DELIVERABLE

Project Acronym:	DCA
Grant Agreement number:	270927
Project Title:	Digitising Contemporary Art

Project Acronym:	**DCA**
Grant Agreement number:	
Project Title:	**Digitising Contemporary Art**

D6.2 Best practices for a digital storage infrastructure for the long-term preservation of digital files

Revision: 1.0

Author(s): **Sofie Laier Henriksen, Wiel Seuskens and Gaby Wijers (LIMA)**

Project co-funded by the European Commission within the ICT Policy Support Programme		
Dissemination Level		
P	**Public**	**x**
C	**Confidential, only for members of the consortium and the Commission Services**	

REVISION HISTORY AND STATEMENT OF ORIGINALITY

Revision History

Revision	Date	Author	Organisation	Description
0.1	09/12/2012	Sofie Laier Henriksen / Wiel Seuskens / Gaby Wijers	LIMA	Text
0.2	06/03/2013	Rony Vissers	PACKED	Internal review
0.3.	23/03/2013	Sofie Laier Henriksen / Wiel Seuskens / Gaby Wijers	LIMA	Amendments
0.4.	21/04/2013	Robert Gillesse	DEN	External review
0.5.	26/04/2013	Sofie Laier Henriksen / Wiel Seuskens / Gaby Wijers	LIMA	Amendments
1.0	06/05/2013	Rony Vissers	PACKED	Final version

Statement of originality:

This deliverable contains original unpublished work except where clearly indicated otherwise. Acknowledgement of previously published material and of the work of others has been made through appropriate citation, quotation or both.

Table of content

1. Introduction

This deliverable focuses on bit preservation, i.e., the physical storage of the data and how today's institutes and companies store their collections of data. It gives an overview of storage media types and systems available for creating a good storage infrastructure. Each storage type is explained in terms of its properties and what advantages and risks the system gives in view of long-term preservation. The best practice can be used as a guideline for those in charge of maintaining contemporary art collections, digital archives, in the event of buying new equipment or when making a preservation plan. Storage is defined in this deliverable as a hardware media on which one can store digital content. In more strict preservation terms a digital storage infrastructure for long-term preservation can be seen as the Archival Storage entity in the OAIS (Open Archival Information System) model. Archival storage is a storage infrastructure, which provides the means to store, preserve and access digital content. The OAIS model was explained more thoroughly in our first deliverable *D.6.1 Guidelines for a long time preservation strategy for digital reproductions and metadata*. The D6.1 emphasis was on the so-called logical preservation of data, i.e., how to preserve the authenticity and integrity of the files. Below is a brief summary of the main points:

- Making sure there is someone who is responsible for taking care of the data, and that his/her role is carefully defined.
- Planning renewal and maintenance.
- Making a decision on which strategy to use and when.
- Having several copies of the same data.
- Having the copies on at least two different types of hardware storage.
- Storing backups at different locations.
- At ingest, making sure there are a few well used formats, well described metadata and persistent identifiers, so that the data can be found and identified.
- That the preservation plan and all the equipment is evaluated at least every five years.

If an institution can be integrated with the different departments in the OAIS model, and it can use different preservation methods such as migration and emulation, it is possible to make sure the files can be understood and accessed years from now. The idea is that the D.6.1 and D.6.2 deliverables complement each other. One cannot make a good plan without considering both the logic preservation and the physical storage aspects.

Most storage systems have been developed with the aim to increase speed of access and file sharing, without the intent of long-term preservation. This should be kept in mind when planning long-term storage and preservation of content, because keeping a collection of data does not necessarily mean needing frequent access. Some properties might be beneficial if one needs a storage and network system for everyday work, but a seemingly “slow” solution might be better for preservation purposes. Therefore it is important to make a distinction between data for preservation and data for access. Always have a preservation master separate from distribution copies and frequently used files. A preservation master can often be a larger, high quality file and a distribution copy is usually smaller and easier to share. Regardless of the type of hardware used for storing data, it is a physical device with the risk of mechanical failure. So in order to maintain the data content the equipment has to be exchanged and checked every now and again, in order to avoid losing access to the data. Errors and loss of data can also be caused by human factors: if someone is not paying attention it can cause major damage. This emphasises the need to deal with redundant data and provide for backup, preferably on different kind of media and on different locations. Also, to prolong the lifespan and lower the costs the hardware will have to be stored and handled correctly.

2. Storage media infrastructures

Infrastructures can have different sizes that are sui(for varying applications and complexities. A small system can be stand-alone and be manually handled; a larger system often needs automated management systems and retrieval functions. Either way the complexity required all depends on the size of the digital collection, the facilities, and available budget and staff. Below, the different hardware types and the technical installations needed to run them most efficiently will be explained.

2.1 Ideal storage media properties

Some types of storage media have come and gone just as fast as the digital formats and programs have been developed. Examples of hardware that are already obsolete are DAT tapes (Digital Audio Tapes) and floppy disks. They went off the market along with their reader/writers (playback equipment), which causes difficulties when information on old DAT tapes, floppy disks or out-dated computers needs to be extracted. Hardware is also prone to physical decay if it is subjected to disasters or not properly stored and handled. This results in mechanical failures, rusting, mould growth, delamination of plastic layers, magnetic layer disruptions etc. In addition, the manufacturing of hardware can sometimes be inconsistent, causing some hardware components not to work as expected. This can all eventually lead to problems reading and retrieving the data. It is therefore important for an institution with a digital repository to consider recommendations for types of physical storage media and how to maintain them for long-term preservation. There follows some advice on preferred properties for storage media.

Preferred properties of storage media:

- **Easy to use**: The more information each medium can contain, the fewer units of media are needed. A small number of media is easier to handle and maintain.
- **Easy to copy**: How fast and easy is it to copy between two media of the same type? Can it be done without the risk of losing data? How easy is it to extract data from the media to new storage technology? The easier it is, the more time and money one saves.
- **Widespread usage:** The more widespread the technology is, the longer it will typically survive on the market. It will be easier to obtain new equipment that can read or utilise the media as well as having updates or repairs done.
- **Low price:** What is the price of obtaining new versions of the same media? What does it cost, and does the media need a reader in addition? How much storage space can you get for your money and how much do you need in the future? Think of the long term, it might be a cheap start up solution, but over time it might become expensive. Another solution might be better.
- **Durability:** What is the life expectancy of the media, before any risk of errors occur? Make sure it is a stable format that has a well-known life cycle. This can be incorporated into a preservation and maintenance plan.

(Ravn-Grove 2010)

It can also be helpful to consider the physical facilities and personnel available when picking the right system and services. Also an institution should be prepared to provide the best environment and storage space available. If a given institution already has an ICT infrastructure[1] and a room that can meet the requirements, maybe it's best to keep the collection in-house. However, if it doesn't have such parameters then it should consider hiring someone or outsourcing the tasks.

The following topics should be considered when investing in storage media:

- **Storage facilities:** Size of room needed, temperature and ventilation possibilities, risk of flood etc.
- **Maintenance:** Personnel to keep track of the hardware, test and check files integrity.
- **Need for expansion:** If new storage is needed is there room for it? Can it be integrated with the system? When developing new equipment, its size tends to decrease, so one might be able to use the same storage space with more data capacity.
- **Environmentally:** Does it make a lot of noise, heat up, or use a lot of electricity?
- **Budget:** What does it cost to set the system up, compared to its maintenance, including the electricity bill and so forth?

These criteria can be held up against the different available solutions in order to find the most suitable one for one's digital content. At the same time it is recommended that one should have more than one storage type, to avoid difficulties retrieving data if any errors occur on specific hardware. Which makes it possible to "mix and match" to best suit the four criteria of any institution:

- **Size of data collection**
- **Budget**
- **Internal staff/expertise**
- **Use of data (for access or preservation)**

2.2 Network systems

Before we explain the different type of storage media, it is useful to know how the storage can be connected to local and external networks, making it possible to access the common data set from various computers. A network system setup is independent of the type of storage; whether it is Hard Disk arrays, tape libraries, optical disk jukeboxes or even Solid State Drive arrays. In order to access files that are stored remotely, i.e., not directly on the computer you are working from (let's call this the host computer), there are different network setups to access storage devices. These networks connect the storage devices with the host computer, making it possible to retrieve the files needed. Different types of network setups will be described in the following.

NAS [2] (Network Attached Storage), a NAS server is a preconfigured file server[3] with internal hard disks, which makes its storage capacity available via LAN [4] (Troppens 2009). This type

[1] ICT= information and communication technology, meaning an IT network and server system set up and maintained by an IT department.

[2] Network Attached Storage, a NAS server is a preconfigured file server with internal hard disks, which makes its storage capacity available via LAN (Troppens 2009, storage network explained).

of server system was developed for file sharing. It is relatively easy to install and expand a NAS server system. Provided it does not become a giant system, with many individual departments connected, it is also easy to maintain making it a cheap solution. However, NAS servers are not suitable for providing storage space for databases that need frequent access, because each file must pass through the internal system of the file servers twice before the files arrive at the computer where they are required (Troppens 2009). This makes the system slow and not suitable for quick access to databases, backup procedures or multimedia applications. It could be a good network solution for digital preservation specifically.

SAN (Storage Area Network) is an advanced network connecting storage devices to servers, making the devices appear like locally attached disks to the operating system (Tate 2006). The network basically "cheats" the file system into believing the disk is locally connected, even though it's not. A SAN consists of a communication infrastructure, providing physical connections and a management layer, which organizes the connections, storage elements, and computer systems to make data transfer more stable (Tate 2006). Usually SAN is the type of network used for a cloud server setup (For more on Cloud storage see chapter 7).

Both NAS and SAN networks can be achieved by using for example Fibre Channel or iSCSI (Internet Small Computer System Interface) technology. Fibre Channel is at present the predominant storage network technology and allows serial transmissions for high speed and long distances, low rate of transmission errors and little delay of transmitted data (Troppens 2009). The transmission through Fibre Channel is only possible if the storage devices have some kind of connection point, called a port. The connection between different ports can be set up in different ways, depending on storage size and the complexity of the network. iSCSI is a form of technology used for serial input and output just like Fibre Channel. Fibre Channel and iSCSI are often referred to as I/O[5] buses (Input/Output), because they transfer data between components in a computer/server.

2.3 Storage media

In the following chapters there will be an introduction to different types of physical storage media and services that follow with them, used in institutions today. There will also be a description of pros and cons for each type of storage media as well as a case study on an institution utilising that type of media. The institutions do not necessarily use only one type of storage media, because different media has different benefits, which will be highlighted at the end of this section on storage media systems section.

[3] A file server is a computer or computer program that manages access to a centralised resource or service in a network.

[4] Local Area Network, a data network with low geographic extension (maximum several tens of kilometres).

[5] Input/Output.

3. Data tapes

The technology of data tapes (magnetic tapes with the capacity of storing digital data) has been around for almost thirty years (Mearian 2012). Since its invention, a large range of different types have been developed, some of them already obsolete (e.g., DAT tapes) while others are still going strong. Today magnetic data tapes are one of the most widely used hardware storage devices for digital files in larger archives, together with hard disk drives. There are two types of tape housing design: the single-hub and the dual-hub cassette (Bradley 2009). Single-hub tapes ensure minimum friction and accurate tape guidance, but the dual-hub is faster, can contain more data on the same amount of tape and is easier to spool through. The most common dual-hub cassettes are the 8mm helical scan tapes by Mammoth and Sony AIT. However, a quick search on the Internet shows that the dual-hub tape technology is no longer being produced; with Sony declaring in 2009 they would completely stop the development of AIT in 2010 (Murai 2009). If an institution has data stored on helical scan tapes it is therefore recommended that it be migrated onto other storage media before equipment becomes too difficult to procure. Single-hub LTO type is the most well known option when choosing tapes and is therefore the focus of this chapter.

3.1 LTO (Linear Tape Open)

The most widely used tape in the field of digital preservation is LTO Ultrium (Linear Tape-Open) (see Figure 1). The word “linear” refers to the way the tape drive reads or writes the tape, i.e., it is read with a fixed head and not with a helically / rotary head (Bradley 2009). The LTO Ultrium was first developed in 1990 as collaboration between HP, Quantum and IBM. It is an open format, in order to provide a wider range of products that might be compatible with the tapes, and also to make it easier for different vendors to make tapes and reader/writers (see Figure 1) with the same properties. When creating a digital archive it is always good to pick open standards, because it there is a lower risk of obsolescence, a wider choice of brand and therefore lower prices.

Figure 1: a LTO-4 tape and a LTO-5 reader/writer.

There are a lot of companies with their own brand of LTO Ultrium tapes, such as:

- IBM MTC
- Quantum
- Sony
- Storagetek
- HP Ultrium

Regardless of the brand, the LTO tapes should work in a similar way; so an institution can have more than one brand of LTO tapes. What is important is that all drives (reader/writers) are compatible with the tapes and the indexing systems one chooses. It can be difficult for the drives to read some types of tapes and it will be costly to invest in many different drives. Since 2000, a new generation of LTO tapes have been made every two-three years and the last generation was released in 2010. This was generation five, and was called LTO-5. According to the manufacturers it has a data capacity of 1.5 TB (terabyte) for uncompressed material and a transfer rate[6] of 140 MB/s[7] (Pease et al 2010, Mearian 2012). LTO-6 was released in December 2012 and has a capacity up to 8 TB and data rates of 525 MB/sec (see Figure 2). However, tests that LIMA has performed show that the actual capacity of compressed video data only has a benefit of 2%, which means it's not really beneficial to use compressed material and that a LTO-6 tape only has a storage capacity of about 4 TB and not 8 TB.

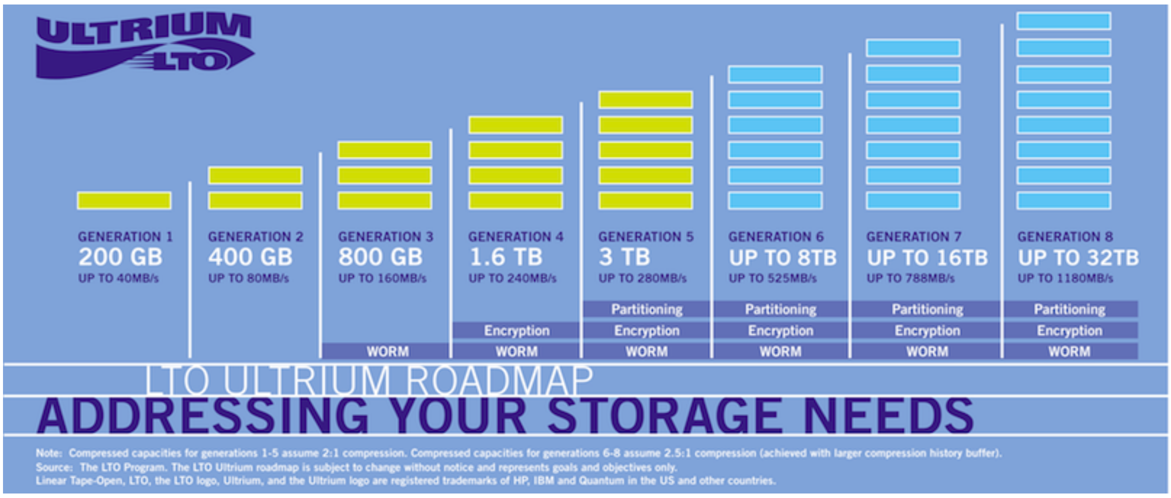

Figure 2: development of LTO tape storage capacity from www.ultrium.com, 2012

Before generation five, tapes didn't have any file indexing system, making it difficult to search and find files on the specific tapes. However, with the release of LTO-5, a system called Linear Tape File System (LTFS), it is now possible to create self-describing tapes (Pease et al 2010). This means that there is a built-in index on each tape with a file system making it as easy to find files as on a USB stick. One should be aware, however, that the LTFS software is hardware specific, i.e., not all readers/writers can understand the software. There is LTFS specifically for HP and specifically for IBM, and they are not compatible. There are some open source LTFS programs, but they are not quite on the same usability

[6] The time it takes for data to be read or written onto the media (whether a tape, hard disk or other storage device), the transfer rate is measured in MB per second.
[7] The transfer rate means the time it takes for data to be read or written onto the media (whether a tape, hard disk or other storage device), the transfer rate is measured in MB per second.

level. Since generation three there has been the so-called WORM (Write Once Read Many)[8], which protects against overwriting data that has already been added, and which made the LTO tape usable for restricted documents and preservation masters that need to be kept unedited.

If a specific brand of LTO is not being produced anymore or there are several new generations of tapes, the existing tapes should be updated. One should also make sure that newer drives are compatible with older tapes so they can still be read. It's important to be aware of the backward compatibility characteristics of the tapes. The LTO Ultrium compatibility is defined with two concepts demonstrating investment protection.

1) An Ultrium drive can read data from a cartridge in its own generation and two prior generations.
2) An Ultrium drive can write data to a cartridge in its own generation and to a cartridge from the immediately prior generation in the prior generation format.

For example:

- An Ultrium format Generation five drive will read and write data on an Ultrium format Generation four cartridge as specified by the Generation four format and read data on an Ultrium format Generation three drive.
- An Ultrium format Generation four drive will read and write data on an Ultrium format Generation three cartridge as specified by the Generation three format and read data on an Ultrium format Generation two drive.

As long as the drivers and readers are available, then in principle one doesn't have to make any changes but there will be a higher risk of unreadable material. Therefore it is recommended converting to a newer generation of tapes every four-five years.

3.2 Autoloaders / tape libraries

For a small collection of data a manual backup procedure will be fine, but an institute with a larger collection of tapes could benefit from so-called automated robotics (Bradley 2009). There are different sizes and complexities to robotic libraries. They are basically closed shelving systems from which each tape can be taken out automatically and put into a supplied drive. The drive can be set to backup any possible changes or additions from a server, as well as check for errors on the tapes, which allows for quicker access and also automatic backup of any changes. A small autoloader can only hold a limited amount of tapes and it will be necessary to manually swap tapes if the tape collection is bigger than the autoloader's capacity. Data should also be copied on tape or other storage devices stored offline, to ensure that any errors occurring during automated backup isn't fatal for the files concerned, although this does require some sort of regular plan and manual work. For a small institution with an increasing amount of tapes it can become an issue to manually handle all the tapes. However, at the same time autoloaders might not be an option because they can be expensive to invest in and use more space than the tapes alone. Also setting the system up requires people with expertise. For institutions with larger storage room

[8] Write Once Read Many, a technology that ensures that data on a HDD, magnetic tape or a DVD/CD cannot be overwritten again, but can still be read repeatedly.

capacity and technical staff it is a valuable help to keep track of tapes and check for errors automatically. A couple of institutions using these are the National Library of the Netherlands and ZKM – the Center for Art and Media in Karlsruhe.

3.3 Preservation issues

Because a tape is read/written by the tape surface running over the read/write heads there is always a risk of the tape being damaged while being played, especially if the tape is read often. It can also take up to several minutes before one retrieves the right information, because the tape has to be loaded into a reader and spooled to the right spot. Therefore it is recommended, no matter how reliable the tape, that HDDs be used for frequently accessed content. Tapes are also sensitive to handling and climate change. Dropping a tape can cause detrimental damage, such as the drivers not being able to read the tape. If tapes are stored under high temperature, high humidity or both they might be prone to the sticky shed syndrome, where the tapes plastic particles become sticky resulting in difficulties when reading. Also as the data tapes are magnetic they are sensitive to magnetic fields and should therefore not be exposed to magnetic fields or electromagnetic pulses, as this will cause the magnetic particles, holding the data, to shift position, thereby eliminating the information. In many ways, a tape is as such more suited for long term, inactive storage of data than a hard disk drive, because the life span is a lot higher (around thirty years) the error rates are lower (Bradley 2009) and it's an affordable technology. Although one should keep in mind that the generation of tape reader/writers change every two-four years, which might mean it will be difficult to read thirty year old tapes, because the reader/writers are no longer are available. The difficulty of testing the integrity of the files on the tapes depends on the type of software and setup that is used. Obviously if tapes are not setup to remote access then one will have to physically handle the tapes and connect them to the software that does the testing.

Solutions for optimum preservation of data on data tapes are as follows:

- Have data redundancy (several copies of the same data).
- Check files on tapes and the actual tapes every six-twelve months or on an automated basis to make sure the files and tapes are still working correctly.
- Due to the previously explained new generation of tapes one should base the replacement of tapes on the backwards compatibility and availability of drives for that generation of tapes.

Pros and cons

Pros:

- Low prices.
- Scalability.
- Energy saving.
- Transfer rates relatively high.
- Low error rates compared to hard disk drives.
- Life span of ca. thirty years (not taking into account the reader/writer equipment).
- After generation LTO-5 the possibility of internal file index for faster access to the right data.

- Supports compression and encryption.
- Has WORM.

Cons:

- Writes in a linear way, meaning that excessive accessing is not recommended, because the tape will be worn faster.
- Access times are long.
- Manual handling unless connected in autoloader.
- Can become an expensive solution if robotic autoloaders are needed.
- The reader/writer equipment is not compatible with future generations and needs to be updated.

(Pease et al 2010, the linear file system)

4. Hard disk drives

Most people know that a computer is usually supplied with a hard disk drive (HDD) to store data and programs (although one can also get computers with solid state drive (SDD)). For laptops these are usually 2.5 inch HDD because they are small enough to fit into a computer and made for low energy consumption. A hard disk can also be in a desktop computer or as an external device, as with a 3.5 inch HDD, which is bigger and craves a more sufficient power supply. The HDDs come with different storage capacities and speeds and for archival purposes there are several hard disk drives to choose from. A basic explanation of a hard disk drive is that it contains a number of hard platters/disks with a magnetic surface that can rotate rapidly with help from a spindle (Figure 3). On these disks the data is recorded or read through the magnetization of ferromagnetic material done by read/write heads that control the level of magnetising (Hoel 1999, p.153).

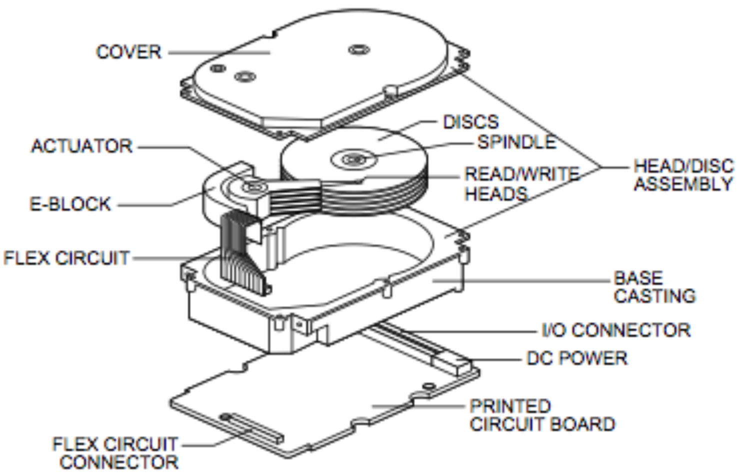

Figure 3: from Anderson (SCSI vs ATA)

The rotation of the disks is measured in RPM (rotation per minute) and has been continuously developed to go faster and faster. According to a study in 2007 HDDs can spin with 15,000 RPM (Anderson 2007). This calls for a stronger motor, and also for features which make sure the data can be read as fast as possible, keeping the temperature down and the mechanical parts in place.

4.1 Connecting an external HDD

There are different methods of connecting an external HDD to a computer:

- USB (Universal Serial Bus)
- Firewire
- SCSI (Small Computer System Interface)[9]
- SATA (Serial Advanced Technology Attachment)[10]

[9] Small Computer System Interface is a set of standards for connecting and transferring data between computers and peripheral devices.
[10] Serial Advanced Technology Attachment is one cable containing a minimum of four wires that creates a point-to-point interface on a hard disk.

- SAS (Serial Attached SCSI)[11]

Usually USB and Firewire connected HDDs are for plugging into personal computers. They are easy to plug in to any computer with that type of connection. SATA is typically the type of HDD built into a personal computer, although today a SATA HDD can also be made for external enterprise/professional use together with SCSI (Bradley 2009, s.109). This would typically be larger storage and server systems.

4.2 SCSI and SATA technology recommended

SCSI and SATA facilitating HDDs are usually said to be the most reliable type of disk drives (Schroeder 2007). It is not recommended to use USB and Firewire connected HDDs for archiving and long-term storage purposes (Hunter 2011). As they are not made to be connected to larger systems they are difficult to monitor and can be unreliable. The functionality and features of the two recommended types of HDD technologies are slightly differentiated, resulting in a price difference. Some claim that the SCSI is more reliable and it is also more often seen in large storage system contexts, with the use of SAS. But HDD with SATA is becoming more and more common, maybe because it costs less, but also because it has been shown that SCSI and SATA HDDs have equal replacement rates, indicating that the failure rates are approximately similar, provided the replacement occurs when the disk actually fails (Schroeder 2007). It has been difficult to retrieve publications after 2007 about the development in these areas, so it is not known at this point whether there has been a significant change in the technology. In practice SCSI is becoming less popular because of the pricing and it often has less capacity than the SATA solution.

4.3 Server systems

Depending on the size of the archive needed, the storage capacity and numbers of HDDs can be regulated. They can all be connected to each other as a unified storage system and linked to servers and individual computers of a whole institution. This allows for quick access and backup, also for HDD stored on other locations. It’s possible, and for some institutions enough, to have a 'bunch' of individual disks that needs to be maintained separately, but they all have to be individually connected to a server. To make this easier they can be put together in a so-called JBoD (Just a Bunch of Disks), which means that the individual disks are put together at one connection point. This can lead to slow access if several disks are activated at the same time, or if one disk fails, and there is no automatic backup of the data on that specific disk. So for larger collections and for safety reasons a RAID[12] (Redundant Array of Independent Disks) is recommended. This is a system that brings together a number of hard disks to form one virtual hard disk (Troppens 2009). Data is distributed across the drives in one of several ways called "RAID levels”, but the server sees it as one big disk. RAID systems are faster and more fault-tolerant than individual physical hard disks because if a HDD fails it can be exchanged without data being lost. The data can be reconstructed with the help of the other disks, because it is distributed and connected between the disks. How this is built up depends on the level of RAID.

---

[11] Serial Attached SCSI is a network that moves data to and from computer storage devices with the help from SCSI technology. It allows the use of longer cables and less bulky connectors.

[12] Redundant Array of Independent Disks, originally Redundant Array of Inexpensive Disks, is a storage technology that combines multiple disk drive components into a logical unit. Data is distributed across the drives in one of several ways called "RAID levels”.

There also exists a system called MAID[13], a storage technology like RAID but with energy management. MAID turns the HDDs off to save energy as soon as the content on the disks is not in use. It has also been said that a MAID could increase the life expectancy of the HDDs, because they would not spin constantly. However, Google published an article in 2007 saying that the usage and higher temperature of a HDD does not seem to correlate with the disk failure rate, making this argument less reasonable (Pinheiro 2007). For the cause of disk errors see page 17. MAID seems to be an advanced system and is not very well known, so it still isn't clear if it increases the life expectancy of HDDs, and as such is not a very widespread system. Most institutes go for a RAID setup, although there could be beneficial energy savings using a MAID system.

In terms of which RAID level is suitable for long-term storage, it is important to focus on the level that provides the highest fault tolerance and lowest risk of data loss, and then on its overall performance and speed. There are many different setups, unfortunately none of them with the perfect, risk free solution. The following levels are some of the most used. They are only mentioned for their overall properties and not explained in full detail.

Table 1 - RAID levels explained, based on Troppens 2009:

RAID level	**Type of performance**	**Overall properties**
RAID 0	Block by block striping.[14]	Increase in performance.
RAID 1	Block by block mirroring.[15]	Increase in fault-tolerance.
RAID 0 + 1 / RAID 10	Combined striping and mirroring.	Increases performance and fault-tolerance. Problem is that all data is written onto the physical hard disk twice, thus doubling the required storage capacity.
RAID 4 or RAID 5	Parity[16] instead of mirroring.	Increases storage capacity, but there is a risk of losing a lot of data at once.
RAID 6	Double parity.	Compromise between RAID 5 and RAID 10 by adding a second parity hard disk to extend RAID 5, which then uses less storage capacity

---

[13] Massive Array of Idle Disks, storage technology like RAID, but with energy management. Instead of the HDD being constantly on, the MAID turns them off to save energy as soon as the disks are not actively accessed.

[14] In striping, physical blocks (disks) of different storage devices are brought together to one volume. This increases the speed of read and write operations, since the load is distributed over several physical storage devices (Troppens 2009).

[15] Storage devices that are connected copy the data on each other, meaning there will be at least two copies of the same data on different devices.

[16] Parity data is an additional digit of information that helps one recover lost data. It only keeps the difference between two datasets, instead of mirroring the whole dataset.

		than RAID 10. Slows down write performance.

If you do set up a RAID do not only rely on RAID 0, because it has no reconstruction feature and even worse; because data is spread over two disks, risk of failure is doubled. If one of the two disks fails, data is lost. RAID 1 and RAID 6 have been seen used in different institutions during the DCA project and should have suitable levels for preservation purposes.

4.4 Preservation issues

The development of new hard disks has mainly focused on increasing storage capacity and access speed. Today you can get one enterprise hard disk drive with up to 3 TB storage capacity and an access speed of 6 Gbit/sec and around 10,000 RPM.

Although the introduction of HDD into a digital archival system can be expensive, the costs have been decreasing constantly since the technology was introduced and it is possible to get an affordable system setup with a lot of storage space. If the size of storage is not too large, price levels for investing in HDDs as magnetic tapes can be almost equivalent (Bradley 2009). There are of course different parts of each technology, which can make one or the other the most costly. For example the magnetic tapes might not cost much, but the drives or autoloaders to keep track of them can be expensive, then again the energy costs of keeping a HDD array running and cool is a continuous expense. When it comes to checking the files' integrity, it is easier to do such automated tests on HDDs if they are internal and set up in a server system. This way one can save a lot of manual handling of tapes and access time will in general be decreased.

It has been shown in different surveys that HDDs, no matter what the type of technology, have an equal risk of failing whether they are constantly in use or switched on and off (Pinheiro 2007). Also, if a new HDD survives without experiencing any errors in the first year of use, it will usually remain stable for the next couple of years. This is called the infant mortality phenomena, which means that if a HDD is prone to errors and failure it will often happen in the first years of use (Schroeder 2007). This emphasises the fact that one should keep a sharp eye on storage and do regular check-ups every three-six months in order to discover any errors. Other typical HDD failures according to Wikipedia are:

- Head crash: The reader/writer head comes in contact with the rotating platters, due to mechanical shock or other reason
- Bad sectors: The magnetic sectors on the platters might become faulty because of magnetic fields.
- Circuit failure: Parts of the electronic circuit might fail
- Motor failure: electric motors might fail or burn out
- Miscellaneous mechanical failures: mechanically moving parts do have a risk of break or fail.

(Wikipedia, last viewed April 2013: http://en.wikipedia.org/wiki/Hard_disk_drive_failure)

Solutions to avoid the loss of data during a failure of any kind of storage media, are data backup, data redundancy (several copies of the same data) and SMART[17]. It is recommended that one change a hard disk after three-five years, since manufacturers have estimated the life span to be 40,000 hours (Bradley 2009, p.105), which is approximately four-six years. It has been remarked that after five years failure rates seem to increase (Schroeder 2007). From personal experience we have seen HDDs failing after three years of intense use (e.g., an email server).

Pros and cons

Pros:

- Random and quick access.
- Contains file index managed by file index systems.
- Multiple files can be opened and used at once by multiple users.
- Portable across platforms and operating systems.
- Scalable.

Cons:

- For scaling content it can be an expensive solution.
- Energy costly.
- SATA HDD systems have been proved to have higher error rates than magnetic tapes[18].
- Life span of only five years and an initial infant failure risk.

[17] A self-monitoring analysis and reporting technology best suited for predicting trends of large aggregate populations of HDD (Pinheiro 2007).

[18] In a study it was said that the disk failure rate and the replacement rate of SCSI HDDs was almost equal to SATA HDDs (Schroeder 2007). This indicates that SCSI systems also have higher error rates than magnetic tapes, but this has not been empirically tested.

5. Solid State Drives

SSD (Solid state drive), also called Flash Drives, are differently built than hard disk drives and optical media. They don't have any moving parts; instead they are based on semiconductor chips[19] (Chen et al 2009). The overall performance and stability is good, because the chips don't heat up and aren't prone to mechanical failures. A SSD has fast reading times, low energy usage and can be extremely small in relation to the amount of storage. Therefore they have become increasingly popular in consumer products, such as MP3 players, mobile phones, laptops and even external storage. The main reason not to implement SSDs for larger storage networks has usually been the high prices, but in 2008 Google announced plans to exchange some of their HDD systems with SSD, mainly due to energy savings (Claburn 2008). Recently the cloud host company CloudSigma has started slowly moving all its data onto SSDs, claiming that they give fewer inconsistencies in performance (Clark 2013). Although SSDs could be a fair competitor for HDDs, not many tests on their durability have yet been conducted. Some say the hardware parts used in SSDs are the same as in HDDs and therefore are sceptical as to the overall hardware stability of SSDs (Marks, 2012). A couple of tests showed that although a SSD at some levels is more effective than a HDD, there are still some elements that need improvement. The SSD handles random writing and reading differently than the HDD. During reading SSDs are very fast, but the writing speed is seriously reduced, increasing the power consumption considerably. Also a SSD can find random files quickly, but it is not particularly fast when the files are searched in a more sequential manner, which is important when searching large databases (Dumitru 2007, Chen et al. 2009). Some solutions are being proposed to deal with such issues, which might improve the function of the SSD, but at the same time increase price.

5.1 Single-Level Cell Flash recommended

There are different ways a SSD can be built up, either through a MLC[20] flash (Multi-Level Cell) or a SLC flash (Single-Level Cell)[21]. MLC basically means that each memory element of the SSD (called a cell) is capable of storing more than one bit of information. A SLC on the other hand can only store one bit per cell (Spanjer 2008). If SSDs are considered for long term preservation of digital content, it is recommended that only SLC flash is used, because it has lower error rates, higher performance speed, higher number of program erase cycles and uses less energy than the MLC flash (Spanjer 2008). Program-erase (PE) cycles are a sequence of events in which data is written, erased and then rewritten to a SSD. A SSD has a limited amount of PE cycles because each cycle causes physical damage to the medium (Kwanghee 2009), eventually causing the SSD to malfunction. Apparently an SLC SSD can handle a higher number of PE cycles, but the exact life span depends on how many times the data is overwritten. The more you overwrite the quicker the SSD gets worn out. Manufactures claim the life cycle of an SSD is five-seven years (Hodges 2011).

---

[19] A conductor is a device that conducts or transmits electricity. Semiconductors are today made of silicon and are the base of any computerised device and can be found in microprocessor chips (Brain M.).

[20] Multi-Level Cell, a technology of solid-state drives that can contain several bits on one cell in a microchip.

[21] Single-Level Cell, is a technology of solid-state drives, where one single bit is stored on one single cell in a microchip.

5.2 Preservation issues

The overall pros would be the speed and ease of access and the low power consumption. Making it a technology that could easily compete with HDD systems, especially as it is said to have a lower rate of errors than HDD (Hodges 2011). At present it is not fully attractive for institutions, mainly because of the price. As a result, one doesn't come across many institutions using SSD solutions for long-term preservation purposes and the uncertainty of the lifespan and issues during writing needs to be improved before SSDs will be a competitive alternative. The main reason for using SSDs today is the speed of use, which is not the biggest priority when it comes to preservation. As for the integrity of files on SSDs it depends on the setup, but it would follow the lines of HDDs.

Pros and Cons

Pros:

- Low power consumption.
- Compact size.
- Shock resistant.
- High performance for random data access.

(Based on Chen et al 2009)

Cons:

- High prices.
- Low performance in data writing (this can be improved with certain built-in software).
- Uncertain lifespan.

6. Optical disks

Commonly known storage media are optical disks such as CDs (Compact Disc) and DVDs (Digital Versatile Disc). A CD can contain 700 MB of information and a DVD maximum o17.08 GB or gigabyte (if it is double sided and has two data layers) (Byers 2003 p. 28). If they aren't pre-recorded at the manufacturer (CD-ROM), they can either be recordable (CD-R, DVD-R) or rewritable (CD-RW, DVD-RW). A typical CD consists of at least three layers: a polycarbonate base, a reflective layer, usually of aluminium or gold, and a protective coating (Adelstein 2009). A DVD has even more complex layers, since it can be both double-sided with one or two layers of data. If the CD is recordable it will also have a laser-sensitive dye coated on the polycarbonate base. There are three commonly used dyes: cyanine, phthalocyanine and azo (Bradley 2009). When a CD-R is recorded a laser burns the dye and leaves transparent pits and non-transparent areas, which cannot be erased (Bradley 2009, Adelstein 2009). The areas will be reflected in the metallic layer when read by a laser in the CD/DVD drive. A Rewritable CD has germanium, antimony or tellurium as the recordable layer instead of organic dyes. When recording on a rewritable disc, the layer is heated up and controllably cooled. This can be done again, hence the name rewritable, but only a limited amount of times. A DVD is usually more stable in terms of chemical and physical decay because the metal and dye layer are embedded at the core of the disk instead of at the surface (Iraci 2012). There are other types of optical disks such as Blu-ray; although they do have a larger data storage capacity they are susceptible to the same errors and problems as CDs and DVDs.

6.1 Blu-ray

In 2002 the Blu-ray Disc Association was formed between Sony, Samsung and Pioneer, amongst others, and the Blu-ray format became available shortly after that[22]. This was to compete with the DVD and the now already lost format HD-DVD (High Definition). Blu-ray ended up outcompeting the HD-DVD format. A two-layered Blu-ray disk can contain up to 50GB of data. Today you can even get a triple layer Blu-ray with a capacity up to 128GB (Hodges 2011). The name comes from the laser, which reads the Blu-ray disk, and is blue whereas the laser reads or writes CDs or DVDs is red and the rills of information are thinner and more compact. As for using Blu-ray for long-term storage of data, WORM technology and playback equipment is backwards compatible with, for example, CDs and DVDs and early Blu-ray formats. The main difference between the disk types is the laser used to produce them: Blu-ray has a larger capacity of data storage and a higher data transfer rate. A rewritable Blu-ray disc can be rewritten up to 10,000 times (Bennet 2011). It is said to be more robust than DVDs, but it's still a disk with layers of materials sensitive to the climate and wrong handling. Whether it stays on the market, only time can show... maybe the optical disk format will end up being replaced by another form of technology all together.

6.2 Preservation issues

Using "stand-alone" optical disks for long-term preservation is not recommended, at least not for larger archives (Bradley 2009, V.D. Hoeven). This is mostly because of the relatively small amount of information each disk can contain compared to a magnetic tape or HDD. In a large repository containing a large amount of information, it would be too time-consuming

[22] http://www.sony.net/SonyInfo/News/Press_Archive/200205/02-0520E/

and space craving to have all the data on optical disks. The CDs and DVDs are also often of varying quality depending on production and type of material. They are susceptible to scratches while handling and errors during writing can mean that the content is unreadable. The formats are also very much independent on the consumer market and it is not easy to say how long optical disks will be around. In 2011 Apple stopped supplying optical disk drives in some of their products, claiming we don't need them (Hope 2011). This is of course mostly seen in a personal computer setting context, as the optical disk drive is still used in many work situations, but there is a potential risk of obsolescence; one need only think of all the CD and DVD record stores that have closed over the last five years. Another thing to consider is the file formats an optical disc can contain. The audio and video formats suitable for CDs and DVDs have to be compressed. If the original digital file is of a higher, less compressed format, it is not advised that it be converted into mpeg formats without keeping the original file. The life expectancy of an optical disk is also very dependent on production quality, type of dye and metal materials as well as what type of jewel casing it is stored in (Iraci 2012). If all these measures are right a CD or DVD can last for about fifty years. At the same time, the risk of playback systems going obsolete, compressed file formats and the difficulty of handling vast amounts of disks makes it a less attractive alternative for long-term storage. If you do use optical disks for preservation purposes you should consider changing formats.

7. Cloud storage

Cloud computing is an emerging technology for easy access and storage of files, software and platforms. To put it simply, using cloud computing means having a third party provider of digital storage, processing capacity, infrastructures, platforms and software programmes (Vaquero 2009). Different levels of using the cloud have specific terms, according to v.d. Hoeven 2012:

- Infrastructure as a Service (IaaS): using storage/computing elsewhere.
- Platform as a Service (PaaS): using predefined remote computer platforms running specific system software such as databases or software development environments (IDE).
- Software as a Service (SaaS): using software tools or even complete suites remotely.

In this deliverable we will focus specifically on the first point, infrastructure services in the cloud (this will also be referred to as storage service or cloud storage) and how it can be beneficial in a digital preservation circumstance. Because the cloud is an external service, it is easier to scale to fit one's needs. A company doesn't have to invest in their own equipment and the service can be made to suit any size of data collection. The data content is still stored on hardware components (usually HDD or Tape systems), the difference is that the access goes through the Internet and the cloud host is responsible for maintaining the hardware. This way the IT department can focus on other issues rather than just maintaining old servers and installing new. If an institution doesn't even have an ICT infrastructure the cloud is an easily scalable and relatively cheap solution for any IT related needs.

Table 2 - Properties of cloud computing, based on Vaquero 2009

Properties of cloud computing	**Explanation**
User Friendliness	Anyone should be able to understand the interface, and access his or her files.
Virtualization	You can access the content of the hardware remotely and see the data anywhere.
Internet Centric	Everything in the cloud is accessed through the internet. Without internet, no cloud content.
Variety of Resources	Choice between different levels of services such as infrastructures, platforms, software and applications. More or less complex solutions.
Automatic Adaptation	The system can be setup to automatically keep all files, calendars and emails up to date, no matter what device is in use.
Scalability	Without having to buy new equipment to integrate with the existing network you can get more storage space in a few minutes.

Pay per use and resource optimization	You don’t have to pay for maintenance or hire people for the IT department, you only pay for what you exactly need.

Often “the cloud” is mentioned as a service where everything is out of your hands and given to the cloud hosting company, but in fact there are several ways to use cloud computing:

Public cloud

- Access over the web by third party vendor.
- Requires less start-up money, because you don’t have to invest in equipment yourself.
- No operational or maintenance staff is required either.
- The amount of storage and services is scalable without the cost of buying new equipment.

Private cloud

- The cloud is set up solely for one organisation, with same accessibility as the public version, but only for people connected to the organisation.
- Gives control over all assets.
- Costs more because you need staff and equipment to make the cloud work, although a third party vendor can also do this.
- Not as scalable, because one can’t expand without buying new equipment.

Hybrid

- For maximum flexibility and same scalability as with the public cloud.
- Combination of some data content in a public cloud, more confidential parts in a closed, private cloud. Also this can be for backup reasons.
- At the same time one can keep control over the content

(HP video, white paper)

7.1 Preservation Issues

Cloud computing and storage is on everyone's lips these days. Many are considering, or already using, cloud services in one way or the other. However, very few are looking at it in a preservation context. Indeed the question is whether it is a viable solution for long-term preservation. The two main issues are the risk of losing data and the security of data being leaked or hacked, but other things also come into play, as one can see in the following table.

Table 3 - Risk of cloud computing in a preservation context, based on v.d. Hoeven 2012:

Risk	Explanation
Less control over data.	As managing digital assets is a core task of a cultural organisation the cloud represents a risk in this regard.
Copyright and legal issues.	Extra security measures are needed if collections contain personal information or

	allow only for restricted access due to copyright constraints. Data in the cloud has a higher risk of being hacked than offline data. Or there could be data not complying with local regulations. If a shared (available for multiple users) cloud is used, it must be defined explicitly at which geographical locations the data is stored to avoid legal issues.
Dependency and stability on a specific vendor.	The extended preservation period increases the likelihood of having to deal with multiple vendors during the lifetime of the data set. Migration of data and adaptation of underlying software is then required.
Loss of data, due to payment issues.	In case of financial problems, not paying the cloud provider will result in direct loss of data hosted elsewhere. This is not the case if you have your own storage infrastructure – which can be shut down while keeping the data on the media carriers.
Increase in access and availability.	Aspects that are contradictory to long-term preservation, where an organisation must be sustainable with a focus on preservation rather than providing short-term access.
Upload and migration issues.	To be able to upload or download large files such as video a fast and reliable internet connection is needed. It will be slower and more time-consuming than a regular in-house network.

Most cloud hosts offer storage, but have no strategies for the long-term access of content. This means an institution that facilitates the cloud has to make sure that its data is usable for years to come, because the host will not. The risk of losing data is unfortunately not something that can be ruled out; the best way is to have a clear deal with a cloud host that has a record of good back-up procedures and focuses on preservation. The risks mentioned in the above table apply for any institution, regardless of the size of its collection. External cloud computing is at the moment not necessary for large institutions that already have an ICT infrastructure that can handle large digital collections. However, it could be a solution to make a private cloud, i.e., internal system accessible via Internet that one has full control over oneself (v.d. Hoeven 2012). For smaller institutions the benefits of using a public cloud host to maintain storage, ingest and access a collection might outweigh the risks, because it is much cheaper than setting up one's own network from scratch. Yet, it is vital to ask any potential cloud service providers what regulations they are under and what is guaranteed if anything out of the ordinary happens for the safety of the digital collection and the copyright and privacy of the information.

7.2 Recommendations when choosing cloud storage

Questions to ask any cloud host company before selection (Findlay 2010):

- How easy / difficult / costly will it be for them to meet any record-keeping requirements specified by your organisation, for example additional metadata fields?
- Whether any additional charges would be levied by the service provider in the event of your organisation seeking to remove information from 'the cloud'?
- If they will commit to storing and processing your information in specific terms that are acceptable to your organisation?
- Whether they will make a contractual commitment to obey privacy requirements on behalf of their customers – both local to the organisation and in the location or locations(s) where the information is stored?
- Ask for an assurance that the service provider retains no copy of the records or information after the termination of the contract.
- Whether they are regularly subjected to external security audit or certification processes?
- How many administrators will have access to your files and details?
- If they have measures such as multiple geographically separated backup sites in place, so that they can do a complete restoration of your records if needed, and how long this would take?
- As well as complete restoration of data, how will they go about finding and restoring particular specified files or sets of files and what timeframes will they guarantee for this? For example, if someone accidentally deleted some files or if some data becomes corrupted.
- When restoring files, can they ensure that the structure of these files (not just the content) and associated metadata is maintained?
- What are the guaranteed service provision parameters offered by the provider (given the greater likelihood of service disruption due to provision over the internet) – and what action will they take in the event of service disruption? Do they offer any recompense?
- Whether service providers subcontract part of their service offering to third parties and, if so, what contractual agreements they operate under?
- Whether there are any relevant standards they are certified as meeting?

Even though you can ask these questions, there are some things that you will have to take their word for, such as how many backups they have. For example you can't possibly know if it is true when they say they have twenty locations. One should aim for a cloud provider that can answer sensibly, has a good reputation and provides a good contractual agreement. There are a few cloud hosts offering specific services concerning preservation, which might be the best alternative. Although incorporating long term access to cloud storage is at its infancy these service providers might have a solution:

- Arkivum: http://www.arkivum.com/ (UK based)

- Duraspace - Duracloud: http://duraspace.org/about_duracloud (Open for using other vendors, makes it possible to use more than one cloud provider)
- Tessella - Preservica: http://www.digital-preservation.com/solution/preservica/ (Based on Amazon cloud, but the only all-in-one solution)

Arkvium is based in the UK, which means it is answerable to European law, adding a certain security regarding privacy, but it is not open source, so once a collection of data is stored with it, it might be difficult to change to another contractor. Duraspace is the only vendor offering open source services, providing the possibility of using more than one cloud provider and reducing the possible cost of changing or migrating from one cloud host to another. Tessella-Preservica is the only all-in-one solution available where you get help from ingestion to storage, which might be very useful for smaller institutions without any ITC. Be aware that they are using Amazon cloud service; for now this means the servers are all located in the US and as such are answerable to US regulation, which could be a security risk, due to political acts such as the patriot act. All of these companies have a specific focus on long-term preservation, which make them unique in the forest of cloud storage providers. It's of course possible to make contract directly with better-known companies such as Amazon, but they don't offer any preservation precautions. Keep in mind that even if the chosen cloud computer host does follow all recommended regulations there is still a risk of data loss. Therefore it is highly recommended to store a backup of the digital collection offline on own terms.

Pros and cons

- Cost/benefits
  - Can be cost saving
  - Reduced pressure on IT department for maintenance
  - Easy access, also outside office premises
  - Better and easier collaboration - geographically

- Issues/Risk
  - Security
  - No preservation plan
  - No control and checking of the integrity of the files
  - Not yet ideal as a long term preservation solution
  - Less control over own content

8. How to choose

In the previous sections different types of hardware for digital storage have been explained with their positive and negative properties. There are some systems more suitable for long-term preservation and others better for frequent access and use (for example tapes for long-term storage and HDD in RAID for accessibility). It is therefore important to make a distinction between data for long-term storage (preservation masters) and data for everyday purposes. Other variables, such as the size of the data collection and the budget available, influence what type of solution is best. A large collection with hundreds of TB of information needs different management systems than a small one with only a few hundred GB. The budget and the type of equipment you need to get started also have to be considered. For example if an institution started out using DVDs it should consider alternative storage possibilities. Also, as soon as a collection begins to grow continuously, it will be beneficial on a practical level to have larger storage systems instead of many small disk or tape stacks that need manual handling (2009). An institution does not have to only use one type of hardware or setup; a mix of solutions is usually the best option to meet the four key variables:

- **Size of data collection**
- **Budget**
- **Internal staff/expertise**
- **Use of data (for access or preservation)**

The following table gives an overview of suitable storage media depending on these four parameters. Obviously this is a simplified view and there are other factors that come into play, but it gives an idea of where to start.

Table 4 - Suitable storage media depending on the four key variables:

Size of data collection	**MB**	**GB**	**TB**
Recommended storage	HDD, LTO tape, cloud	HDD, LTO tape, SSD	HDD, LTO tape, SSD
Size of budget	**Small**	**Medium**	**Large**
Recommended storage	HDD, LTO tape, cloud	HDD, LTO tape	HDD, LTO tape, SSD
Internal IT staff	**0**	**1-5**	**5<**
Recommended storage	Cloud and remote backup on HDD by external vendor	LTO tape or HDD server systems	LTO tape, HDD and SSD server systems
Use of data	**Daily access**	**Backup (refreshed)**	**Preservation archive**
Recommended storage	HDD, SDD, cloud	LTO tape, HDD, SDD, cloud	LTO tape, HDD

8.1 Recommended brands

The following brands or companies (see Table 5), offering different storage media or services, are recommended specifically for preservation purposes. This overview is based on brands preferred by institutions as well as the focus of the production company, if they have a special emphasis on preservation and have solutions specifically for archival purposes. These are not necessarily the only brands one can use, just what is considered top of the line and best practice. As long as one takes precautions, cheaper solutions can work as well; the most important is to have several copies on several devices, to keep data from getting lost. Keep in mind that the technology changes quickly and brands tend to come and go; try to stick with those that are well known and regarded as open standards.

Table 5 - A few brands of storage media that institutions use and are recommended, from a preservation perspective.

Storage media type	**Brand and company**
HDD	Silent Cube by Fast LTA (recommended by the National Library of the Netherlands)
LTO Tape	LTO tapes and reader/writers of the newest generation (at the moment level 5) by IBM or HP
SSD	Intel or Vertex (with SLC technology)
Cloud storage	Arkivum (http://www.arkivum.com/) Duraspace (duracloud) (http://duraspace.org/about_duracloud)

9. Maintenance

This section is about how stored data can be maintained. Digital preservation is about continuously checking data to keep it accessible and understandable. This means that physical parts as well as software parts have to be maintained, as is nicely summed up in the following quote.

"It is also critical to note that any current technology, data tape or otherwise, will require migration to a newer storage media in the next 5-10 years, otherwise the data tapes themselves will become inaccessible through either tape failure or format obsolescence. "

Bradley 2006 (UNESCO CD-DVD)

There is always a risk of data being lost or inaccessible. This can happen in a number of ways, but one of the most common is hardware failure (Figure 4). The stability of hardware has probably been improved since this chart was made, but hardware are still mechanical devices that will degrade and at some point fail to work.

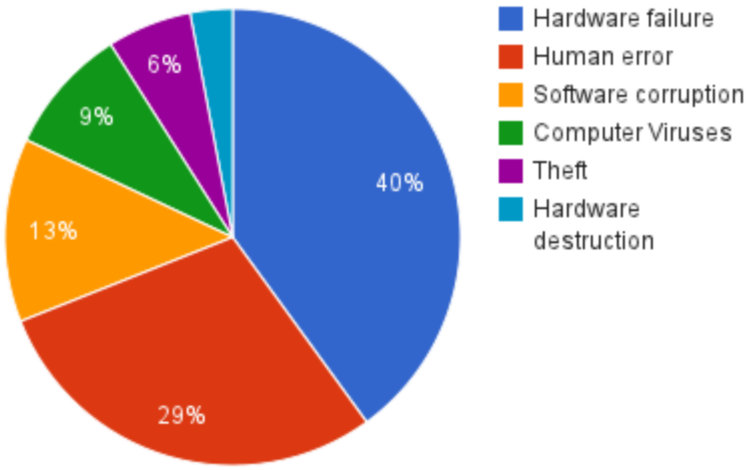

Figure 4: Reasons for losing data. Source: The insurance Agency Inc.(2001) & Ontrack Data International, inc. (2003).

It is interesting to note that the second most common reason for data loss is human error. Either way it is vital to have a backup plan and a maintenance procedure to stay ahead of any possible issues. Preservation planning was described in *D.6.1 Guidelines for a long-term preservation strategy for digital reproductions and metadata.* In short it entails:

- Making sure there is someone who is responsible for taking care of the data, and that it is defined what this actually means.
- Planning renewal and maintenance.
- Making a decision on which strategy to use and when.
- Having several copies of the same data.

- Having the copies on at least two different types of hardware storage.
- Storing backups at different locations.
- At ingest, making sure there are a few well-used formats, well-described metadata and persistent identifiers, so the data can be found and identified.
- That the plan and all equipment are evaluated every five years.

When evaluating the data and hardware there are some things to consider, in order to safeguard the proper content and also save money and time.

9.1 Evaluation of content

When it is time to evaluate the content for preservation consider how much the collection has grown since the last time, and if there have been any specific files that have been accessed more than expected or vice versa. In the following table (Table 6) the different topics worth considering are explained. That way you make sure only the necessary material is preserved and that the storage infrastructure is up to date and suitable for your institution.

Table 6 - Based on Oswalt et al 2010

Evaluation of content	**Considerations/modifications**
What is being stored? Is the data that is already stored still needed? (Is there new material, is some information out-dated?)	If the data format is becoming out-dated and it is still worth preserving, it is necessary to consider migration or other methods for keeping the data accessible.
What is the data's growth rate? (How much and how often is new data added to the collection)	Depending on how quickly the collection size is growing, more storage space is needed. It could be helpful using cloud hosting services to minimize the workload of maintaining larger and larger storage facilities or investing in an in-house ICT infrastructure.
What are the trends in its growth? (Will more or less data be added in the future?)	If it is only a temporary growth, it can be enough to invest in a few extra drives or tapes, but if you know it will increase a lot over time, compare your present solution with other options to make daily maintenance easier.
How long do we retain this data?	Is there any of the content that is not meant to be stored for long term that can be handled differently? Or are some data only meant to be kept for legal purposes, in which case make sure they are not kept longer than needed.
How often is the data accessed?	If the data is on tapes, but needs to be accessed often, it should be considered if HDDs are a better option. If you have preservation master data on HDDs, but don't use them often, consider making

	access copies for the HDD and invest in tapes for storing the master data.
Can some of the data be taken offline from the network to lower the cost of archival storage?	If a lot of the data is not accessed often, but needs to be stored for long term, it can be more efficient to keep it offline on tapes at a remote place.
Is the data accessible easily and quickly?	If there have been incidents where data needed couldn't be easily reached, maybe it is better to use a HDD server system or a cloud hosting service.
How often is the data backed up?	To make sure that the backups are up to date, they should be created every time any changes are made. Also onto remotely stored data.
Where are the backups stored?	It minimizes the risk of lost data if the data is stored on several media, at different locations. If you manually handle for example tapes, whenever there is an update, the tapes have to be moved back to their location, for safety precautions.

9.2 Error checking procedures

To avoid any errors in the data, the storage devices in use should regularly be checked. This can be done with different codes and algorithms that the storage devices can utilise automatically. A common error is when a bit in a sequence of data changes, a so-called "bit-flip". This can occur when a single bit changes, or when multiple bits in a row change their value. The reason for this can't always be explained, but it can be any number of reasons such as temperature changes, voltage, humidity, pressure, vibrations, power supply fluctuations, electronic interference, ground loops, cosmic rays, loose connections, aging components, resistive or capacitive variations and noise in the system (Alchesky).

Error detection - In *D6.1 Guidelines for a Long Time Preservation Strategy for Digital Reproductions and Metadata* the term checksums and how to check them were described. In short, this is a way to detect any discrepancies in the bits (a 1 has changed to 0 and so on). Just detecting the error is not enough, since a correction of the fault is necessary to maintain the data.

Data scrubbing - error correction reduces the risk of accumulating single correctable errors. Data scrubbing is the automatic reading of disks and checking for defective blocks.

ECC memory - Stands for Error Correcting Code memory, and is a protocol that can detect not only single but also multiple bit errors, and correct the errors on the fly. To benefit fully it is a good idea to set up the operating system to receive notifications when errors have been fixed. The errors could indicate that hardware needs to be changed or that there is something else wrong.

9.3 Backup procedures

When it comes to backing up procedures it is important to make a distinction between collections of data that are stored for long-term preservation and data that needs quick and easy access on a daily basis. The latter needs to be backed up, but not necessarily stored long term. It is recommended to have a copy for easy access available on a different server system than the actual preservation masters. It is also necessary to have several copies of the preservation masters (two or three) to reduce the risk of losing data. A more automated backup function that keeps an institution's data accessible and up and running even with minor operating problems would be suitable for larger collections. This would include the RAID system setup, as explained in chapter 4.3. RAID is a way of insuring that if one disk crashes the data on that disk can be reconstructed from the other disks. This backup system is made for quick recovery (which is the basis of backup) but benefits a long-term preservation system too. Robotic automation systems for tapes can also be used for easy and automated backup procedures and can be connected to a server, ensuring any changes are copied to a new tape. Regardless of whether the system is automated or not, the important thing is that there is a distinction between backup copies and preservation masters, because the latter should not need to be backed up very often. Each preservation master and also distribution copies should be copied and stored so there are at least two. These two versions should be placed at two different locations. In any preservation plan a procedure for dealing with errors or malfunction should be written down.

9.4 Preventive conservation: how to store the storage?

When storing data, it is worth thinking about the environment that the hardware storage devices will be exposed to. By storing them properly one can increase data life expectancy. An ideal room for storing any archival records (both analogue and digital) is in the middle of a building, without any windows. This provides the most stable climate, without too great a fluctuation in temperature and UV light. A top floor location often means problems with hot summer days and cold winter days, and there is also a risk of roofs leaking. In a basement or bottom floor location the risk of water or insect damage and humid air are often high. It is of course possible to choose such locations, but keep in mind it will require more equipment and time spent maintaining the temperature and humidity level stable. The room must be ventilated, especially with high spinning HDDs, as they emit a lot of heat. There should be no main water pipes near the selected room, because if the pipes are broken it can cause damage to the storage room.

The ideal storage environment:

- Relative humidity maintained between 35%-40%.
- Temperature maintained between 15-21°C (depending on hardware type).
- Temperature and humidity monitoring.
- Fire alarm.
- Internal, non-water, zero-residue fire suppression system.
- Restricted access to protect media from theft.
- HEPA air filtration.
- Magnetic shielding (specifically for tapes).
- Security cameras and alarms.
- Backup power to facility.
- UV shielding on lights and any windows.

- No water pipes near or above storage area.
- Storage area should not be made of concrete (increases humidity to near-100% during fire).
- Storage area should be insulated with material that can maintain temperature within storage area for several hours in event of fire outside storage area.

(Recommendations from Hunter, Guideline 3 on Digital Storage, Appendix 5, 2011)

Of course it is not always feasible for an institution to meet all ideal standards. However, being aware, and adjusting towards them will be beneficial in the long run.

HDD

The ideal storage facilities for HDDs requires well-ventilated rooms, an average temperature of 20°C and a maximum of 50% RH. It has been shown that a high temperature does not significantly cause HDDs to fail (Pinheiro 2007).

Data Tapes (LTO)

LTO tapes are actually quite robust and are said to last for thirty years. However, often they are kept on similar conditions as older, analogue video formats or books usually are stored. This is typically around 19°C and with a relative humidity of max 30-40%. This might not have to be so strictly retained, but it is important to keep the humidity in check in order to avoid any mould growth. Tapes are also sensitive to magnetic fields and bad handling, because of loose parts such as the spool of tape; therefore careful and minimal use is essential.

Overview

Table 7 gives an overview of preservation issues for each storage media, ideal storage climate as well as disaster management tips.

Table 7 - Preservation issues and tips for different storage media.

Digital media storage type	**Magnetic Tape**[23, 24]	**Hard Disk Drives**[24]	**SSD**[25]
Preservation issues	Avoid high intensity magnetic fields. High humidity can cause mould growth. Temperature and RH fluctuations can cause binder degradation, sticky tape, magnetic shedding and layer separation.	Regular maintenance checks are essential to detect potential wear-out and parts failure. Machines should be used in an air-conditioned environment.	USB memory devices should never be used for anything other than the transfer of records from one system to another. MLC (Multi-Level Cells) flashes should also never be used to store digital content for a longer period of time. The SLC type of

[23] Recommendations from Adelstein, IPI Media Storage, 2009.
[24] Recommendations from Pinheiro, Eduardo, Wolf-Dietrich Weber, Luiz André Barros, 2007.
[25] Recommendations from Hunter, Guideline 3 on Digital Storage, Appendix 5, 2011.

			solid-state drives is recommended.
Ideal storage temperature and humidity	11°C - 50% RH. A range between 11-23 °C with corresponding 50% or 20% RH can be accepted as long as there are few fluctuations.	Avoid temperatures below 20°C.	Ideal temperature is still being determined. A SSD has been seen working between -40 to 55 C.
Recommended time of renewal	Four-eight years (after one or two tape generations it is considered best to renew the tapes, although a single tape can last for about thirty years). All new tapes should be of the newest generation.	Five years.	Five years. Should retain information for a decade powered down, but recommended to refresh more frequently.
Handling and disaster Management	Never stack tapes high to reduce the chance of damage from tapes falling. Before loading check sticky labels on cartridges and never load a tape with a peeling label. Never touch the tape (either the recording or non-recording surface). If using an automated tape library, ensure that the library supports partitioning[26] so that older and newer generation media can both be supported.	Use RAID set up to cope with simultaneous failures on two disks. If in servers, they should be mounted in racks in machine rooms to make it easy to find and replace any disks.	Periodically verify data integrity (ideally write whole memory, then verify and then store).

[26] The division of a hard disk or a tape to create an appearance of having separate hard drives or tapes in one. Normally used for file management or multiple users. A partition is created when you format the hard disk or tape. On a personal computer different drives will typically be nominated with a letter followed by colon, such as *C:* or similar.

10. Price examples

Although having some kind of error correcting procedure and automated systems are essential to keep the data fit and well, it does cost time, money and increase the risk of damaging hardware. Also it is recommended to have several copies, preferably on different media, which may lower the risk of data loss, but at the same time increases the costs. Therefore one should check and correct errors regularly, but not too often. One should keep at least two copies, but having more than three becomes a hassle in terms of storage space and maintenance. It is always difficult to estimate the happy medium and PrestoPRIME (a European consortium for the preservation and digitisation of audio-visual content) has therefore developed a tool to calculate the cost versus risk, regarding the number of preservation copies and the frequency of checking and scrubbing the data (PrestoPRIME 2009). It is recommended to make a visualisation with an individually made model with the help of their tool, but in Figure 5 you can see an example.

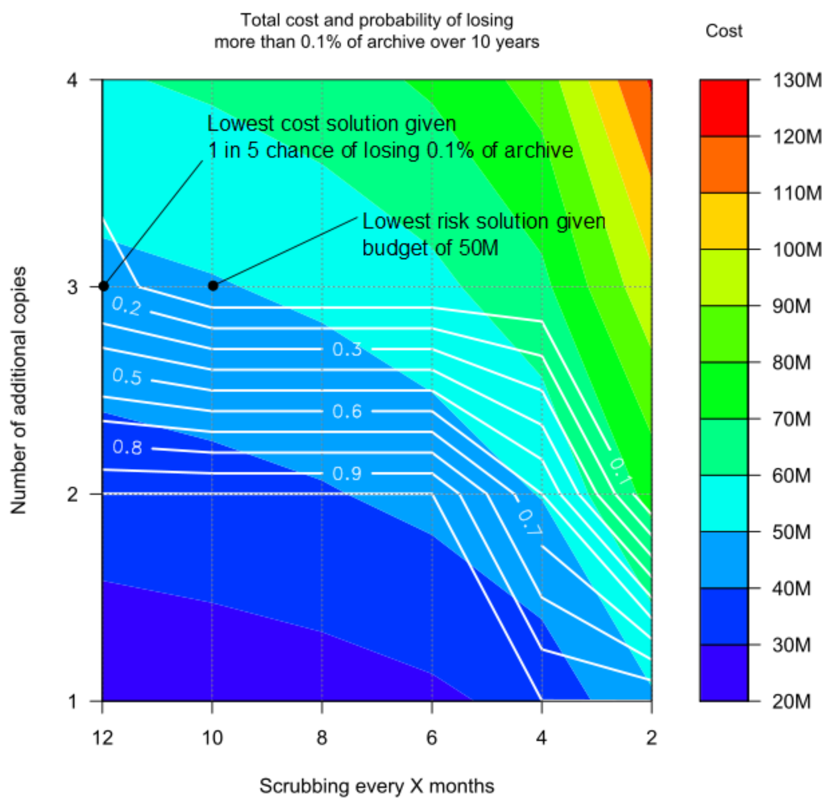

Figure 5: from PrestoPRIME 2009

"...an institution with a budget of 50 million (Euros), the strategy with the lowest probability of loss of more than 0.1% of the archive over 10 years is to keep 3 additional copies of each asset and to check the integrity of the files every 10 months."
PrestoPRIME (2009).

Another factor is the price for different types of storage infrastructure solutions, not only including the storage media, but also the costs of energy and maintenance. For this deliverable we try to make price examples, however, the market prices on all factors change inadvertently and the prices estimated in this chapter should only be seen as an indication of the present market and not as a precise value. An overview of storage media solutions and their price and capacity can be seen in the following table (Table 8), based on Hodges 2011 (prices are from the American market) and Bradley 2009.

Table 8 - Price estimates for different types of storage media

	HDD	SDD	LTO Tape
Media cost per TB	€ 91	€ 2361	€ 36
Max capacity of one unit	3 TB	2 TB	3 - 6 TB
Archive life according to manufacturers	Three - five years	Five - seven years	Seven - ten years
Total cost of ownership (TCO) [27]	High	Very high	Medium
Maintenance per year for <100TB	Medium	-	High
Maintenance per year for >100TB	High	-	Medium

It's important to take into account not only the initial price of buying the storage media, but also the cost of managing equipment (readers, autoloaders etc.), energy costs and maintenance/renewal. This is called the total cost of ownership (TCO) and can be difficult to estimate precisely; therefore it's only an indication in Table 8. SSDs are still the most expensive solution for storage, and although the prices do drop over time there is still a gap between HDD and SSD prices (Mearian 2009). Normally a monthly payment would be necessary to accustom the needs and data collection size, including energy consumption.

The International Association of Sound and Audiovisual Archives (IASA) has compared prices and maintenance costs of tapes and HDDs in its guideline from 2009 (Bradley 2009). Surprisingly their tables show that if you have a data collection of less than 100 TB, investing in HDDs can be a cheaper solution than using tapes (Bradley 2009). Then again the initial price of tapes fall according to the more data capacity you need, whereas the price for HDDs stays the same. Which means that if you have under 100 TB of data it will be cheaper to invest and maintain HDDs, but if you have over 100 TB of data the price of buying tapes will be almost half the price, and the cost of maintaining the tapes will become approximately 30-35% lower than HDDs. There are different factors coming into play here; one is the initial cost and the other is the maintenance over time. When investing in tapes it depends on how many tapes you need. If it is less than 300, it is not necessary to buy robotic tape libraries and the initial cost of buying equipment as well as the energy consumption is lowered, whereas you will have to pay for the hourly labour costs to manually handle and store the tapes properly. With HDDs you are provided with a whole system setup from the beginning, but the prices per unit and for maintenance don't drop the more data capacity you need, so the prices remain high, whereas the cost of the tapes drops the more you have.

[27] Total Cost of Ownership, the cost not only of purchasing equipment but also energy consumption and renewal.

10.1 Cost model for digital preservation

Many organisations have tried to come up with a tool for calculating the costs of digital preservation over time, covering not just the price of hardware but also people payment, time for developing software and risk of migrating/converting files from one format to another. There are some usable forms around, still being improved, although one should be aware they require expertise not only in digital preservation but also administration and finance. The cost of digital preservation will always be difficult to estimate, because many parameters change and depend on the individual collection. There is no easy model or tool that will give an instant result, but we do present a few below that could be of help. LIFE3 is a British developed tool and one of the simplest Excel tools made for calculating life cycle costs (www.life.ac.uk/tool). You pick the years you want to calculate for, the format you have (printed documents, digital material, sound material), how many items and how large your institution is. From these parameters the tool calculates the total costs. Obviously the results are only an estimate and a lot of parameters are not taken into consideration. For example it doesn't take into account staff working hours you would need and there is a limit to the type of material you can choose (for example there are no videos). The Danish National Archive and Royal Library have made a more extensive Excel cost model for each department in the OAIS model (explained in D6.1). The results should cover the calculation of costs for three stages, starting with digital preservation planning and migration, followed by ingest and ending with archival storage cost estimates. See more on www.costmodelfordigitalpreservation.dk.

10.2 Tiered storage

Regardless of what storage media is the cheapest, it is recommended to use more than one type of media, i.e., investing in both cheap and expensive systems. Most large institutions will have some form of hierarchical storage management (HSM), using several tiers defining different levels of storage quality. This means the data is categorised in different types of storage media in order to reduce total cost of ownership (Rouse 2005). The categories can be based on different needs such as level of performance, frequency of access and so on. For example the National Library of the Netherlands has four tiers, which it in common terms call gold, silver, steel and bronze. Gold and silver are the highest levels and contain data that must be readily available. To achieve this, data is stored on HDD in a SAN. Steel and bronze stands for durability, meaning backup and preservation copies, which in this case are on NAS and tape. The reason for this categorisation is that prices for high availability HDD systems are usually higher than for high capacity HDD and tape solutions. It can therefore be financially beneficial to keep only the most frequently used data on the expensive system and the more long-term backup and archival records on cheaper solutions.

10.3 Cloud cost

The cost of using cloud storage has not really been mentioned in the above, because it involves a whole other way of calculating the costs. Using cloud storage usually means having an agreement with an outside vendor who takes care of energy and maintenance (unless you are making a private cloud). Therefore they typically offer different subscriptions that can be changed to suit your needs and also you cannot account for the costs of energy as you do if you have to maintain the hardware yourself. The question is if this is economically viable for long-term preservation. The prices are not dropping like the prices for hardware; some say they only drop 3% annually (Angevaare 2012). Also prices are not based on costs but on market value and this is partially regulated by Amazon's dominant market position. Regardless of this, the main issue which people seem to be most concerned about is the time and cost of moving data from one service to another, if this becomes necessary. Some companies might go bankrupt, change policy or increase their prices, so changing storage services every now and again might not be an unrealistic scenario. However, at present this is both technically difficult and expensive, and so having a solid agreement and choosing a host with digital preservation in mind comes before the price.

11. Case studies - examples from real life

This chapter gives examples of how institutions of varying sizes and with different types of data collections manage their storage infrastructure. The institutions focus on preservation, but their solutions reflect their differentiating budgets, expertise and facilities.

11.1 Large data collection: National Library of the Netherlands

The National Library of the Netherlands (Koninklijke Bibliotheek - KB) is a fairly big institution with a yearly budget of € 45 million and around 300 employees. The maintenance of data storage and IT is not outsourced because they have a fully equipped IT department and ICT infrastructure, including 12 staff members dedicated to keeping everything up and running, including maintenance of storage facilities, servers, websites and so forth. It also has an IT department for innovation and development working on new applications and databases. Its digital collection contains text and image files of mainly Dutch newspapers, books, magazines, scientific publications and lately also archived websites. Everything, except the scientific publications, is available via the web. This is due to the agreement between the KB and international publishers stating that access to scientific publications of these publishers can only be acquired at the library's premises.

Data and Storage media

The digital collection is constantly growing, but is today about 300 TB of data (excluding backup copies). The data is stored on a variety of media, regularly backed up in-house and on one off site location.

Table 9 - Type of storage media at KB

Storage media	Capacity and main use	Brand
HDD - SAN setup	350 TB, access to data via web, database storage and indexing.	HP
Magnetic tape	500-600 TB, double backup and archive, one in-house, one offsite[28]	IBM LTO (generation three, four, five)
HDD - NAS setup	50 TB, new archiving system being implemented, more durable than SAN, faster than tape and optical.	Silent cube (by Fast LTA)
Optical disk	20 TB, current long-term preservation system consisting of Ultra Density Optical (UDO) storage. This solution is end-of-life and currently being replaced by Silent Cube NAS.	IBM, Plasmon

[28] Any digital content that is stored on storage media that is not connected to the main headquarters of an institution.

HDD - SAN and NAS: The hard disk system set up as SAN is used for accessibility purposes and needs to be fully running at all times. It provides a fast system (HP EVA) with quick access times, and with the use of RAID 1 and 6, which makes sure the data can automatically be reconstructed from other disks when there is a disk failure. There is an automatic notification system that alerts the staff when a disk malfunctions, the system will then also shut down the failing disk until something can be done. Within four hours a new disk should arrive from the provider. In case of power failure there is also an extra power supply that will kick in. The disadvantages of HDD in this setup are low durability and risk of data loss. Therefore all data is backed up on tape. The other HDD setup is NAS. The KB has acquired several Silent Cubes, NAS-based storage solutions with extra long-term preservation features. Because data stored on Silent Cubes are written with four-time redundancy, the chance of losing bits is decreased significantly. Furthermore, it follows the principle of Write-Once-Read-Many (WORM), which makes it impossible for operators to delete data from the storage solution. The silent cubes could be compared to a MAID system.

Magnetic Tapes used are LTO tapes produced by IBM, generations three, four and five. Generation three is being migrated onto five. Tapes are used for backup and archive purposes, therefore there are two collections with the same data, one collection in-house and one at a remote storage facility. A robotic library connected to a server system facilitates both collections. Every significant server in-house has a program that checks for any changes and every night new data is written onto a tape at both sites. This data written is compressed beforehand. If any errors occur the tape system will notify by email that something is wrong and action can be taken. The tapes are regularly spooled through to check for malfunction or corruption of bits, but it is not known how often they are checked.

Optical disks (UDO) are used for long-term storage of some digital collections, but are end-of-life and currently being replaced by Silent Cubes. The reason UDO is not found suitable for long-term storage anymore is because it has become a rare solution, out of production, costly and has a slow access time (> 45 seconds).

SSDs are considered but not implemented mainly because the prices are still too high. In the future they might be reconsidered for high availability of databases, indexing and web services if prices have dropped to a suitable level.

Cloud storage services are not utilised for different reasons. First of all, the library already has an in-house IT infrastructure, so it doesn't need a third party provider. Secondly the cloud is still seen as a somewhat unstable platform, with a risk of losing data and a low security and privacy level. This is not in line with the policies and regulations of data protection and long-term preservation defined by the KB.

Storage climate

The HDD SAN and NAS systems are stored under strict environmental control, with regulated temperature and humidity to avoid overheating. The in-house data system has a lower oxygen level, in case of fire, so it cannot spread and backup power supply (UPS) for electrical failures. The tapes are stored in the basement level, in conditions similar to books: low moisture and regular temperature of 19°C. Both areas are secured with alarm and lock systems.

Preservation plans

When asked if the institution has to apply to governmental rules and regulations the answer was that apart from budget restrictions and a rule of having a double back up, the library can form its preservation plan as its sees fit. Concerning migration plans or regular error checking procedures, the SAN system is regularly migrated and updated so it does not need a specific migration plan. NAS records are only migrated if an error occurs, not after a specific time frame. Tapes are regularly refreshed, spooled through and migrated from older generations to new. Some specific projects are implemented to lower costs and improve storage media lifespan, such as converting file formats or changing to new brands of storage media. It has, for example, been decided to convert TIFF files to JPEG2000 to free up storage space, as data loss or change during conversion is minor. All incoming data is processed regardless of the format, because the KB is not allowed to dictate any formats to suppliers. For rare or difficult file formats there are normalisation procedures to make sure the files can be read later on. Error checking of the bits is at present not a regular procedure, but a system for checking checksums is soon to be implemented. However, checksums are always generated for each file or package and ingested into the archiving system. The checksums are stored in a separate database. The system will then automatically compare the checksum of the file and the one saved in the database, and if any discrepancies occur an error will be reported. Any form of data scrubbing or use of ECC[29] memory does not seem to be in use. The storage media used by the Silent Cubes are WORM (Write once, read many), to avoid any accidental overwriting of the original material. There is an important distinction between what is preserved for backup and archive and how the data is classified as one or the other. Backup is data that can be overwritten after a year, to free up storage space, while archival data should never be deleted. From previous incidents it has been seen how important it is, with correct data management and classification of the data, to make sure the right data is labelled as archival and not only as back up.

11.2 Medium data collection: LIMA

LIMA, the author of this deliverable and one of the participating partners in the DCA project, is an institution dedicated to research and preservation of media art. Experts from the Netherlands Media Art Institute (NIMk, discontinued as of 1 January, 2013) established the LIMA Foundation. In the years to come LIMA will profile itself as the international service platform for media art. With active, online distribution of its own collection, a shared service structure for the storage and conservation of media art, and a lab devoted to the permanent preservation of such works for makers, producers and museums all over the world, LIMA is striving toward lasting access to media art. LIMA provides digitising and permanent digital storage for the media artworks by more than 500 artists, including Marina Abramovic, the Vasulkas, Nan Hoover, Stansfield/Hooykaas, Servaas and Jan Fabre, and for more than 15 collections, including those of the Kröller-Müller Museum, the Van Abbe Museum, and the Netherlands Cultural Heritage Agency (RCE). LIMA is now a small institution comprised of 6 employees, whereas NIMk was a medium sized institution The staff has many years of experience with preservation and digitisation of video art and 2 staff members are in charge of digitising, storing and preserving data.

[29] Stands for Error Correcting Code memory, and is a protocol that can detect not only single but also multiple bit errors, and correct the errors on the fly.

LIMA's data collection comprises mainly video art material. Budget, equipment and preservation plans are determined on own terms and is not influenced by outside restrictions, but of course is dependable on best practice in the field as well as what kind of funding is available.

Data and Storage media

NIMk/LIMA has two main data collections. The preservation masters are stored on LTO tapes of generation five, on location and with one off-site backup. Distribution copies are stored on HDD, maintained by an external provider, Toltech, an open source company. Also the HDDs are on location and with one off-site backup. These copies can be accessed via the server on location. There are about 16 TB stored on hard disks, incl. 32 TB backup. There are about sixty LTO tapes (ca. 90 TB); including the backup this accounts for one hundred twenty tapes and 180 TB of data). Previous results of video digitisation were stored on digital Betacam, but in the last two years the data from the digital Betacams have been transferred to LTO generation five. These tapes are therefore fairly new and have not needed any update or renewal. Every tape is checked immediately after writing, and each file on the tape gets a MD5 hash tag (a checksum) embedded within the file metadata. The LTO tapes are not of any particular brand; HP, Sony and IBM are all in use. Usually the choice depends on the given price at the time. The data on the tapes are indexed with a self programmed software based on TAR (Tape Archive), a file format and program made for finding and extracting files on tape. This means that when a tape is put into a reader/writer TAR can find and extract the files needed (not to be confused with the LTFS indexing system). It is very important to document what file formats and what lengths each file on a given tape has, because a tape, not utilising LTFS, cannot be read like a HDD. One has to know which program to use in order to extract the information on the tape in the right way and that is what the TAR system is used for. NIMk/LIMA does not use a robotic autoloader system because it is still cheaper and easier to insert the tapes manually. It is however considered that if the collection grows to above five hundred tapes it will become necessary to have some kind of automatic system, but this will increase the cost of energy and maintenance hours. When choosing storage media, NIMk was considering LTO tapes or Blu-ray optical discs. The choice fell on LTO tapes because video files require a high amount of data capacity and if you have a collection that's larger than 20 TB it becomes difficult to handle using optical discs. By not having the tapes connected in an automated system it can be very cost effective. Also, the tapes are rarely accessed, because there are distribution copies accessible via the HDD system. The tapes will only be used in specific circumstances, e/g/, when a video is needed in another format than the distribution copies. The data on the tapes are neither compressed nor encrypted, for preservation reasons. Compressing videos only gives up to 2% more storage room on each tape so it was considered unnecessary. Encrypting data can be very risky, for example if the encryption is old or can’t be identified it will be difficult to decrypt, leaving the data on the tape inaccessible.

Storage Climate

LTO tapes on location are kept in a vault together with other video formats at 19°C and a relative humidity (RH) between 35-55%. The humidity is kept in check with the help of a dehumidifier and the temperature and humidity is regularly checked. The room is always locked, and is on a middle floor, away from any windows and main water pipes. The off-site backup LTO tapes are not kept under any specific climate conditions, although there has been discussion that they might be stored at the storage facilities of the Stedelijk Museum (Amsterdam). The HDDs are placed in a locked room on the ground floor, with a regular air condition system, to avoid overheating.

Preservation plan

NIMk previously had an eight-year cycle plan. Every seven-ten year the data on the tapes have been migrated to a new, (at the time), best practice storage media. The last renewal was finished in 2012. But since most new media art is now digital the workflow is changing. Instead of a specific cycle, the artworks will be copied and migrated when LIMA receives them, and then regularly checked and migrated with automated systems. There is a new generation of LTO out every second year and it will be necessary to migrate onto the newest type of LTO and also invest in new reader/writers regularly. The reader/writer used today is backwards compatible and can read LTO generation three and four and can write generation five, but it will not be able to read or write generation six to ten. As long as it is possible to acquire generation five readers it will not be necessary to migrate within the next eight years, unless the writer itself fails for some reason. During this time the tapes are manually put into the reader and a program is started to verify and automatically check the MD5 hash tags. If there are any discrepancies between the MD5 files and the previously recorded MD5, the specific file will be compared to the back up, and overwritten if the backup is correct. This procedure has so far not been needed. The whole procedure and file index is documented in a database.

11.3 Small data collection: Royal Museums of Fine Art of Belgium

The Royal Museums of Fine Arts of Belgium, Brussels, host a collection of approximately 22,000 artworks (paintings, sculptures, works on paper, decorative art), covering a time period that ranges from the 15th until the 21st century. The institution includes five museums: the museum of modern art, the museum of ancient art, the Magritte museum, the Constantin Meunier museum, and the Antoine Wiertz museum. The missions of the MRBAB-KMSKB include the conservation of the collections, conducting research projects, and organising exhibitions. They have been part of the DCA project, with the aim of digitising their collection of contemporary art. The task of storing and preserving digital collections is a new initiative, because the museums' main focus has been on the physical collections. Following the DCA project, there is now a collection of images around 1 TB in size. There is also the daily maintenance of the museums' IT infrastructure.

Data and Storage media

The collection of 1 TB is stored on Hard disk drives, tapes and DVDs. To find and locate the files the museums use DAM (Database Activity Monitoring). This is a technology for analysing activities in databases. Hard disk drives are set up in cabinets and use SAS technology. The server system used is SAN. The HDDs are set up in a RAID 6, for preservation and speed purposes, like the National Library of the Netherlands. It seems that the museum has a good procedure of backup and also stores HDDs in different locations. The tapes used are LTO Ultrium generation four, and the reader/writer machines of generation three. The tapes seem to be the main backup. The museum uses robotic tape libraries, which can hold about sixteen tapes at a time. The robotic tapes are connected to the overall server system so they can be accessed without manual handling. The reader/writer has not yet been changed, but there has been an incident when a tape could not be read or had another error. As for the storage climate, the museums don't have any specific procedures or parameters for this.

Preservation plan

The data is regularly accessed, which might indicate there is no distinction between preservation masters and copies for distribution. Every five years the HDDs are changed and checked, but there is no procedure in place for checking the files' integrity and function. This is a big risk for long-term accessibility. The institution does not particularly recommend using LTO tapes as a preservation storage type, and highlights one good feature - that they can be stored off-site and offline. The tapes are changed every two years. The museum also stores images on DVDs, which are stored in regular room temperature and handled manually. When data is written to the DVDs they are not checked for errors, but in this context it is not known if the DVDs are used for distribution purposes or long-term storage. There are some good starting points, but it does seem that the focus has not been on long-term storage. This could be for a number of reasons; one of them being that the institution has not had a digital collection for long. Secondly, it is a fairly small collection of data. With few adjustments the museum could be on its way to a good long-term preservation plan. The main point is making a good logical preservation plan. Such a plan should include the regular checking of the files' integrity, by integrating error-checking software. Regarding the physical storage it would be wise to move anything stored on DVDs or other optical media onto a more durable storage medium, especially if the DVDs contain original files.

11.4 Small data collection: Macedonian Museum of Contemporary Art

At present the MMCA's permanent collection is made up of 1,800 works by Greek and foreign artists. Apart from the permanent display, the MMCA has put on over hundred exhibitions of works by Greek and foreign artists such as Fluxus artists, Robert and Sonia Delaunay, Matta, Viallat, Beuys, Uecker, Greenaway, Hockney, Ernst, Barlach, Beckmann, Zervos, Kaniaris, Varotsos and many others. The museum's library is made up of 2,500 books and reviews on art. All catalogues published by the MMCA, audio-visual material and the museum's archive on artists are available to the public.

Data and storage media

The collection of data comprises 3 TB stored on HDDs and DVDs and USB flash drives which are used for distribution and backup. The collection is mainly made up of video files and maintenance is all in-house. Two types of HDD connections are used: SCSI and USB. The HDDs are set up in a HP Proliant G5 server and RAID 5 backup system. The hard disks use the vendor specific HP Advanced Memory Protection Technology. The data is checked with online ECC software every two months. The material is also backed up and refreshed every two months. The content is only accessed if there is an update. There are no governmental restrictions or regulations for preservation so the museum follows its own protocol for controlling the data. The hardware storage is at 24°C and 55% relative humidity in a ventilated room with an installed air cleaning system.

Preservation plan

There is a regular and automated system to take care of the HDDs and backups with several disks in different locations. The museum also does regular checks for bad sectors and blocks, as mentioned above, with ECC software being used on a two-month basis. If an error occurs the RAID 5 system will alert the IT staff. During migration every file is checked at the end of the procedure with maintenance software. The museum seems to have good control of their HDD system. Improvements could be made to avoid using DVDs, unless it’s for distribution purposes only. The museum claims the main reason for still using DVDs is that they're easy to use. There is also a small risk when using hard disk drives with USB connections. For a growing collection of data the museum's on to a good start, especially if it's made a written plan for the future handling of its collection.

Terms and abbreviations

An institution's digital collection will in this document be referred to as data, digital content, files, data collections or digital collections. Hardware will be referred to as storage media or hardware. There follows a list explaining different abbreviations and words used throughout the document.

I/O - Input/output

JBoD - Just a Bunch of Disks, see more in chapter 4.3

LAN - Local Area Network is a computer network that interconnects computers in a limited area such as a home, school, computer laboratory, or office building using network media.

MAID - Massive Array of Idle Disks, storage technology like RAID, but with energy management. Instead of the HDD being constantly on, the MAID turns them off to save energy, as soon as the disks are not actively accessed.

MLC flash - Multi-Level Cell, a technology of solid-state drives that can contain several bits on one cell in a microchip.

NAS - Network Attached Storage, a NAS server is a preconfigured file server with internal hard disks, which makes its storage capacity available via LAN (Troppens 2009, storage network explained).

Offsite - Any digital content that is stored on storage media that is not connected to the main headquarters of an institution.

Onsite - Any digital content located in-house, that is accessible from any computer at a workplace.

Partition - The division of a hard disk or a tape to create an appearance of having separate hard drives or tapes in one. Normally used for file management or multiple users. A partition is created when you format the hard disk or tape. On a Microsoft personal computer different drives will typically be nominated with a letter followed by colon, such as *C:* or similar.

RAID - Redundant Array of Independent Disks, originally Redundant Array of Inexpensive Disks, is a storage technology that combines multiple disk drive components into a logical unit. Data is distributed across the drives in one of several ways called "RAID levels”.

SAN - Storage Area Network is an advanced network connecting storage devices to servers, making the devices appear like locally attached disks to the operating system.

SAS – Serial Attached SCSI is a network that moves data to and from computer storage devices with the help of SCSI technology.

SATA - Serial Advanced Technology Attachment, is one cable containing a minimum of four wires that creates a point-to-point interface on a hard disk.

SCSI - Small Computer System Interface, see more in chapter 4.

Server - A computer or computer program that manages access to a centralised resource or service in a network.

SLC flash - Single-Level Cell, is a technology of solid state drives, where one single bit is stored on one single cell in a microchip.

TAR - Tape Archive, file format and software for writing and extracting files from magnetic tapes. TAR can also be used for indexing tapes, but it is not a strong part.

TCO - Total Cost of Ownership, the cost not only of purchasing equipment but also of energy consumption and renewal.

Transfer rate - The time it takes for data to be read or written onto the media (whether a tape, hard disk or other storage device), the transfer rate is measured in MB per second.

WORM - Write Once Read Many, a technology that ensures that data on a HDD, magnetic tape or a DVD/CD cannot be overwritten again, but can still be read repeatedly.

References

Adelstein, P., Z. (2009). *IPI Media Storage, quick reference, 2nd edition.* Image Permanence Institute, Rochester, US, p. 1-8.

Alchesky, L. C. (?). *Single Bit Error - A Memory Module Supplier's perspective on cause, impact and detection* (Online). Smart modular technology, white paper. http://www.smartm.com/files/salesLiterature/dram/smart_whitepaper_sbe.pdf

Anderson, D., Dykes, J. & Riedel, E. (2003). *More than an interface — SCSI vs. ATA. In:* FAST '03: 2nd USENIX Conference on File and Storage Technologies, p. 245-257.

Angevaare, I. (2012). *The wiggly wobbly jelly: predicting costs of digital archiving* (Online). Blogpost about the conference "Screening the Future 2012". http://www.ncdd.nl/blog/?p=1849

Bradley, K. editor (2009). *Guidelines on the Production and Preservation of Digital Audio Objects, 2nd Edition*. IASA – International Association of Sound and Audiovisual Archives, p. 8 – 137.

Brain, M. (?). *How Semiconductors work* (Online). http://electronics.howstuffworks.com/diode.htm

Byers, F. R. (2003). *Care and Handling of CDs and DVDs – A Guide for Librarians and Archivists*. National Institute of Standards and Technology, Maryland, USA (Online). P.42. http://www.clir.org/pubs/reports/pub121/pub121.pdf

Clark, J. (2013). CloudSigma spins-up all-SSD storage cloud. The Register, posted under section Cloud, 4th April 2013 (Online). http://www.theregister.co.uk/2013/04/04/cloudsigma_launches_ssd_storage/

Findlay, C. (2010). *Advice on managing the recordkeeping risks associated with cloud computing, version 1.0* (Online). ADRI (Australasian Digital Recordkeeping Initiative) and Council of Australasian Archives and Records Authorities. http://www.adri.govt.nz/

Hodges, K. (2011). *The Landscape of Data Storage - Cost Advantages per Terabyte of Data* (Online). Hie Electronics. http://www.hie-electronics.com/wp-content/uploads/2012/06/cost_advantage_per_terabyte.pdf

Hoel, I. A. L. (1999). *Dokumenter med informationsbærende lag. Bevaring af dokumenter* – Det teoretiske grundlag. Nordinfo publikation 42. P.148-175.

Hoeven, v.d., J. (2012). *Cloud! Unless... - Pros and cons of cloud technology* (Online). The Dutch National Library. http://www.ncdd.nl/blog/?p=2347

Hope, D. (2011). *Disc deathbed: Is the optical drive obsolete?* (Online). Tech & Gadgets, msnbc. http://www.msnbc.msn.com/id/43883426/ns/.../disc-deathbed-optical-drive-obsolete/#.T7Nh-591BTc

Hunter, A. (2011) *Guideline 3 - Digital Storage*. Public Record Office Victoria PROS http://prov.vic.gov.au/government/standards-and-policy/storage

Iraci, J. (2012). *The Effect of Jewel Cases on the Stability of Optical Disc Media.* Restaurator 2012, Vol. 33, p. 17-47.

Kwanghee Park, Dong-Hwan Lee, Youngjoo Woo, Geunhyung Lee, Ju-Hong Lee & Deok-Hwan Kim (2009). *Reliability and Performance Enhancement Technique for SSD array storage system using RAID mechanism.* In: 9th International Symposium on Communiations and Information Technology 2009, ISCIT.

Marks, H. (2012). *Are SSD Based Storage Arrays Doomed?* Network Computing, for IT by IT. May 2012 (Online). http://www.networkcomputing.com/storage-networking-management/are-ssd-based-storage-arrays-doomed/240000637

Mearian, L. (2009). *Review: Hard disk vs. solid-state drive,- is an SSD worth the money?* (Online). Computerworld, 18.06.2009. http://www.computerworld.com/s/article/9134468/Review_Hard_disk_vs._solid_state_drive_is_an_SSD_worth_the_money_

Mearian, L. (2012). *As 60th anniversary nears, tape reinvents itself* (Online). Computerworld, 29.03.2012. www.computerworld.com/s/article/print/9225514/As_60th_anniversary_nears_tape_reinvents_itself

Muray, T. (2009). Letter from Sony concerning the stop of AIT Drive sales (online). http://sony.storagesupport.com/files/SONY%20AIT%20EOL%20second%20notice%2010-12-09_Fin.pdf

Pease, D., Amir, A., Villa Real, L., Biskeborn, B., Richmond, M. & Abe, A. (2010). *The Linear Tape File System* (Online), p. 1-8. http://storageconference.org/2010/Papers/MSST/Pease.pdf

Pinheiro, E., Weber, W. D. & Barroso, L. A. (2007). *Failure Trends in a Large Disk Drive Population.* In: FAST '07: 5th USENIX Conference on File and Storage Technologies. 1-13p. http://static.usenix.org/event/fast07/tech/full_papers/pinheiro/pinheiro.pdf

PrestoPRIME (2009). *Visualising cost versus risk of loss* (Online). University of Southampton IT innovation centre. http://prestoprime.it-innovation.soton.ac.uk/imodel/visualisation/

Oswalt, B. & Turner, R., L. (2010). *The changing world of records management: Requirements of digital storage.* Association of Business Information Systems 2010 Refereed Proceedings. P. 57-60.

Ravn□Grove, J. (2010). *Hvilket lagringsmedie er bedst?* (Online). Statsbiblioteket, Copenhagen. http://digitalbevaring.dk/lagringsmedier/

Schroeder, B. & Gibson, A. G (2007). *Disk failures in the real world: What does an MTTF of 1,000,000 hours mean to you?* In: FAST '07: 5th USENIX Conference on File and Storage Technologies, p. 1-16.

Rouse, M. (2005). *Tiered storage - definition* (Online). SearchStorage. http://searchstorage.techtarget.com/definition/tiered-storage

Tate, J., Lucchese, F. & Moore, R. (2006). *Introduction to Storage Area Networks. (Online)* IBM Redbooks. https://www.redbooks.ibm.com/redbooks/pdfs/sg245470.pdf

Troppens, U., Müller-Friedt, W., Wolafka, R., Erkens, R. & Haustein N. (2009). *Storage Networks Explained - Basics and Application of Fibre Channel SAN, NAS, iSCSI, InfiniBand and FCoE, Second Edition*. John Wiley & Sons, Ltd, United Kingdom.

Vaquero, L. M., Rodero-Merino, L., Caceres, J. & Lindner, M (2009). *A Break in the Clouds: Towards a Cloud Definition.* ACM SIGCOMM Computer Communication Review, Vol. 39, Nr. 1, 2009. P. 50-55.

Interviews

Jeffrey van der Hoeven MSc
Project manager Digital Preservation, Innovation & Development
National Library of the Netherlands
11. July 2012