www.ichim.org

Les institutions culturelles et le numérique

Cultural institutions and digital technology

École du Louvre
8 - 12 septembre 2003

IMPLEMENTING CONSERVATION-QUALITY DIGITAL MEDIA SYSTEMS AND DISTRIBUTED ARCHIVES

Peter Cornwell, Juergen Enge and Hartmut Joerg,

ZKM - Center for Art and Media Karlsruhe, Germany

« Acte publié avec le soutien de la Mission de la Recherche et de la Technologie du Ministère de la Culture et de la Communication »

Abstract

Since 2000, the convergence of new low-cost fast intra-networking technology and very large storage components has enabled completely new solutions to conservation of, and widespread accessibility to, cultural media. In a project commencing during mid 2002, ZKM has investigated the mid- and long-term implications of these technologies in the context both of archival and conservation and also of access by researchers and the public. This work has centered upon the construction and management of very large disk storage systems by non-specialist personnel in the museum environment, and builds upon ZKMs experience with largescale automated CD archives. While the advantages of disk-based storage for media archive is well reported, this project addresses the reliability of such systems for repository of high quality masters, enabling the complete disposal of conventional tape media for the first time. The ZKM Storage Area Network (SAN) project will reach approximately 100Tb (Terabyte) capacity by mid 2004 and will eventually house the audio and moving imagery collection of the ZKM | Media Library as well as mirroring several other collections. This paper presents specific experiences in loss-less digitisation processes and the automatic generation from digital masters of a range of outdated, current, and future compressed data formats for internet and mobile wireless users. The interoperability of different vendor storage products and open systems-based implementations is reported and the personnel training and management procedures necessary to sustain and guarantee the integrity of such archives over very long periods is also addressed.

Introduction

Several fundamental problems face organisations holding collections of valuable media titles, especially film and video, because of the deterioration in quality accompanying each access and the limited lifetime of conventional media. Computer-based solutions have enabled widespread accessibility to the research community and public through location-based media libraries and increasing performance of the internet. However, the compression techniques employed to make manageable the huge file sizes resulting from

the digitisation process using economically available computing equipment have led to reduction in quality of the content. While this is acceptable for much usage such as entertainment and documentation for researchers, it does not address the long term conservation of unique masters and moreover, the CD, DVD or digital tape media employed have little better life expectancy than the magnetic tape that they replace.

The scale of many media collections poses another difficulty for long term security of valuable titles. Conventional tape media are accessible only at real time play speed and so archives of thousands of hours of moving imagery, of which there are many, cannot, even with large staffs and huge investment in equipment, transcribe titles fast enough to keep ahead of media expiration. Many organisations hold hundreds of thousands of hours of material that is increasingly irretrievable.

The ZKM holds such collections of unique media art work and documentation, and since 1997 its Media Library has addressed accessibility both by the scholarly community and the public, using robotic media management, state of the art browsable location-based viewing systems and a comprehensive internet service.

However, rapid growth in collection size, together with the risk to quality of title masters from media deterioration has required a far-reaching and urgent review to be undertaken, as well as research into new technologies. Of particular concern has been the impact to cost of ownership – capital and consumable equipment, operating cost, maintenance contracts, personnel, facilities – and security of irreplaceable titles, together with the development of infrastructure for new conservation techniques. In turn, this has led to pilot projects to evaluate technologies in detail and to carry out service deployment trials to assess the personnel issues for management of installations and the success of new interfaces for users. In this paper we report trials with several specific technologies that have resulted in the development of a sustainable architecture for largescale disk-based digital archive. We then describe the implementation of new computing and networking infrastructure at the ZKM to support such an installation, as well as a project to evaluate new means of access to largescale archives by the user community using wireless-connected personal computers.

1. Strategies for Media Management

The underlying approach taken in this work is simple - we have concentrated on the needs of conservation and dissemination of moving imagery using digital techniques by developing capture, secure storage and dissemination methods for the very highest quality data that can be obtained from analogue masters. This makes the techniques that have been developed applicable equally to still imagery, audio and geometrical information of conservation quality because, although differently specialised digitisation equipment may be required, the size of resulting data objects is small by comparison with moving imagery. The most significant benefit of this "loss-less" digitisation approach, which admittedly involves the management of exceptionally large datasets representing each title, is more far-reaching - and it is an important goal of the project to determine its practical feasibility. The use of open systems data representation standards, together with the capacity for very high resolution gives such an archive architecture the potential for "future-proofness". This arises from an intersection of separate developments in data storage and communication technology which make it possible for the first time to manage storage of the scale of hundreds of Tbs (Terabytes) economically and safely.

The fundamental property of digital representation - that data-sets can be cloned to produce identical multiples - now enables the intrinsic limitation of access time of traditional media to be overcome. Loss-less digitisation and copying of valuable collections on multiple disk systems at different locations within an installation offer a level of fault-tolerance equivalent to the reliability of existing tape media. Automatically maintaining identical copies of one installations' titles at another geographic location - a technique known as mirroring - gives unprecedented security. The breakthrough in this approach is that the fast geographic cloning enabled by new communication and storage technology overcomes the lifetime limitation of individual physical media. A disk has a life expectancy in continuous service that is lower than that of a magnetic tape (although an analogue tape looses quality with each access and no tape can be accessed continuously - on the other hand, a disk is a non-contact device and its contents do not degenerate with access. If disks are "spun down" when not in use their life expectancy can be of the same order as that of tape archives). However, its contents can be cloned in the

order of a couple of days, compared with a large tape-based collection – which it might not be possible to copy before media deterioration renders some titles irretrievable. Apart from the risk to content, the huge cost in human intervention of copying a tape collection is prohibitive. Cloning a media collection on disk also results in "regeneration" - the making of an identical multiple with exactly the same quality - the same is true of digital tape, wheras copying analogue tape or film results in loss of information.

Moreover, there are side effects of disk-based archive methodologies, which actively contribute to the potential accessibility of collections. Firstly, the geographic mirroring of archives enables collection sharing as well as ensuring survivability in the face of disaster at one or even several sites. This is simply not possible with large collections based on tape based media because of transcription time. Materials in a disk-based archive can also be protected by access permission or even by encryption on disk, individually, so that copyright protection can be administered on a title-by-title basis.

Another, no less important side effect of direct-access disk-based archives, which actually arises from the intrinsically distributed implementation of practical installations, is the transformation processing that can be achieved automatically as background computation. A storage area network having the bandwidth necessary to playout datasets even of filmic resolution (we will consider in more detail below the ratio between processing and storage to achieve this) functions inherently as a collection of interconnected computers which can function as a parallel processing resource. Computation to search and access titles for users and to manage the consistency of geographical mirrors does not exhaust the capacity of this architecture. Either in "its spare time" or by scheduling through load-balancing, the storage network can also accomplish other computationally demanding tasks, such as the transformation of titles into new format standards or the production of a variety of compressed formats. This is an important indicator of the sustainability of digital media systems, because data representation standards are almost as fluid as information technology itself and the possibility of migrating a collection from one lossless representation method to another, or as significantly from one metadata structure to a new emerging standard is crucial to the future-proofness of the archive. Significantly, this can be achieved in the background - even, if necessary, producing a completely new version of the collection in parallel with the original.

Another aspect of this, investigated in a pilot project reported in more detail below, is where small hand-held devices are used to access ZKM | Media Library titles, and versions of collections formatted for different access methods can be produced automatically or even "on-demand". This is particularly important as internet performance improves and users expect increasing quality of compressed materials to be delivered from web sites.

A strategic view such as this is necessarily decoupled from the practical feasibility, cost and organisational issues that face real installations - the underlying approach might be simple but the delivery of a long term service is very different. In examining the last sixteen years experience of development and maintainance of the ZKM | Media Library, the evolution of technology as well as archive methodologies has been crucial to the formulation of a sustainable architecture for a new generation of digital archive. Similarly, much is written about performance and cost trends in the information technology sector and in particular, the growth forecast in communication performance and storage capacity has informed our work in the same way as that of other institutions and corporations. However, there is a third element, which has influenced the work reported here profoundly, determining whether to pursue incremental improvements or seriously to examine the more radical alternatives described above - and that is the Open Source movement. Together with the Internet Engineering Task Force's Storage Area Network standards (www.ietf.org) - the organisation that sets Internet standards (Simpson, 2003) - the impact of open systems on the sustainability of practical disk-based media archives is far reaching because it insulates installations from dependence upon individual computer vendors over the very long lifetimes significant of archives. Consequently, the ZKM Storage Area Network project has concentrated on installation management and sustainability issues arising from an open systems strategy - on the reality of operating very large installations with low risk, commensurate with the long term security of national collections. Important also, has been the establishment of strategic collaborations with other institutions and corporations for development of mirrors of collections and of new standards.

2. Historical Background

The ZKM | Center for Art and Media in Karlsruhe is a unique institutional model, responding to rapid developments in information technology and today's changing social structures and providing a forum for research, production, presentation and discourse in the fields of art, science and media technology. The ZKM comprises the ZKM | Media Museum, the ZKM | Media Library (Mediathek), the Exhibitions Department, the Museum for Contemporary and six research institutes: the Institute for Visual Media, the Institute for Music and Acoustics, the Institute for Net Developments, the Insitute for Basic Research, the Institute for Film and the Insitute of Media and Economics.

Since 1989 ZKM has been collecting broadly-based media art as well as contemporary art employing more traditional techniques. The different collections are made public via temporary art exhibitions in diverse in-house exhibition spaces, in the Media Museum and the Media Library as well as via travelling exhibitions, conferences and a comprehensive internet presence.

2.1 Media Museum

The ZKM | Media Museum is one of the world's first museums for interactive art, having a thematic spectrum extending from interactive film to cyberspace simulation technology and current internet software applications. The Media Museum defines itself as ZKM's show case, making projects, works and developments of ZKM's research institutes and visiting artists available to a wide range of audiences. In this way, the Media Museum not only ensures transparency, but at the same time documents the diverse activities of the individual departments. The permanent collection of the Media Museum is unique, having an international reputation due to the breadth of its holdings and comprehensive early media art works. The main task of those working at the Media Museum is the storage, restoration, presentation, and mediation of media art, especially interactive media art.

2.2 Media Library

The ZKM | Media Library comprises audio and video collections and a library. Occupying floor space amounting to some 600 sq.m., the entire Media Library offers a comprehensive collection of international video art, contemporary music, and literature

relating to 20th century art. A central database (http://biblio.zkm.de) allows search requests covering the inventory of about 1,300 digitised art videos, about 12,500 audio works, with a focus on electroacoustic music as well as about 30,000 books, 120 reviews and about 500 CD-ROMs and DVDs.

2.2.1 Target groups

The target group consists of visiting academics, users in the ZKM institutes and the State Academy for Art and Design Karlsruhe as well as the interested public. The diverse composition of the Media Library's users requires flexible means of access: as a service institution for the scientific staff and the exhibition team as a place for research and teaching for the State Academy for Art and Design and as an educational centre on media art for the interested public.

2.2.2 Collections

The video art collection of 1,300 digitised titles includes central positions of video art, the video magazine "Infermental" from the 1980s as well as work awarded the International Media Art Award. In addition, the collection contains documentary material of great importance to the research community.

The audio collection with its 12,500 titles emphasizes contemporary music. Of special importance is the International Digital Archive for Electro-Acoustic Music (IDEAMA) from the medium's beginnings until 1970. Besides significant modern works and the archive of the Deutsche Gesellschaft fuer elektroakustische Musik (DegeM), the collection also hosts Blues, Jazz, Rock, Pop and movie soundtracks, as well as one of the largest soundscape collections worldwide. Users can also listen to works that have been awarded the Karl-Sczuka-Award for radio play and radio art (donated by the Suedwestrundfunk since 1955).

2.2.3 Digital Storage

Since the opening of the new ZKM building in 1997, the Media Library has concentrated on innovative presentation and delivery forms for media art. These include a directly accessible central database that stores information for each title, and a fully-automated jukebox retreival system for video-CDs and audio-CDs. The video-CDs are stored as

MPEG and the audio as Wave. This jukebox system and individual viewing userstations allow for individual perusing of the moving image collection without time-consuming tape loading cycles. Stand-alone systems are provided for viewing DVDs. The jukebox installed in 1997 - a Terastore CDR76 - supports 2064 CDs and it is currently 98% utilised. The core components of the system are a catalogue server for the database and a media control server that manages requests and channels imagery and sound in response to users at the public terminals. This includes the choice of a video or audio title as well as the functions of playback such as Play, Forward, Rewind, Stop, and Pause. The relevant information concerning these user choices and play functions is linked either to a MPEG2 Cache with a capacity of 50Gb or a CD (Audio CD or CD-ROM for video) from the Jukebox. In both cases the data stream is decoded before being transported via a crossbar to the user stations in analogue form. (Frieling, Joerg , 1999)

Figure 1: CD-Jukebox

2.2.4 Accessibility

The user stations are situated in the public area of the Media Library. A special exhibition design for ten individual stations equipped with two seats each was created in 1997 by the Franco-Canadian media artist Luc Courchesne, and four historic listening booths, designed by Professor Dieter Mankau. In order to preserve the autonomy of the art works, the online catalogue is detached from the actual viewing system and can be seen on a separate screen. The monitor displaying retrieved titles is reflected by semi-mirrored glass. This effect is similar to a gallery of floating pictures; the work of art is no longer viewed in the usual context of a monitor presentation but as a frameless image, freeing perception from connotations of the television. Additional equipment, seminar rooms and

the ZKM lecture theatre are available for researchers.

Figure 2: ZKM | Media Library user stations

3. Research into Digitisation Methods and Linux-Based Archives

The Media Library collaborates with the Institute for Visual Media and the Institute for Music and Acoustics in research into conservation of archive materials in the media library. A grant from the Bundeskulturstiftung has been awarded for digital restoration of early video material using an initial sample of about fifty titles, including video materials of the Media Museum. Another research initiative has been the uncompressed digitisation of the International Digital Electro-Acoustic Music Archive (IDEAMA) target collection on 140 CD-Rs, completed in 1997, which is now completely stored on hard disc. Other programmes include the digitisation of the Documenta5 archive by Karl Oskar Blase dating from 1972 and that of early works by Ira Schneider and Woody Vasulka. Work in progress includes the digitizing of historic interviews with philosophers and artists conducted by Gerhard Johann Lischka originally mode on analogue cassettes. This work is based in a digital video editing suite in the Media Library.

During 2002, in a collaboration with the Institute of Contemporary Art in London and U.K. company Street Vision, loss-less digitisation techniques were developed for video

which have since resulted in a Linux-based very high resolution workstation for video and audio. The project has digitised a sample of the ICA's collection of historical lectures from the last 20 years by key figures such as Baudrillard and Derrida and performances and debates, which were held on magnetic media near to end of life. Completion of this work will secure the future of the ICA material, which otherwise faces irretrievable loss, and promises to make accessible this material to the scholarly community for the first time. It has also proven that secure disk-based archives are now robust enough for deployment in a museum service. Government grant applications developed and submitted jointly between the ICA and the ZKM have formulated justification for mirroring of collections for security and shared accessibility. This research has been used in the MetaPlex project - part of the Future Cinema exhibition at ZKM in 2002 (Weibel, 2003)- in which a virtual museum comprising a large number of conservation-quality films, videos and interactive artworks was installed as a navigable virtual environment.

In a related project, a mixed sample of film and video holdings of the Japan Animation Association (JAA) has been digitised using the same loss-less techniques. By contrast with the ICA project, the higher size of dataset per second play-out time of film is a significant challenge to disk-based playout systems - requiring upwards of a Gbit per second throughput to deliver filmic quality imagery in real time. However, the selection of animation titles was partly based on their generally shorter length than features - which was commensurate with disk capacities economically available at that time (2001-2002), and the fact that this collection also has numbers of titles endangered from media deterioration. It is significant that loss-less digitisation, which effectively arrests deterioration in title quality due to media failure, enables future application of restoration software not even yet designed.

4. Constraints Arising in the Existing Media Library Installation

Several shortcomings have become increasingly clear during extended operation of the Media Library, which arise partly through the need to commission equipment of necessarily fixed functionality in order to mount the collection. This is obstensively no more than the simple imperative facing any IT installation: that of being unable to wait interminably for the next technological leap. However, it belies a deeper problem of lack

of scalability in storage products until the last 12 months, to which we will return later. Other limitations are intrinsic to the medium employed for primary storage - the recordable CD.

During six years operation the CD jukebox has been filled as library personnel digitised titles from tape. The machine has operated at 98% of its 2064 CD capacity for some time and the new collections that have been addressed are now stored on DVD and hard disk. The possibility of upgrading some of the player stations of the jukebox to handle DVD has proved impractical and so the maximum capacity remains fixed at 1.3Tbytes instead of almost 10Tbytes using DVD. However, even this limitation, especially for a machine more than five metres long represents unacceptable density now that 250Gb disk drives, of a few centimetres in overall size can be purchased in quantity for a few hundred dollars - or less in quantity.

These are not the principal difficulties, though. Of far greater relevance, given the expectation of current archive users, and moreover the needs of sustainability and preservation of titles, is the limited functionality of the CD itself and its automated handling by the machine. Firstly, the mechanical transport requires many seconds to retrieve a new title and any level of random access is far too slow. This means that play-out is, in practice, linear only and that titles are cannot exceed the size of the individual CD without significant delay (and in fact modification of the installation's software). This mode of operation is reflected in the structure of the jukebox, in that its mechanism reliability is such that repeated disk exchange, such as would be required to off load all of the datasets at once - for example, to clone the collection to disk - leads to repeated failure, the remedy of which requires manual intervention. This means that the jukebox is exceedingly difficult to move it has to be decommissioned over several months even to refurbish the facility that it occupies. Worse, there is only one copy of each disk and a significant overhead is incurred if a CD in the jukebox fails. As we will discuss below, the total lifetime of a CD is not much greater than magnetic tape, and the statistical mid-life failure rate of a two thousand CD collection is not insignificant. There remains the cyclic and huge task of copying entire collections in advance of deterioration due to limited media lifetime. Borgmann (1999) estimates media lifetime as follows:
'the lifespan of an optical disk is thirty years (...). On magnetic tape, information would

have lasted one year, on a videotape two years perhaps, and on a magnetic disk five to ten years.'

5. The ZKM Storage Area Network (SAN) project

The contrast between media and equipment having expected lifetimes only of a few years on one hand, and the very long term horizons of preservation and monotonic growth of storage requirements of digital media collections represents alone a 'grand challenge'. Together with declining budgets for cultural institutions under pressure in both local and national government, we must clearly consider radical solutions to avoid the situation of the broadcasting companies - resigned to watch their historical titles evaporate.

The technical feasibility of building huge disk archives has been demonstrated by companies such as Google (Shread, 2002), which routinely operates tens of thousands of computers in its caching of search keys of the internet. Also experience with the Terastore jukebox at ZKM has seen the 5m machine of six years ago overtaken by a few disks costing five hundred dollars and fitting into a shoe box - and delivering much higher performance and functionality at the same time.

5.1 Digital Media Archives in the 21st Century

Building and operating such systems within the budget constraints of the modern archive is another matter, but the key lesson from the experience of the last sixteen years is that the sustainability of new digital media archives must be a paramount concern. Condensing this experience, it seems that just a few issues fundamentally affecting sustainability also reveal potential solutions to controlling the cost of ownerhip over extended periods - both of equipment and the equally important issue of retaining generations of personnel able to operate such installations with the reliability necessary for the entrusting of title masters.

Firstly, loss-less digitisation must be employed if the tyranny of transcription accompanying end of life of media is to be broken. This immediately causes a significant escalation in the storage capacity required, since the dataset equivalent to one hour of broadcast quality video (far less film) is about 200Gb - compared with a 650Mbyte CD of MPEG. By loss-less, we mean any digitisation method which, in the case of moving

imagery for example, records all of the information available in the original and stores it using a digital representation that preserves that information frame by frame. Compression - or any other transformation - might be applied provided that it is reversible - i.e. that the original high resolution digital master can still be regenerated.

Secondly, the cloning of whole collection datasets is essential not only for security in the face of equipment failure or natural disaster, but in order economically to overcome the limited lifetime of individual disks. In order to achieve this, the access performace of the storage and the communications interconnecting an archive with a mirror installation must be high enough to enable the building of an identical multiple, potentially of hundreds of Terabytes, within a few weeks. Again this escalates the storage capacity requirements, since at least three mirrors are required to ensure the security of title masters.

Thirdly, storage equipment must support open - and widely adopted - data communications and filestore standards so that, even if standards evolve, it is possible to migrate collections from old representations to new, without loss of title information. Interoperation of products from different vendors, as part of a storage area network is essential as the long-term security of a collection cannot assume the long term survival of any single computer company. Arising from this issue is the increasing need to perpetually upgrade installations into the distant future, rather than expect to discard machines such as the CD jukebox at a definite end-of-life. It is in this are that the project has conducted specific trials with open systems, which are discussed shortly.

Fourthly, the actual location at which a title is stored - even the specific physical disk - will become increasingly unknowable as installation capacities (and complexities) escalate. The copyright and mirroring requirements necessary to protect both the intellectual and material property of titles must become the subject of automated tools of considerable sophistication if such archives are to be operated without the constant attention of highly skilled, expensive personnel and in-house expertise, which will be increasingly difficult to maintain for very long timescales.

5.2 The Contribution of Open Systems

The original motivation and dynamics that have driven the early sentiments of Open

Source Software and later, the Linux operating system, from a self-organised independent developer community to become one of the major driving forces in the software industry are now intricately entwined in a myriad different factions and corporate spin, and certainly beyond the scope of this text. However, the direct sentiment of the Free Software Foundation at http://www.gnu.org/philosophy - "The enemy is proprietary software"- remains a rally call of a global phenomenon which has now reached commercial maturity and proven worthy of forming the basis of installations of international significance (Stallman, 2002).

The contribution of Open Systems to the development and sustainability of disk-based archive systems at the ZKM is three-fold: 1. precipitation of international standards for data communication, storage formats and management protocols such as mirroring; 2. reliability and maintainability of large systems over very long periods, through widespread technical review and the ability to fix code defects and customise software to specific project needs, which is unachievable using the products of any individual company; and 3. the potential for significant reduction in costs of purchase, operation and maintenance and enhancement. For this reason we briefly review a sample of major programmes that have considered or adopted an Open Systems strategy in order to gauge, more than anything else, the security of archives implemented in this way.

5.2.1 OSS Uptake in the Public Sector

Preference legislation in favour of open-source software (OSS) or Free and Open-Source Software (FOSS) spans the globe geographically, politically, and economically. Governments in Europe (Denmark, France, Sweden, UK), the US (Oregon, Texas, Delaware), Asia (China), and South America (Argentina, Colombia, Peru, Brazil) have either already adopted or are contemplating open-source preference laws.

United States

Oregon is the first state to consider making it mandatory for agencies to evaluate using OSS in cases where it is feasible. If they are already using proprietary, software they must be able to justify its use over open source.

Australia

It is speculated that the Commonwealth Government might enact a bill that effectively mandates the use of OSS at government level. Recent statements by South Australian officials indicate that they are far from reaching concensus on the matter, but some government departments have already migrated to Linux, such as the Bureau of Meteorology (Kidman, 2003; also Sterlicchi, 2003).

European Union

In 2002 the European Commission released a study, conducted by Unisys, that advocated the pooling of OSS between EU administrations. In May 2003, the Munich City Council announced that it would migrate all of its desktop and notebook PCs (of which there are approx. 14,000) from Microsoft Windows OS to Linux. Interestingly, a spokeman for the Council cited the same reasons for switching to Linux that politicians in Oregon have cited against OSS (Galli 2003).

5.2.2 OSS in the Private Sector

Unilever

Consumer product giant Unilever announced in January 2003 that it will migrate all of its systems - in over 80 countries - from proprietary Unix to the Linux system platform. In providing an enterprise test bed for Linux and as a new member of the consortium OSDL (Open-Source Development Lab), Unilever hopes to contribute to Linux OS' uptake in private sector applications. In the more immediate term, however, Unilever has stated that it expects to cut hardware and software costs as well as improving performance. The company estimates total savings to be of the order of hundreds of millions of dollars. (McCue, 2003)

Amazon.com

Hewlett-Packard has helped Amazon.com build a Linux system architecture to support its international on-line sales business. Amazon's stated reasons for using Linux were primarily for its increased flexibility and security (according to the advertising campaign currently running in London Underground stations), as well as decreased operating expenses (cut by roughly 25%) (Shankland, 2001).

Silicon Graphics Inc.

Of more direct relevance to this text, many special effects and film and video post-production companies in Europe and Hollywood, such as Moving Picture Company, have been using Linux clusters to achieve affordable supercomputer performance for several years. In particular the principal computer graphics supercomputer vendor since the 1980s, Silicon Graphics, has used Linux in its mainstream product range for several years and supports Linux systems with its popular Maya 3D software suite.

5.3 Linux-Based Storage Networks

The ZKM research institutes as well as the IT Department (EDV) have operated large disk-based storage systems since completion of the new building for video editing, for the projects of visiting artists in the Media Museum and the Visual Media and Music and Acoustics Institutes, as well as for backup of some 600 desktop machine connections throughout the ZKM. The storage technologies employed range from proprietary enterprise servers to third party SCSI and Fibre Channel disk systems and also Linux-based RAID installations assembled in-house from components for specific research projects. The ZKM network infrastructure also comprises three different fibre optic installations which provide very high interconnect capacity within the building, capable of operating at Gigabit speeds. In the Media Library we have already discussed the digitisation of the Digital Electro-Acoustic Music Archive, which is held entirely on disk.

Development from these resources towards a storage area network serving the whole organisation is a largescale task which will not be complete until 2005, when it is planned that generic fault-tolerant disk storage will be distributed throughout the building. This will enable both load-sharing and mirroring to service very high bandwidth applications, such as the non-linear delivery of filmic resolution imagery into the Media Theatre, as well as a high level of automation of backup for research and administration and security for the Media Library collections. Network security and access protection schemes are of the highest importance in these plans.

Better utilisation of the existing heterogeneous storage, in which individual installations have been confined to specific projects and services, is being achieved by upgrading of the underlying network infrastructure, in particular through the installation of more copper

Gigabit switching and additional fibre capacity. The most significant development, however, which forms the storage platform that will replace the CD jukebox and eventually enable transition of the Media Library collections from magnetic tape to disk, is the development of open systems-based network-attached storage. This approach differs from the existing structure, in which storage is either part of, or attached primarily to specific servers or research computers, which in turn are networked. Instead, clusters of Linux-based computers have been installed which are not tailored to specific computational needs, such as film post production (although they are efficient for that role). Rather, they have high bandwidth connections to the building network using channel-bonded Gigabit ethernet and offer a variety of logical connections for client applications. The latter range from Apple-based DVD authoring workstations to iSCSI-attached film streaming and loss-less video digitisation systems, as well as flexible NFS storage space to augment IT services on an "as-needs" basis.

The architecture and management of these cluster storage systems has been developed in order to address the very large capacity and high security needs envisaged in the Media Library, as well as to explore the sustainability and cost of ownership of very large disk-based archives. Interoperation trials with proprietary SAN solutions, as well as with the wide variety of PC, Apple, Sun Microsystems and Silicon Graphics systems within ZKM, are ongoing. The support, in particular of Internet Engineering Task Force (IETF) partner Eurologic, Inc. has been particularly valuable.

The basic element of the network attached storage cluster is a fast Xeon-class processor with four 250Gb serial ATA (SATA) disk drives and a parallel ATA system disk. In quantity such drives cost less than 200 Euros and we have test experience of previous models operating continuously under high load for 24 months. The choice of mainboard, which has two on-board Gbit ethernet connections and multiple internal data busses, enables the Redhat Linux operating system to deliver the full disk bandwidth to the network. Sixteen of these storage elements are housed in a 19" rack system designed by Street Vision in the U.K., which also contains network switching and air handling. These systems are relatively easy to construct in-house as they require only plugging together components "out of the box". The largest commissioning cost is testing of compatible network and disk drivers with different Linux kernel configurations and soak testing of

assembled hardware. Mechanical packaging of the cluster, however, represents a significant departure, which has direct bearing on the long term sustainability of the installation. The 19" enclosure system and associated power and air handling is designed for very long lifetime, including periodic relocation and frequent upgrade of computing and disk. It is planned that excepting failure, disk drives be replaced within two years and mainboards at least every four years and the enclosure system is designed for maximum ease of access and has flexible fixtures for most different and planned board configurations. An important side effect of this strategy - the disk and mainboard components are both purchases below 1000 Euro - is that the storage capacity of the installation doubles every two years because of advances in technology. Computational performance grows at a slower rate, but is already at the level of many supercomputer installations.

The computational support for these installations includes comprehensive console logging to remote management systems and instrumentation of fan speeds, power supply voltages and chip temperatures so that preventative maintenance can be automated as much as possible. Diagnostics lead to pro-active notification of maintenance personnel by SMS and email to ensure maximum availability of the installation and security of data. The design of the enclosure system employs a number of new materials from the aerospace industry which, in particular, protect the individual computer storage elements from handling and shock. This means that an installation can be re-configured or even shipped with only moderate precautions (Patterson 2003).

On-going development in the project includes formalisation of management procedures and training to enable secure operation of the installation by personnel joining the organisation on contract, so that we can evaluate the level of specialised knowledge in which ZKM must make a long term investment in order to sustain these systems for very long periods. Current research in the project is concerned with the development of mirroring protocols to enable more automated management of security. There is work ongoing also in the field of new copyright protection mechanisms and in the automated production of different presentation formats for accessiblity of Media Library collections using the internet and also, for visitors to the ZKM, using handheld computers and digital assistants.

6. New Accessibility to the Media Library

While the acquisition and sustainability of largescale digital media storage is a key infrastructure that the ZKM must develop in order to ensure the future of its Media Library collections, the accessibility of those collections to the public and to researchers using widely available desk and internet-based and mobile equipment is also of great concern. In particular, the rapidly changing formats and performance of both browsing interfaces and communication channels must be addressed if media collections are to remain widely and efficiently accessible. Consequently, another research project running at the ZKM in parallel with the storage area network development concerns the archive infrastructure - the transformation of master datasets into compressed formats suitable for delivery to a multiplicity of viewing environments (archive client system) - and a service trial using wireless LAN and handheld personel digital assistants (PDAs).

It is not proposed that location-based (such as the viewing stations already described) or desk-based browsing of the archive from intra-net connected computers be entirely replaced by mobile systems. Some users actively prefer stationary access to mobile devices, for example researchers who require a desk for their studies. Moreover, the current performance of wireless LAN makes the 100Mb or even Gbit intranet highly preferable for browsing sound and image data. In contrast to PDAs, such as the PocketPCs employed in the trial, whose small display only allows a limited part of the information to be shown, desk-based stations will give access to high resolution audio-visual material as well as metadata. The PocketPCs will thus serve as an introduction and first contact with the audiovisual content.

Due to the limited transfer rates of wireless technology and the limited processing available on mobile devices, the compression rate for digitisation has to be carefully chosen. Wireless data rates are faster than most domestic internet connections and so yet another set of constraints must be addressed. All of this contrasts with more general archive principles concerned with preservation, in which it is important to digitise material without compression to the extent that storage space is available. Eventually, a

multiplicity of different digitised versions will be necessary in order to meet the demands of combinations of different communication channels and client viewing systems. The equation is complicated still further by the fact that specific clients only run specific application software and so the archive interface must not only recognise the compression level appropriate, but also determine the permissible client application interface.

6.1 Real world example

In the Media Library mobile access trial, two prototype players have been commissioned using the IDEAMA collection (electroacoustic music) can be received via PocketPCs and network based MP3 players.

The PocketPC solution works in a similar manner to the normal web client, but the interface design is adapted to the size of the small screen. This playout is connected via wireless LAN or Bluetooth. At present the audio files are streamed through a windows media server at the highest bitrate commensurate with the IPAQ processor (Strong ARM, 200Mhz). An Xscale processor version is planned, for which a flash interface will be developed to optimize usability.

Figure 3 Handheld Media Library viewer

The SLIMP3 solution delivers distributed playout under the control of a server. This device is used to play archive content at locations where computers are inappropriate. It is possible to employ limited control of the SLIMP3 using its infrared remote interface.

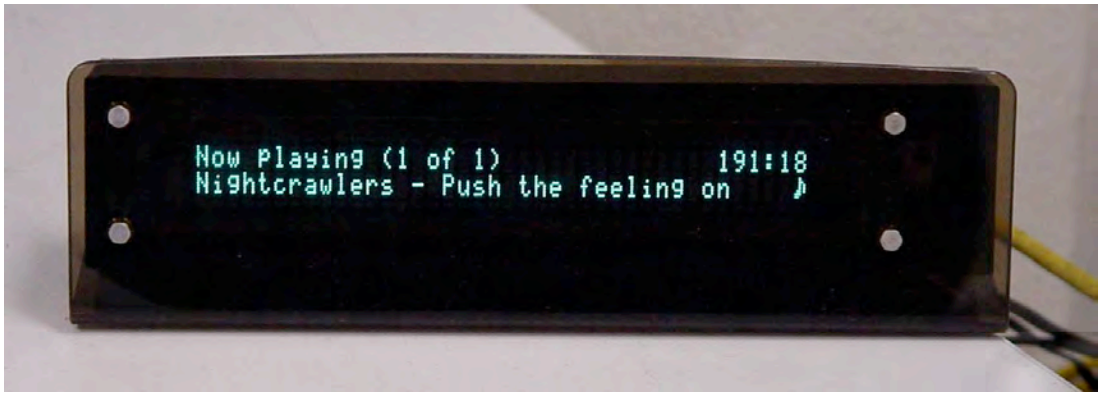

Figure 4 SLIMP3-player

SLIMP3 playout (only) client (http://www.slimp3.com). A further audio device, which is currently being evaluated is a memory card mp3 players. The user collects the music through a web interface and rents a player, which has titles pre-copied to the memory card (to prevent copying, the players are sealed).

Nearly all of the servers delivering this service are Linux-based excepting only application software not supported on Linux: a windows streaming server is operated, which is connected to a Linux fileserver to get the required media files. The database web servers are currently resident on the same machine because the load is not currently heavy enough to justify dedicated equipment. The main streaming server is also a dual Xeon-class machine, which runs a quicktime streaming server, a real server, an MPEG-3 radiostreamer and a videolan server. Due to high license costs, the realserver will be decommissioned and a migration made to quicktime, windowsmedia and MPEG (1,2,4).

Only the audio collection will initially be made available via PocketPC, since the data transfer rate required for sound is much smaller than for video works. Subsequently an equivalent presentation form will be implemented for video data. One alternative to address the problem of small screen size of the PocketPCs would be the application of TabletPCs which combine a larger screen with the advantage of mobility.

6.2 Documentation

The Media Library contains extensive information on all collection titles, which is updated on a regular basis. In addition to standardised bibliographical descriptions (artist, work title, year, duration, location), most works are linked to extensive descriptions and

artists' biographies, in three languages: German, English and French. On the PocketPCs, the bibliographic description can only be shown in a condensed form, because texts that are too long cannot be read on the small screens. Therefore two versions of the same information have to be created and made available: the short form with the essential bibliographic specifications such as the artist's name and work title and the full bibliographic information, which can be provided at a desk-based interface or via the internet for research purposes.

One solution to providing broader information, such as biographies and work descriptions, could be to transfer them into audio files, which could then be played back by the PocketPCs in the same way as acoustic works. However, this would require the narration of more than a thousand texts. A better solution, already employed by intelligent browsers like Opera, is to dynamically pre-process content for display on hand-held devices on small screens.

6.3 Copyright issues

The CD jukebox system is physically located at ZKM and the Media Library viewing stations are connected to it with a link that would be difficult to remote. Consequently the use of the archive has been limited to the ZKM building and therefore, operated in accordance with current copyright regulations (that allow in-house use). Since the new circumstances will grant access regardless of location and even to wireless users outside the ZKM, copyright infringements need to be prevented. The system will therefore have different copyright levels, which are automatically selected by access type. These type include the location and character of the client system as well as the authentication data. The authorisation subsystem will provide a profile of circumstances in which the media object can be delivered. In the case of complete rejection due to the copyright restrictions, copying will be prevented and a request generated for pay access. A special type of trusted clients is being investigated, which uses a sealed hand-held device to be lent by ZKM and recognized by the authentication subsystem.

Conclusions

The contrast between the security and persistence required of an archive and the shifting sands of information technology is one of the most important issues for media collections. Safe-keeping of individual title masters used to be a concern only of loss of quality upon copying from one analogue medium to another. With individual collections now comprising thousands of titles, it is no longer possible to employ storage technologies that require manual intervention, whether it be human or robot. The role of information technology to digitise titles, to maintain their safe storage as identical multiple copies at different locations and to protect owner's copyright, is crucial.

Experience with the ZKM SAN project indicates that there is not however a unique technical prescription for the installation and operation of digital media storage systems. Rather, a highly flexible architecture is important, employing open systems to insulate collections, metadata and retrieval systems from the impact of evolution in storage technologies, encoding and representation standards and the financial stability of individual manufacturers. For example, although it is forecast that growth in magnetic disk storage will slow before 2015 and that holographic storage will become commercially available during 2004, we expect to be able to employ these technologies on a plug-in basis as they come and go. The important thing is to be able to move collections freely and quickly between installations.

Similarly, experience with new museum data access technologies indicates that being able to convert titles and metadata into various different compressed forms to match evolving user interfaces and communication channels is more important than highly refined permanent library installations. In both of these approaches the control of conservation master datasets is paramount.

References

Borgmann, A. (1999). Holding on to Reality: The Nature of Information at the Turn of the Millenium. Chicago: Chicago University Press

Frieling, R. and H. Joerg (1999). Direct access to the media : the Mediathek of the Center for Art and Media in Karlsruhe in: Document numérique Vol. 2 No. 3-4. 1998. Les bibliothèques numériques Paris : Hermes Science Publications, 1999 ISSN 1279-512, ISBN 2-86601-738-2

Galli, P. (2003). Why Munich Dumped Microsoft for Linux. In eWeek. Posted 08-May-2003. http://www.eweek.com/article2/0,3959,1110809,00.asp

Kidman, A. (2003). No Open Source Mandate: Aust Govt. In Zdnet Australia. News & Tech. OS. Posted 19-Feb-2003. http://www.zdnet.com.au/newstech/os/story/0,2000048630,20272191,00.htm

McCue, A. (2003). Unilever Moves to Linux for Savings. In Web Active Magazine. Posted 30-Jan-2003. http://www.webactivemagazine.co.uk/News/1138361

Patterson, D. (2003). A Conversation with Jim Gray. In ACM. Queue. Tomorrows Computing Today. Storage. Vol. 1 No. 4. June 2003.
http://www.acmqueue.org/modules.php?name=Content&pa=showpage&pid=43

Shankland, S., Kane, M. and Lemos, R. (2001). How Linux saved Amazon millions. In . CNET News.Com. Posted 30-Oct-2001, 5:20 PM. CNET News.Com. http://news.com.com/2100-1001-275155.html

Shread, P. (2002). Google Experiments With Distributed Computing. In ITManagement. Posted 27-March-2002. http://itmanagement.earthweb.com/netsys/article.php/999551

Simpson, D., (2003). iSCSI gains a toehold in SAN market. In i-tech. Intelligent Storage Solutions, April 2003.
http://www.i-tech.com/Storage%20Info%20Center/indnews153.php

Sterlicchi, J. (2003), Rann's Man Slams Open Source. In Austrialian IT. Posted 15-Jul-2003. http://australianit.news.com.au/articles/0,7204,6751841%5E15306%5E%5Enbv%5E, 0.html

Stallman, R. (2003), Why 'Free Software' is better than 'Open Source'. In gnu.org. http://www.gnu.org/philosophy/free-software-for-freedom.html

Weibel, P., Shaw, J. (2003). Future Cinema:The Cinematic Image After Film. Cambridge/Mass.: MIT Press