

Automatic processing of CERN video, audio and photo archives

Michal Kwiatek

CERN, Geneva, Switzerland

Michal.Kwiatek@cern.ch

Abstract. The digitalization of CERN audio-visual archives, a major task currently in progress, will generate over 40 TB of video, audio and photo files. Storing these files is one issue, but a far more important challenge is to provide long-time coherence of the archive and to make these files available on-line with minimum manpower investment.

An infrastructure, based on standard CERN services, has been implemented, whereby master files, stored in the CERN Distributed File System (DFS), are discovered and scheduled for encoding into lightweight web formats based on predefined profiles. Changes in master files, conversion profiles or in the metadata database (read from CDS, the CERN Document Server) are automatically detected and the media re-encoded whenever necessary. The encoding processes are run on virtual servers provided on-demand by the CERN Server Self Service Centre, so that new servers can be easily configured to adapt to higher load. Finally, the generated files are made available from the CERN standard web servers with streaming implemented using Windows Media Services.

1. Introduction

CERN has a 50-year old tradition of divulgation and preservation of research results. Already the CERN Convention of the 1954 states that “the Results of its experimental and theoretical work shall be published or otherwise made generally available” [1]. Furthermore, over the years, many distinguished speakers visited CERN and presented the results of their work. Audio recordings of the early period and the later video recordings have been stored in the CERN audio-visual archives, but until very recently they were not available to the general public. Digitalisation of these archives is currently in process, but the digitalised video files are far too large to be easily accessible over the Internet without further processing.

CERN’s institutional repository has been available to everyone from CDS, the CERN Document Server [2], supported by IT-UDS group¹. The software used to manage this repository, CDS Invenio [3], has been developed at CERN and is licensed under the open source General Public License. As of February 2006, CDS managed over 500 collections of data, consisting of over 800,000 bibliographic records in the field of particle physics and related areas, covering preprints, articles, books, journals, photographs and more [4]. Multimedia capabilities were gradually added, but CDS was never supposed to be a multimedia server *per se*. In consequence, although some multimedia content was available from CDS, human interaction was necessary to prepare this content for presentation on the

¹ User and Document Services Group of CERN IT Department

web. Also, CDS lacked advanced streaming capabilities like adaptation of streaming bit-rate to the available network bandwidth.

At the same time, the IT-IS group² has been providing general web services to all CERN users[5], with advanced streaming functionalities based on Windows Media Services. IT-IS has also been providing the Server Self Service Centre[6], where virtual servers can be configured on demand to perform the CPU-intensive encoding of multimedia files. Last but not least, IT-IS operates a general file storage system based on DFS[7]. This last service has been suggested as the storage system for the digitalised media archives, but it was evident that by using the aforesaid services together, CERN media archives could be put on the Internet with minimum manpower investment.

This paper explains how the existing CERN services have been used together to publish CERN audio-visual archives on the Internet. It also presents the Media Archiver software developed at CERN to provide the link between these services and process CERN video, audio and photo archives.

2. Archiving workflow

The archiving process, presented on Diagram 1, starts when members of Audio-Visual team or Photo Lab copy video, audio and photo *master*³ files to *Media Archive*⁴ disk space on DFS. The same people are responsible for putting the description of the files in CDS. New *master* files are discovered automatically by the Media Archiver software and as soon as the *metadata*⁵ is available from CDS, *slave*⁶ files are registered for future processing according to conversion profiles.

The Media Archiver software runs in parallel on multiple worker nodes. *Slave* files are chosen for generation based on predefined priorities and current load on the worker nodes. Generating a *slave* file effectively means encoding the *master* file into the lighter web format. Both *master* and *slave* files are stored on DFS.

When Media Archiver completes generation of a *slave* file, it notifies CDS and the new files becomes available to the end users. The users discovers the new *slave* file by browsing CDS, but the file itself is served from DFS by means of Media Archive Web Server.

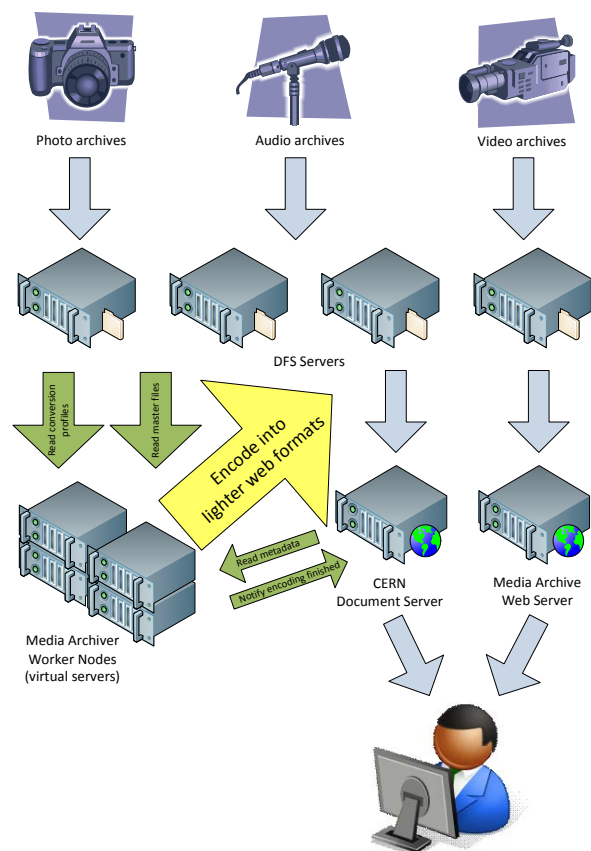


Diagram 1, Media Archive workflow.

² Internet Services group of CERN IT Department

³ *Master* files are the high-quality, large-size files stored in the archive in order to ensure their long-term preservation.

⁴ *Media Archive* will be used throughout this paper to describe the combination of services used together in order to store digitally and publish on the Internet CERN's audio-visual archives.

⁵ *Metadata* is the structured description of the files stored in the archive and contains: title, text description, author, copyright information and event date. *Metadata* for video files additionally contains frame rate, aspect ratio and size in pixels.

⁶ *Slave* files are the lightweight versions of the *master* files created and stored for presentation on the Internet. *Slave* files are the result of conversion of *master files* according to conversion profiles.

2.1. Synchronisation of changes

Media archive of an institution of CERN's size changes constantly. New files are added (both new material and old material only recently digitalised), existing files are updated and moved within the archive and descriptions are updated. In order to have an archive where valuable material to be found in an easy and reliable way, these changes have to be synchronised.

The Media Archiver software constantly monitors changes within the archive:

- It discovers new, updated or deleted *master* files.
- It is notified when *master metadata* is updated.
- It discovers new, updated or deleted conversion profiles.

Any change is propagated to the *slave* files. For example, if a *master* file is deleted, all its *slaves* are deleted as well.

3. The storage

Three factors have determined the structure of *Media Archive* storage: physical capacity of the servers, ease of use and finally the potential to grow in the future.

CERN DFS service has been chosen to accommodate *Media Archive* files. This service is based on Distributed File System[8] solution of Microsoft Windows Server 2003: multiple *shares*⁷ of numerous servers are linked into one logical tree structure. Users of computers running Microsoft Windows see this tree as a unified disk space; it is in addition mounted locally as an additional disk drive. Users of other operating systems connect to each of the *shares* separately: it is less convenient but still functional, since the structure of the tree is dynamically documented. The size of a single *share* is limited by the size of disk partition on which it resides. In 2006, when the structure of *Media Archive* was defined, the biggest available servers fitting CERN infrastructure had 14 data disks of 465GB each. These were grouped into two RAID5 partitions of 5 and 7 disks with a hot spare for each partition. It resulted in two partitions of 1.81 TB and 2.72 TB respectively. Four such servers have been acquired resulting in the total disk space of 18 TB divided over 8 partitions.

From the point of view of the Audio-Visual and Photo Lab teams, who work on Linux or Mac operating systems, it would be most convenient to have all of its data on a single partition. This was possible for audio and photo files, but even the existing video files would not fit on a single partition. Moreover, each of the three categories of video files (Movies, Conferences, Rush) was supposed to grow above the 2.72 TB limit within a year and would therefore need to span several partitions.

Finally a hierarchy of type/category/year has been chosen together with the opportunistic space allocation policy: whenever files of a certain type (for example video files) were not fitting on a single partition, DFS links were moved to the level of category. When this was not enough, because for example movies occupied more than 2.72 TB, links were moved to the level of year so that different years could be placed on different partitions. The Audio-Visual team confirmed that no category would grow over 2.72 TB in years 2006 and 2007 and later we could expect arrival of bigger disk servers. In conclusion, this organisation of storage enabled maximum ease for the users without the necessity to reorganise the folders periodically.

It is also worth mentioning that the disks of CERN DFS servers are backed up on tape within the general backup and restore service[9] provided by IT-FIO⁸

4. File formats

The choice of formats for the files stored in the archive, called the *master* formats, has been a very important one. The ideal *master* format should:

- natively support the recording hardware, so that *master* files have the same quality as the original footage,

⁷ Windows *share* is a directory that has been made available on the network and can therefore be shared among different computers.

⁸ Fabric Infrastructure and Operations group of CERN IT Department

- be supported and easy to read on the widest range of operating systems,
- have good chances to stay supported and easy to read on the same operating systems in the future,
- be recognised by external customers of the archive: television companies (for video files) and newspapers and publishing houses (for photo files),
- be the format in which the teams work and save their work,
- be easy to process automatically with the available, well-supported tools.

4.1. Audio *master* format

WAVE PCM[10] has been chosen as the *master* format for audio files. It is a good example of the ideal *master* format, as it satisfies all of the criteria mentioned above.

4.2. Photo *master* format

The choice of the photo master format has been more difficult because many photographers, including those of CERN Photo Lab, work in the raw format, which is the native format of the camera they use. The problem of the raw format is that the files can only be read using professional photo-editing software and that the format varies from camera to camera. It is also not understood by general purpose programming libraries.

In order to be sure that the files can be easily read in the future and at the same time to let the photographers continue to work in the way they that they are used to, both raw and original-size, high-quality JPEG[11] files are stored in the archive. Only the JPEG files are used as the basis for generation of *slave* files.

4.3. Video *master* format

The case of video files is the most difficult one because of the different formats natively supported on Mac and Windows computers, the complexity brought about by the high definition video and the proliferation of video codecs. This last issue is particularly important, because even though the newer codecs are more effective, the number of different codecs used in the archive should be as small as possible. Use of more than just 2-3 well chosen codecs would make it difficult to ensure that the archived files will be easy to read and process in the future.

The primary aim of the Media Archive project has been to ensure the storage and easy processing of historic video material, previously stored on tape. For this purpose, we have chosen the AVI container format and the DV codec[12]. The DV codec is the native codec of many standard definition digital video cameras and the DV AVI files are easy to read on both Windows and Mac computers. This ensures that the support for DV files will not disappear with the new version of the codec and that tools (including programmatic tools) are and will remain available. The same would not be true for some newer and more effective, but not so well established codecs. The tape archives are therefore digitalised directly into DV AVI format.

Some of the newer videos in the archive come from DVD discs. Although the DVD standard allows higher compression and therefore lower quality than DV, it is equally popular and widely supported. DVD uses the MPEG-2 format[13] with some restrictions related to allowed resolutions, aspect ratios, frame rates and bit rates [14]. To enable compatibility between files stored in the archive and the DVD discs, the MPEG-2 DVD-quality has been chosen as the second *master* format.

Two problems remain unsolved. The first one is related to the fact that the Audio-Visual team is used to work on Mac computers. As a consequence, they save their work in the Quicktime Movie (MOV) container format[15]. Unfortunately, MOV files are not easy to play on Windows computers, for example DV MOV files played with Quicktime software on Windows show a low-quality stream instead of the original DV stream. At the same time, support for Mac computers at CERN is limited and the management considers stopping it completely to reduce costs. MOV files are also difficult to

process automatically on Windows computers. All this makes the MOV container a bad choice for the *master* format.

This problem does not really affect the archives (which are digitalised from tape directly into the DV AVI format), but is relevant to the recent recordings, often edited by the Audio-Visual team. These, however, are recorded in high definition, which brings us to the second unsolved problem.

There are currently several competing formats of high definition video files. The Audio-Visual team alone recorded in two high definition formats: HDV and DVCPRO HD, each specific to the camera the team used at the given time. These formats are not easy to process automatically and they may well become obsolete with the recent introduction of AVC, which is based on the H.264 MPEG-4 codec [16]. As already mentioned, presence of so many different *master* formats in the archive is undesirable, as it complicates the management of the archive and makes it hard to ensure that the files will be easy to read in the future. We have therefore decided to postpone the introduction of support for high definition video until one of the formats establishes itself as the de facto standard, which can be measured by the availability of tools and ease of use on various operating systems. The current workaround is to export the HDV and DVCPRO HD files into the MPEG-2, DVD-quality format and to store these files along with the original footage.

4.4. *Slave* formats

Slave files are the result of conversion carried out to enable presentation of the archives on the Internet. The main criteria for *slave* formats are that the files are easy to read by the majority of users and easy to generate automatically and serve on the Internet with well-supported tools. They also need to be small enough to transfer over the Internet.

According to these criteria, Windows Media Video (WMV) [17] has been chosen for video *slave* format. For each video *master* file, WMV files of three bit-rates are generated for download: 128 kbps, 512 kbps and 1000 kbps. In addition, the fourth file containing three streams of different bitrates is generated to enable optimized streaming: depending on the available bandwidth the appropriate stream is chosen so that the user can see the best possible quality without delays. Also, 9 frames of the video are extracted and saved in the JPEG format for fast preview of the video file contents.

Similarly, Windows Media Audio (WMA) [18] is used as audio *slave* format, but in addition MP3 and Vorbis OGG formats are used due to their popularity and availability of reliable tools.

For photo *slave* files, reduced size JPEG files are used; the generated files are either A4 at 144 dpi or A5 at 72 dpi. Icons of 180 by 120 pixels are also created.

5. Media Archive web server

Media Archive web server is the front component of *Media Archive*, as it serves the *slave* files to the end-user. The configuration of the server is based on Windows Server 2003 operating system with Internet Information Services and Windows Media Services[20]. The web server reads the files directly from DFS storage.

Internet Information Services (IIS) is the traditional HTTP server: it serves files over the HTTP protocol. This is well suited for download, but it does not guarantee good presentation of video content that is watched as the file is downloaded [21]. Windows Media Services (WMS) are used instead to enable more efficient use of network bandwidth with the RTSP or MMS protocols over UDP-based or TCP-based transports. Finally, to enable smooth cooperation with clients behind firewalls, streaming over HTTP is enabled with the WMS HTTP Server Control Protocol plug-in[22]. Since both IIS and WMS use port 80 for the HTTP protocol, two IP addresses are assigned to the server.

6. The Media Archiver software

The basic role of the Media Archiver software is to carry out the conversion of *master* files into *slave* formats. But Media Archiver does in fact much more: it creates and maintains the link between the individual services already existing in CERN IT department to provide the user with the functionality of a unified Media Archive.

The four services with which Media Archiver directly interacts are presented on Diagram 2. Instances of Media Archiver software are run every minute on the Worker Nodes, the virtual machines configured within the Server Self Service. Media Archiver compares the state of *master* files on DFS servers with the state of the database. New, modified or deleted files are discovered and new *slave* files are registered for future generation. Alternatively, CDS may send a web service (WS) call with information that the metadata for a given *master* has changed. This call will be registered in the database and acted upon as soon as the next instance of Media Archiver becomes available.

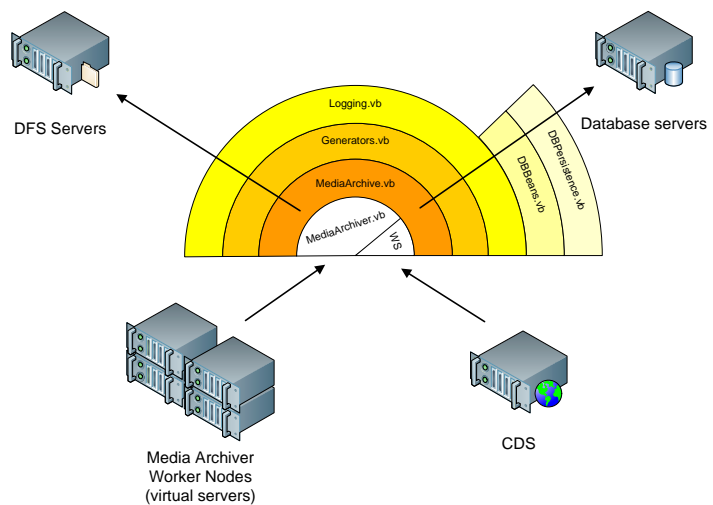


Diagram 2, Media Archiver software modules.

6.1. Software modules

The software is modular; each module is presented on the diagram as a circular layer. When a new instance of the software is run, the flow of control is started in the main method of MediaArchiver.vb – the entry point of the application. The flow of control will normally go through all layers of the software:

- MediaArchiver.vb, which synchronises the files in DFS filesystem with Media Archiver database and registers files for future parallel processing,
- Generators.vb, which implements conversion of the files into specific *slave* formats,
- Logging.vb, which logs all of the operations depending on the configuration and severity level of the message: to text files on DFS, event log on the worker nodes, database table or sends the message by email,
- DBBeans.vb, which maps database tables and records to classes and objects and implements operations on these objects,
- DBPersistence.vb, which implements basic operations on database records.

The names of the layers correspond to the names of VB.net source files.

6.2. Conversion

The actual conversion, or encoding of *master* files into *slave* formats resulting in generation of *slave* files, is implemented within the Generators.vb layer and depends on the format. Clear separation between the layers makes it easy to add support for additional *slave* formats.

6.2.1. *Video conversion.* Generation of video *slave* files in WMV format has been implemented using Windows Media Encoder 9 Series SDK[23]. The implementation is based on the code sample provided with the SDK [24]. Extraction of frame images from DV AVI files has been implemented with calls to DirectShow COM objects through a strongly typed interop assembly: Interop.DexterLib.dll. Frame images of MPEG-2 videos are extracted using external ffmpeg open source video encoder[25].

6.2.2. *Audio conversion.* Generation of WMA audio *slave* files has been implemented with the above-mentioned Windows Media Encoder 9 Series SDK code sample. MP3 files are generated using the LAME external encoder [26] and OGG Vorbis files are generated with the oggenc2 encoder[27].

6.2.3. *Photo conversion.* Processing of JPEG files has been implemented using Microsoft.NET APIs.

7. Conclusions

Automatic processing of CERN video, audio and photo archives have been implemented within the Media Archive project described in this paper. Synergies created thanks to the use of existing IT services allowed us to complete the project with minimum manpower investment and to ensure that future total cost of ownership will be low.

Apart from the original objective of making CERN audio-visual archives available on the Internet, this project achieved a less visible, but equally important goal. Strict rules of storage and description of audio-visual *master* files, enforced within the automatic conversion process, ensured that the future generations of CERN users will inherit a well organised, coherent and useful media archive.

Acknowledgments

I would like to warmly thank Alberto Pace, who initiated the *Media* Archive project, provided a proof of concept and introduced me to the world of video encoding. This project would not have been possible without the involvement of members of CERN Photo Lab and Audio-Visual team: Maximilien Brice and Jacques Fichet and members of IT-UDS group: Christian Montarbaud, Thomas Baron, Jean-Yves Le Meur and last, but not least, Diane Berkovits, who enabled the integration between Media Archiver and CDS. Separate thanks go to Tim Smith, whose remote supervision ensured that the project did not drift away from users' needs.

References

- [1] CERN Convention of Member States, Article II, www.cern.ch/dsu/ls/conventionE.htm
- [2] Pepe A, Le Meur J-Y, Simko T, Computing in High Energy Physics (CHEP-2006) I 467
- [3] CDS Invenio, www.cern.ch/invenio/index.html
- [4] Op. Cit.
- [5] CERN Web Services, www.cern.ch/WebServices/
Lossent A, CERN Web Services 2005, Post-C5 presentation for CERN IT Department, www.indico.cern.ch/conferenceDisplay.py?confId=a045813
- [6] Ormancey E, Using Virtual Servers for the CERN Windows infrastructure, Post-C5 presentation for CERN IT Department, www.indico.cern.ch/conferenceDisplay.py?confId=4149
- [7] Overview of the Distributed File System Solution in Microsoft Windows Server 2003 R2, microsoft.com/downloads/details.aspx?familyid=5e547c69-d224-4423-8eac-18d5883e7bc2
- [8] Idem
- [9] Backup of Computer Server Nodes at CERN, www.cern.ch/service-backup/
- [10] WAVE PCM sound file format, ccrma.stanford.edu/courses/422/projects/WaveFormat/
- [11] The JPEG photo file standard, www.w3.org/Graphics/JPEG/itu-t81.pdf
- [12] DV Data in the AVI File Format Specification, download.microsoft.com/download/1/6/1/161ba512-40e2-4cc9-843a-923143f3456c/DVAVSPEC.RTF
- [13] MPEG-2, en.wikipedia.org/wiki/MPEG-2
- [14] Idem.
- [15] QuickTime, en.wikipedia.org/wiki/.mov#QuickTime_file_format
- [16] H.264/MPEG-4 AVC, en.wikipedia.org/wiki/H.264/MPEG-4_AVC
- [17] Encoding Audio and Video with Windows Media Codecs, www.microsoft.com/windows/windowsmedia/howto/articles/codecs.aspx
- [18] Idem
- [19] OGG, en.wikipedia.org/wiki/Ogg
- [20] Streaming from a web server, www.microsoft.com/windows/windowsmedia/howto/articles/webserver.aspx

- [21] Web Server vs. Streaming Server,
www.microsoft.com/windows/windowsmedia/compare/webservvstreamserv.aspx
- [22] Firewall information for Windows Media Services 9 Series,
www.microsoft.com/windows/windowsmedia/forpros/serve/firewall.aspx
- [23] Windows Media Encoder 9 Series SDK,
msdn2.microsoft.com/en-us/library/ms991204.aspx
- [24] Automating the Encoding Process,
www.microsoft.com/windows/windowsmedia/howto/articles/AutomatingEncoding.aspx
- [25] FFMPEG, ffmpeg.mplayerhq.hu/
- [26] The LAME Project, lame.sourceforge.net/index.php
- [27] OGG Vorbis encoder, www.rarewares.org/ogg-oggenc.php