

In this issue:

storage technology
file formats and standards

media asset management tools
digitisation systems

February 2014



PRESTO
CENTRE

AV Digitisation and Digital Preservation
TechWatch Report

#01



PrestoCentre TechWatch Reports are published twice to three times a year except for special issues. Statements of fact and opinion are the responsibility of the authors alone and do not imply an opinion on the part of the officers or other members of PrestoCentre.

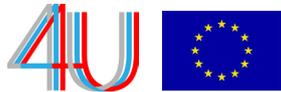
Published by the Presto4U project, 2014. All rights reserved.
Materials may not be reproduced without written permission. For information and contributions, contact us at +31 20 894 3570 or +1 347 404 5337; editor@prestocentre.org.

PrestoCentre Foundation
PO Box 1060, 1200 BB Hilversum, The Netherlands
Telephone: +31 20 894 3570 or +1 347 404 5337
Editorial offices: editor@prestocentre.org

Editorial Team
Marius Snyders, Mirella Bulsink

Contribution to this edition
Werner Bailer (Joanneum Research), Simon Factor (Moving Media), Peter Schallauer (Joanneum Research), Paul Walland (University of Southampton IT Innovation Centre)

Presto4U is a two-year project supported by a core network of 14 PrestoCentre members.
The Presto4U project receives funding from the European Commission's Seventh Framework Programme (FP7).



Foreword

PrestoCentre's TechWatch Reports are about identifying technology trends and business issues which exist in digital AV archiving and finding a way to bring clarity in a language that is accessible to non-specialists.

PrestoCentre wants to be a strategic partner who connects the preservation challenges and needs of audiovisual archives to the solutions available from industry and research, guiding audiovisual media owners through change and helping them shape up for the future. We also help vendors in the domain by assessing the impact of solutions on archives at different stages of development, by analysing relevant standardisation activities, and identifying areas for new work in research and innovation.

This first TechWatch Report has been written by members of PrestoCentre involved in the Presto4U project and was compiled through meetings they had with specialist technology vendors and researchers late 2013. This TechWatch is the first of what will become a regular report, published twice to three times a year. In the next report we will be reviewing some of the technology on view at the NAB show in April. We will also be providing an update on the evaluation of AV Archiving tools being assessed through the Presto4U project in the first half of 2014.

Thank you for your interest in PrestoCentre's TechWatch. Please feel free to follow us on Twitter @prestocentre and provide any feedback or thoughts you may have.

Table of contents

1. Storage Technology	05
2. File Formats and Standards	07
3. Media Asset Management Technology	08
4. Digitisation Systems	09

1. Storage Technology

During the past decade, much debate has been entered into regarding the choice of tape or disk for long-term storage of file based audiovisual assets. The debate began with cost comparison, but as disk decreased in price this became less of an issue. However the primary reason for choosing disk remains around speed of access, which, for most archives excluding 'production' archives, is not an issue. The raw cost of tape versus disk does not consider the cost of the software and services required in utilising data tape, and this has been an issue for many users in recent years. The proprietary nature of how software applications such as Media Asset Management systems (MAMs) wrote to tape meant that users were tied to the Independent Software Vendor (ISV) who supplied the MAM or other software system that interfaced with the storage, thus creating a whole new set of challenges for archives. This resulted in additional cost and reliance in the event of the system (or vendor) failing, or when it was time to upgrade the application stack / condense the archive as newer higher capacity tape generations were released.

The LTO forum was set up in the early 2000s by a number of LTO manufacturers to create an industry platform for LTO development. At that time there was a wide range of data tape formats available like SAIT, DLT and DDS. The forum was created to allow many companies to provide the same tape cartridge. Now, a decade later, the LTO forum has adopted LTF5 within the programme, which is good news for archives. LTF5 is an open standard for writing data to LTO tape through a standards POSIX interface; it opens up a range of possibilities to archives and media distributors as to how we can use LTO. With LTF5, LTO becomes a distribution format: We can write and send tapes to other facilities without having to worry about the software systems used to read or write the tape.

In the Presto4U project we have been testing the LTF5Archiver software; a copy of the software is available on our website for you to download.¹ It's free. The goal of the software is to effectively manage the storage of generic files and it is optimised for working with multimedia (big, order of several GBytes) files. In our evaluation, we gave the tool a TRL² score of 4-6, which generally means it is available in demo form but will require some integration and development to run. LTF5Archiver has reached good scoring in nearly all the characteristics considered. Its principal point of strength is its simplicity of use and maintainability, its good interoperability via use of the LTF5 standard and low consumption

of resources. LTF5Archiver has been developed with the support and input of a team at one of our member organisations, the RAI Technology Centre in Torino in Italy, who will be using the system in an upcoming project which will see the digitisation of a very large collection of BetaCam Video tapes and storage of the resulting files using the LTF5Archiver software.

So data tape (perhaps now more than ever, with the addition of LTF5) offers a value proposition that is very relevant to the AV Archiving sector through low cost and high bandwidth access, but not necessarily random access. According to the LTO forum representatives we spoke to at the International Broadcasting Convention 2013 (IBC),³ they plan to release new generations of LTO for the foreseeable future, in keeping with the 2 year refresh rate they have been on since the early 2000s; they do not see any headroom issues in terms of the capacity and data rate of future product releases.

It does seem that the boundaries of storage density for spinning disk is starting to plateau at this time, and storage technology companies are beginning to try new ways of increasing the density of disk storage to stay on track with Moore's prediction. Among those we have come across in 2013 are HAMR which stands for Heat Assisted Magnetic Recording, which augments the write head on a disk drive with a laser allowing for higher density storage. Seagate have been particularly active in R&D in this area and in March 2012 became the first hard drive maker to achieve the milestone storage density of 1 terabit per square inch using HAMR technology. Another approach on trend in 2014 with storage resellers is SMR – Shingled Magnetic Recording which involves writing to a disk sequentially (ironically similar how we write data to tape!), which allows for improved linear bit density and can increase the storage available by up to 15%.

1. <https://www.prestocentre.org/library/tools/ltf5-archiver>

2. NASA defines Technology Readiness Level [Mankins, 1995] as: "Technology Readiness Levels (TRLs) are a systematic metric/measurement system that supports assessments of the maturity of a particular technology and the consistent comparison of maturity between different types of technology."

3. IBC is an annual event for professionals engaged in the creation, management and delivery of entertainment and news content worldwide. Attracting 50,000+ attendees from more than 170 countries, IBC combines a highly respected and peer-reviewed conference with an exhibition that exhibits more than 1,400 leading suppliers of state of the art electronic media technology and provides very good networking opportunities. <http://www.ibc.org/>

Somewhere on the outskirts of the great tape versus disk debate we find interesting alternative solutions like Sony's Optical Disc Storage and Disc Archiving Corporation's ALTO system. These alternative systems take a slightly different approach to storage and each has its own unique features and value proposition.

The ALTO II system from DAC is a hybrid storage system which uses disk drives to store data but only spins these disks up when a media file is required for retrieval or when the system is performing health checks. By keeping the disk array dormant most of the time, the system does not draw the energy or require the environmental management of 'spinning' disk. Essentially, ALTO II treats disk like tape — it is designed to store large data files for infrequent access.

The operational benefit is increased speed of access to large files within a media production environment. DAC promote their products as taking just 3 seconds to spin up the disk that the required data is on and that their system can support many concurrent users. Within a system using LTO, the amount of tape drives available in the system being less than the amount of concurrent users can be a bottleneck.

The main benefit of this approach when compared to disk is the lower running cost of the system. This approach has been tried in the past by companies such as COPAN Systems who were an early pioneer of this approach, which they called MAID Technology (Massive Array of Idle Disks) in the early 2000s but there was not enough demand for them to build a business and the assets of the company were acquired by SGI in 2010.

Sony have in the past year introduced a range of upgrades and new components to its Optical Disc Archive System. This uses a proprietary Sony cartridge that contains an optical disc to store data. Sony promotes their Optical Discs as having a 50-year lifespan and the latest generation of cartridges launched in 2013 can be read at a rate of over 1Gbps. The cartridges can be loaded into a desktop reader drive or kept within a robotic library system; Sony's library system is called 'PetaSite'. The value proposition here is for long-term reliable storage combined with high-speed access from a library. No LTO manufacturer will promote its media as being good for 50 years, however a long term digital archiving strategy is unlikely to be reliant upon a single carrier for such a long time.

The move to the 'cloud' has had a lot of lip service in the past year but the general feeling among vendors and archive data users is that the term 'cloud' is too broad and can mean a wide range of things depending on the application. Consumer cloud products are not really relevant to long-term data storage — they are more to do with on-demand computational horsepower for running web applications. Cloud storage offers have emerged from companies such as Amazon with their Glacier product, however the business case is not clear for large AV archives. In most cases, cloud storage as offered by some vendors is more of a managed data storage service. Think of it as a safe corner in a data centre with your name on it. This sort of storage is also still mainly to support online delivery of 'product' files, for example streaming media to websites or via download services, not for storage of preservation quality master files.

Matthew Addis of Arkivum wrote a good post on cloud archiving and Amazon's Glacier specifically.⁴ And while you're at it, why not read his other posts on how storage always involves compromises between safety, cost, performance and other factors.⁵

4. <https://www.prestocentre.org/blog/glacier-causes-global-warming>
 5. <https://www.prestocentre.org/blog/203>



Exhibit 1: One thing we can say about the ALTO II system is it looks pretty fetching in bright Yellow!

2. File Formats and Standards

Lossless compression of audiovisual data is a very relevant technology for long-term preservation, providing storage savings over uncompressed storage while still being able to reconstruct bitwise identical data. JPEG2000 lossless is one of the most common lossless compression technologies for video data, but has a rather high computational complexity, especially for the encoding.

Recently, alternatives have emerged that have lower computational complexity but achieve compression rates comparable to JPEG2000 lossless. One of them is FFV1, which has been created as part of the FFmpeg open source project. It has already been successfully used in preservation projects — see the extensive report in the December 2013 issue⁶ of AV Insider. A similar technology is the TICO codec from intoPIX, which can be scaled from lossless to visually lossless, having low computational complexity and decoding latency. While all these emerging codecs have very interesting features, one should not forget that there is a risk over well-documented formal standards (such as JPEG2000), which is the lack or incompleteness of information. As paradoxical as it may seem, this even holds for open source software. As long as an active developer community is around, it is easy to add support for new features and to support new platforms, but working on someone else's sparsely documented code can take significant time just to make a small fix.

Lossless compression of audiovisual data is a very relevant technology for long-term preservation, providing storage savings over uncompressed storage while still being able to reconstruct bitwise identical data.

FIMS, the Framework for Interoperable Media Services, is an effort by the Advanced Media Workflow Association (AMWA) and the European Broadcasting Union (EBU) to define standard service interfaces for media processing systems. FIMS is already becoming widely adopted in industry and will improve interoperability between components from different vendors and reduce the risk of vendor lock-in. There are some projects within FIMS which are quite relevant for preservation workflows: repository access, transform (i.e. transcoding or transwrapping of media files) and, most

recently, quality analysis. For quality analysis, FIMS is also working on a standard report format for quality analysis results and is collaborating with the EBU Quality Control (QC) group on the definition of quality checks.

At IBC2013, EBU QC group showed the first version of their so-called periodic table of QC elements. This set of elementary video QC tests and their parameters have been presented for the different layers of video QC, as there are wrapper, bit-stream and baseband. The goal of EBU QC group is to define a complete set of QC tests relevant for the use cases: ingest, exchange and delivery, and archive migration/digitisation. The periodic table is available online.⁷

At IBC2013, file based video quality control (QC) received a lot of attention. A major trend in the area of QC is that the integration of different QC engines. Companies like Amberfin and Cube-Tec showed the integration of multiple QC engines, with the goal of having the results of the different complementary engines in a single user interface, allowing for fast and efficient QC decisions. The solution offered by Cube-Tec integrates the VidiCert QC tools specialised on film and video migration/digitisation QC. File-based video QC is currently a hot topic in standardisation, as well as workflow integration for all typical steps in the video lifecycle, from production, delivery and exchange up to preservation.

6. <https://www.prestocentre.org/library/resources/av-insider-4-film-files-formats-and-future>

7. <https://tech.ebu.ch/webdav/site/tech/shared/qc/IBC2013qcperiodictable-final.pdf>

3. Media Asset Management Technology

Media Asset Management (MAM) is an on-going challenge for archives.

The process of digital migration is bringing some of the underlying challenges to light. There seem to be plenty of technology providers who have some or all of the solutions available, but from discussions at IBC2013 it appears that whilst some MAM providers include archive support in their thinking, not all of them realise (or care) that their technology is also useful for archives.

The intention here is to look at the key issues that need to be addressed in the process of storing, maintaining and successfully retrieving files. Firstly, the actual storage. There doesn't seem to be much of a problem finding space for all the data anymore, even though many people will quote storage as an issue, citing the amount of space needed to store a 2-hour, native format 4k movie (several terabytes, the exact value being a function of the point being made). As we have seen in the storage technology review in this same report, there will be plenty of storage options available. What is of more interest here is the reliability of the storage, the accessibility of the content, the metadata handling available and the search and retrieval capability.

Media Asset Management (MAM) is an on-going challenge for archives. The process of digital migration is bringing some of the underlying challenges to light.

Metadata is absolutely crucial. There are plenty of vendors who are interested in finding ways of annotating your ingested data files, mostly through some kind of ingest workflow solution which they offer. The choice is between an open standard (such as MXF) and a closed proprietary standard. The advantage of a complete package is that installation and support is simple, the disadvantage being that you are 'locked into' a single vendor solution, and export of metadata is difficult other than to other systems using the same schema. Many MAM suppliers are offering MXF based workflow solutions, which have the advantage of conforming to a widely accepted standard and allowing migration or interfacing with other storage systems. Since these workflows often incorporate quality assurance and

content analysis components it is beneficial to use metadata schemas that conform to open standards, since this provides both compatibility between systems and longevity. It's a sad fact of life that proprietary systems can disappear or change in a way that doesn't (or shouldn't) happen with standardised implementations. Of course, there will still be challenges of standards migration and ensuring adherence to standards, but there is at least a historical record to support regression for recovering past files.

Several vendors talk of using more complex profiles of MXF for content management. At present, the simplest profiles (called OP1a and OP-ATOM) are used, however MXF profiles such as OP1b and OP2a/b can be used in the future to aggregate more of the different components which comprise a media package. This is something which is being actively considered by workflow providers, but care will need to be taken to ensure backwards compatibility with older files.

Another big issue under discussion is emerging formats that will need to be stored in archives as media production standards evolve. The most commonly anticipated change is to Ultra-high definition television (UHDTV), which will demand more storage space, higher frame rates and more efficient compression formats. Archives will need to balance the demands on file size against the degree of compression, with mathematically lossless compression being a desirable, but perhaps unaffordable option. Compression standards giving 'visually lossless' results, such as JPEG, H265 or even the newly proposed TICO format from IntoPIX will need to be considered, and should be supported by any archive workflow.

4. Digitisation Systems

While MXF would seem to be a natural choice for archive applications, there seems to be so many variants of the standard that things need to be constrained. Application specific or constrained specifications are important to avoid future access issues. Digital Rapids provide a range of video ingest and encoding tools which include their StreamZ hardware encoders, which have been a popular choice for video digitisation projects and service providers for many years. At IBC2013 we could see their product development roadmap and hear about issues they see around the archiving space. While much of the digitisation of physical media archives requires hardware for the ingest of the video signal, Digital Rapids see a significant demand from archives who have previously digitised collections and are now seeking to update or migrate these digitised collections from an older digital format, e.g. one that may have presented as a good option at the time but has since become obsolete. An example of this would be an AVID OMF format; there is a demand to migrate files in this format to a format that they can better marry metadata to and that has a longer shelf life. JPEG2000 in an MXF wrapper is by far the most popular request that vendors are getting in terms of archival video formats. The main motivation is the mathematically lossless compression mode available and the fact that it is an intra frame format and therefore temporal artefacts are significantly reduced. It is a well-defined standard and while there may be nuances there is a sense that it is going to be useful in 10 years. MXF is popular as a wrapper due to the richness of metadata support that it offers — however the variance of how the metadata is written means that they want to see more standardisation in that regard.

AVC Intra is becoming popular and IMF (Interoperable Mastering Format) is playing into this as post-production companies and content producers adopt it. IMF is a work in progress with many of the US Movie studios at present, gaining traction as a “Mezzanine” production format. Here, the rationale is that they could use IMF as a long-term archival format so the standard is being developed with this in mind. IMF is popular because it boils down the metadata and compression format choices to a formula that includes production metadata from upstream, and allows interchange during the post-production and mastering process. AS11 was an initiative by the DPP (Digital Production Partnership) that sought to create a constrained standard for MXF. In this they can define video stream types for both SD and HD content, AVC Intra for HD and D10 for SD.

Another area of interest is what processing needs to be done around the digitisation of media, such as up-scaling or de-interlacing. A process such as de-interlacing can have a huge impact on the quality outcome of a digitisation process. With this in mind, Digital Rapids have put some effort into creating new algorithms for advanced de-interlacing to maximise visual fidelity.

One company that have been operating at the forefront of the audio digitisation space since 2000 is NOA, well known for their audio digitisation and workflow systems. This year they have expanded to the video space with the launch of their Video Ingest System.

Some of the interesting things about the NOA system include use of FFV1 Lossless Video Encoding and integration with VidCheck automatic video QC tools. The 2-channel system is designed for digitisation of SD formats; the system uses an SDI input, so requires a A/D converter to present an SDI signal to the ingest system. They have 9-pin control to the VTR for control and monitoring of the video machines including native RF readout to monitor signal quality. A plug-in from Harmonic (Carbon Coder) delivers multi format transcoding within the system for a range of video file format outputs.

NOA are using Lossless FFV1 as a video encoding format in their products for the master file capture format. They like this because it is an open source codec and free of cost, it is supported by a community and has free available player technology in VLC. They have tested the FFV1 format and have found that it has comparable compression ratios to the more expensive codecs such as JPEG2000 and, performance-wise, it is superior in its encoding speed (V3 can be multithreaded). This is their choice for a mathematically lossless archive format. FFV1 V3 was launched three weeks before the IBC2013 show. NOA want to see wider adoption of open source standards for video encoding. They feel that MXF could be a problem for some archives, as the range of attributes a file could have may mean that the format is not easy for a small archive to decode in the future, if the structure of the file is somewhat proprietary. NOA like the fact that LTFS is gaining traction because it represents an easy and cheap way to provide access to LTO storage. Cube-Tec, another company traditionally known for its archive solutions in the audio domain, are also working on developments in the video digitisation space. Cube-Tec



Exhibit 2: Apparently NOA now stands for 'Not Only Audio!'

have pioneered in audio since 1990 and have been at the forefront of the industry in recent decades, leading many standards and technical innovations. One such new technical innovation they were showing at IBC is their Video RF card, developed in association with the Fraunhofer Institute. The company also offers a video digitisation system, Quadriga Video. This latest innovation takes an entirely different approach to the traditional methods of getting a video signal from tape. Apart from the problem of decaying tape, we have the problem of decaying machines (and the folks who operate them!). This is the problem that to some extent Cube-Tec are trying to solve with 'RF Direct Transfer'. Their approach is a physical modification to an existing videotape machine, where a card is inserted into the machine which takes the RF Video signal directly from the video playback head and sends this out to a software process which then synthesises the signal path of the 'perfect' VTR.

The objective of this system is to make the video player circuitry redundant and provide an optimum video signal from tape that is not subject to any degeneration from the ageing tape playback machine's internal signal path. Clearly this is a novel and innovative approach, which is very interesting, but does it work and is it terribly expensive? From what we could gather, the hardware side is fully functional however until the tape machine signal path modelling systems are optimised in speed and quality

domains, a full commercial product is not on the market. However, we could see that this approach will become important as time moves on and we are forced to 'virtualize' ever more of the signal chain, removing the reliance upon ageing and hard to maintain legacy playback technology. The main area of R&D is in the modelling of the tape machine signal path. This development work is still ongoing and the two companies are working together to deliver a solution that will provide a playback result superior than that possible through an ageing VTR.

Cubetech's MXF Legalizer provides file-based, large-scale QC and repair of files that are damaged in the container or the essence stream, and software to automate this process. The company is solving issues regarding standards compliance and compatibility for the MXF file container structure. They do not so much look at the image quality or detect faults in the image, however they do go as far as to correct image and audio quantisation rates if this is a requirement.

PrestoCentre's TechWatch Reports are about identifying the technology trends and business issues which exist in digital AV archiving and finding a way to bring clarity in a language that is accessible to non-specialists.

ISSN: 2352-2941