-process on file-formats for long-time-preservation of photographics the quality- pro's and -con's of the different formats and codecs has been discussed; in addition the learn more about the robustness of the different formats and codecs, a small comparative test has been performed by violating a sample-photo-file in

different formats and qualities by deleting and altering parts of the file. For the editing of the files a simple Hex-Editor (HxD V1.7.7.0) has been used; the violated files has been controlled and monitored with two different standard graphic programmes (Adobe® Photoshop CS4; IrfanView 4.10); for the graphic examples in the following lists the same programmes has been used to produce the thumbnails and details.

The results of three of the tests will be shown here:

3. Files violated by deleting a big coherent block of x bytes (16/32/64/128 bytes)
4. Files violated by periodically deleting a block of 3 bytes x times (4/8/16 times)
5. Files violated by randomly deleting 1 byte x times (4/8/16 times)



The original file produced by shooting a RAW-pictures with a semi-professional SLR-Camera (Canon ® 40D) and using the RAW-file (Canon CR2) as basis for producing the following files to be tested:

Format	Details	File-Size in KByte
Canon RAW2	Original RAW Data from SLR	9.709
TIFF	16 bit, uncompressed	59.076
TIFF	16bit, ZIP-Compression	47.934
BMP	8bit, uncompressed	29.525
PNG	8bit	22.723
DNG	compressed	8.749
JPEG	Q100	5.021
JPEG	Q100, progressive	4.726
JPEG	Q50	557
JPEG	Q50, progressive	275
GIF	interlaced	4.664

General results:




- The length and position of the data loss, and the encoding used, including compression, all have a bearing on the affects of data loss.


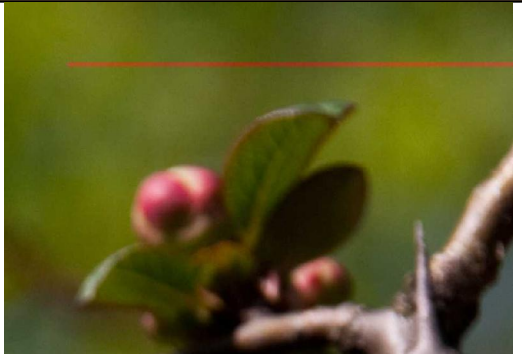


As with the audio data loss experiments presented above, experiments on data corruption (altering, but not deleting) bytes has not yet been done and is an important area for future work. Data loss is typical when transmitting data, e.g. packet loss over networks, but data corruption is more likely in storage systems.

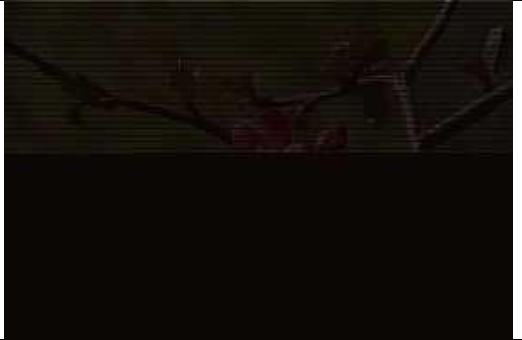
The results in

No artefacts visible
File can't be opened
Syntax:

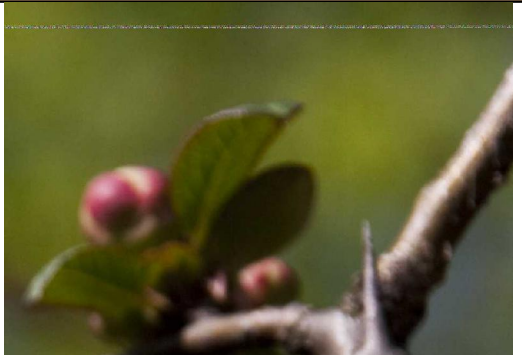
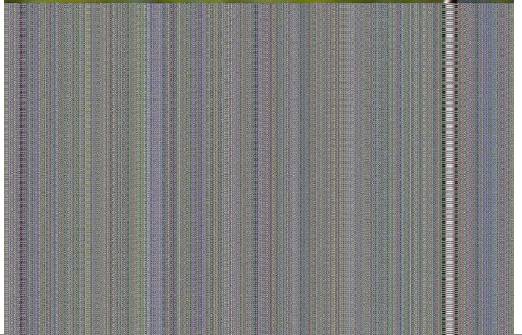

A) 1 big block (16bytes long) deleted at start position hex 00040000






Format	Details	16Byte Block-Error / Start	Result
CR2		00040000	Can't open file
TIFF	16 bit, uncompr.	00040000	
TIFF	16bit, ZIP-Compr.	00040000	
BMP	8bit, uncompr.	00040000	


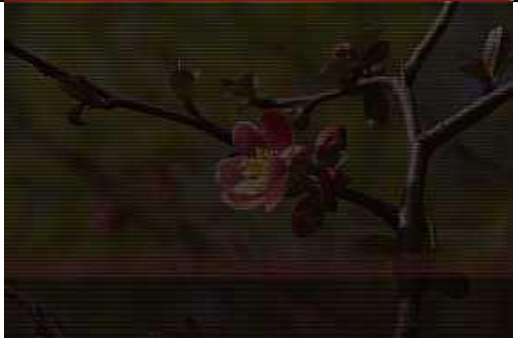
PNG	8bit	00040000	
DNG	compressed	00040000	Can't open file
JPEG	Q100	00040000	
JPEG	Q100, progr.	00040000	
JPEG	Q50	00040000	
JPEG	Q50, progr.	00040000	

GIF	interlaced	00040000	
-----	------------	----------	--

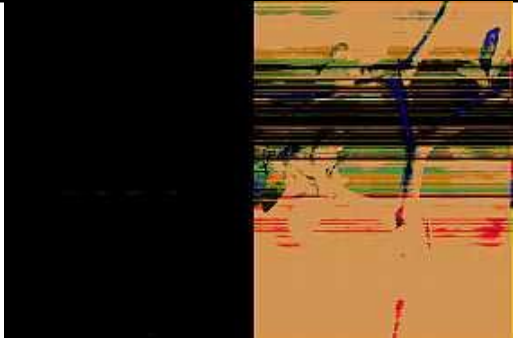
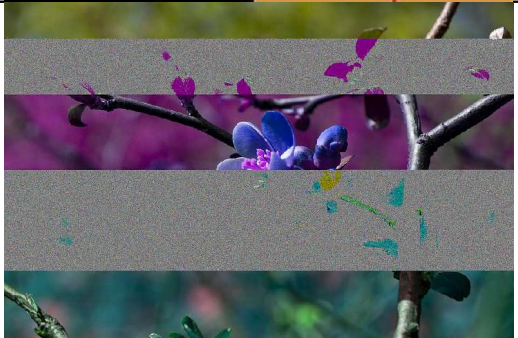
B) Periodical deletion of block (3bytes long); 4 times, step hex 00007000


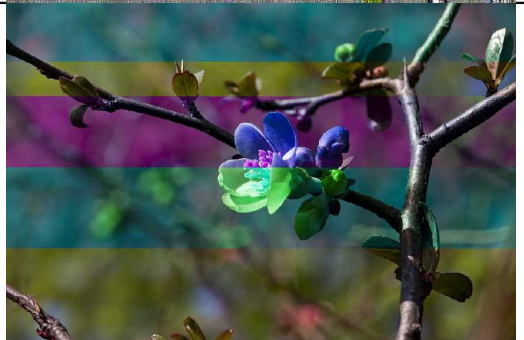
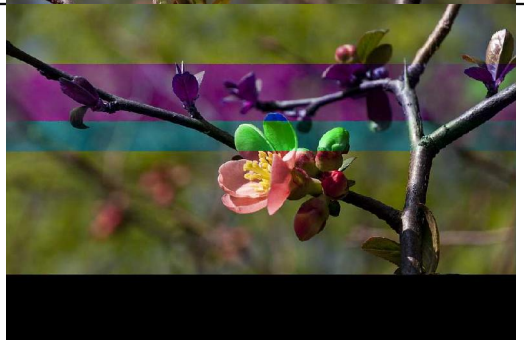

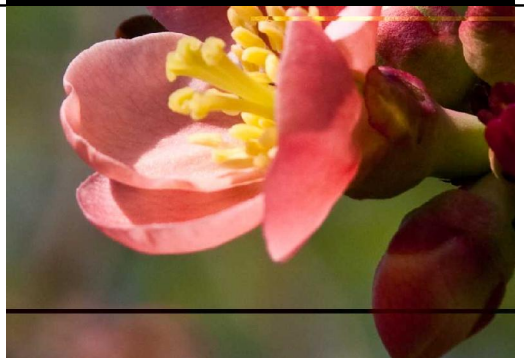
Format	Details	Periodical 3Byte Block-Error / Start	Result
CR2		00100000	Can't open file
TIFF	16 bit, uncompr.	00100000	
TIFF	16bit, ZIP-Compr.	00100000	
BMP	8bit, uncompr.	00100000	




PNG	8bit	00100000	
DNG	compressed	00100000	
JPEG	Q100	00100000	
JPEG	Q100, progr.	00100000	
JPEG	Q50	00040000	

JPEG	Q50, progr.	00001000	
GIF	interlaced	00100000	

C) Random deletion of block (1byte long); 4 times

Format	Details	Periodical 3Byte Block-Error / Start	Result
CR2			
TIFF	16 bit, uncompr.		

TIFF	16bit, ZIP-Compr.		
BMP	8bit, uncompr.		
PNG	8bit		
DNG	compressed		
JPEG	Q100		

JPEG	Q100, progr.	
JPEG	Q50	Only thumbnail-picture left ...
JPEG	Q50, progr.	
GIF	interlaced	

4.2 Summary

Data corruption does happen when using IT storage with catastrophic effects on AV files.

- Data corruption causes major problems for AV content in file format. This applies to video, images and audio alike. In the worst case, which is not infrequent, files simply can't be opened or played by their respective applications.
- Considerable further work is needed to investigate how data corruption of preservation video formats translates into visible artefacts.
- The use of encoding, in particular compression, can massively amplify even low levels of data corruption, or result in whole files becoming useless.
 - For single images (and intra-frame encoded video) a single byte of corruption to a compressed image can render the whole image completely useless. The sensitivity to data corruption is not correlated to the level of compression, e.g. lossless JPEG2000 is just as sensitive to data corruption as lossy compression.

- For corruption of intra-frame encoded video there is at least a possibility to use concealment techniques, e.g. interpolation between adjacent frames, to correct the effects of data corruption – provided that the number of frames affected in a sequence is low. The same is unlikely to be true for inter-frame encoded video due to the temporal propagation of errors between frames.
- For audio, major audible artefacts can be generated and persist well beyond the temporal location where data loss first takes place. Files compressed with a variable bit-rate are most vulnerable. Files compressed by constant bitrates are more stable and robust to corruption. PCM-files are extremely vulnerable to data loss, resulting in unusable content in almost all cases.
- More investigation is needed for both audio and video formats, in particular video that uses inter-frame encoding. The expectation is that all compressed formats in common use are not likely to be at all robust to data corruption, even at low levels of data corruption.
- Compressed formats are in general much more sensitive to data corruption than uncompressed formats. Due to the ‘amplification’ effect that compression has on data corruption, the percentage saving in storage space is often much less than the percentage increase in the amount of information that is affected by data corruption.

5. File Quality-control in current use

The quality control processes that are currently in place in a range of audiovisual archives were surveyed to build a picture of current practice. This indicates the areas that archives are currently finding the most important to concentrate on and perhaps in areas where little QC activity takes place it can indicate a lack of awareness of some of the issues. Full details of the archives surveyed can be found in Annex 3: Quality Control case studies

- The BBC uses no fully- or semi-automated processes so far; nearly all the quality checks and controlling is done manually. These include quality review by manual inspection of the results of the production process as well as the use of technical checks, e.g. peak programme meters or for problems such as Harding flashing which can be dangerous to viewers with photo sensitive epilepsy. In the BBC's D3 project, automated error logging is used as part of the transfer process to flag up areas where there may need to be subsequent manual inspection of the programme material.
- INA employs quality control checks during digitisation, ingest and migration. The measures used include content checking using tools, e.g. for defect identification, and manual spot-checks by operators. Documentation is checked for completeness and structure. Importantly, checksums are used during the migration process to detect corruption (INA have performed several large-scale migrations at the petabyte scale between tape robots).
- B&G does most of its quality control checks against metadata, in particular MXF datastructures, during content ingest. The content itself, having typically been already broadcast, is not checked directly for quality problems. The MXF structure is checked for consistency with the essence it wraps by using tools that compare files to be ingest against templates, e.g. for D10-30/50 or XDCamHD.
- In ORF's planned new content management and storage system, a wide range of Quality Control checks are planned. These include syntactic checks of file formats, e.g. wrapper or encoding compliance to standards as well as checks on the content itself, e.g. blockiness or audio silence. Checks are done at various stages of the process, e.g. production prior to ingest, during ingest and then again during any migrations that take place within the archive. The plan is to automate the QC process as far as possible.

Comparing the content QC processes of the broadcast and archive partners in the PrestoPRIME consortium reveals that there is not a common set of tools, processes or techniques in place. Each archive has its own approach, with some more developed and automated than others. Much of the QC processes that are in place focus on ensuring the quality and standards of content admitted into the archive, for example identification and checking of file formats (wrappers, metadata, video, audio) against standards at the syntactic level and also checking content (often manually) for visual or audible quality problems (e.g. during digitisation, transfer or format migration) of the essence. Less attention is paid to proactively monitoring data integrity when the content is within the archive. This is partly because many archives still operate a 'items on shelves' model and the bulk of their content is not yet in digital file form. However, this is changing and several

of the archives surveyed recognise the need to review and further develop their QC processes in this area.

Annex 1: OCTAVE Allegro worksheets

Allegro Worksheet 1	RISK MEASUREMENT CRITERIA- REPUTATION AND CUSTOMER CONFIDENCE		
Impact Area	Low	Moderate	High
<i>Reputation</i>	Reputation is minimally affected; little or no effort or expense is required to recover.	Reputation is damaged, and some effort and expense is required to recover.	Reputation is irrevocably destroyed or damaged.
<i>Customer loss</i>	Less than _____% reduction in customers due to loss of confidence	_____ to _____% reduction in customers due to loss of confidence	More than _____% reduction in customers due to loss of confidence
<i>Other:</i>			

Allegro Worksheet 2	RISK MEASUREMENT CRITERIA- FINANCIAL		
Impact Area	Low	Moderate	High
<i>Operating Costs</i>	Increase of less than _____% in yearly operating costs	Yearly operating costs increase by _____ to _____%.	Yearly operating costs increase by more than _____%.
<i>Revenue loss</i>	Less than _____% yearly revenue loss	_____ to _____% yearly revenue loss	Greater than _____% yearly revenue loss
<i>One-Time Financial Loss</i>	One-time financial cost of less than \$ _____	One-time financial cost of \$ _____ to \$ _____	One-time financial cost greater than \$ _____
<i>Other:</i>			

Allegro Worksheet 3	RISK MEASUREMENT CRITERIA- PRODUCTIVITY		
Impact Area	Low	Moderate	High
<i>Staff Hours</i>	Staff work hours are increased by less than _____% for _____ to _____ day(s).	Staff work hours are increased between _____% and _____% for _____ to _____ day(s).	Staff work hours are increased by greater than _____% for _____ to _____ day(s).
<i>Other:</i>			

Allegro Worksheet 4	RISK MEASUREMENT CRITERIA- SAFETY and HEALTH		
Impact Area	Low	Moderate	High
<i>Life</i>	No loss or significant threat to customers' or staff members' lives	Customers' or staff members' lives are threatened, but they will recover after receiving medical treatment.	Loss of customers' or staff members' lives
<i>Health</i>	Minimal, immediately treatable degradation in customers' or staff members' health with recovery within four days	Temporary or recoverable impairment of customers' or staff members' health	Permanent impairment of significant aspects of customers' or staff members' health
<i>Safety</i>	Safety questioned	Safety affected	Safety violated
<i>Other:</i>			

Allegro Worksheet 5	RISK MEASUREMENT CRITERIA- FINES AND LEGAL PENALTIES		
Impact Area	Low	Moderate	High
<i>Fines</i>	Fines less than \$_____ are levied.	Fines between \$_____ and \$_____ are levied.	Fines greater than \$_____ are levied.
<i>Lawsuits</i>	Non-frivolous lawsuit or lawsuits less than \$_____ are filed against the organization, or frivolous lawsuit(s) are filed against the	Non-frivolous lawsuit or lawsuits between \$_____ and \$_____ are filed against the organization	Non-frivolous lawsuit or lawsuits greater than \$_____ are filed against the organization.
<i>Investigations</i>	No queries from government or other investigative organizations	Government or other investigative organization requests information or records (low profile).	Government or other investigative organization initiates a highprofile, in-depth investigation into organizational practices.
<i>Other:</i>			

Allegro Worksheet 6	RISK MEASUREMENT CRITERIA- USER DEFINED		
Impact Area	Low	Moderate	High
<i>Other:</i>			

Allegro Worksheet 7		IMPACT AREA PRIORITIZATION WORKSHEET
PRIORITY	IMPACT AREAS	
	Reputation and Customer Confidence	
	Financial	
	Productivity	
	Safety and Health	
	Fines and Legal Penalties	
	User Defined	

Allegro Worksheet 8		CRITICAL INFORMATION ASSET PROFILE
(1) Critical Asset <i>What is the critical information asset?</i>	(2) Rationale for Selection <i>Why is this information asset important to the organization?</i>	(3) Description <i>What is the agreed-upon description of this information asset?</i>
(4) Owner(s) <i>Who owns this information asset?</i>		
(5) Security Requirements <i>What are the security requirements for this information asset?</i>		
<input type="checkbox"/> Confidentiality	Only authorized personnel can view this information asset, as follows:	
<input type="checkbox"/> Integrity	Only authorized personnel can modify this information asset, as follows:	

<input type="checkbox"/> Availability	This asset must be available for these personnel to do their jobs, as follows:	
	This asset must be available for _____ hours, _____ days/week, _____ weeks/year.	
<input type="checkbox"/> Other	This asset has special regulatory compliance protection requirements, as follows:	

(6) Most Important Security Requirement

What is the most important security requirement for this information asset?

<input type="checkbox"/> Confidentiality	<input type="checkbox"/> Integrity	<input type="checkbox"/> Availability	<input type="checkbox"/> Other
---	---	--	---------------------------------------

Allegro Worksheet 9a	INFORMATION ASSET RISK ENVIRONMENT MAP (TECHNICAL)	
INTERNAL		
CONTAINER DESCRIPTION	OWNER(s)	
1.		
EXTERNAL		
CONTAINER DESCRIPTION	OWNER(s)	
1.		

Allegro Worksheet 9b		INFORMATION ASSET RISK ENVIRONMENT MAP (PHYSICAL)	
INTERNAL			
CONTAINER DESCRIPTION		OWNER(s)	
1.			
2.			
3.			
EXTERNAL			
CONTAINER DESCRIPTION		OWNER(s)	
1.			
2.			

Allegro Worksheet 9c		INFORMATION ASSET RISK ENVIRONMENT MAP (PEOPLE)	
INTERNAL PERSONNEL			
NAME OR ROLE/RESPONSIBILITY		DEPARTMENT OR UNIT	
1.			
2.			
3.			
EXTERNAL PERSONNEL			
CONTRACTOR, VENDOR, ETC.		ORGANIZATION	
1.			
2.			

Allegro Worksheet 10		INFORMATION ASSET RISK WORKSHEET			
Information Asset Risk	Threat	Information asset			
		Area of concern			
		(1) Actor <i>Who would exploit the area of concern or threat?</i>			
		(2) Means <i>How would the actor do it? What would they do?</i>			
		(3) Motive <i>What is the actor's reason for doing it?</i>			
		(4) Outcome <i>What would be the resulting effect on the information asset?</i>	<input type="checkbox"/> Disclosure	<input type="checkbox"/> Destruction	
			<input type="checkbox"/> Modification	<input type="checkbox"/> Interruption	
		(5) Security Requirements <i>How would the information asset's security requirements be breached?</i>			
	(6) Probability <i>What is the likelihood that this threat scenario could occur?</i>	<input type="checkbox"/> High	<input type="checkbox"/> Medium	<input type="checkbox"/> Low	
		(7) Consequences <i>What are the consequences to the organization or the information asset owner as a result of the outcome and breach of security requirements?</i>		(8) Severity <i>How severe are these consequences to the organization or asset owner by impact area?</i>	
			Impact Area	Value	Score
			Reputation & Customer Confidence		
			Financial		
			Productivity		
			Safety & Health		
		Fines & Legal Penalties			
		User Defined Impact Area			
Relative Risk Score					
(9) Risk Mitigation <i>Based on the total score for this risk, what action will you take?</i>					
<input type="checkbox"/> Accept		<input type="checkbox"/> Defer		<input type="checkbox"/> Mitigate	
For the risks that you decide to mitigate, perform the following:					
<i>On what container would you apply controls?</i>		<i>What administrative, technical, and physical controls would you apply on this container? What residual risk would still be accepted by the organization?</i>			

Annex 2: OCTAVE Allegro Risk Analysis Example

This section contains a hypothetical example of an organisation whose business is to provide archive hosting services.

Imagine that this business hosts AV content owned by others and that maintaining the integrity of the content over the long-term is one of the main values of the service to the customers.

The objective of this section is to show the process used for Allegro risk assessment in the context of data integrity.

The organisation first prioritizes impact areas according to its business strategy. This is done by numbering the areas from the least important (1) till the most important (n) as the following:

Allegro Worksheet 7	IMPACT AREA PRIORITIZATION WORKSHEET
PRIORITY	IMPACT AREAS
5	Reputation and Customer Confidence
4	Financial
2	Productivity
1	Safety and Health
3	Fines and Legal Penalties
n/a	User Defined

For each critical asset the following risk analysis steps are then required.

The “Critical Information Asset Profile” contains the description of the asset and the security requirements. Finally, only the most important security requirement(s) is considered to be analysed.

Allegro Worksheet 8		CRITICAL INFORMATION ASSET PROFILE	
(1) Critical Asset <i>What is the critical information asset?</i>		(2) Rationale for Selection <i>Why is this information asset important to the organization?</i>	
Digital AV Data Repository		This is all the digital audiovisual data deposited in the archive by our customers.	
(3) Description <i>What is the agreed-upon description of this information asset?</i>			
SafeStore			
(4) Owner(s) <i>Who owns this information asset?</i>			
Us (archive service providing organisation)			
(5) Security Requirements <i>What are the security requirements for this information asset?</i>			
<input checked="" type="checkbox"/> Confidentiality	Only authorized personnel can view the assets in the repository.	Only customers who have supplied AV data and have paid for its preservation should	
<input checked="" type="checkbox"/> Integrity	Only authorized personnel can modify the assets in the repository.	Only repository administrators should be able to modify data.	
<input checked="" type="checkbox"/> Availability	This asset must be available for customers to do their jobs, as follows:	Customers must be able to access their data at any time of the day or night with the exception of Sunday which is maintenance day	
	This asset must be available for 24 hours per day, 6 days per week, 52 weeks per year.		
<input type="checkbox"/> Other	This asset has special regulatory compliance protection requirements, as follows:		
(6) Most Important Security Requirement <i>What is the most important security requirement for this information asset?</i>			
<input type="checkbox"/> Confidentiality	<input checked="" type="checkbox"/> Integrity	<input type="checkbox"/> Availability	<input type="checkbox"/> Other

In the three following sheets the organisation identifies the containers of the information asset. OCTAVE Allegro considers the containers where the asset is stored, transported and processed. The containers may be technical, physical containers or people. These containers may be internal or external at another organisation.

This analysis enables identifying the boundaries of the threat environment and the scope of the risk assessment.

Allegro Worksheet 9a		INFORMATION ASSET RISK ENVIRONMENT MAP (TECHNICAL)	
INTERNAL			
CONTAINER DESCRIPTION		OWNER(s)	
1.The data in the repository is stored on a local primary storage server, and replicated on an external storage provider. The customers get their data from the primary server.		Us	
2.Archive Service Platform: The Platform offers services for discovery and accessibility of data stored in the repository.		us	
EXTERNAL			
CONTAINER DESCRIPTION		OWNER(s)	
1.Internet: The data is accessed online. Data is sent on the internet.		unknown	
2.Storage space at an external provider is used to store copies of the data for safety and disaster recovery.		Amazon S3	

Allegro Worksheet 9b		INFORMATION ASSET RISK ENVIRONMENT MAP (PHYSICAL)	
INTERNAL			
CONTAINER DESCRIPTION		OWNER(s)	
1.Tape copies of the data are created on a regular basis and used for local backup.		Us (IT dept.)	
EXTERNAL			
CONTAINER DESCRIPTION		OWNER(s)	

Allegro Worksheet 9c		INFORMATION ASSET RISK ENVIRONMENT MAP (PEOPLE)
INTERNAL PERSONNEL		
NAME OR ROLE/RESPONSIBILITY	DEPARTMENT OR UNIT	
1. Archive service managers: responsible for the services made available to the customers	Customer service	
2. Technical staff: responsible for managing the platform and ensuring the storage has sufficient capacity.	IT dept.	
EXTERNAL PERSONNEL		
CONTRACTOR, VENDOR, ETC.	ORGANIZATION	
1 Customers supply the data stored in the repository	Various	
2. Storage service providers license space to us for storing remote copies of the data	Amazon S3	

After analysing the different containers the organisation now has an idea about the environment where threats may originate. In the “information asset worksheet”, the next step is to identify the threats and the consequences.

Severity parameters (section 8 of the sheet) are calculated for each impact area. Severity (high, medium, low) semantics are identified in worksheets 1, 2, 3, 4 and 5. The estimated severity value (high=3, medium=2, low=1) is multiplied by the organisation impact area priority (worksheet 7) to obtain the score. This relative risk score enables to classify risks and prioritize them in the context of organisation’s mission and business objective.

For example, if reputation is most important to an organisation, risks that have an impact on the organisation reputation will generate higher scores than risks with equivalent impacts in another area. For each area of concern a sheet should be filled.

Allegro Worksheet 10		INFORMATION ASSET RISK WORKSHEET			
Information Asset Risk	Threat	Information asset	Digital AV Data Repository		
		Area of concern	Data is altered when unauthorised individual gains access to the service.		
		(1) Actor <i>Who would exploit the area of concern</i>	Current Employees		
		(2) Means <i>How would the actor do it? What would they do?</i>	Access the remote copy at the storage service provide over the Internet. Access the primary copy using internal systems when at work.		
		(3) Motive <i>What is the actor's reason for doing it?</i>	Accidentally (e.g. screw up a maintenance process) or deliberately (e.g. because of a dispute with management).		
		(4) Outcome <i>What would be the resulting effect on the information asset?</i>	<input type="checkbox"/> Disclosure	<input type="checkbox"/> Destruction	
			<input checked="" type="checkbox"/> Modification	<input type="checkbox"/> Interruption	
		(5) Security Requirements <i>How would the information asset's security requirements be breached?</i>	Only specific authorised staff should be able to modify the data.		
	(6) Probability <i>What is the likelihood that this threat scenario could occur?</i>	<input type="checkbox"/> High	<input checked="" type="checkbox"/> Medium	<input type="checkbox"/> Low	
	(7) Consequences <i>What are the consequences to the organization or the information asset owner as a result of the outcome and breach of security requirements?</i>	(8) Severity <i>How severe are these consequences to the organization or asset owner by impact area?</i>			
			Impact Area	Value	Score
	The reputation of the organisation is badly affected.	Reputation & Customer Confidence	High	15	
		Financial	Med	8	
	This modification will have impact on our relation with the content provider. This may include fines and legal penalties.	Productivity	Low	3	
		Safety & Health	Low	1	
It requires effort to get the data back to its original state	Fines & Legal Penalties	Med	4		
	User Defined Impact Area	n/a			
Relative Risk Score				31	

The next step is to indicate a mitigation approach and process.

At each contain, the organisation indicates the controls to mitigate the indicated risk. We normally choose controls to reduce the likelihood that a threat happens but we should not ignore other measures which reduce the impact severity level as well. Finally, residual risk may still exist. However; this should be of an acceptable level for the organisation.

(9) Risk Mitigation	
<i>Based on the total score for this risk, what action will you take?</i>	
<input type="checkbox"/> Accept	<input type="checkbox"/> Defer
<input checked="" type="checkbox"/> Mitigate	<input type="checkbox"/> Transfer
For the risks that you decide to mitigate, perform the following:	
<i>On what container would you apply controls?</i>	<i>What administrative, technical, and physical controls would you apply on this container? What residual risk would still be accepted by the organization?</i>
Primary Storage Server	Access control should be enforced to only allow authorised users having the archive manager role to modify data. Logging is required for accountability and recovery purposes.
Remote Storage Services	Access control should be enforced to allow only authorised users from the organisation to access and modify the data at the storage provider. Logging is required for accountability and recovery purposes.
Tape copies	Tape copies need to be protected and kept in a safe place. Any security code or keys for physical access should only be available for authorised staff.

Annex 3: Quality Control case studies

BBC

BBC uses no fully- or semi-automated processes so far; nearly all the quality checks and controlling is done manually.

Final Quality Checks (Technical Review) at BBC-Scotland⁸⁴

For example, the instruction manual used at BBC-Scotland, where file based working has been established supposedly, says on the matter of checking studio-recordings on files:

“There is no fool-proof way of checking that a high resolution version of any studio recording is being created. It's the same, whether you're recording on file or tape. Tapes offer some reassurance that a recording is taking place (but that's all). Nor can you check the high resolution at your desktop.

There is no real reason why a file recording would not be created, but until the technology is proven and while confidence in file-based working builds, a back-up tape recording will be made in Media Central of all studio output.”

BBC-Scotland has a dedicated Multi Media Review Area (MMRA) for carrying out technical reviews of completed programmes in high or standard definition. Programmes can be reviewed in the MMRA as files which have been created on the CPS or as tapes:

For Programmes delivered on file, a craft editor or dubbing mixer copies the completed sequence to the MMRA folder on an Avid Interplay system. A Post Production Operator takes the file in the MMRA folder and replays it, via an Airspeed machine, to assess for technical compliance using a variety of high quality monitoring and assessment tools:

- High definition widescreen picture monitors for visual assessment of picture quality
- Dolby 5.1 surround sound system for aural assessment of sound quality
- Waveform monitor for colour line up, luminance etc.
- Peak programme meters (PPM) for sound levels and phase coherence
- Safe area indicator to ensure captions etc. are viewable on equipment of various types and aspect ratios
- Harding flashing pattern analyser (FPA) to identify sequences categorised by OFCOM⁸⁵ as dangerous to viewers suffering from Photo Sensitive Epilepsy (PSE)

During the tech review a tape backup copy of the programme will be made (on digital betacam), along with a DVD and/or VHS copy if required for subtitling. The digibeta backup will be spot checked and retained in the Tape Vault by Media Management as a physical backup of the programme.

⁸⁴ Pacific Quay, Glasgow

⁸⁵ UK communications regulator; <http://www.ofcom.org.uk>

If the programme passes its tech review, and the editorial review is also flagged, the file will be moved into a "Transmission" folder on Avid® Interplay⁸⁶. This triggers the transfer process into the Digital Library, under the control of the Media Manager, and, ultimately, onto Transmission Payout.

A programme which fails its MMRA will be moved into a "Failed" folder.

Although the production of a programme is fully file-based, the complete quality-checking is done in the common "programme-clearing"-process.

Spot Checking of Audio-Productions

Also Audio-files used in BBC's radio-broadcasts are not checked by dedicated automated or semi-automated routines (beneath those used by the different storage-units; like ECCs⁸⁷ on Harddisks, SSL⁸⁸ and IPsec⁸⁹ on Networks, etc.); all further checks are done on the content only by spot-checking the signal via listening.

Identification of critical areas by using D3⁹⁰-replay logs

The only major investment BBC made in automatic checking, is in adapting the hardware of the D3 players, so that all the read errors in the D3 machine can be detected and logged. That log is then used to guide the manual checking of the resultant file.

The QC-part of the D3-Preservation is described as follows:

- Takes place after the transfer of the D3-signal to MXF⁹¹-OP1a⁹² on LTO3⁹³
- MXF-Files are extracted from LTO3 to local storage (Harddisks)
- D3 replay error logs (from serial port of D3-playback unit) are used to identify critical parts/spots
- Areas around those parts/spots are checked manually
- Finally a full visual check is performed

ORF

Specifications for QC in P-CMS and on Storage-entry⁹⁴

The specifications for the new P-CMS & Storage-system and the incorporated workflows at ORF have a very elaborated part on QC; below the main specifications are listed:

⁸⁶ Production Asset Management System by AVID®

⁸⁷ Error correction code

⁸⁸ Secure Sockets Layer; cryptographic protocol providing security for communications over networks

⁸⁹ Internet Protocol Security; protocol suite for securing IP-communications

⁹⁰ D3 digital composite video tape-cassette format (8bit uncompressed 4fsc PAL video / 4 x 20bit 48kHz digital audio)

⁹¹ Material eXchange Format

⁹² Operational pattern of MXF; SMPTE 378M

⁹³ 3rd generation of LTO (Linear Tape-Open; magnetic tape data storage technology)

⁹⁴ Production Content Management System

- Primary Goal: all stored files have to match predefined specifications; a check for output to miscellaneous target-systems is possible.
- Therefore templates for all wrappers/codecs in use are existing and the according use is done automatically; new templates can be generated.
- Standard Point of time for QC: Input/Output of file(s) into/from the P-CMS
- Additional point of time for QC: transmutation, transcoding
- Syntactical Control:
 - Wrapper-Control
 - MXF (SMPTE 377M) compatibility
 - Generic Container (SMPTE 379M) compatibility
 - Operational Pattern
 - SDTI-CP⁹⁵ (SMPTE 326M) compatibility
 - Clip or Frame based wrapping
 - Status of streamable flag (2nd flag of Byte 15 in the OP UL)
 - Syntax of KLV⁹⁶ structure
 - Status of MXF Files (open/closed, complete/incomplete)
 - Valid "Duration" in Metadata
 - Elements in the Header Partition
 - Elements in the Body Partition
 - Constancy of the Edit units per Body Partition
 - Number of Body Partitions
 - Regularity of Body Partitions
 - Number of Essence Tracks
 - Validity of Timecode Tracks (EBU Rec.122)
 - Validity of Index Tables
 - Valid Header Metadata Repetitions
 - Updated Metadata in Footer Partition
 - Validity of RIP⁹⁷
 - Read-Out of DMS-1 Data
 - Read-Out of File, Picture, Sound and Data Essence Descriptors
 - Codec-Control (Video & Audio)
 - D-10⁹⁸ compatibility (SMPTE386M + SMPTE356M + EBU D94-2002)
 - XDCAM⁹⁹ HD-422¹⁰⁰ compatibility (ISO/IEC 13818 + SMPTE 381M)
 - AVC-I¹⁰¹ compatibility (RP2008 + IOS/IEC 14496-10)
 - JPEG2000 compatibility (SMPTE422M + ISO/IEC 15444)
 - VC-3¹⁰² compatibility (SMPTE2019-4 + 2019-1)
 - Bitrate
 - Maximum of coded Frame Size
 - AES¹⁰³ (SMPTE 382M) compatibility
 - 8-channel AES (SMPTE 331M) compatibility
 - BWF¹⁰⁴ compressed & uncompressed (SMPTE 382M) compatibility

⁹⁵ Serial Data Transport Interface Content Package

⁹⁶ Key-Length-Value; data encoding standard

⁹⁷ Routing Information Protocol

⁹⁸ MPEG-2 based video compression format

⁹⁹ tapeless professional video system

¹⁰⁰ 3rd generation of XDCAM, using 4:2:2 profile of MPEG-2

¹⁰¹ AVC-Intra; video codec for production quality HD

¹⁰² SMPTE standard

¹⁰³ Standard for digital audio

¹⁰⁴ Broadcast Wave Format

- Video-Control
 - Videoformat
 - legal Videolevel
 - legal Colour space
 - VBI¹⁰⁵ and Ancillary Data (SMPTE 436M)
 - AFD¹⁰⁶ informations (SMPTE 2016)
- Video-Control
 - Validity of Syntax (MPEG PES¹⁰⁷, ES¹⁰⁸ or Transport Streams on ETR 101290 or Standard-conformity)
 - Dolby-E¹⁰⁹ compliance (for Dolby E-streams)
 - Sample Rate
 - Channel-togetherness (Mono, Dual Mono, Stereo, Joint Stereo)
- Semantic Control
 - Video & Audio
 - Black Frames (configurable black level for “x” % of frame-content)
 - Blockiness (configurable amount of block-dimension for “x” % of frame-content)
 - Audio silence (configurable signal-threshold for “x” Audio-channels)

The checks will be used in this basic QC-rule-framework:

- Complete check
 - During or prior to Production-review-process (including complete manual-visual checks by technical supervisor and production-master)
 - For newly ingested material (including raw-material, external productions, etc.)
 - On entry to P-CMS
 - On entry to Storage System
 - For migrated files (from DiMi-System¹¹⁰)
 - On entry to Storage System (depending on load-surveys, this may be partly shifted into the DiMi-System)
- Basic checks (parts of Wrapper- and Codec-control)
 - For rushes (only in use since complete checks can be done in (near) real-time)
 - On entry to P-CMS
 - For “reviewed” and broadcasted files/productions
 - On entry to Storage System

The whole QC-process will be fully automated (watch-folder, etc.); for faulty files certain thresholds will be introduced to decide on the further proceedings (e.g. all data below thresholds = proceed; some/one data above threshold = decision needed on further proceeding (manual intervention by content-manager); many data above threshold = abort and decline proceeding (info to order-system / manual intervention needed).

¹⁰⁵ vertical blanking interval

¹⁰⁶ Active Format Description

¹⁰⁷ Packetized Elementary Stream; specification in MPEG-2

¹⁰⁸ Elementary Stream; usually the output of a av-encoder

¹⁰⁹ Audio encoding and decoding technology

¹¹⁰ Digital Migration System-framework at ORF

B&G

At sound and vision several checks are done during ingest.

An application (MXF-Checker) written by Technicolor (former NOB) does this check and shows compliancy with our format yes/no During the building of the archive we certified broadcast to deliver files based on test files delivered by this broadcaster. (so we certified broadcaster A delivering files coming from software/platform X) At this moment, these broadcasters are still certified and we do not check the files during ingest. Main reason for this is the fact that these files are broadcasted before they are automatically added to the archive.

Any file problems will be found within the playout environment and be solved before playout can occur. That way we are sure the files entering the archive are compliant to our standards. However, other files (for example, generated by images of the future) are checked one by one.

The base of this application is 'IRT Analyzer Pro' which will generate a XML file with all sort of information regarding the MXF and its contents. This output XML is compared to a profile within the MXF-Checker application and the output of the application tells us that the file is compliant with our archive yes or no The profiles contain information about video format, colour depth, metadata, audio format etc. which has to be on the right place in the MXF file to ensure compliancy. At this moment we have profiles for D10-30/50 and XDCamHD.

INA

The following checks are done at INA

1) Digitisation :

Control of digitisation results:

- Audio : all files are verified using CubeTec Dobbin/Quadriga; reported errors are verified by an operator, and correction/comparison with original/re-digitisation is required if appropriate.
- Video : Random sampling and verification is made by an operator.
- Film : Random sampling and verification is made by an operator. This is done on physically repaired films, and on telecine tapes (Digibeta).

2) Ingest :

Encoding (Digibeta -> MPEG2 files):

- Every day samples are taken and verified by an operator, main objective is to detect problems (head-clogging, malfunction) before physical media are returned to the remote vaults.

File ingestion :

- A specific software tool is used to verify that a set of technical and documentation parameters are set right : file names, all required files present (browse, broadcast),

same length, right technical parameters, accessibility (a documentary record has to exist).

3) Maintenance :

When the files are in the system, no specific verification is made, other than implicit maintenance by the HSM, and reaction to problems when happen. A more proactive approach is being considered.

4) Migration :

There was once a complete migration from the previous data tapes and robots to a new one (it took 6 months to migrate half a Petabyte). Checksums were used to detect corruption.