

In this issue:

What's new for film scanners
Video ingest systems at scale

Microservices and open source
Storage services

October 2016



AV Digitisation and Digital Preservation
TechWatch Report

#04



PrestoCentre's TechWatch Reports identify the business issues and technology trends that exist in digital audiovisual archiving and bring clarity in a language that is accessible to non-specialists.

PrestoCentre TechWatch Reports are published up to twice 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 behalf of other PrestoCentre members.

Published by PrestoCentre, 2016.



The text of this work is licensed under the Creative Commons Attribution 4.0 International License. To view a copy of this license, visit <http://creativecommons.org/licenses/by/4.0/> or send a letter to Creative Commons, PO Box 1866, Mountain View, CA 94042, USA.

PrestoCentre Foundation
PO Box 1060, 1200 BB Hilversum, The Netherlands
Telephone: +31 35 677 1813
office@prestocentre.org

Contributed to this edition
Simon Factor (Moving Media Ltd.), Erwin Verbruggen (Netherlands Institute for Sound and Vision).

Series editors
Johan Oomen, Aubéry Escande, Anna van der Meulen.

Foreword

The IBC, or International Broadcasting Convention, is Europe's largest professional broadcast fair. Held annually in September in Amsterdam, the trade show saw a record year in 2016, with 56,000 in attendance over the five days of exhibition and conference activities. The buzz at the show was that things are picking up again after the disruption of the past decade, where the financial turbulence and technological shifts exerted great pressure on the audiovisual business sector across the world.

Big numbers were not just a feature of conference attendance: they were to be found in many of the new products on show, with new benchmarks being set in terms of production video file formats, scanner imaging resolution and storage capacities. Happily, smaller numbers were also to be found, with prices continuing to fall in key product areas. At IBC in 2016 we saw companies continuing to innovate and deliver interesting solutions to support archives in leveraging digital technologies for both preservation and access. There was evidence of technology vendors actively engaging with open source tools and microservices in their products and workflows to limit costs, deliver greater efficiency, and increase flexibility and control.

For this fourth report in the TechWatch series for audiovisual archiving, we took a close look at film scanning developments, systems for video ingest, the potential of microservices and storage providers. The scale of the IBC show and volume of exhibitors means that it is not possible for us to provide an exhaustive list of vendors who were exhibiting new products and innovations in the specific categories we cover in this edition. Certainly, there are many other companies who offer great film scanning, video digitisation and storage technologies who were in evidence at the show and we look forward perhaps covering their products in future issues!

Table of contents

1. What's new for film scanners	05
2. Video ingest systems at scale	08
3. Microservices and open source	10
4. Storage services	11

1. What's New for Film Scanners

Over the past decade many scanner manufacturers were under extreme pressure due to the fundamental changes in their market. Post-production companies became less reliant upon film, as production moved to digital and therefore what were once the biggest ticket items in the post-production toolkit became surplus to requirements. The bottom fell out of the scanner market while many leading companies were restructured or sold off. Other manufacturers who counted scanners among their wider product range scaled back their activities in this area. One such respected and longstanding company, Cintel, founded in 1927 with great pedigree in this field, was acquired by Blackmagic, who sought to take the knowledge and capabilities of this established business and configure it towards delivering a disruptive, low-cost product. Which in theory is a boon for archives, offering a low-cost route to acquiring scanning hardware and the possibility to bring some of this expensive process in house.

Another company on a path to disrupt the film digitisation business is Vintage Cloud, who at IBC announced the news that they had acquired the well-known film editing equipment company Steenbeck. The reason for this? Vintage Cloud is taking a different approach to working with archives

to preserve audiovisual heritage. They see the hardware used for film scanning as a component within a broader digital business model. The film scanner they have developed uses Steenbeck as a platform, providing an easy to use system that facilitates digitisation but is ultimately connected to the storage cloud where assets can be analysed enriched, found and monetised.

Lasergraphics

One of the innovations present at IBC was the demonstration of a 10K scanning system by Lasergraphics from California. In recent years, the company has built a range of features into its products that make it an interesting option for archive film projects, and their new 10K Director product was positioned as a world first at the show. The flagship 'Director' is a sprocketless capstan-drive scanner and therefore tolerant and gentle with older items that may be subject to shrinkage. Their latest model includes their 'Automatic Failed Splice Recovery System' which picks up where the scanner left off after a splice repair during a scanning process. While pushing the boundaries on resolution in their premium products, Lasergraphics continues to develop the 'Scanstation' range of lower cost, easy to use scanners designed for lower budget / higher throughput applications. At NAB in 2014 we reviewed the original Scanstation, in 2016 the company added a new product to the line-up: the 'Scanstation Personal', which provides a more constrained feature set at an even lower price.

Blackmagic Cintel

After several years in development and much anticipation of its launch, the Blackmagic Cintel scanner finally began shipping in spring 2016, and according to Blackmagic



there are now over 100 units sold worldwide. Blackmagic announced the scanner in 2014, with a promised price tag of less than \$30,000, however at that time it was unclear when the product would enter full production.

Shipping of the Blackmagic Cintel has now commenced

and overall the package has many admirable features in addition to its relatively low price tag. The Cintel Scanner is handsome, compact and therefore more portable than most scanners. The configuration is straightforward, as is the case with most Blackmagic products: simply plug and play over Thunderbolt 2, for the Apple users among us. Thunderbolt 2 supports data transfer rates of up to 30Gb/s which allows the Cintel to run 4K scans in real time. The transport on this scanner is sprocket driven, so may not be suitable for more delicate items that are subject to shrinkage. The product ships with DaVinci Resolve, so out of the box the package provides a complete solution for scanning, grading and delivery. The Cintel can handle optical audio, however there is currently no solution for magnetic tracks. Blackmagic have a magnetic audio track reader in development and told us that they hope to launch this before the end of the year as an optional extra.

A 16mm film gate can be purchased for \$1,495 as an optional extra to the standard 35mm unit which is priced at \$29,995.

Vintage Cloud

Over the past four years, Vintage Cloud has developed its digitisation technology using Steenbeck machinery as the foundation, as these leading film handling and editing tables are well designed, tried and tested. Vintage Cloud has brought them into the digital age by adding a very elegant user interface that consists of a multi-screen array to replace the original viewer, and simple multi-function buttons and controls on the tabletop surface. Steenbeck

machinery always inspired confidence in users as a robust and solid piece of equipment for working with film, and Vintage Cloud has done a remarkable job of evolving these machines for use in digitisation applications, combining what makes the Steenbeck great with cutting edge technology and really well-designed tactile controls.

As a piece of equipment the Vintage

Cloud scanner is a flat deck, so different in its orientation to the majority of upright scanners. It will handle optical and magnetic audio tracks and, as it also includes two audio plates, can be used to work with separate optical and magnetic audio elements for 16mm, which is a standout feature that just a few other companies such as Sondor offer. As is the case with the original Steenbeck hardware, you can buy a 16mm or 35mm version. It is not possible to swap out the gate to use the same transport as you can do with some upright scanners.

The camera streams uncompressed SDI and can be used to capture 4K at 25 frames per second real time or 2K at 50 frames per second double speed.

The light source uses white, red, green and blue LED's and it is possible to control the intensity of each to vary light and throughput of the source. This allows for colour correction and modification based on the light source and can be especially useful when working with older film types such as discoloured Kodak stock from the 70's, where the lights source can be adjusted to compensate by boosting the blue and green channels.

Facial recognition is applied to identify individuals and analyse sentiment based on facial expressions

Once the sound and image have been captured, the Vintage Cloud system comes into its own and we see why this innovative solution is in a different category to what we typically see as film scanning equipment. The system is equipped with a range of smart indexing tools, applying third party API-based computer vision services to analyse key frames from the scanned images to detect objects. Facial recognition is applied to identify individuals and analyse sentiment based on facial expressions. All of this information is recorded as frame- and asset-level metadata. The audio signal is passed through an Automatic Speech Recognition system (again via a third-party API) and a text transcript of the audio is also generated to supplement the computer vision-generated meta tags. The fact that the system is open architected to leverage cloud-based third-party analysis tools





via APIs means that these things remain flexible and can be swapped out when improvements happen in the future. This automatically generated metadata is harvested for each digitised film to create a searchable index of descriptive contextual information describing the

contents of the film asset, which is hosted online or 'in the cloud', opening up a new world of commercial possibilities.

Vintage Cloud's vision for this is to bring the metadata and the key frames for these digitised film assets online and support the development of a footage marketplace that is easy for film collections to participate in, due to the combined ease of operating the digitisation workflow and automating the creation of the most important metadata to drive re-use.

Vintage Cloud operates a flexible pricing model; their machines can be purchased or rented. The 16mm machine is €179,000 and 35mm is €199,000 for purchase or can be rented for €6,000 or €6,500 per month respectively. A 35mm system is currently installed and in use at ITV in Leeds and another at Danske Film Services in New York. Both 35mm and 16mm machines are available for viewing at Vintage Cloud's facility in Copenhagen.

Digital Vision

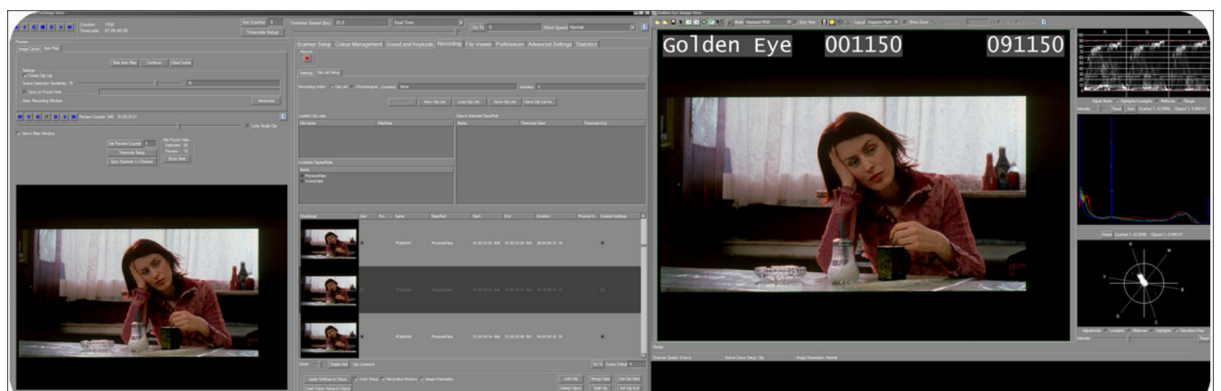
Digital Vision's Golden Eye is a multi-format scanner that can work with film gauges up to 70mm. The scanner is designed

for working with a wide range of archive materials and like many high-end scanners it uses a sprocketless drive so can handle old and delicate film. The Golden Eye also uses two cameras: the first is a non-Bayer 4K RGB camera which captures the scanned images and the second is a technical camera that supervises the film as it is processed measuring sprocket holes while detecting and reading the optical audio tracks. The Golden Eye also has a magnetic audio track reader for combined magnetic elements. The machine will scan 2K at up to 24 frames per second and 4K at up to 11 frames per second.

The current version of the Golden Eye is the Mk4, of which Digital Vision told us there are currently 20 units shipped and in use around the world. The machine is in use in a number of national film archives, including the Dominican Republic and the Philippines, with units in Hong Kong, Australia, three in London and two in the US.

Digital Vision offers its scanner with the Phoenix software, which they present as an end-to-end finishing tool. The software covers a range of workflow stages - from grading and restoration to finishing and output to, for example, DCP or ProRes. The Phoenix file-based restoration modules include automatic scratch and dust removal. The full, end-to-end Phoenix system is €28,000.

Pricing for the Golden Eye scanner starts at €140,000, upgrade to 4K with a range of gates is an additional €20,000. Other options such as a wet gate are also available.

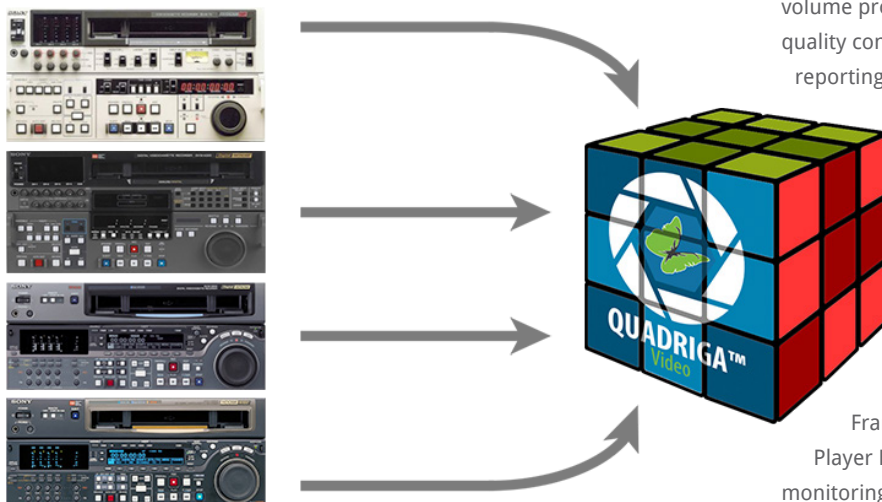


2. Video Ingest Systems at Scale

Systems that support high-volume video tape ingest were on show this year from both Cube-Tec and NOA. The objective with all of these video ingest systems is high-volume multi-stream ingest per operator with a wide range of Quality Assurance tools being used to ensure high quality outcomes in large scale video digitisation projects.

Quadriga

Cube-Tec has been working on a range of video signal quality analysis tools for some time and this year at IBC they were demonstrating the new Quadriga Video multi-stream digitisation system. Quadriga will be a product familiar to many who have previously worked on audio migration projects and the audio product has been around for over a decade. The video version is new for 2016 and the company has integrated many of the RF signal analysis tools that we reviewed at the 2014 IBC show¹ into the Quadriga Video product. Quadriga video can support up to 4 video players per system via 4 x SDI inputs. It uses 9 pin machine control to remote control the video transport and the RS 232 for collecting machine operating and tape condition data from the Sony Interactive Status Report (IRS) feed for each video machine.



Quadriga Video uses GPU-based video encoding to deliver four simultaneous streams of JPEG2000 lossless video encoding. Users can also choose from linear 8-bit or 10-bit encoding, depending on the input format and output as an MXF file.

The Quadriga Video system can import a range of XML metadata types, including descriptive metadata, to pass through the process. Time location metadata can be used to control the system to extract particular sequences or multiple items from a single tape or stitch together longer items from across multiple tapes - thus avoiding the need to manually set in and outpoints for the digitisation pass.

Pricing for the Cube-Tec Quadriga Video for the full 4 channel version is €30,000. Single or dual channel systems are also available upon request. The optional real-time multi-stream JPEG2000 lossless feature requires an upgrade to the GPU at around €1,000 per system. Pilot systems are currently deployed at a number of archive project locations including Deutsche Rundfunk.

Framelector & PICO

The NOA Framelector system has been chosen for use at ORF in their current large-scale video digitisation project, which is being delivered by ATOS. The project includes the digitisation of 600,000 video items with a total running time of 300,000 hours. Assuring quality of output and consistency in this high volume project was a key objective and now four levels of quality control are used. In addition to the interactive status reporting available via the Sony RS232 output on a Digital

BetaCam Video Player, NOA developed a new RF measurement tool to strengthen the QC process and provide additional signals to detect problems in tape sources.

Their RF Logger sensor is soldered onto the Digital BetaCam Player board in order to measure the RF output from the heads.

The signal is processed and amplified by a small appliance and delivered to the Framelector system via USB. This modification to the Player Hardware enables an additional layer of remote monitoring of the RF signal quality.

In the ORF project, ATOS is also using the VidiCert QC tools from Joanneum. They use it to perform a baseband analysis of the digitised video files and combine this with the NOA system outputs to surface problem files to operators for manual quality control of the digitised video.

NOA is a big proponent of the use of the FFV1 video format, for which the Internet Engineering Task Force (IETF) have

¹ See https://www.prestocentre.org/system/files/library/resource/techwatch_report_01_2014_0.pdf



The ATOS Installation at ORF using NOA Framelector multichannel video ingest systems

begun a standardisation project, announced in Berlin during the summer of 2016. Currently they have 2 broadcast archive projects underway using FFV1/MKV in eastern Europe and recently the format was chosen also for use by Indiana University for a high profile project in the US. NOA chose to provide FFV1 as a native option for encoding within its Framelector products, as they see it has many benefits as a lossless, open source file format that is easy to use, has low computational overhead and is growing in adoption.

Another development from NOA at IBC this year is the introduction of their Framelector Pico. This is a stand-alone version of the NOA video ingest system suitable for low-volume, single channel applications. The Pico includes many of the powerful QC capabilities of the full Framelector system but is not designed for high volume multi-channel applications. As it does not include some of the Job database and network features of the full Framelector systems, it does not require any additional servers or network infrastructure beyond the single workstation, which makes it suitable for archives looking for a self-contained low-volume video ingest solution.

The Pico offers a range of features that make it stand out from production-oriented video systems that can be used to capture video. In addition to the RF analysis and QC features

the NOA system delivers, this data can be output as a METS or PREMIS format metadata sidecar that can also contain a range of information about the system used for digitisation. Prices for a single-channel Pico system start at around €13,000.

3. Microservices and Open Source


MediaInfo has been gaining popularity in recent years as a very useful tool to create technical metadata and for many of the companies exhibiting at IBC it is a component of their own products. We met with Jérôme Martinez, the developer of MediaInfo, to ask him about the standards his tool supports and the rise of microservices in the archiving community and of organisations building their own solutions using open source tools.

MediaConch

MediaConch is a prototype tool that has been developed and launched recently by Jérôme's company MediaArea, funded by the European Commission through the PREFORMA project and with input from Dave Rice, who is well known for his work with the Bay Area Video Coalition in the US. The tool is used to detect different versions of Matroska files and to apply a rule-based filter system to detect files that are not compliant within a file-based workflow. The tool is being used in conjunction with MediaInfo within a microservices workflow to detect when files that do not match a particular standard or profile enter the workflow. The prototype tool has been adopted by Artefactual for use in its Archivematica products. MediaArea will make a more advanced version of MediaConch available before the end of 2016 once some planned developments relating to how a user can edit file policies within a workflow are complete.

MediaInfo

MediaInfo provides a range of options to output technical metadata according to industry XML standards including FIMS and EBUCore. In recent months, Jérôme has been working to integrate this feature into a range of microservice workflows. The range of metadata that can be extracted and structured by the MediaInfo application is growing beyond the typical video and audio stream settings metadata. It also covers new metadata that can be extracted from MXF files, other administrative metadata or acquisition metadata, such



as details of the type of camera and settings used. MediaInfo updates that will enable the extraction and output of this type of rich MXF metadata to FIMS- or EBUCore -compliant XML records will be coming soon.

MediaInfo has been widely adopted in recent years. Many companies such as Tektronix use MediaInfo as a tool to generate technical metadata natively within their products or as a microservices block within their products. It seems that many companies are opening up to the benefits of enabling the integration of microservices into their products to allow flexibility and customisation within user configurations. As production and archive workflows become more advanced and user specific, companies who continue to take a closed approach to integration with open source tools may find their products are less attractive.

MediaArea has recently been working on a number of projects to bring more open source tools into the market, for example using FFV1 video and Matroska file wrapping. It seems that while there are open source standards emerging there is not always the knowledge base or skills within an archive to develop solutions using these open source components or to build their own microservice-driven workflows. This is where companies such as MediaArea come into play: their business model is not in the creation of the open source software but rather in its implementation, modification and customisation to specific use cases.

Open source development does not mean free development. If people want open source software to continue to develop and open source standards to progress and remain current, they can either invest their own time and knowledge into the development of the resource or they can sponsor the development through companies like MediaArea. When one considers the scale of some of the companies exhibiting at IBC and the scale of resources behind them it is fantastic to see a company like MediaArea flying the flag for open source with a microservice product that is downloaded 5,000 times per day and is used daily by over 100,000 users around the world!

4. Storage Services

DTernity

On the storage front, Fujifilm launched a new product at IBC this year, called DTernity – ‘Data for Eternity’, described as ‘Archiving as Managed Service’. It fundamentally provides tape-based offsite storage of very large volumes of data. The key selling point for this service is that Fujifilm have developed what they are positioning as fully open source TAR-based standard for writing data to LTO tape. Fuji maintains that retrieval of the data from the tape should never be a problem for a user who requires portability. The main target is clients who have a long term data archiving requirement. It is not a backup service or a cloud storage service and is based upon a minimum 3-year contract. They create two copies of all data using data tape. The second copy is sent to a bunker at a discreet location close to the Dutch border - which all sounds rather James Bond!

One key aspect to this approach is that the data is not connected to any network or the internet and therefore has the security of an ‘air gap’. Fuji jokingly call the service ‘NSA proof’ and for this reason the service has gained some traction in the medical, legal and financial sectors where confidentiality and security is of high importance. Most of the current clients for the service fit into this category. Since launching the product, Fujifilm has identified the broadcast and audiovisual archiving sectors as potential customers

and is so far working with one large post-production company in the UK. Their key selling point for audiovisual clients is that they are offering a ‘migration free archive’: the client sends them the data on whatever data carrier they use and the Fujifilm team will ingest all of the data, regardless of what system was used to write the original tapes, to the DTernity platform.

Clients will see an online index of all of the data (but will not have network access to the files themselves) in the DTernity archive and where required can order any files they need to be delivery upon a data storage drive, data tape or USB stick. The Fujifilm team also offer data migration services to archives to move large data payloads from legacy data tape formats to new.



Fuji jokingly call the service ‘NSA proof’ and for this reason the service has gained some traction in the medical, legal and financial sectors, where confidentiality and security is of high importance.

1c per GB per month is the base price for the storage service. There is no ingest cost based on the minimum 3-year contract. The data recall price is said by Fujifilm to be competitive compared to cloud services, however is typically in minimum 2.5 TB blocks.

Storage DNA

Other interesting storage technology was on show from Storage DNA, who have been working on a range of interesting storage applications using Sony Optical Disc Storage and LTO-7. Storage DNA are using Sony Optical Disc in a way that differs from other applications that use it as an archive and restore platform: their software enables users to mount and access the Sony Optical Disc in a way similar to spinning disc. Storage DNA has been using the LTF5 standard since the 5th generation of LTO in its products. The LTO consortium is currently about to launch LTO-8, which should have about 12 TB storage per tape. The read/write speed for the current generation 7 LTO is about 300 MB/s, which is faster than standard spinning disks. This feature of LTO-7 is key to Storage DNA's products, as they leverage this high-speed read/write access using tape to unlock performance benefits in audiovisual production and archiving workflows.

Storage DNA explained to us that its products are very useful in applications where there is a requirement to read or write data quickly at very high volumes. As most systems that use tape go from source to disc to tape, effectively using the disc as a stepping stone, the speed of the disc is the limiting factor. With Storage DNA's product, they go straight from source to tape at 300 MB/s. With multiple tapes in parallel they can write extremely large data streams in real-time, which has applications in the very high resolution capture and transfer workflows that are becoming the norm in 4K production.

The Storage DNA turnkey system runs on HP LTO Hardware, while its software will work with many other LTO hardware systems and robotics, e.g. from Spectralogic, Quantum or Oracle. The range covers single-drive starter systems with 8 slots, providing 48 TB of storage, going up to 3.4 Petabyte Libraries using LTO-7 (with 540 tape slots and up to 42 drives) with an LTO-8 version with double the capacity due for release in 2018. Storage DNA quote the American and British Film Institutes, the National Archives of Fiji and the Olympic Council among their audiovisual archive customers. Pricing starts at US\$10,000 for a single LTO Drive software license while a turnkey 288 TB system with a two-drive, 24-slot LTO-7 Robot is \$130,000.

PrestoCentre's TechWatch Reports identify the business issues and technology trends that exist in digital audiovisual archiving and bring clarity in a language that is accessible to non-specialists.

ISSN: 2352-2941