A NEW APPROACH TO AUDIOVISUAL DIGITAL ARCHIVING

Matthew Addis, Richard Lowe, Lee Middleton

University of Southampton IT Innovation Centre, Southampton, UK

OVERVIEW

There are over 100M hours of audiovisual material in Europe's archives. The data volumes are huge (hundreds of Petabytes in total) and will double within two years. Most new material is born digital. This material needs to be kept safety, securely and with high levels of content integrity for 50 years or more.

At the same time, digital audiovisual archives are becoming 'embedded' as services within wider networked infrastructures and content-centric processes. The business models and processes surrounding the storage, preservation and access to digital assets are evolving fast; access to archive content now takes place electronically and across organisational boundaries. This is partly fuelling a nascent but growing market for outsourced archive hosting as a service.

This paper presents work in the UK AVATAR-m project on how to specify and govern federated archive services that involve both local and remote storage. AVATAR-m is a UK collaborative R&D project supported by the Technology Strategy Board where the IT Innovation Centre, BBC, Xyratex and Ovation Data Services are developing an innovative approach to large-scale long-term digital archiving within distributed storage infrastructures

Our focus is how to achieve high levels of data safety when multiple storage services are combined. Techniques we use include AV asset decomposition and replication across storage locations and the use of proactive data integrity monitoring. This together with simulation and modelling techniques allow trade-offs between the level of data safety (availability over time expressed as a risk of loss) and the cost (storage, networking, processing, management, maintenance etc.) to be investigated.

Our tools allow organisations to profile the generation and consumption of archive assets including the requirements for safety, security, longevity and accessibility. These profiles then allow storage provision to be planned in terms of long-term access, ingest and retention and technical specifications created ready to be matched against storage solutions or managed services.

BACKGROUND

Current projections are that over 90% of all new information is digital and that the volumes generated over the next two years will be larger than the total volume of all information ever created previously in human history (see Figure 1). This is as true for audiovisual content as it is for other types of digital information. For example, YouTube is growing at about 20 percent every month, which equates to over 300% per year. In the professional audiovisual (AV) archive world, UNESCO estimates there are 100M hours of content in existence. Broadcast archives project that this will grow at 5M new hours every year, which given increases in frame rates and resolutions can mean a data volume doubling in as little as 18 months.

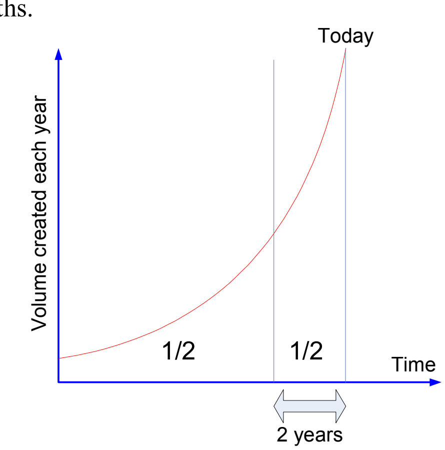

Figure 1 Exponential growth in AV archives

Audiovisual content presents demanding challenges for digital preservation, especially given the preservation ideal of storing content uncompressed. Standard Definition digital video has an uncompressed data rate of about 270 MBit/s and even when stored with compression, e.g. 50MBit/s DV, multiple Petabytes of storage are required for a typical broadcast archive. HD requires five times as much space. In digital cinema, 4K requires up to 30 times the data rate of SD and for 3D cinema with twin data streams at up to 144 fps the volumes are truly vast. This presents a real problem. The cost of maintaining this content is uncertain where estimates range from 'half the price of analogue' [6] to nearly 'twelve times higher' [7].

When considering all the new devices, techniques, services and business models with which to create, distribute and use audiovisual media, it is essential to consider how this content will be archived and how it can be maintained in a way that allows it to be easily accessed and used for years to come. The creation, consumption and archiving of audiovisual content are inseparable topics, yet archiving tends to be an ‘after thought’ and is often neglected compared to the more immediate concerns of how to best take advantage of new forms of content and the user experiences they bring.

Mass storage technology, e.g. disk arrays or tape robots, appears to be an obvious solution to storing large volumes of content and making it easily accessible. However, the long-term safety of content in these technologies is far from assured. There is a widespread assumption that bit-level preservation using mass storage technology is a solved problem, e.g. using RAID disk and offsite tape backup. However, the reality is that for large data volumes (e.g. the petabyte level) data corruption or loss can be caused by failures in hardware, bugs in software, and human errors. Field studies of large disk-based systems, e.g. by CERN [24] and [26], reveal data corruption taking place silently without detection or correction, including by 'enterprise class' systems that are explicitly designed to prevent data loss.

Storage manufacturer metrics of MTTF (Mean Time To Failure) or MTTDL (Mean Time To Data Loss) for media or systems are meaningless as a measure of the ability to preserve data [23]. They deal with the case of complete and catastrophic loss of all the data in a system. They neglect that data loss can actually take place incrementally and in a way where the corruption of just a few bits of information can render large parts of a video file unusable due to the way the content is encoded. Furthermore, manufacturer data on MTTF are based on their own models or tests and often don't match observations in the wild by Google [25], NetApp and others.

The implication for digital archiving when using mass storage technology is the need for a continuous activity of data integrity checking and repair, which in turn requires copies of content to exist in multiple storage systems in different locations. Recognising that mass storage technology can’t be relied upon is a key feature of our approach and we take the view that no storage technology or storage service provider should ever be relied upon to maintain data integrity no matter what their MTTDL or Service Level Agreement might state.

STATE OF THE ART

Whilst there is intensive interest in preservation strategies for digital content [1][2][3][4], in general there is little work on practical implementations tailored for the needs of audiovisual content. This includes archive integration into production, distribution and consumption workflows, with dynamic preservation processes required as a consequence. For example, the OAIS Reference Model [5] defines some of the processes required for long-term preservation and access to information objects, but does not specify how to monitor audiovisual objects or the systems they are stored in, identify when migration should take place or to what an audiovisual object should be migrated to.

More widely, different archive implementation models need to be considered including value chains and business models delivered through multiple service providers or organisations (e.g. outsourced services, federated preservation across organisations etc.). These value-chains and business models are liable to evolve rapidly over time because of the relative rates at which storage, networking, processing are evolving [8], e.g. as evidenced by the explosion in online services such as Amazon S3, EC2 and SQS [9].

The economies of scale, power, cooling and staff costs that can be achieved by organisations like Google [10], mean that as network costs continue to fall, in-house solutions will become increasingly expensive compared with outsourced or federated models.

Different approaches will be applicable depending on the type and volume of content or the need for access across organisational boundaries, and the use of mixed models is likely considering robust preservation strategies typically involve multiple copies of content in multiple locations to mitigate against technical obsolescence or content loss.

Whilst audiovisual archives typically use dedicated in-house systems for storage and processing (e.g. transcoding) of their assets, various technologies exist to support data federation and remote data services in distributed environments. Many have emerged from the Grid community, including storage services and high-performance data transfer tools, e.g. GridFTP[11], SRB[12] and RFT[13]. These are used as part of Data Grid Management Systems[14] to support the needs of large-scale scientific applications e.g. High Energy Physics Experiments at the CERN LHC. iRODS[15] is one that has already been used for digital library applications, persistent archiving, and real-time data systems, where management policies (sets of assertions that these communities make about their collections) are characterised in terms of rules and state information.

Remote access to archive hosting services is yet to emerge in the broadcast industry, although there are services for remote access to data for distribution, e.g. VIIA from Ascent Media[16] and data transfer within the enterprise, e.g. DIVAGrid from Front Porch Digital [17]. The digital library community has meanwhile been busy creating software frameworks for implementing preservation environments. These include open source solutions, e.g. DSpace[18] which provides standard services for ingestion and access and is ported to run on top of SRB for managing distributed data, Fedora[19] which associates display functions with each data type, allows relationships to be imposed on records, and maps semantic labels on records to an ontology, as well as simple, off the shelf systems such as Greenstone[20], and commercial systems, e.g. ExLibris[21] – however none are designed specifically with the challenges of AV content in mind.

In summary, the technology state of the art is one of fragmentation where individual communities, e.g. broadcast and digital libraries each provide pieces of the puzzle. The challenge is one of integration and adaptation to the specific challenges of audiovisual content.

AVATAR APPROACH

Storage in AVATAR-m is heterogeneous, reflecting the broad range of storage types that an archive may want to utilise, both in-house in locally managed IT systems and remotely through storage provided as a service.

The emphasis of our solution is on networked storage, such as spinning disk or media jukeboxes, e.g. in a SAN or as NAS. Additionally, online remote storage provided as a service is also supported to allow archives to make use of third-party storage services such as Amazon S3 or use storage in remote locations over protocols such as ftp or http. Our approach is to combine these disparate storage types and locations, so they are aggregated together into a single storage solution (Figure 2).

Adapters are used for each storage type that the storage aggregator interfaces to, but since most operations are done at the file system level additional adapters are only required for storage services, which offer different APIs (Figure 3).

Rather than assigning each AV asset to a specific tier, available storage locations are ranked dynamically using a cost function based on factors such as the current and average read/write rates and availability. This model can be applied to storage services too, e.g. by making use of performance monitoring services such as CloudStatus [22].

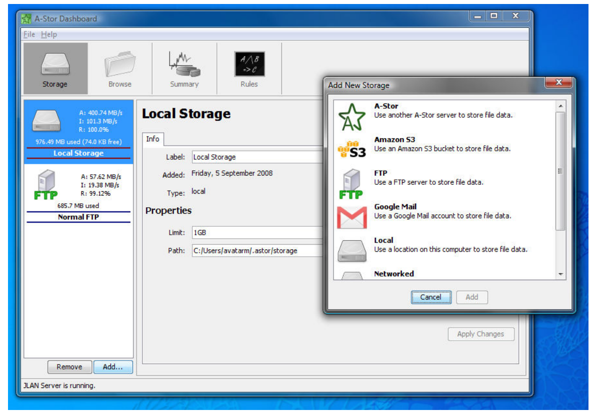

Figure 2 Adding storage locations to an AVATAR archive

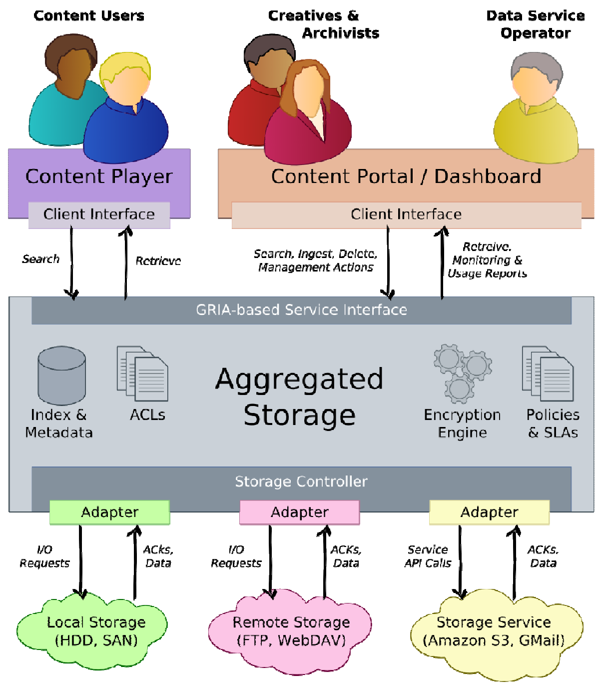

Figure 3: AVATAR-m aggregated storage

The use of the storage is also monitored ensuring that content that is accessed frequently is made available from higher-ranked (and therefore faster) locations, whilst content that is not accessed often is moved to slower storage (Figure 4)

The rules that determine what gets moved can be modified through management policies that can be assigned to specific items or classes of items, such as all files of a certain type or belonging to a certain user or project.

The interface to AVATAR-m from a user perspective, i.e. someone who produces content that needs to be subsequently archived, or someone who needs access to content in AVATAR-m, is simply a network accessible file share.

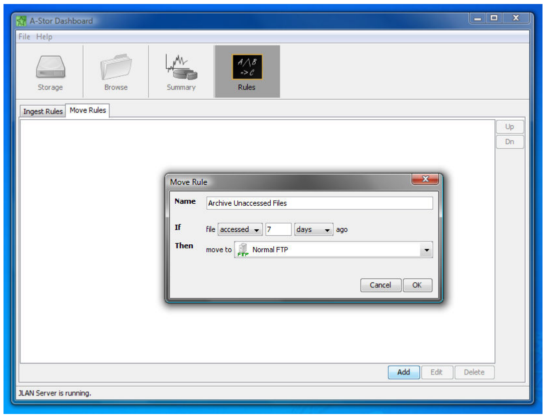

Figure 4 Specifying data move rules in an AVATAR archive

The details of where a particular file is stored, or how it is accessed and retrieved, are completely hidden from the user. All they see is the ability to access files as if they were on a local file system (Figure 5)

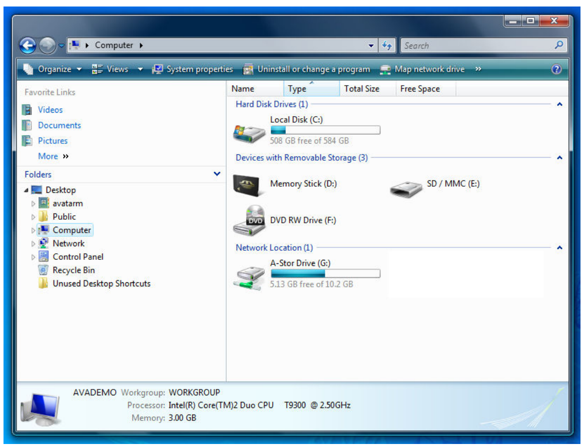

Figure 5 AVATAR archive exposed as a file share (G:)

Our approach is in some ways similar to hierarchical storage management (HSM) systems, but with the advantage in our case of being able to utilise third-party storage services as well.

The key benefits of this approach are:

- Direct control is not needed over the storage systems used at individual storage locations, which allows the approach to be used with storage service providers or in federated archiving.
- Storage of an AV asset can be optimised on an asset-by-asset basis, i.e. assets need not all be treated in the same way simply as large data files.
- Multiple copies of an asset can be stored and managed transparently to the users of the asset.

In AVATAR, audiovisual assets are broken down into ‘chunks’ that are distributed and replicated across the various storage locations in the system. Each chunk is stored as a file and AVATAR maintains an index of where the chunks are and how they should be reassembled to create the AV asset. All chunks and their parent asset are MD5 hashed to create digests for use in integrity checking, both when the file is retrieved on request of a user and periodically as part of active integrity monitoring and repair.

The allocation of chunks to storage can be done based on rules, for example to select the storage locations with highest performance (Figure 6).

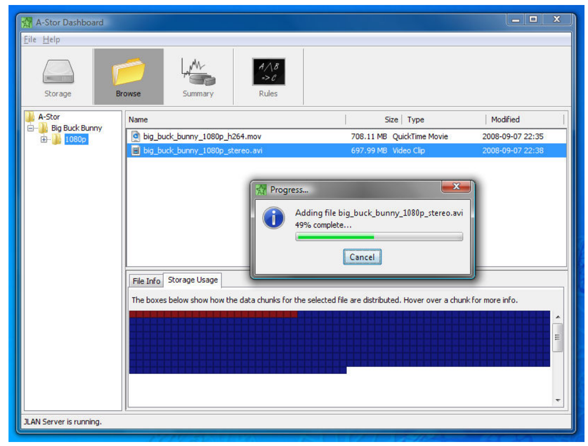

Figure 6 AVATAR file splitting and distribution

The decision on how to chunk a file (e.g. how many pieces, what size, what type) can be made on an asset-by-asset basis according to rules applied to that asset and the context in which it needs to be archived. This allows sophisticated archiving policies to be enacted, e.g. to allow varying levels of replication to be applied to different parts of an asset (e.g. video, audio, metadata, key frames, header, index etc.). This recognises that not all parts of an AV asset are ‘equal’ when it comes to preservation and indeed not all collections of items should be treated equally either. Architecturally, we use plug-ins to be added that can understand files at an application level and then decide how best to chunk these files, e.g. by disassembling an MXF asset into component pieces, or extracting key frames from an MPEG.

Experience with an initial implementation has found that there are often significant differences between the parameters with which storage services are defined (storage capacity, access latency, delivery bandwidth etc.) and the level at which archive operators characterise their archive (rates and volumes for ingest and access, retention scheduling to encapsulate value, preservation priorities and asset safety). To address these differences, we developed a storage planning tool that allows archive requirements to be specified using parameters (e.g. data volumes and data i/o) that are both application and technology implementation neutral. The tool can be used by an archivist, external service provider, or in-house IT manager to define

SLAs in archivist terms or to interpret resource-level SLAs. Through a series of screens, the user can specify one or more collections of assets and the associated ingest, access and retention profiles. For example, a collection might be born digital content of a particular genre or it might be a particular type of analogue carrier being migrated into digital form in a preservation project. The ingest profile specifies the rate at which items are put into the archive and can be expressed in various ways, e.g. items per month or terabytes per year (Figure 7).

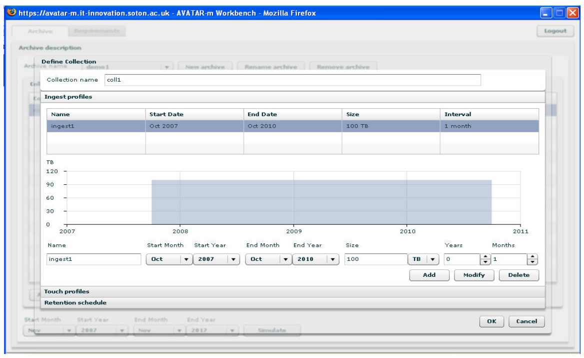

Figure 7 Example ingest profile

The access profile specifies how often material is likely to be accessed and can be expressed as an average rate or as a periodic activity. The retention schedule specifies how long each item of content needs to be retained before it is re-appraised and includes an estimate of how much content is likely to be retained after that point (Figure 8).

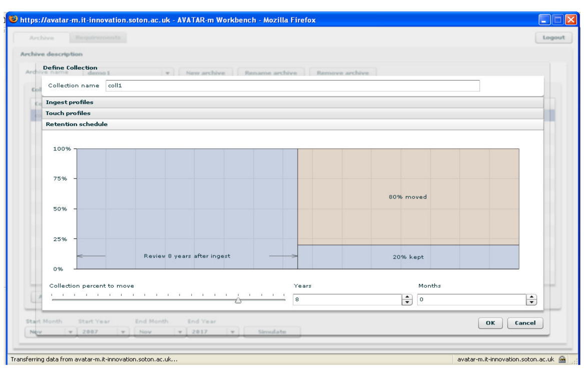

Figure 8 Example retention schedule

Ingest, access and retention profiles are aggregated across the collections to define the overall needs of the archive. The tool allows simple storage solutions to be simulated (e.g. tape libraries) using technology roadmaps (e.g. LTO data tape) to profile investment and migration and find deviations from the archive needs, e.g. resulting from device contention during concurrent migration and access (Figure 9).

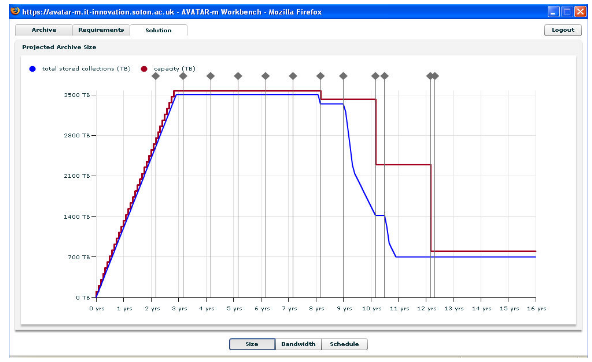

Figure 9 Example storage plan

The ability to separate the concerns of archiving at the 'business level' and storage at the 'technology level' along with our system that allows storage to be combined from both local and remote sources, means that we can support new digital archive value chains, for example archiving as a service (Figure 10).

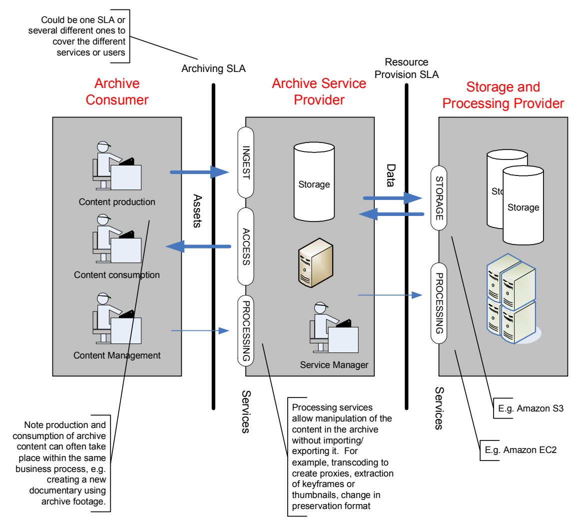

Figure 10 Archive value chain enabled by AVATAR

In the next phase of the project, we plan to develop a combination of process modelling and statistical techniques to calculate the workloads placed on an archive from the processes that involve the archive, including ingest, access, transcoding and maintenance (e.g. through migration). This will be combined with models of the failure modes (Figure 11) for storage in the system and how these can be monitored and content repaired.

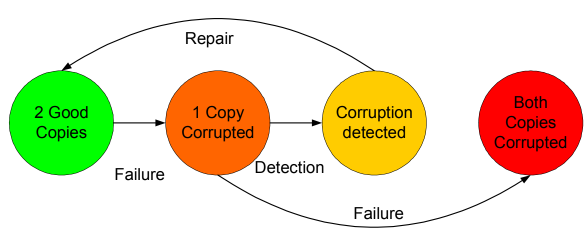

Figure 11 Failure, detection and repair in a two-copy system

This will allow archives to plan out how to use storage in a way that is both cost effective and results in the required level of safety for the assets stored within it. More advanced requirements estimation will form the basis of round-trip capacity planning, SLA definition, archive service provisioning, and service usage auditing and reporting, including the case where archive hosting is outsourced.

We expect to extend our use of automated archive rules, for example so we can support policies for who can access what. We also plan to allow rules to be applied to different types of content, e.g. for transcoding video content on ingest to create proxies for access, or to be triggered in reaction to system events so we can react to drops in availability of storage locations, e.g. if a copy is on Amazon S3 and it goes offline then replicate one of the remaining copies to another location to maintain safety.

Our ultimate objective is to demonstrate a decision support tool (dashboard) for planning, monitoring and managing archiving using distributed storage infrastructures in a way that allows suitability, flexibility, scalability and cost to be investigated, trade-offs to be explored, and best-fit solutions to be chosen from the perspective of both the consumer of the services and the provider of the services.

CONCLUSIONS

In this paper we argue that audiovisual archiving is becoming an integral part of content production, distribution and consumption processes. The use of service oriented models, including the delivery of archive hosting through third-party services, provides the key to integrating archiving activities into wider media-centric environments in a way that still allows the archivist to achieve their primary mission – the safety and longevity of their assets.

However, using digital mass storage technology, be it within local systems or through storage as a service providers, is not without risk of data corruption. Therefore, active integrity checking techniques are essential and techniques need to be used that make no assumption that any part of the overall system will ever be 100% reliable.

In order to combine heterogeneous storage into an archive solution that is still usable for producers and consumers of content, techniques are needed to allow the combined solution to be presented transparently to the user, e.g. as a network file system, so that archive storage can be directly integrated into modern production, post-production and distribution environments and workflows.

AVATAR-m addresses these challenges through the use of aggregated and federated storage, a service oriented infrastructure to access and manage this storage, and user interface tools to help with capacity planning and decision support. This allows archive owners to concentrate on the long term management of their content in a secure, safe, and cost effective manner.

ACKNOWLEDGEMENTS

AVATAR-m is a £3M UK Technology Strategy Board supported collaborative R&D project that started November 2006 and runs until 2010. The project involves the BBC, Ovation Data Services, Xyratex and the IT Innovation Centre. For more information please contact avatar-m@it-innovation.soton.ac.uk

REFERENCES

[1] Digital Preservation Lifecycle Management for Multi-media Collections. Arcot Rajasekar, Reagan Moore, Fran Berman and Brian Schottlaender. Book Chapter. LNCS http://www.springerlink.com/content/p3m497wt5452012u/.

[2] CLIR pub106 Building a National Strategy for Preservation: Issues in Digital Media Archiving http://www.clir.org/pubs/abstract/pub106abst.html

[3] Evaluating preservation strategies for audio and video files. Carl Rauch et al. http://www.ifs.tuwien.ac.at/~andi/publications/pdf/rau_digrep05.pdf

[4] A Foundation for Automatic Digital Preservation, Miguel Ferreira, Ana Alice Baptista and José Carlos Ramalho Ariadne Issue 48 July 2006. http://www.ariadne.ac.uk/issue48/ferreira-et-al/

[5] CCSDS 650.0-B-1: Reference Model for an Open Archival Information System (OAIS). Blue Book. Issue 1. January 2002. This Recommendation has been adopted as ISO 14721:2003. http://public.ccsds.org/publications/archive/650x0b1.pdf

[6] "Archiving Movies in a Digital World". David Cavena et al. SUN Microsystems report, VERSION 2.1, June 8, 2007; 29pp David.Cavena@Sun.COM http://wikis.sun.com/display/SunMediaSpace/2007/11/05/Archiving+Movies+in+a+Digital+World

[7] "The Digital Dilemma: Strategic Issues in Archiving and Accessing Digital Motion Picture Materials" Academy of Motion Picture Arts & Sciences, 2007; 74pp available from the academy: http://www.oscars.org/contact/council.html

[8] For example, see Gilder's law: http://www.netlingo.com/lookup.cfm?term=Gilder's%20Law. Edholm's law: http://ieeexplore.ieee.org/iel5/6/29070/01309810.pdf?arnumber=1309810. Kryder's law: http://www.mattscomputertrends.com/Kryder's.html.

Moore's law:
http://en.wikipedia.org/wiki/Moore%27s_law
[9] Amazon Web Services:
http://www.amazon.com/AWS-home-page-Money/b/ref=sv_hp_4?ie=UTF8&node=3435361
[10] IEEE Spectrum: Google Watches Its Watts
http://spectrum.ieee.org/jul07/5382
[11] GridFTP
http://www.globus.org/grid_software/data/gridftp.php
[12] Storage Resource Broker
http://www.sdsc.edu/srb/index.php/Main_Page
[13] Reliable File Transfer
http://www.globus.org/toolkit/docs/4.0/data/rft/
[14] DGMS
http://www.sdsc.edu/srb/index.php/DGMS
[15] https://www.irods.org/index.php/IRODS:Data_Grids%2C_Digital_Libraries%2C_Persistent_Archives%2C_and_Real-time_Data_Systems
[16] http://www.ascentmedia.co.uk/
[17] http://www.fpdigital.com/
[18] DSpace http://www.dspace.org/
[19] Fedora http://www.fedora.info/
[20] Greenstone http://www.greenstone.org/
[21] ExLibris digital library solution:
http://www.exlibrisgroup.com/
[22] CloudStatus http://www.cloudstatus.com/
[23] Bit Preservation: A Solved Problem? David H Rosenthal, Stanford University. In proceedings of iPRES 2008: The Fifth International Conference on Preservation of Digital Objects, British Library, London. http://www.bl.uk/ipres2008/index.html
[24] Silent Corruptions, KELEMEN Péter. CERN IT. LCSC 2007, Linköping, Sweden.
[25] Failure Trends in a Large Disk Drive Population, Eduardo Pinheiro, Wolf-Dietrich Weber, Google. Proceedings of the 5th USENIX Conference on File and Storage Technologies (FAST'07), February 2007
[26] Are Disks the Dominant Contributor for Storage Failures? A Comprehensive Study of Storage Subsystem Failure Characteristics. Weihang Jiang, Chongfeng Hu, and Yuanyuan Zhou, University of Illinois at Urbana-Champaign; Arkady Kanevsky, Network Appliance, Inc. FAST '08 pp. 111–125 of the Proceedings.