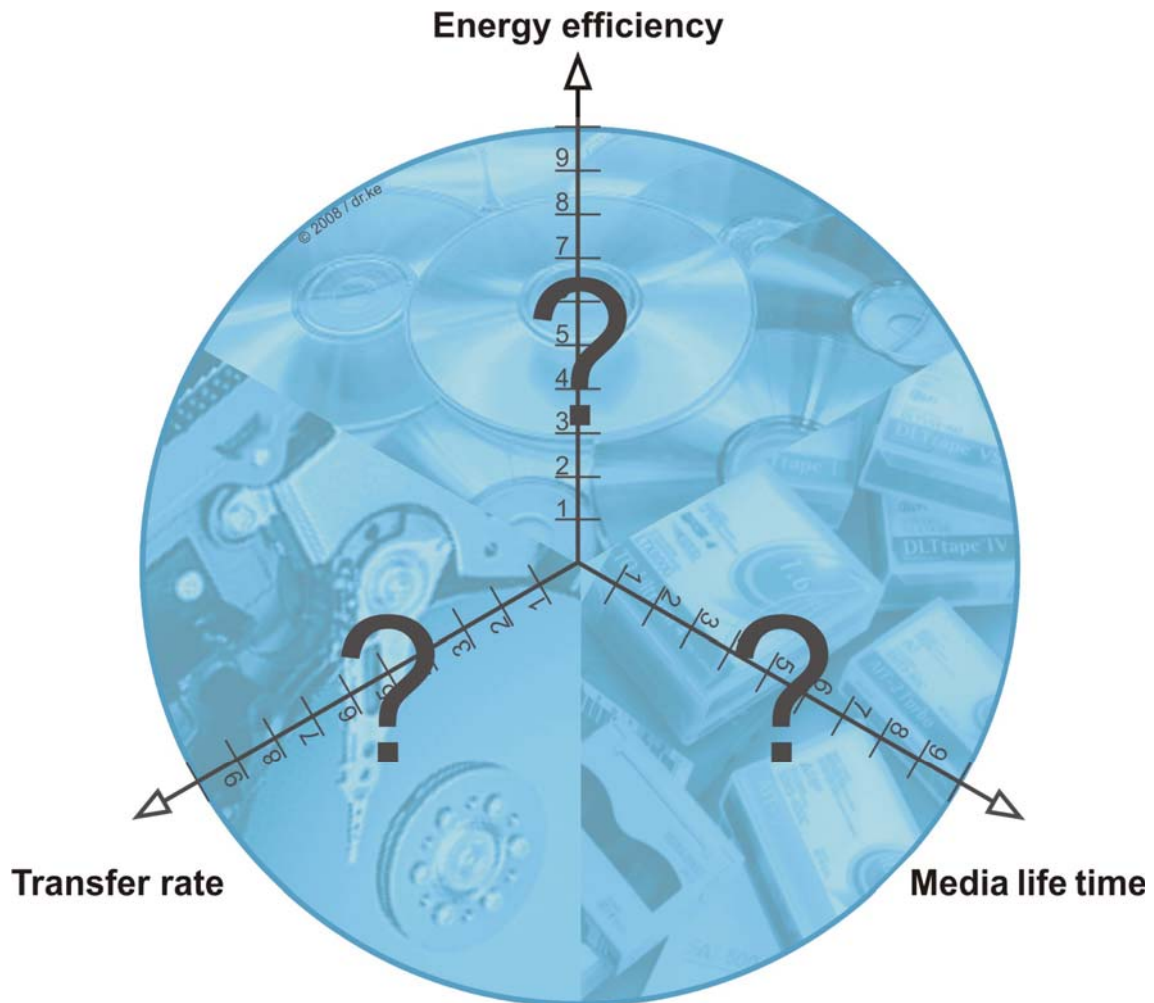




Secure Data Storage

- An overview of storage technology -



**Long-term archiving from extensive data supplies requires more than only big storage capacity to be economical.
Different requirements need different solutions!
A technology comparison repays.**

Author: Dr. Klaus Engelhardt



Secure Data Storage

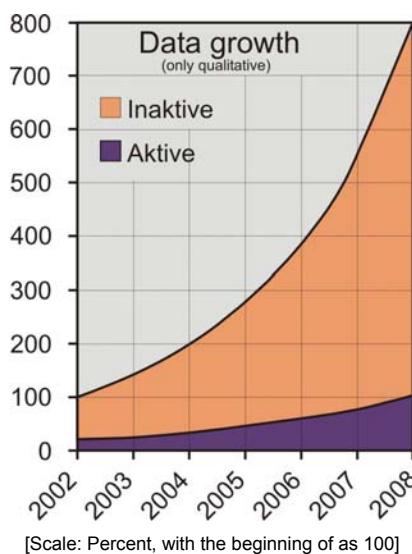
- An overview of storage technology -

Author: Dr. Klaus Engelhardt

Audit-compliant storage of large amounts of data is a key task in the modern business world. It is a mistake to see this task merely as a matter of storage technology. Instead, companies must take account of essential strategic and economic parameters as well as legal regulations. Often one single technology alone is not sufficient to cover all needs. Thus storage management is seldom a question of one solution versus another, but a combination of solutions to achieve the best possible result.

This can frequently be seen in the overly narrow emphasis in many projects on hard disk-based solutions, an approach that is heavily promoted in advertising, and one that imprudently neglects the considerable application benefits of optical storage media (as well as those of tape-based solutions). This overly simplistic perspective has caused many professional users, particularly in the field of long-term archiving, to encounter unnecessary technical difficulties and economic consequences. Even a simple energy efficiency analysis would provide many users with helpful insights. Within the ongoing energy debate there is a simple truth: it is one thing to talk about 'green IT', but finding and implementing a solution is a completely different matter.

1. INTRODUCTION



Fact 1: The amount of digital data stored in companies doubles every one to two years, depending on the type of company and sector. Thus the biggest challenge is to meet the demands of this ongoing growth, which is set to continue and even gain momentum in future. This can be accomplished not only by expanding the capacity of storage media, but also, in view of the new and increasingly complex needs of the market, by devising a practical and comprehensive storage strategy.

However, contrary to the common opinion, a strategy is not just limited to controlling the investment costs for the acquisition of storage media – it should also take account of all the parameters that are involved in a complete solution that ensures long-term economic success and above all provides an answer to the challenge of data storage.

A vital point to consider is the life cycle of the different types of data and documents (key terms here are active or passive data; ILM, fixed content etc. – see Glossary).

Fact 2: In determining the cost-effectiveness of a particular storage solution, it is not sufficient to simply examine the acquisition costs, as is often the case in practice, or to attempt to gain accurate estimates by taking personnel and storage space factors into consideration on a pro-rata basis. A professional approach is to add to these costs the aspect of the energy efficiency of a particular solution, including all other (hidden) expenditure. The key term is TCO – Total Cost of Ownership. It simply refers to the total cost calculation of all the direct and indirect costs of an economic asset throughout its entire serviceable life.

Special attention should be paid to the high energy costs involved, a factor that many users underestimate. For many years the impact of computer centres and data storage



solutions on energy efficiency was hardly considered, not to mention taken seriously. However, since late 2006, this all too wastrel attitude has begun to change, which in view of the ongoing climate debate is long overdue.

Fact 3: Many users do not place enough importance on the issue of information management or storage management. Information management is a strategic corporate task, and one that can determine the success or failure of a company.

The fundamental decisions that need to be made concerning the handling and management of electronically stored corporate data, including data security and archiving, are ultimately management tasks. Thus, in the event of a disaster – which can be triggered by something as innocent as a lost e-mail – it should not be IT administration alone that is held responsible. Clearly, the issues of information and storage management must be given a great deal of consideration by the very highest levels of management.

Fact 4: The current market for comprehensive solutions for large-scale data access, storage and security does not place enough emphasis on application-oriented approaches. A neutral observer will notice how many providers of mainly hard disk-based solutions monolithically propagate their own technological and practical strengths without devoting much attention to their weaknesses or indeed to the benefits of alternative technologies. Such a one-sided approach offers little assistance for companies looking for the best overall solution for the right storage strategy. This in turn leads to widespread confusion and insecurity among many potential users.

Fact 5: In the market for new optical storage solutions based on blue laser technology, the final battle in the war over the format to succeed DVD in the consumer market was fought and won in early 2008. Warner Bros announced that they would cease to support HD-DVD, making them the fifth Hollywood studio to decide to only release films from now on in the Blu-ray[®] format.

The conflict between the two competing Blu-ray and HD-DVD formats was hard-fought, and ultimately damaged the parties on both sides. As a result of this battle, confidence in the industry was lost, causing the market to practically stagnate up to and including 2007, despite the initial successes of both formats. Consumers are used to bearing the brunt of format conflicts, but in this case the manufacturers went too far. This phase should now be over, despite the fact that the HD-DVD consortium believes that they can still turn the tide by offering huge discounts on HD-DVD drives.

The professional market's initial uncertainty as to whether to decide for or against one particular format was mainly due to the lack of a standard format for both Blu-ray Disc (BD) and HD-DVD (High Density DVD). Ultimately, however, the professional market has not been overly impacted by the hard-fought format war. At a very early stage it began to focus primarily on BD technology, and this format has now become the successor to the DVD in all major professional storage applications. The DVD, in turn, still has enough areas in which it retains a very strong economic position as a data storage medium.

The applications of BD cover both the classic data storage market (medium to large amounts of data) and the multimedia sector, as developments have shown. In particular, the most disc drives available and sold on the market are in BD format. The BD camp will now try to ensure that the market begins to settle down.

It should be emphasised that in the professional sector BD is the benchmark for large-scale data storage when DVD-based solutions are no longer sufficient. BD-based storage solutions for large amounts of data have long since undergone their baptism of fire and now offer users obvious long-term benefits. Due to the enormous data capacities that can be integrated into individual solutions and to the advantages that this technology offers in terms of functions, energy efficiency and cost-efficiency, the professional



sector holds a great deal of promise for BD. Ultimately all these technologies are set to profit from the enormous increase in data and from the tightening of national and international legal regulations (Basel II, Sarbanes-Oxley Act, Turnbull Report, etc.), which will create more and more need for storage solutions. It is only as a result of these trends that manufacturers have pushed development forward. This has been of benefit both to BD as well as the second UDO generation (UDO2), which has successfully been launched onto the market.

2. DATA STORAGE IN COMPANIES

At first glance there seems little distinction between compliant long-term archive technology and high performance magnetic disk storage media. Yet on closer inspection it can be seen that this perception is a result of intensive marketing campaigns by the proponents of hard disk-based storage solutions, with little consideration of what the storage media are used for. If one looks at developments from a user's perspective, the actual situation becomes clear. There are obvious benefits to matching different storage media to specific application and business requirements. Often, this entails searching for the right combination of tried and tested technologies and not simply hoping to strike it lucky by deciding on one technology only, particularly with regard to strategic solutions in large companies.

Quite some time ago, the accepted wisdom was to automatically store long-term and audit-compliant digital data on optical media. Data was cached on hard disks and to ensure system reliability backup and mirroring solutions were established. The issue of system stability has often been solved by using complex tape-based storage solutions. Yet these approaches made it very labor-intensive to adapt solutions to growing data sets and cope with the recurrent need for data migration, particularly if proprietary solutions were no longer feasible.

Hard disk-based approaches are also increasingly making headway into the field of long-term archiving, spurred on by the enormous fall in unit prices (euros/ capacity unit). Black box concepts are being developed that use digital signatures to deal with the data immutability, will offer scaling, security and migration, and integrate WORM tapes. It is unclear and indeed doubtful in some cases whether all the practical requirements for comprehensive long-term storage and especially for archiving can be met by such solutions. In any case, it has become considerably more difficult to critically assess the particular demands of a comprehensive storage solution as a result of the growing range of solutions on offer and the increased pressure from advertising.

Conclusion: It has become significantly more difficult for potential users to find optimal solutions for comprehensive systems for the storage of and rapid access to large amounts of data. Therefore the requirements for storage models need to be compiled and analyzed in advance with much more care and attention than was the case several years ago. Aspects of corporate strategy must be carefully considered given the legal regulations and the idiosyncrasies of each market sector. It is no longer sufficient to be educated on only one possible technological solution, even if the advertising strategies from some providers would like to create this impression that there are no credible options.

Instead, it is essential to be aware of the strengths and weaknesses of a range of possible technologies before making a decision. A key challenge is to be able to combine different approaches to create a solution that merits being described as "the best possible" or "an economic necessity". Thus all three basic technologies should be given due consideration without prejudice and with the requirements of the system in mind. Optical storage solutions should be given as much consideration as tape-based or hard disk-based options. And often the solution to the puzzle is a hybrid application that combines technologies to complement each other.



3. APPLICATION SYSTEMS

Throughout the entire sphere of company-wide information management (including public and scientific organisations), terms and concepts are often used that people understand in different ways. Sometimes the concepts for which different current terms are used overlap, sometimes different terms refer to practically the same thing, other times several concepts are used to describe the very same thing, and often the terms that are used are little more than buzzwords.

To avoid causing any more confusion, the following will be terms will be used in connection with the content of this White Paper (with a focus on long-term data storage / storage solutions and the management of large amounts of data): **DMS** (**D**ocument **M**anagement **S**ystem), **ECM** (**E**nterprise **C**ontent **M**anagement), e-mail management, **ILM** (**I**nformation **L**ifecycle **M**anagement), **BPM** (**B**usiness **P**rocess **M**anagement).

3.1 Storage and Archive Management

Data-intensive applications, such as e-business or more specifically CRM (**C**ustomer **R**elationship **M**anagement), e-mail and a wide range of database applications (data warehousing, data mining, groupware etc.), create much of the data that needs to be processed in companies.

An additional factor is the requirements for broad-based applications such as DMS or ECM. It is not possible to get a grip on the amounts of data in a company simply by increasing the available storage capacity. Instead, it is necessary to use intelligent and economical storage management to control data and to ensure that it is accessible over time and in accordance with business requirements.

Storage systems are an essential part in the overall organization of information. Comprehensive storage management goes far beyond backup and recovery. Other key tasks include capacity planning and coping with the particularities of data availability on the one hand and the requirements of long-term archiving and legal security on the other. Here ILM is becoming increasingly important and now overlaps to a great extent with information and storage management systems.

A definition of ILM: ILM is a storage strategy. Its goal is to use optimized technology, rules and processes to transfer information and data to the most economical storage medium available, depending on the value of the data in question. The control mechanisms of this transfer process are based on the priority, value and cost associated with specific information or data. (Thus ILM is not a software product but a combination of technology, storage media and processes. Its primary aim is to ensure that information is constantly available at the right time, at the right place and in the right form, and at the lowest possible cost.)

DMS, ECM, ILM and similar systems serve the common purpose of administrating all data in a company and to ensure that it is constantly available. The key focus of such systems is managing data that is essential to business operations. Such systems make use of a wide range of criteria, classifications and data structures to guarantee access to the data when required. Data that has been archived but that is no longer essential or can no longer be located is of no value for a company and in some cases can create additional corporate risk.

Important criteria for all applications include the currency (age) of the data, their role in business processes, the requirements of corporate strategy and legal regulations such as immutability and the retention periods (fixed content, compliance, etc.). The general task of information and storage management is to incorporate all criteria and other company-specific requirements into a viable solution.



Depending on the sector and type of company involved, fixed content (i.e. stored data that is no longer altered) is reliably estimated to constitute up to 80 percent of corporate data, with this figure set to increase in future. One reason is that information is becoming increasingly important so must be retained for longer periods of time. This older data seldom changes or must be specifically retained in its original form. Companies want to maintain access to more and more data for longer and longer periods of time. This can often mean retaining business records for decades, placing new challenging demands on hardware and software, and particularly company-wide storage management.

To meet this challenge, comprehensive solutions that attempt to deal with this trend must pay particular attention to the long-term properties of potential technologies. Here it already becomes clear that there are key business requirements that can best be fulfilled by optical storage media rather than hard disk-based solutions.

3.2 Storage networks / storage organisation

Important aspects of storage networks include: unlimited access to the mass storage media, shared and centralized use of high-performance systems, balanced data load between the available backup devices, central hardware monitoring and configuration, improvement of the amount of data that can be administrated and managing load on local area networks.

DAS - **D**irect **A**ttached **S**torage: DAS refers to a mass storage device that is directly attached to a server (without a storage network). They may be individual internal or external hard disks or an external disk array consisting of several hard disks. The server and storage systems normally communicate with each other via dedicated SCSI or fibre channel interface. If storage resources need to be expanded, the DAS approach is relatively complex; in addition, as all the data flows through the servers, which may form a performance bottleneck in the system.

NDAS - **N**etwork **D**irect **A**ttached **S**torage: NDAS refers to storage media (mainly hard disks) that can be connected directly to a network without using a PC or a server and which appear on the target system as a local data storage device.

NAS - **N**etwork **A**ttached **S**torage: NAS refers to mass storage units connected to the local area network that are designed to expand storage capacity. NAS is usually used to avoid the complex process of installing and administrating a dedicated file server.

SAN - **S**torage **A**rea **N**etwork: In the field of data processing, SAN refers to a network that allows hard disk subsystems and storage units such as tape libraries to be connected to server systems. SANs are dedicated storage networks that connect and decouple servers and storage systems via broadband networks such as fibre channels. A SAN differs from a local area network (LAN) in that it forms a network between servers and the storage resources they use.

iSCSI is a protocol that describes SCSI data transport via TCP/IP networks. Thus an SCSI device (such as a hard disk, RAID system, optical devices or tape libraries) can be connected to a network (Ethernet LAN) via a bridge. On the server side communication takes place via a network card, while the storage system is “seen” as being connected locally via SCSI. This allows storage systems and computers to be located at great physical distances from each other.



4. TECHNOLOGIES, SYSTEMS, MEDIA

This section deals with the storage of very large amounts of data, and provides a concise overview of the possible technologies or media that can be used, from CDs through DVDs to BD and UDO, as well as, hard-disk storage to tape.

4.1 CDs, DVDs and BDs

Many years ago things were much simpler. Documents that had to be stored for a long period of time for legal or policy reasons were archived in their original paper form and were then shredded when the retention period expired. As these collections increased in size and required ever larger archiving space. Sometimes the records were photographed and placed on microfilm or microfiche, which could then be archived instead of the paper documents, thus saving a great deal of space. (There have always been certain documents that had to be stored in the original form for legal reasons, which required solutions involving film-based and paper-based archives, both of which had to be maintained.) As reading and re-magnification technology gradually improved, these solutions were used for to cope with the constantly increasing frequency of documents access in long-term storage. However, this type of solution has mostly outlived its usefulness.

As storage media continued to develop and new formats evolved, improved technology and efficiency provided the possibility to store very large amounts of data, regardless of the type of requirements involved. The transfer rate and capacity of hard disk and tape-based technology, for instance, developed at a dramatic pace.

However, the biggest step forward in the flexible and standardized long-term storage of large amounts of data did not occur until the invention of the **optical disc (CD - Compact Disc)**. The magic word associated with the optical disc is **WORM (Write Once Read Many)**, or more precisely "**True WORM**". With this development, the ideal means was available to store documents over a long period of time in a cost-effective manner, permanently, securely and immutably (audit-compliance).

For writing and reading CDs a laser beam with a wavelength of $\lambda = 780 \text{ nm}$ was used (1 nanometre = 10^{-9} m). The CD's capacity has remained at 650 MB.

Yet developments did not stand still and attempts were made to considerably improve the performance parameters of the CD by using laser light of increasingly shorter wavelengths.

A particularly important step was an increase in the CD's storage density and a reduction of access times, as these factors ultimately determined whether technological developments would continue or come to a halt. This must be understood in the context of the enormous increases in capacity per medium that were achieved in hard disk technology and the subsequent falls in media unit prices in €/GB. However, this factor alone is no guarantee of an economically viable storage solution, particularly not for long-term storage.

The second generation of optical discs based on the CD is referred to as **DVD (Digital Versatile Disc)**. DVDs offer a capacity of 4.7 GB per layer/side, which allows for a total of 9.4 and 18.8 GB respectively for dual-layer and dual-layer double-sided discs. Red laser light with a wavelength of $\lambda = 650 \text{ nm}$ is used for reading/writing.



DVD

The development of the third generation of optical media and drives, originally based on the CD and then on the DVD, is a particularly important milestone in increasing both capacity and read/write speeds. This technological leap forward also demonstrates that optical media continues to have a great deal of potential in terms of being able to increase their capacity.



The breakthrough was achieved by using laser light of an even shorter wavelength, as had been the case in the development of UDO. In contrast to the DVD, the wavelength for the new medium is a mere 405 nm. This allows for very high recording densities on the discs coupled with a considerable reduction in access times.




Fig. 4.1.1: Structure of a BD dual layer

Parameter	Blu-ray Disc (BD)
Diameter	120 mm
Protection	Hard coating, no cartridge
Capacity	single layer 25 GB dual layer 50 GB next generation 100 GB
Laser wavelength	405 nm
Numerical aperture (NA)	0.85
Write transfer rate	4.5 MByte/s
Read transfer rate	9.0 MByte/s
Layers/sides	1/1 or 2/1 or (shortly) 4/1


Tab.: 4.1.1: Key parameters of the BD

The thickness of the Blu-ray layer is approx. 100 μm . The cover layer L1 is 75 μm thick, and the spacer layer is 25 μm thick. Substrate thickness is approx. 1.1 mm.

The new technology is known by the abbreviation “**Blu-ray**®” and was supported by a consortium encompassing over 100 companies, including the founders Sony and Matsushita. The **Blu-ray Disc (BD)** has a recording density that is 6-7 times higher than that of the DVD. As a single layer format it offers 25 GB of storage capacity, as dual layer 50 GB. The maximum data rate of a first-generation BD is 9 MB/s at single reading rate. Drives with four-times (4x) reading rate are now available on the market, as are the corresponding media, and the upwards trend is set to continue.

The Pioneer Corporation announced that now a 20-layer optical disc is feasible. Every layer can store 25 GB of data by which 500 GB capacity would be available on a Disc. This developing result, including a temporary specification, was presented at the international symposium for optical data memories (ISOM) in July, 2008 on Hawaii. The market availability of the 20-layer Disc is calculated for 2011. The method of the multi-layer is compatible to the Blu-ray technology and therefore sure for the future. Important side effect: The production is substantially easier than, for instance, those of holographic memory medias.

Technology of the CD, DVD and BD formats



Legend: NA = Numerical aperture, λ = wavelength of the laser light in nanometres

Fig. 4.1.2: Overview of the different capacities



Several Blu-ray formats are available on the market: **BD-ROM** as a distribution medium, **BD-RE** for HD-TV recording in the film and television industry, and **BD-R** as a once-only recordable storage medium. The format often referred to as **BD-RW**, a rewritable PC storage medium, is actually the BD-RE format.



Although the rival **HD-DVD (High Density DVD)** format was used until the early 2008 alongside **BD** to a certain extent in the electronic entertainment industry (see the remarks on this point under chapter 1./ Fact 5), the storage market has chosen **BD** (Sony / Matsushita) and **UDO** (Plasmon) for professional optical storage applications very early.

The following diagram shows the trends in developments for storage capacity for UDO and the next generations of CD-based media (the roadmap for UDO and BD):




Fig. 4.1.3: Roadmap for Blu-ray Disc and UDO

4.2 UDO - Ultra-Density Optical



UDO was developed for professional use (long-term archiving requirements) by Plasmon as the successor to **MOD** (Magneto Optical Disc). UDO is based on phase-change technology.

The UDO medium is enclosed in 5¼"-format cartridges. Three variants are available for archiving solutions: **True Write Once** (once-only recording, undeletable), **Compliance Write Once** (once-only recording, data can be destroyed if necessary using a shred function) and **Rewritable** (rewritable up to 10,000 times).

Both the UDO and BD technologies use laser light of the same wavelength. Yet they differ in one key aspect: in the first-generation lens system that they use. Whereas Blu-ray uses a 0.85 NA lens, the numerical aperture of the UDO1 lens was 0.7. Numerical aperture is a measurement of the level of precision with which an optical system can depict details (see also Fig. 4.1.2). As a result of its slightly higher NA, BD also automatically achieves a higher storage capacity of 25 GB per layer on one side (currently 50 GB as a dual layer), while UDO1 has a capacity of 15 GByte/layer on both sides. Second generation UDO2 uses a 0.85 NA lens to provide 60 GB (2 layers of 15 GB each on two sides).



UDO drive with medium [Plasmon]



As the developmental trend in Table 4.2.1 below shows, Plasmon offers potential users investment security. This is also visible in their commitment to making the generations backward read compatible.

Parameter	Generation 1	Generation 2	Generation 3
Diameter	130 mm	130 mm	130 mm
Protection	Cartridge		
Capacity	30 GB	60 GB	120 GB
Laser wavelength	405 nm	405 nm	405 nm
NA (aperture)	0.7	0.85	0.85
Write transfer rate	4 MByte/s	8 MByte/s	12 MByte/s
Read transfer rate	8 MByte/s	12 MByte/s	18 MByte/s
Layers/sides	1/1	1/2	2/2

Tab. 4.2.1: UDO performance parameters and roadmap

UDO achieved broad market acceptance in 2004 when Ecma International Information and Communication Technology (the industrial association for the standardisation of ICT systems) declared that the ISO (International Organisation for Standardisation) and the IEC (International Electrotechnical Commission) had officially recognised Plasmon's UDO media format standard ECMA-350. This standard specifies the mechanical, physical and optical properties of 30 GByte UDO media with rewritable and genuine write-once technology. A large number of hardware and software providers, including POINT Software & Systems, IBM, HP, EMC/Legato, Symantec/Veritas and Verbatim (Mitsubishi), support the UDO standard.

4.3 Optical Storage Devices

Today there is a very broad range of optical storage media available for use in a wide spectrum of storage concepts. The misconceptions surrounding optical media, namely limited storage capacity, long access times, sensitivity to external influences, and high maintenance and handling costs, should be seen as little more than modern-day myths.

Fact: Optical storage is used today across virtually industry including manufacturing, public administration, the health sector, banks and insurance companies. If large amounts of data need to be stored for long periods of time in a reliable and immutable form that conforms to legal requirements, then an overall solution involving the use of optical storage media is almost always the best choice. It is also true that the development of optical storage technology is by no means over and that the costs for the use and maintenance of optical storage solutions are practically negligible.

It is not just the misguided debate regarding the supposedly high support costs for optical storage solutions that shows the urgent need for a more objective treatment of the issue, yet the debate is a particularly illustrative example of common misconceptions. This is particularly true regarding cost, namely the expenditure on software and hardware support, a professionally installed optical solution has considerable advantages over a hard disk-based system, especially when comparing the lifespan and stability of the respective media. (Which manufacturers can offer a 30-year guarantee for hard disks, as is the case for DVDs, or even 50 years for the now widely used UDO and Blu-ray Disc?)

On the other hand, hard disk-based solutions offer unbeatably fast access times for the storage of data with short life cycles. Thus in many practical situations the most obvious choice is to consider a combination of the strengths of both technologies as a possible solution. There are already very successful appliances that make use of an optical jukebox for long-term, complaint storage with capacities greater than 10TB with integrated RAID storage for high performance. There are also comparable hybrid solutions that make use of magnetic tape storage.



Despite all the practical advantages of tape or hard disk-based solutions, thought should also be given to the limitations of these technologies. However, the scope of this White Paper unfortunately does not allow for a more detailed discussion of the security costs and considerations of each system (such as replacing the tapes every two to four years, frequent integrity check, migration problems, etc.).

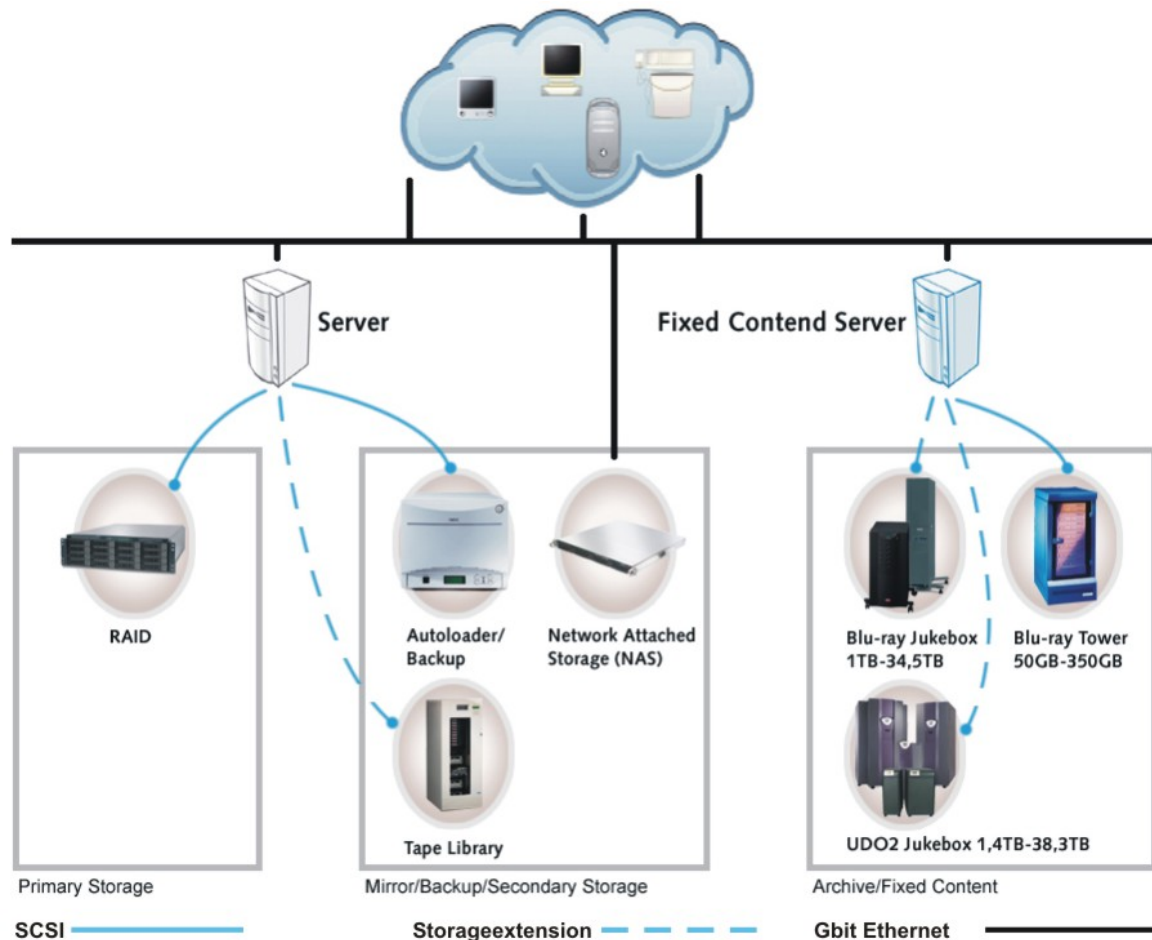


Fig. 4.3.1: Example of application-based storage structure [INCOM 2008]

Fig. 4.3.1 shows an example of how a range of storage components based on various technologies (hard disks / RAID / tape libraries / SAS / Blu-ray tower / Blu-ray jukebox / UDO jukebox) can be integrated into one system.

Similar consideration should be given to the factors involved in long-term storage on hard disk. In order to achieve anything approaching the stability of optical solutions, a hard disk-based solution must have very high redundancy, especially in terms of hardware, which in turn leads to increased energy costs and system complexity.

Thus, if a potential user is faced with the necessity of selecting the most economical system available for given storage tasks, they should not only take the investment costs into consideration, but should also give equal priority to operating costs. As there are several distinct differences between the technologies, particularly regarding cost, special attention should be paid to the lifespan and stability of the solutions themselves. This includes the investment protection that can be expected from a solution in addition to the media used. In view of this, it often makes sense to consider an intelligent combination of a range of technologies from the outset, and not to see one technology as being exclusive of all others. A hybrid architecture is frequently the solution of choice, or should be given serious thought from a broad economic perspective.



Interim conclusion on optical solutions

Given the tangible strengths of optical storage solutions when examined from the viewpoint of long-term archiving, it is worthwhile to list the main practical advantages of such methods. This is especially true when you consider how many potential users are actually unaware of these benefits (not least because their perspective is blurred by the overly simplistic marketing campaigns of a number of providers of hard disk-based solutions):

- ✓ Optical solutions when used properly are particularly cost effective. This is true of the installation costs and particularly the long-term operating costs.
- ✓ Optical solutions are highly energy-efficient compared to RAID systems which is one reason why their operating costs are so low.
- ✓ Optical solutions are transparent and manageable. This is true both in terms of the technology used and the overall configuration of hardware and software.
- ✓ The key benefit of jukeboxes is their high level of reliability. The mechanics are robust, access is precise and operation is reliable.
- ✓ In everyday applications, jukeboxes are very easy to use. Optical solutions, including the storage media themselves, are extremely user-friendly and simple to operate.
- ✓ It is extremely simple to add a jukebox to a storage environment, even ones that make use of the latest generation of drives/media since there is no need to change the administration software. As such, concerns over system capacity are easily addressed.
- ✓ The maintenance costs of optical solutions are relatively low. Data migration is only rarely necessary due to the long lifespan of the media. In cases where migration is required, it can be easily carried out by using software-based mirroring in an updated system.
- ✓ Modern jukebox solutions based on BD or UDO already provide capacities of tens of Terabytes. Additional jukeboxes can be rapidly connected and integrated to the network.
- ✓ The number of drives has no impact on capacity, only on access speeds, which makes it particularly attractive for long-term archiving.
- ✓ The CD started out with 650 MB. The DVD has a capacity of up to 18.8 GB. Currently, the BD provides 50 GB and UDO2 offers 60 GB per disc. The technology is undergoing constant development, and there is no end in sight (see also Figure 4.1.3).

To IDC (2008) optical memory technologies will assert themselves in the future very well, in 2010 with a (estimated) interest in the installed capacities of just 29% (compared with hard disk of 58% and tape of 12%).




Fig. 4.3.2: Installed storage capacity in dependence of technologies (2008-2010 estimated)
[Source: IDC 2008]



4.4 Hard disk and RAID


4.4.1. Hard Disk

The **hard disk** (HD – or **hard disk drive** – HDD) is a magnetic storage medium. In 1973 IBM launched the renowned Winchester Project, which led to the development of the first drive containing a built-in hard disk (Winchester disk drive).

Interfaces

Parallel interfaces such as ATA (IDE, EIDE), or SCSI were initially used to allow hard disks to communicate with other computer components. ATA was mainly used in home computers, whereas SCSI was used in servers, workstations, high-end PCs. Nowadays serial interfaces are more wide spread.

For technical reasons, there are limitations to data transfer rates; SCSI Ultra-320, for example, is limited to a maximum of 320 MByte/s. Initial efforts were made to develop the Ultra-640 (Fast-320) with a rate of up to 640 MByte/s, but were eventually abandoned.



Hard disk read/write head
 Fig. 4.4.1.1: Read/write head of a typical hard disk drive, 3.5" format [Source: Petwoe]

To all intents and purposes, even the 320 had reached the limitations of what is electronically possible. Today, Serial Attached SCSI (SAS) is the most commonly used system. This initially had defined transfer rates of 300 MByte/s. The next stage of development, currently in use, offers up to 600 MByte/s, with a further stage in development offering up to 1,200 MByte/s.

The first serial interfaces for hard disks in widespread use were SSA (Serial Storage Architecture) and fibre channels (in the form of FC-AL/Fibre Channel-Arbitrated Loop). Where as SSA disks are now virtually unknown, fibre-channel hard disks are still used in large storage systems. (The fibre channel protocol is an interface for an electronic interface, not, as the name suggests, for an optical one.)

New interfaces have since become established, such as Serial ATA (S-ATA or SATA). Advantages over ATA include a higher data transfer rate.

External hard disk drives can also be connected via universal interfaces, such as USB or FireWire, although the hard disk drives themselves are equipped with traditional interfaces (predominantly ATA or S-ATA). Communication between the hard disk and other components is even higher with a fibre channel interface, and if used with glass fibres is above all suited for use in storage networks (such as SANs). Even in this configuration, communication with the disk is not carried out directly but via a controller. All types of hard disk interfaces may be used, such as FC-AL, SCSI or Serial ATA. If IP networks are also included, a rival approach may also be used – iSCSI.

Data security / reliability

Due to the technology they are based on, hard disks are subject to reliability risks and thus may require extensive measures to be taken to improve operating stability depending on the application in question. Section 5.5.2 deals with the risks in more detail; only the key points are listed here.



The main causes of security risks up to and including catastrophic failure are:

- ⇒ Overheating (see also Energy efficiency)
- ⇒ Mechanical strain from vibrations, dust etc. leading to increased wear in the vicinity of the read/write head – so-called head crash.
- ⇒ External magnetic fields (dynamic and static)
- ⇒ Electrical faults or use in spaces with excessive humidity
- ⇒ High-frequency access, leading to high load alternation in the read/write head (head crash).

The reliability of hard disks is also measured by the average number of operation hours a disk runs before it fails. Thus the unit of measurement of hard disk failure is time (hours of operation). As this is ultimately a statistical problem, hard disk manufacturers only give approximations of the expected lifespan. MTBF is the abbreviation used to refer to the average time until reparable failure, MTTF for the average time to irreparable failure (see also Section 7 and Glossary).

Remarks on the security of hard disk systems in practice

Any serious attempt to guarantee an appropriate level of reliability for applications that make use of hard disk systems usually involves a considerable amount of effort. It is by no means sufficient to claim that hard disk-based systems can be used everywhere and for all requirements with no risk of failure. Potential users should always ensure to check which storage systems will be used for which purposes. Hard disk systems are unbeatable for specific applications, yet there are enough problems in practice (such as long-term archiving) where other technologies have very clear advantages, particularly compared to hard disk systems. It is very useful to illustrate this with a remark made by an acknowledged expert on hard disk systems:

- *“Probability of a data loss in the first five years of a 100-TB enterprise storage system using SCSI disks and RAID 5: 24 %!”*

From a presentation entitled “Highly-reliable disk storage: Improving disk-based storage reliability together with RAID” given at SNW Europe 2007 in Frankfurt, by the head of the “Advanced Networking and Storage Software” project conducted by IBM Research’s Zurich Research Lab.

This remark is solely intended to raise awareness among potential users of the reliability issues involved, and should be seen in the context of the above-mentioned presentation as a whole.

The specifics of hard disk-based solutions

As a form of online media, hard disks are specially designed for constant and rapid read/write access. This is where the strengths of hard disk-based storage solutions lie – for dealing large quantities of high-access active data.

Yet in the case of the audit-compliant long-term storage of large amounts of data, the limitations of this technology quickly become clear. A brief examination of the time needed to read large volumes of data from IDE disks, as required in data migration, speaks for itself. The results of a specially conducted practical test have shown that reading 50 TB of data in a migration scenario can take up to several weeks; which is not a trivial matter. Of course, the time needed can be reduced using particularly elaborate RAID systems, but even then the migration of very large amounts of data still requires days to complete, which in itself is far from the final stage of a long-term archive.



4.4.2. RAID


RAID refers to **R**edundant **A**rray of **I**ndependent **D**isks

This background to this system was research conducted at the University of California, Berkeley, in 1987 in an attempt to find a cost-effective method of operating hard disks jointly (in an array) as a single logical hard drive. The aim was to try to move away from the large and extremely expensive hard disks used at the time. However, using several hard disks at once increased the risk of failure among the individual components. This problem was tackled by deliberately storing amounts of redundant data (which was then distributed logically to the hard disks). The researchers named the individual arrangements RAID levels. The term RAID was originally derived from the title of a study “A Case for Redundant Array of Inexpensive Disks (RAID)”.

Thus the organizational logic behind a RAID system is to combine two or more physical hard disks (of a computer) to form one logical hard drive. Compared to physical hard drives, the new RAID drive often has considerably higher data throughput (depending on the type and number of hard disks involved), as well as a large storage capacity. The deliberate generating of redundant data leads to a major increase in the reliability of the entire system, so that if individual components (or even individual hard drives) fail, the functionality of the RAID system is not affected. In modern systems it is therefore also possible to replace individual hard disks (even several disks, depending on the RAID level) during operations.


Consequently, RAID systems offer the following overall advantages:

- ⇒ Increased data security (via redundant data storage)
- ⇒ Increased transfer rates (improved performance)
- ⇒ Possibility of constructing very large logical drives
- ⇒ Replacement of damaged hard disks during system operations, and addition of hard drives to increase storage capacity during operations
- ⇒ Reduced costs via the use of several less expensive hard drives
- ⇒ Rapid improvement of system performance if needed



Keywords: Striping/Improved performance without redundancy. In terms of the modern definition of RAID (see above) level 0 is not strictly a RAID system, as there is no redundancy. However, RAID 0 provides a higher transfer rate. Several hard disks are combined, and write operations are carried out in parallel via a cache. This also increases the performance of the reading process.


Statistically speaking, the key advantage of RAID – increased reliability – is actually turned on its head in a level 0 system. The probability of failure of a RAID 0 system consisting for example of four drives is four times higher than that of a single drive. Thus level 0 only makes sense in cases where data security is of minor importance. The most common applications are in the field of music and video, where the main concern is to store large amounts of data rapidly.



Keyword Mirroring: The RAID 1 array system usually consists of two hard drives (or pairs of two or more drives) that contain the same data, a method known as mirroring. RAID 1 thus offers complete data redundancy, which entails, however, that system capacity is limited to the storage size of the smallest hard drive in the array.



RAID Level 5



Keywords: Performance and parity. RAID 5 provides both the advantages of high reliability through redundancy and improved performance. In addition, RAID 5 is a particularly cost-effective variant for the redundant storage of data on at least three hard drives. Level 5 is thus used in the majority of applications, although special consideration must be given to how write-intensive an application may be.

From the user's perspective a logical RAID drive does not differ from one hard disk in isolation. The arrangement of the individual drives (such as redundancy) is determined by the RAID level. By far the most widespread levels in practice are RAID 0, RAID 1 and RAID 5.

There may nevertheless be considerable reductions in performance in both random and sequential write access. If this is area of particular concern, a combination of levels 0 and 1 would make the most economic sense.

Less common RAID levels

RAID 2 is no longer of relevance for modern practice. This method was only ever used in mainframe computers.

RAID 3 (striping with parity information on a separate hard drive) is the precursor of RAID 5. In RAID 3 redundancy is stored on an additional hard drive. It has since disappeared from the market and has predominantly been superseded by RAID 5, where parity is distributed equally over all hard drives.

In **RAID 4** systems, parity information is also calculated and then written onto a dedicated hard drive. The disadvantage of traditional RAID 4 systems is that the parity drive is involved in all read and write operations. As each operation involves the use of one of the data drives and the parity drive, the parity drive is subject to more frequent failure. Due to the dedicated parity drive, RAID 5 is almost always used instead of RAID 4, an exception being that Network Appliance uses RAID 4 in its NAS systems.

RAID 6 (advanced data guarding / redundancy over two additional hard drives) functions in a similar way to RAID 5, but can tolerate the failure of up to two hard drives. RAID 6 systems calculate two instead of one error correction values and distribute them over the drives so that the data and parities are located in blocks on separate drives. This means that data stored on n hard drives requires $n+2$ hard drives in total. This provides a cost advantage compared to RAID 1 (single mirroring). However, its disadvantage compared to RAID 5 in particular is that the computing power required of the XOR processors used in this system needs to be much higher. In RAID 5, the data from a data row are added together to create a parity bit (and, if re-synchronization is required, the data from a data row are re-synchronized via addition).

In RAID 6 systems, on the other hand, the parity bit needs to be calculated over several data rows – re-synchronization then requires complex computational models involving matrices and inverse matrices from linear algebra (coding theory), in particular if two hard drives fail. A RAID 6 array requires at least four hard drives. RAID 6 is hardly ever used today.

CAS - Content Addressed Storage

To complement the discussion of RAID, this section will briefly discuss CAS, a further system used to increase data security.



CAS is a special hard disk-based storage method that allows direct access to individual objects and that is designed to allow for the immutability of the stored information.

In a CAS system, information is accessed within the database according to its content by having an upstream computer calculate a hash value which is used as both a storage and access key. This makes CAS systems considerably slower than pure RAID systems and closed, proprietary systems, which generate large dependencies particularly in the case of long-term data storage. CAS is normally used for storing and accessing statistical data, i.e. primarily for fixed content, such as business documents, patient data, etc.

The first commercially available CAS system was marketed by the company EMC on its Centera platform, and is now typical of a CAS solution. The goal of the system was to store fixed content on fast hard drives.

4.5 Tape – technology and media

Tape has been declared to be extinct many times, but this technology is particularly illustrative of the fact that reports of the death of a technology are often greatly exaggerated. Due in large part to the unforeseen and constantly growing flood of information, coupled with new technological possibilities, providers have made extraordinary efforts to further the development of tape as a storage medium. The considerable improvements achieved in the performance of tape and tape drives have since led to an ongoing renaissance of tape-based storage solutions.

The strengths of tape-based storage lie in applications where large and very large amounts of data have to be stored and secured in a cost-effective way that uses as little physical space as possible. From a data storage perspective there are basically two main uses of tape-based solutions – backup and large-scale data storage. Depending on the data volume and the application environment, either single drives or autochangers and libraries are used. The storage capacity of the latter can in effect be expanded limitlessly, i.e. into capacities of hundreds of terabytes or petabytes.

The technology has its roots in magnetic tape, also referred to for audio storage as cassette tape. It consists of a long and narrow plastic film coated with a magnetic (single or double-sided) material. The tape and the reels are usually enclosed in cassettes. The main types of data storage format as reel-to-reel, single-reel cassette and double-reel cassette.

As is the case for all magnetic recording methods, a magnetic medium (in this case the tape) is fed under a magnetic head, causing the magnetic particles in the tape to align in the direction of the current. The current alternates according to the information to be recorded, so that each change of direction stands for a “0” or a “1”. This presents the first shortcoming of tape-based systems: to read or write data they must always be in contact with the read/write head, which leads to increased wear and tear, and a shorter lifespan of the media themselves. In some solutions this problem is alleviated by leaving an air cushion between the tape and the magnetic heads, which reduces wear on the tape (as in the IBM 3590 tape).

In principle there are two methods of writing data onto the magnetic tape: start-stop mode and streaming mode, the latter being the more modern solution. The data are transferred to the recording head (from a cache memory, for instance,) and are continuously written to the tape as long as data flow is ensured, allowing high write transfer rates to be reached. However, the tape can be stopped to record file marks.



A selection of tape formats

Different type widths are used depending on the particular technology and the specific requirements involved: ¼ inch, ½ inch, 4 mm, 8 mm.

Technology	Sample version	Manufacturer	Format
3590	3590H	IBM	½"
9840	9840C	StorageTek	½"
9940	9940B	StorageTek	½"
ADR	ADR ²	OnStream	8 mm
AIT	AIT-3	Sony	8 mm
DAT-72	DAT-72	HP, Seagate	4 mm
DDS	DDS-4	HP, Seagate, Sony	4 mm
DLT	DLT-8000	Quantum	½"
DTF	DTV-2	Sony	½"
LTO	LTO-4	HP, IBM, Seagate	½"
Magstar	3570	IBM	½"
Mammoth	M2	Exabyte	8 mm
S-AIT	S-AIT-1	Sony	½"
S-DLT	S-DLT 320	Quantum (formatted)	½"
SLR	SLR-100	Tandberg Data	¼"
VS	VS-160	Quantum	½"
VXA	VXA-2	Sony	8 mm

Tab. 4.5.1: Tape technology and manufacturers

DAT – Digital Audio Tape: There are two different variants of this tape format, **S-DAT** and **R-DAT**, although only the latter established itself on the market. As a result of the self-centered interests of international organizations (particularly in terms of the influence of the International Federation of the Phonographic Industry), the consumer market was left out in the cold, and DAT has primarily only been used in the professional market due to its high reliability. Hewlett-Packard, for example, has used DAT as the basis for the DDS backup format.

DLT – Digital Linear Tape and S-DLT: DLT was developed by the former Digital Equipment Corporation as a storage medium for backup purposes. In principle, its demise as a storage format has already been proclaimed (Quantum: "DLT technology has reached the end of its lifecycle."), yet it is still alive. It is capable of high transfer rates. DLT libraries are often used for the backup of large data volumes in high-end networks. Table 4.5.2 lists some DLT variants.

Medium	Tape capacity	Transfer rate
DLT Type IV (black)	20 / 35 / 40 GB	2 / 5 / 6 MB/s
DLT-4000	raw/compressed: 20 / 40 GB	Raw/compressed: 1.5 / 3 MB/s
DLT-7000	35 / 70 GB	5 / 10 MB/s
DLT-8000	40 / 80 GB	6 / 12 MB/s
S-DLT 220 (green)	110 / 120 GB	11 / 22 MB/s
S-DLT 320	160 / 320 GB	16 / 32 MB/s
S-DLT 600	300 / 600 GB	32 / 64 MB/s
S-DLT 1200	600 GB / 1.2 TB	40-80 / 80-160 MB/s
S-DLT 2400	1.2 / 2.4 TB	100 / 200 MB/s

Tab. 4.5.2: DLT variants

Laser technology with an optical servo-track on the reverse of the tape allows the read/write head to be positioned very precisely on one of up to 448 tracks (S-DLT 320).



Today tapes with capacities up to 1.2 TB are already under development, with transfer rates of up to 128 MB/s. These developments are set to become reality within the next 3-4 years.

AIT – Advanced Intelligent Tape / S-AIT: This technology, which was developed exclusively for backup purposes, is the successor to DAT. The format uses tapes with a metallic coating (Advanced Metal Evaporated), which allows storage densities of up to four times higher than those of DAT.

Medium	Tape capacity	Transfer rate
AIT-4 (available since 2004)	200 GB	24 MB/s
AIT-5 (available since 2005)	400 GB	48 MB/s
AIT-6 (scheduled for 2008)	800 GB	96 MB/s
S-AIT-2 (available)	1 TB	60 MB/s
S-AIT-3 (scheduled for 2008)	2 TB	120 MB/s
S-AIT 4 (scheduled for 2010)	4 TB	240 MB/s

Tab. 4.5.3: AIT variants

LTO – Linear Tape Open: LTO, initiated by IBM, HP and Seagate, is a standard for magnetic tapes and the accompanying drives. LTO is an open standard (the only tape standard that is not controlled by a single provider), allowing the media, drives, auto-changers and libraries of a range of manufacturers to be compatible with one another. LTO is able write data to one generation prior to its own and two generations subsequent to its own (see also Table 4.5.4).

Initially two sub-variants were planned and developed: LTO Ultrium and LTO Accelis. However, the latter variant was abandoned before it was launched onto the market. The Ultrium tapes were intended for backup purposes, whereas the Accelis tapes were designed for archiving, which is why the latter were planned to have considerably shorter access times.

There are three LTO Ultrium variants on the market. Similarly to AIT tapes, they are fitted with a memory chip (4 KB). This chip is used to store special data such as the tape serial number and the user log files of the last 100 tape mounts. The following list gives an overview of the different LTO Ultrium variants:

LTO variant	Capacity		Transfer MB/s	Compatibility		Status
	Native	Compressed		Read	Write	
LTO-1	100 GB	200 GB	7.5	LTO-1	LTO-1	available
LTO-2	200 GB	400 GB	35	LTO-1+2	LTO-1+2	available
LTO-3	400 GB	800 GB	80	LTO-1-3	LTO-2+3	available
LTO-4	800 GB	1,600 GB	120	LTO-2-4	LTO-3 +4	available

Tab. 4.5.4: LTO variants

The development of LTO was particularly designed to allow for automated backup. This was taken into account by constructing the tapes in a wedge-shaped design and inserting special notches to enable them to be grasped more easily by robots. LTO libraries are available for large and very large amounts of data, even up to the petabyte level (1 petabyte = 10^{15} bytes = 10^3 TB). Libraries can be connected to the host computer via SCSI or fibre channels.

Considerable progress has also been achieved in the data transfer rate of LTO systems. For example, HP announced the launch of its StorageWorks LTO-4 Ultrium 1840 model for August 2007. According to the manufacturer, the model has transfer rates of 240 MByte/s for compressed (and 120 MByte/s for uncompressed) data. In practice this would permit up to 1.6 TBytes of compressed (800 GBytes of uncompressed) data



to be stored on a cassette in two hours. If these specifications are confirmed in practice, this would be the fastest LTO drive on the market (until the competitors catch up, that is.).

4.6 Hybrid solutions (combined applications)

There are many practical instances in which the advantages of optical jukeboxes or tape libraries are combined with those of RAID systems. The optical component guarantees audit-compliance and the long-term availability of the data, while the RAID system integrated into the other solution ensures very rapid access to specific data. This combined storage package also often assumes a real-time backup function using the cache and management software.

The fact that a considerable number of jukeboxes equipped with discs are also used as pure backup systems as an alternative to tape libraries has been met with astonishment in the ongoing debate within hard disk industry. Yet it should not be surprising that applications of this kind are often found in small and medium-sized companies, as in these companies issues such as cost-effectiveness (including energy efficiency) and the lifespan of solutions have always been a key deciding factor.




Fig. 4.6.1: Backup technologies: a survey of small and medium-sized companies
 [Source: Fleishman-Hillard / speicherguide.de 2007]



5. Background information

5.1 Compliance

In recent years the long-term storage of many types of documents has become subject to new and increasingly encompassing directives and legal regulations. The globalization of many markets also makes it necessary to take account of international legislation. This combination of factors leads to a considerable number of requirements that long-term archiving solutions must fulfill. The following diagram illustrates the wide variety of legal regulations in international markets.

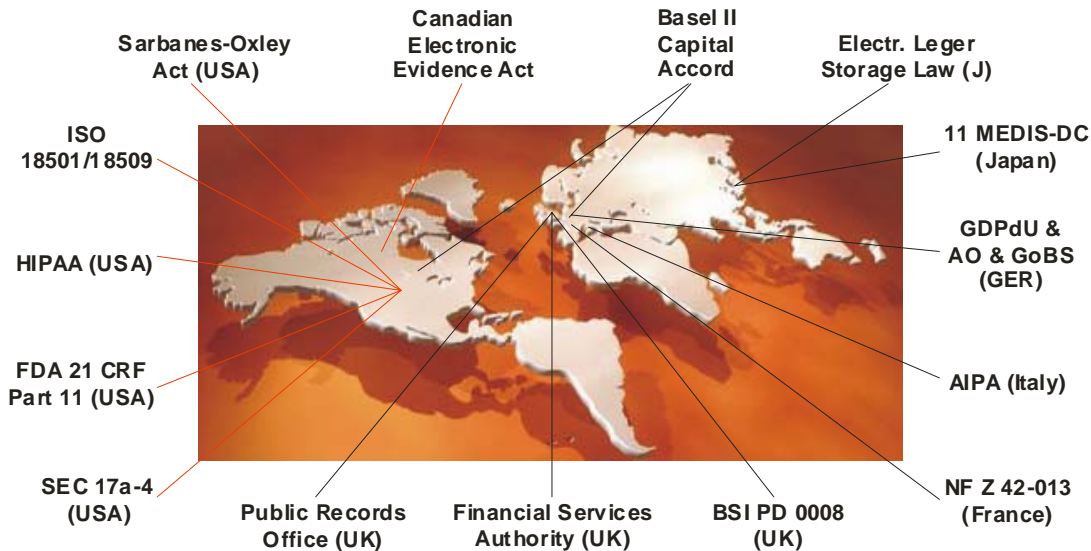


Fig. 5.1.1: Legal regulations on compliance (selection only)

Important: Legislation does not stipulate a specific technology as such, but “only” that due diligence must be undertaken to comply with the above regulations.

5.2 Total Cost of Ownership - TCO

Although the economic consequences of installing IT applications have often been investigated, these consequences have not always been given the same amount of attention in strategic decisions regarding the choice of particular technology. Even today it is frequently the case that companies place the most emphasis on investment costs when selecting IT systems or applications, while failing to give adequate consideration to operating costs, which regularly has major consequences. It is only in the wake of the increasingly heated debate on energy efficiency that some users have realized that there can be much more important cost factors than simply acquisition costs. The keyword is TCO – Total Cost of Ownership over time.

A short general definition of TCO: TCO is the calculation of the total cost of all the direct and indirect costs of an economic unit throughout its entire useful life.

In discussing TCO in the IT sector, all cost-related factors must be taken into consideration, from the costs of choosing and acquiring hardware and software, through customization, maintenance and training costs, to the rental of storage space and energy costs, to mention only the most important factors. Additionally, if TCO is understood in its strictest sense, the additional costs of disposing of components or entire systems at their end of their useful life must also be taken into account. A thoroughly conducted TCO analysis provides the best basis for comparing a range of solution packages for particular IT problems before making a final decision. During the lifespan of systems or applications it is also very helpful to document individual factors, particularly for those IT components or systems that influence energy consumption.



Many mistakenly believe that optical storage solutions are always expensive and complicated. Regardless of how this misconception originally came onto the market, a neutral comparison of potential solutions is always the most appropriate way to determine how systems operate in practice, while of course bearing in mind the specific requirements of each scenario.

The Enterprise Strategy Group (ESG) published a cost analysis based on a concrete case study in the financial services sector requiring storage capacity for an archive of 12 TB. The daily increase in data to account for was 8 GB, with 2,500 archive queries per day (document access).

This comparative analysis is based on an operating period of three years, taking all costs from list prices. Both the investment and operating costs were taken into account. The operating costs include software and hardware maintenance, as well as the direct and indirect energy costs. Additionally, the costs of renting storage space for housing the storage solution were also taken into consideration.

The following diagram shows a comparison of costs in the scenario described above for three hard disk-based solutions, four optical solutions (Blu-ray, two UDO systems, DVD) and for a tape-based solution (LTO-3).

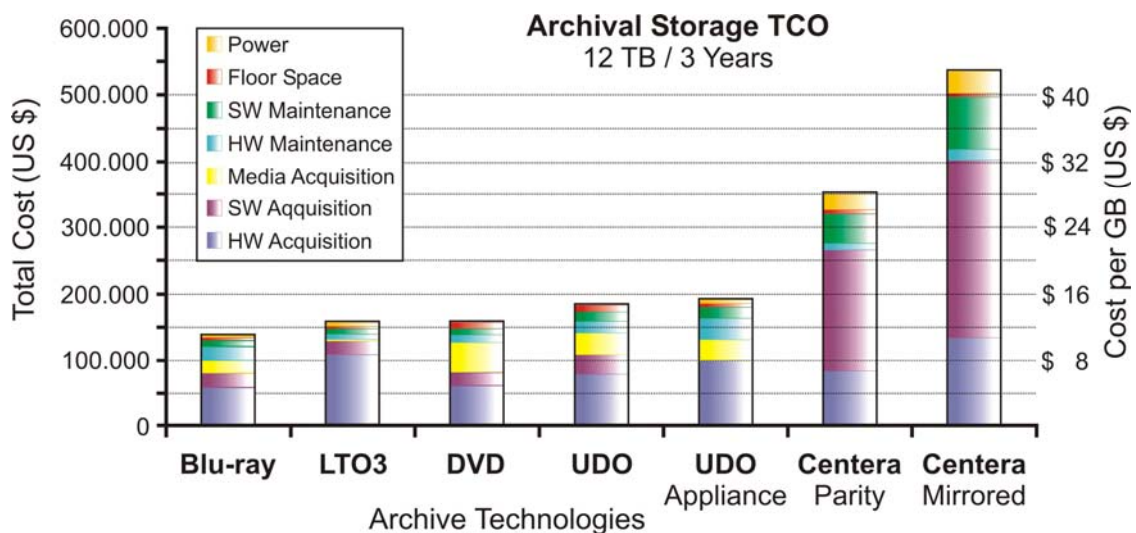


Fig. 5.2.1.: TCO comparison of costs for selected technologies with identical storage capacity (12 TB calculated over three years) [Concept and data: ESG 2006 / Blu-ray data: K.E. 2008]

Conclusion 1 of this ESG case study analysis (additional Blu-ray figures from the author's own database) is that, in this particular scenario, the overall costs of a hard disk-based EMC Centera archive solution were around 295% higher than those for one optical storage solution (the Blu-ray-based solution), approx. 245% higher those for a DVD-based solution, and almost 200% higher than a UDO solution.

Conclusion 2: Always compare like with like. In this particular case, the optical storage solution is without doubt the most economical, and thus the right corporate strategy. A customer should not even begin to attempt to find a hard disk-based solution that fulfils all possible requirements in this given scenario. This is the point of the comparative study documented here: there are appropriate and economically viable applications for every technology. No provider of a specific storage technology should allow themselves to ignore this fact or to sweep it under the carpet, strictly for marketing considerations. This is true for providers of hard disks and of optical discs or tape. A different set of requirements, such as online storage with high access rates, would have led to a hard disk-based solution being selected, even given such a large amount



of data. The study shows how important it is not to purchase off-the-rack storage solutions and not make decisions based on the all-in-one solutions providers may propagate. Instead, it is far better to concentrate solely on the concrete requirements of the application at hand. This is what this White Paper aims to emphasize, with a view to ensuring that the wide-ranging application possibilities of optical storage solutions, and to a degree those of tape-based solutions, are given the consideration they deserve. If the properties of Blu-ray and UDO technology are taken into account, with their enormous storage capacity and high-performance access times, it becomes clear that optical solutions are the ideal answer to many requirements in the fields of long-term archiving, DMS, and ECM, etc., as Figure 5.2.1 clearly shows.

5.3 Energy efficiency: computer centres / storage solutions (“green grid”)

For many years the impact of computer centers (CCs) and data storage on energy efficiency was hardly taken into account, or even taken seriously at all; however, since 2006, this wasteful attitude has begun to change. As is so often the case, the companies involved have not always actually looked for a solution to this serious problem. Instead they have announced in their advertising campaigns that they have recognized the problem or even have it under control. The practical steps taken by some providers have simply been limited to proclaiming that they have developed particularly energy-efficient solutions. Yet often there is a lack of substance to these claims, which are frequently based on inaccurate comparisons.

Some providers of hard disk-intensive solutions, for instance, have advertised their new hard disk models as being particularly energy-efficient, yet this claim rarely says anything about the absolute energy consumption of storage solutions or computer centers full of hard disks. For simplicity's sake, companies have tended to choose the previous model as the benchmark of efficiency, allowing them to appear in a favorable light, when in fact their results were rather modest. Currently it has become fashionable to compare the individual components of a computing set-up, as this allows auxiliary IT power units and auxiliary components (NCPI = network critical physical infrastructure), which consume large amounts of power, to be left out of the equation. However, a storage solution, not to mention an entire computer center, does not merely consist of one hard disk or one drive. This is of course also true of optical and tape-based storage solutions, yet there are considerable differences depending on the features of each system.

However, energy efficiency is not exclusively of importance for the professional IT market. Particularly in the consumer market, two trends have come together to make energy efficiency, or more precisely the lack thereof, more tangible. More and more PCs have reached the performance category of heating appliances, mains adaptors with 850 watts of power consumption have become fashionable, and flat rates for the internet have caused many users to keep their computers running around the clock. As an analogy, nobody would ever think of having their electric heater on 24 hours a day.

Yet a real awareness of the issue seems to have established itself. Although the USA (and China) is the world champion in wasting energy, the IT sector there clearly sat up and took notice of the mood of the times earlier and more consistently than in Europe. For example, the US government's Environmental Protection Agency (EPA) has conducted a detailed study of the issue of computer energy consumption, and has proposed immediate counter measures in view of the worrying results of the study.

The Environmental Protection Agency has estimated that in the year 2006, the servers and storage systems installed in American computer centers, including peripherals (thus including the required cooling equipment), consumed around 61 billion kilowatt hours (61×10^9 kWh, i.e. 61 million MWh!) of electricity. According to the EPA, this figure is twice as high as in the year 2000 and, if the data is correct, makes up a sizeable 1.5



percent of national electricity consumption. Unless huge countermeasures are taken, the EPA has predicted that computer centers will require 100 million MWh by 2010.

Note: The author considers these figures to be somewhat high, as 61 million MWh would be equivalent to around 8,000 MW of installed electrical power, thus requiring eight large-scale nuclear power stations (or other types) for computer centers and PC use in the USA. Yet the American Power Conversion Corp (APC), a leading company in the field of electricity supplies and cooling equipment, has estimated the energy consumption of all computer centers worldwide to be 40 billion kWh (not including the private PC market), i.e. 5,000 MW of electrical power, which is still an alarmingly high figure.

Regardless of the issue of the actual energy consumption of the world's computer installations, it is beyond doubt that the problem has taken on almost unimaginable proportions, especially since the market segment in question is undergoing the world's fastest growth rates. The issue of energy efficiency is thus not a laughing matter or a trivial issue in the IT industry, but one that is relevant for the global climate, on a scale that represents a considerable percentage of the world's energy needs.

It is therefore no wonder that an increasing number of providers of high energy consumption IT components and solutions are developing good arguments in their own defense in the face of increasing environmental awareness among potential customers.




Fig.5.3.1: Important computer centre-related issues from the perspective of CIOs [Source: IDC]

In any case, it is true for both providers and users alike that that company-wide energy efficiency is a strategic field and thus a matter for senior management, and that IT and areas related to it must also be included in energy policy.

This fact is being acknowledged by a growing number of companies. The "EMEA Power and Cooling Study 2006" conducted by IDC showed that chief information officers (CIOs) see energy supply and the actual amount of energy used in computer centers as the most important issues (see also the diagram on the left).

Storage solutions have a major role to play within the debate on energy policy in the IT sector, as the continual growth of data that needs to be stored leads to increasingly high-consumption solutions and systems being installed.

Although other sections of this White Paper discuss the advantages, disadvantages and application-specific idiosyncrasies of various storage solutions from a technical, task-related and economic perspective, the topic of energy efficiency has not yet been included as a parameter for differentiating between potential storage solutions. This field is still in its infancy in terms of the universal comparability of the measured parameters and the methods used.

For this reason this section will only give a mainly qualitative and compact overview of fundamental issues, with some suggestions as to possible approaches. Nevertheless, some substantial remarks can be made on the various storage technologies (hard disk, optical discs such as DVD, Blu-ray, UDO and tapes) that are of relevance for concrete



practical decisions. This topic is given particular attention and will be dealt with in more detail in the next version of this White Paper.

There are as yet no standardized methods for directly comparing the energy efficiency of different IT systems that fulfill similar or identical tasks. However, it is important to note that at least some thought is being given to this issue, which has occasionally resulted in concrete action being taken.

As the term energy efficiency implies, the question at the heart of the debate is how efficiently the energy made available for IT-based components or systems is used. To answer this question, a precise knowledge of all types of energy systems is required.

To illustrate this more clearly, we can consider the example of a single IT component, the processor inside a computer. In order for the processor to be able to fulfill the tasks for which it is designed, such as conducting mathematical operations, it has to be supplied with power. The power is supplied in the form of electricity, the amount of which (measured in watts) depends on the technical specifications of the processor, i.e. its performance specifications. All users know that processors generate heat when they are in use. No one needs this heat, which is why it can be described as “lost heat” from a pure IT perspective.

Thus it is always important when discussing the energy efficiency of a physical unit to know beforehand how exactly this unit is used. The familiar example of a light bulb is particularly illustrative. A traditional light bulb (with filaments) is normally used as a light source, but from an energy-efficiency perspective it is incredibly ill suited for its purpose. Depending on the exact technology used, only 5% of the energy supplied is put to the intended use, namely generating light. The lion's share of around 95% of the energy is transformed into lost heat. The energy efficiency of a light bulb of this kind in its originally intended field is thus shockingly poor, and only merits an efficiency quotient of $\eta = 0.05$. If used as a heat source, the light bulb would achieve a respectable efficiency quotient of $\eta = 0.95$ (not taking too strict a view of the physics involved). This example shows how important it is to compare like with like, and not to compare chalk (light) with cheese (heat).

In the field of IT, this means that, prior to any comparison or study of the cost-efficiency of a particular solution, it is essential to know precisely which IT-related tasks a component, system or an entire solution package (hardware and software) has to fulfill in order to arrive at any conclusions regarding its energy efficiency.

Definition: For any IT system or computer center, the measure of energy efficiency is the ratio of the amount of energy that is supplied exclusively to IT-related components (to fulfill the original IT tasks) to the amount of energy that is supplied to the entire system / computer center.

This can be applied to the example of the processor above: one part of the electricity supplied is transformed into heat, which has to be removed from the processor, as otherwise it would die of “heat exhaustion”. Many PC users know that the IT performance of a computer drops if it is not sufficiently cooled. The less lost heat a processor of a given size generates, the more energy efficient it is. This is why manufacturers have been making great efforts to develop increasingly efficient processors, which is particularly important in all mobile devices, such as notebooks. The battery capacity of many notebooks has not managed to keep pace with the energy consumption of the built-in components, which leads to shorter battery life.

Yet in a notebook it is not only the processor that transfers electricity into (lost) heat. Many components also contribute. This is bad enough from an energy efficiency perspective. It becomes even worse when one considers the efforts that have to be made to remove the heat through the installation and operation of cooling systems for com-



puter installations. Thus it should be no surprise that cooling and temperature control are the main causes of poor energy efficiency for practically all components, systems or computer centers.

In computer centers the biggest “energy guzzlers” (in non-scientific terms) are usually the cooling and temperature control units. But attention should also be paid to lighting, electrical security systems or switching systems. In short, everything in the computer center that serves to maintain operations and needs to be supplied with energy must be considered.

As the main cooling systems in computer centers use air as a coolant, it is virtually impossible to economically reuse the heat generated. This is in stark contrast to the direct water cooling of processors, which is envisaged to be used in the computer centers of the future. Via heat exchangers with acceptable levels of energy efficiency, this technology will allow heat energy to be reused and is set to become established in practice.

Fig. 5.3.2 below gives an overview of the key energy paths in a computer center. The scenario depicted below can also be seen as typical of other physical systems. If this scheme is used for the range of storage technologies available, the enormous differences in energy consumption immediately become evident.

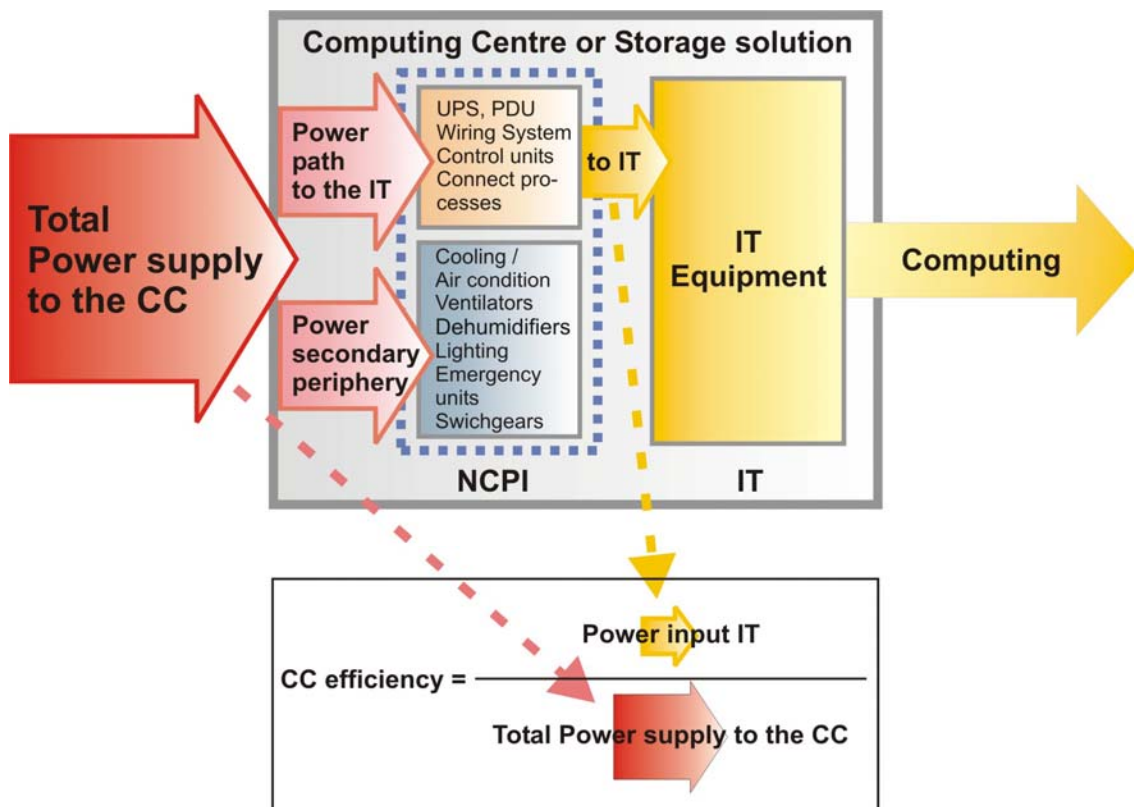


Fig. 5.3.2: Energy paths for a computer centre and efficiency quotient [Source: APC & K.E. 2007]
(Legend: NCPI = data center's networkcritical physical infrastructure, UPS = uninterruptible power sources)

If we take a complex storage solution as an example – such as a collection of data that has to be archived for a long period of time, i.e. for well over five years – with a suitably large number of servers, and compare the energy needs of this solution, involving hard disks that are in continuous operation (24 x 365 hours), with a storage solution based on optical discs, it does not take a mathematical genius to realize the enormous difference in their energy requirements.

A further concrete example is given here. The following table illustrates the situation for a computer installation that was primarily designed for a complex archiving task:



Auxiliary components	[%]
Cooling (ventilators etc.)	35
Dehumidifiers	3
Cold air compressor	9
Junction box	5
Uninterruptible power sources	18
Lighting	1
Switches/generator	1
Total loss	72

Tab.: 5.3.1: Energy requirements of auxiliary components (example only)

The figures:

In this case 72% of the energy provided is required for the auxiliary components alone. This means that a mere 28% of the total energy supply is available for the original task, namely to allow unlimited access to the digital archive and a few other IT tasks, which is why the energy efficiency here is only 28%. This is not a satisfactory result, yet it is not a particularly uncommon one either.

An energy efficiency quotient of 0.28 represents practical scenarios within the IT sector that are by no means a rare occurrence.

Even more alarming examples from everyday practice are well known. Recent research has shown, for instance, that the power required for computer centers has increased more than ten times over the last decade. In some cases more than *three quarters* of the total energy consumption is used solely for cooling the computer center. Thus energy efficiency is becoming an increasingly important issue for IT managers. This is not mainly due to environmental protection considerations, but to energy costs and not least of all due to the risk of failure of many important IT components (predominantly hard disks), which increases dramatically as temperatures rise.

The current relevance of the issue of IT energy efficiency for both providers and users can be seen in the wide range of serious initiatives that have been launched in business and politics, although in Europe the pace of developments leaves much to be desired.

This section discusses two examples (A and B) of high-profile initiatives by IT providers who have paid serious attention to the issue.

A. The Green Grid: This is the name of a newly founded global (non-profit) industry consortium that aims to improve energy efficiency in computer centers and business computing ecosystems. Immediately after its launch The Green Grid published three White Papers that were drafted by the initiative's technical committee. The papers are directed at CIOs, computer center administrators and facility managers, and discuss energy efficiency with recommendations as to how it can be improved in computer centers.

The consortium's board of directors includes representatives of AMD, APC, Dell, HP, IBM, Intel, Microsoft, Rackable Systems, SprayCool, Sun Microsystems, and VMware. The full list of active members is longer and a selection of which is provided below:

1E, 365 Main, Active Power, Affiniti, Aperture, Azul Systems, BT plc, Brocade Communications, Chatsworth Products, Inc., Cherokee International, Cisco, Cold Watt Inc., COPAN Systems, Digital Realty Trust, Eaton, Force10 Networks, Netezza, Juniper Networks, Pillar Data Systems, Panduit Corp., QLogic, Rackspace Managed Hosting, SGI, SatCon Stationary Power Systems, Texas Instruments, The 451 Group, Vossell Solution and Novell.




Fig. 5.3.3: Energy currents (purely qualitative)
[Source: Green Grid]



In view of the combined power of The Green Grid it is tempting to expect great things from this initiative, given the seriousness of the issue. These companies will work together on new practical metrics, methods and best practices for the management of computer centers throughout the world, and will give interested providers and users long-term strategies for energy efficiency in the field of IT.

- B. Climate Savers Computing Initiative - CSCI:** This initiative is composed of leading companies from the IT and energy sectors, as well as environmental organizations, and is headed by Intel and Google. Their goal is to help energy-saving technologies to become established on the market and thus to contribute to climate protection by reducing greenhouse gases. The concrete objectives of the CSCI clearly distinguish it from many other activities and initiatives, a fact that the CSCI emphasizes. It intends to raise IT-relevant energy efficiency to 90% (an efficiency quotient of 0.9) by 2010, which is indeed a very ambitious aim. According to calculations thus far, this initiative will result in a reduction of 54 million tons of greenhouse gases, thus saving approximately US\$5.5 billion in energy costs.

In view of the core and peripheral technologies used in comprehensive storage solutions today, it is easy to determine the potential for saving energy discussed above, yet the time span until that objective is reached may turn out to be too ambitious. There are many computer centers and smaller computing installations that could easily compete with thermal power stations in terms of the megawatts they produce. In order to change this, other partners have joined the two organizations mentioned above, such as AMD, Dell, EDS, HP, IBM, Lenovo or Microsoft, and, importantly, the WWF. This environmental organization, with its climate savers program, is regarded as the actual instigator of the CSCI. In addition to energy providers such as Pacific Gas, client companies such as Starbucks have also joined the initiative. The group members have obliged themselves both to build energy-efficient technologies, if they are manufacturers themselves, and also to implement energy saving strategies within their own companies.

The fundamental aims of the CSCI are modeled on the “Energy Star” guidelines set down by the EPA, yet go beyond these in some aspects. For example, the EPA has an energy efficiency target of 80 percent, whereas the CSCI initiative aims to reach 90 percent by 2010, as mentioned above, and even 92 percent in the server sector.

Outlook

Despite the numerous efforts made to implement energy-efficient IT and storage solutions, it is still true that the growth rate for installations of new computer capacity will remain considerably higher than that of progress in energy-saving technology for quite a while to come. This is due in part to the widespread practice of continually maximizing the computing power of newly installed computer systems, which in practice means setting requirements too high, even if the utilization ratio of the installed computing capacity only rarely reaches anything like cost-effective dimensions (advantage that has taken on extreme proportions in the home computing market, but is also indisputably the case in the professional sector too).



Case study 1 (IBM)




Fig. 5.3.4: Disk vs. tape – electricity and cooling costs calculated over 6 years [IBM]

The above figure shows the energy consumption of disc arrays compared to a tape-based solution. The course of the curve for the tape costs is a little higher than would actually be the case here for illustrative purposes – based on the scale for the figure, the tape line would normally run in a strictly horizontal fashion.

The cost relations illustrated here also hold true for a hybrid solution using optical discs, as, like tapes, these only operate when the data is actually accessed (the hard disks have to be kept rotating 24 hours a day, 365 days a year, which leads back to the energy consumption discussed above).

Case study 2 (IBM)

In 2006 IBM installed **3,116 PB of tape storage**

- ⇒ The equivalent disk-based solution would require 432 MW for electricity and cooling.
- ⇒ A nuclear power plant generates around 1,000 MWeI of power.
- ⇒ To replace the above tape capacity by hard disks would require a new nuclear power station approximately every two years.
- ⇒ As IBM does not install nuclear power plants, they will stay with their tape-based solution.

Note: For individual cases, i.e. for individual storage solutions, a systematic analysis of the energy requirements for the installed optical disc capacity would reach a similar result.

5.4 Storage technologies comparison

Hard disks and tape compared to optical devices

In the past it was always the case that any comparison of the costs of using hard disks or tape for backup purposes would come out clearly in favor of a tape-based media. Yet with the advent of the latest SATA drives, a cost comparison must be undertaken in more detail. This also entails comparing more than just the capacity of the individual media, as major progress has also been achieved in across all storage technologies, and there is no end in sight. While a simple comparison of different types of media using a cost per GB analysis will still favors tape-based solutions. However, a comparison of this kind is misleading in that the price of hard disks includes the price of the drive, and tape drives require additional technology to be used effectively. This means that in order to correctly contrast hard disk and tape-based solutions, it is more accurate to compare a RAID system with a tape autoloader or library.



In all cases, it is essential to compare possible solutions with a more comprehensive perspective. For example, a SATA RAID system can be more cost-effective than an LTO system for at a specific capacity point while LTO may be better solution for another capacity. If the system has to handle very large volumes of data in an archiving application, it may be necessary to remove volumes from a library. In this case, a tape-based solution may be more cost-effective (albeit with delayed in access time).

In order to obtain specific figures for the costs of different technological solutions and media that reflect actual practice, it is not sufficient to only consider pure investment costs and to calculate similar capacity figures. A comprehensive approach must be taken that also includes factors such as energy efficiency. The complete operating costs must also be adjusted to reflect the media being used in order to calculate accurate figures. The next version of this White Paper will explore these specific cost considerations in more detail.

5.5 Lifespan of storage media

Before examining the lifespan of individual media or of the entire storage solution, the nature of the specific storage media must be considered in the context of the total environment.

5.5.1 Retention periods and access frequency

Data are subject to legal and policy requirements regarding how long they must remain accessible. Particularly the legal requirements for retention periods and the type of storage has become more prescriptive in recent years, in part, as a result of the business globalization (see also 5.1: Compliance).

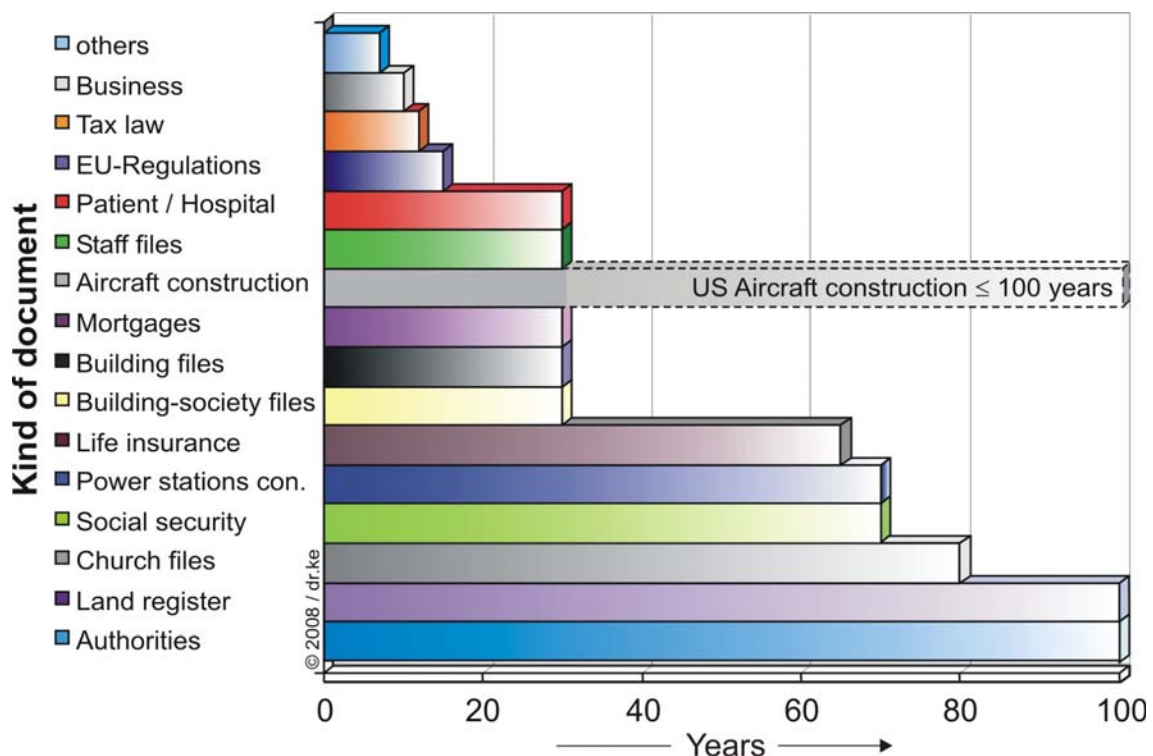


Fig. 5.5.1.1: Practical/legal retention periods for documents [Dr. K.E.]

The selection of business sectors and industries depicted in Figure 5.5.1.1 gives an overview of both legally binding retention periods for documents as well as ones that have become established as best practice over many years.



Whether retention periods are based on practical or legal requirements, they cannot be examined in isolation when considering which technology should be used for long-term archiving. The pragmatic considerations must include the fact that applied retention periods often considerably exceed those stipulated by law. The situation is equivalent to opening a deposit account for a child when it is born and only paying in a sum once, despite the fact that humans can live to the age of 100 or more.

When considering the right choice of storage solution it is important to consider the expected or required lifespan and type of data, the frequency of access throughout the retention period. This is important because of the special requirements of long-term archiving. The following figure (a qualitative depiction only) illustrates a typical pattern of access frequency for data in long-term storage:

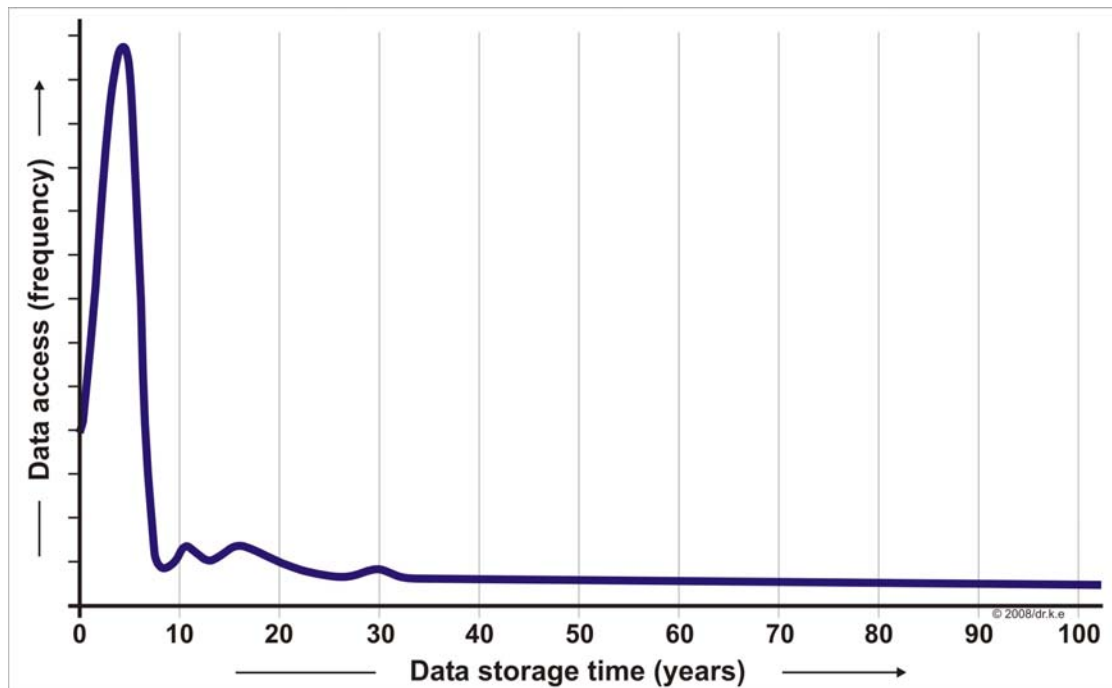


Fig. 5.5.1.2: Access frequency for data in correlation to their lifespan [K.E. 2008]

- Illustration is purely qualitative for a mixture of applications -

It is frequently the case that the migration of data for long-term storage is carried out on a yearly basis, regardless of all cost and security considerations. In such cases, the lifespan of the solution and the media becomes irrelevant and a technology choice can be made without detailed review. Yet in order to find a strictly economical and secure solution to the problem, it is necessary to examine both the storage requirements for the data and the expected frequency of access to the data during the different phases of its lifespan. Not to do so represents seriously irresponsible behavior that can put a company's reputation and business at risk.

5.5.2 Media: error sources and lifespan

Nothing that is man-made lasts for ever, and this holds true for all storage media available on the market. This is particularly true for media that is used widely in everyday practice. That said there are considerable differences in the lifespan of different storage media. The trigrams used in ancient times and other storage formats of former millennia are good examples. However, for modern storage problems the advantage of longevity offered by some ancient media cannot be put to any practical use. The one exception could be paper, which can still justifiably be used in very specific applications for particularly long retention periods. But modern applications in require considerable



effort to ensure the preservation of paper records. The irretrievable damage that is caused when historical libraries burn down is just one example of how valuable record can be lost forever.

If a storage solution is needed for long-term archiving that calls for rapid (i.e. immediate) and constant access to all data for all users than only a digital solution can meet this demand. The storage systems used are as varied as the individual requirements of particular scenarios. Depending on the needs of the application, the solution may contain hard disk systems, or tapes and optical storage components (hybrid solutions). All the media used in these three technologies have different life spans with significant variation.

The life span of hard disk and of most tape-based solutions is only measured in terms of several years. Optical discs and some specialized tapes have far longer life spans. Life span naturally does not refer to vague notions of how long a hard disk will continue to rotate or a tape can be loaded in a drive, but to the length of time a manufacturer guarantees that very precise data and media-specific error frequencies will not be exceeded when the medium is used.

In technical terms life span is measured by the average number of operating hours until a hard disk fails. If the hard disk cannot be repaired, this period is defined as “**Mean Time To Failure**” (MTTF). If it can be repaired, the number of operating hours is defined as “**Mean Time Between Failures**” (MTBF). All these measures of the longevity of medium are purely statistical.

Sources of error in the use of hard disks (selection)

With hard disks, errors can occur both on the media and in the actual disk drive that is directly connected to the media. The following list covers some sources of error that have been observed in practice. Manufacturers of hard disks use both software and hardware methodologies to keep the impact of such errors to a minimum (see also Chapter 4.4):

- ⇒ A very frequent source of error is total failure and disk errors due to heat-related problems. This is a particular risk in systems that have high rotation speeds (see the note in the section on energy efficiency in Chapter 5.3).
- ⇒ Mechanical failures can occur at the read-write head. The dreaded head crash is still a very frequent problem due to environmental influences such as dust, humidity, etc. During operations the read-write head hovers over the disk and is only prevented from touching the disk by an air cushion created by the air circulation formed through the rotation of the disk.
- ⇒ Electrical faults in control systems and inevitable mechanical wear and tear can also lead to failures.
- ⇒ External magnetic fields can irreversibly destroy the sectors on the hard disk. Deletion of data by a magnetic field makes more modern hard disks unusable.
- ⇒ Long periods of inactivity can also make the disk drive’s lubricants become more viscous. This can cause the mechanics to get stuck or to become so stiff that the hard disk cannot even begin to rotate (so-called “sticky disk”). However, modern lubricants are now so good that errors of this kind only occur very rarely.

In conclusion, it can be said that it is impossible to make accurate predictions for the life span of a hard disk. It remains a theoretical problem, which is why MTTF and MTBF are only statistical measures. The actual life span only becomes apparent when a hard disk crashes.



In addition to the sources of error listed above, there are other factors that make it difficult to give precise estimates of the life span of a hard disk:

- ⇒ Vibrations from external sources and impacts can lead to increased wear and tear and thus reduce the life span of the mechanics of the system.
- ⇒ Frequency of access has a direct influence on the lifespan of a hard disk. As the frequency of the head movements increases, so does the probability of error, especially in mechanical parts.
- ⇒ Energy efficiency remains an important issue. If a hard disk is operated at too high an ambient temperature, overheating problems can quickly arise leading to total failure. Thus the heat-related parameters of the environment have a direct impact on life span. The maximum operating temperatures stipulated by manufacturers are generally seen to be reliable as hard disks rarely encounter overheating problems below these temperatures.
- ⇒ The statistical basis of life span estimates makes it difficult for manufacturers to make precise statements when introducing new models. It is impossible to make inferences from one model of hard disk to another, and thus with every new model manufacturers have to gradually begin to determine the expected life span. A statistically sufficient number of hard disks must to be observed in practice before being able to make reliable claims regarding their longevity.
- ⇒ Hard disks in notebooks are subject to particular mechanical strain due to their transport frequently. Although specially cushioned hard disk mountings are often used, these hard disks have a shorter MTTF than desktop hard disks.

Generally speaking, server hard disks are designed to have higher MTTF than is the case for normal desktop hard disks. Thus it should be possible to expect them to have a corresponding higher lifespan. Yet the practical factors listed above often prevent this from being the case, as these servers are in constant operation and usually have to withstand considerably higher access frequencies. This leads to mechanical wear and tear (at the read/write head), affecting the hard disk system as a whole. Thus it comes as no surprise that the life span of server hard disks only amounts to a few years. There have been many cases in practice where hard disks have been replaced as a precaution to avoid the risk of data loss, even though they could have continued to be used. But ultimately what counts is practical experience and responsible maintenance procedures.

It should come as no surprise that manufacturers cannot give any precise figures for the lifespan of their hard disks, only statistical estimates. Reliable estimates range from 5 – 10 years, depending on the type of use involved even if the disks are replaced at an earlier stage.

Sources of error in the use of tape-based systems

Modern magnetic tapes are technically complex and at the same time extremely sensitive. If they are not used properly, irretrievable data loss or the total failure (destruction) of the medium is the inevitable consequence. Tapes cannot cope with heat or cold, are particularly sensitive to dust, humidity and magnetic fields (even weak ones). There are many places in practice where tapes ought not to be used, yet IT managers are often not aware of the risks.

The location for a tape library or an autoloader should always be cool, dry and dust-free. Tapes are predominantly used for backup purposes. If data security is to be taken seriously, the tapes must be physically transported and distributed to different locations. During transport they can be easily damaged, leading to complete data loss if not transported in special protective containers. The tapes should be subjected to as little



vibration as possible during transport, and the move should be as rapid as possible to avoid large temperature fluctuations. When in its final storage location the tapes should be placed in an air-conditioned, dust-free safe.

Most errors result from improper use and environmental influences usually combined with a lack of awareness of the sensitivity of tapes and their drives.

Conclusion: Tape-based media should be replaced at regular intervals. These intervals have little to do with the actual lifespan of the media, but primarily dependent on the duty cycles of the media. The key factor is the frequency of use (read/write access patterns). An additional key variable is wear and tear on the tape through stop and go operations. If the data stream is relatively moderate, information is usually not written to the tape at full speed, which is why drives repeatedly stop to wait for additional data. Fortunately replacing the tapes is no longer a matter of gut feeling since some modern tape-based media indicates when it is time to be replaced by measuring wear and tear. Only by replacing tapes on a frequent basis can the required level of data security be guaranteed.

Sources of error in the use of optical discs

Data is written to and read from optical discs without any contact to the storage medium, avoiding any form of mechanical wear and tear. Nevertheless, like hard disk and tape-based systems there are guidelines regarding storage conditions for optical media. This includes both the environment in which it is used and the quality of the media itself. Dust, high humidity, strong vibrations during access and extreme temperature fluctuations should all be avoided. However, the limits for all of these parameters are higher than those for hard disks and tapes due to the robust nature of the recording technology and media construction.

To protect the physical media care should be taken to avoid scratches on the data layer of the disc. However, complex error correction methods are used to allow discs to be read even if they are scratched. The two most important error correction methods are parity bits and interleaving.

Parity bits are information that is added to the data to be able to verify or correct the status of the data. For instance, a parity bit may store information as to whether the sum of the digits of a byte is odd or even. If several parity bits are used that have been generated in different ways it provides redundancy for the system. This enables checks to be made as to whether a byte or a frame has been read correctly. With increasing redundancy it is possible to detect errors and correct them.

Often a special process is used to interconnect the data across several data packets with the aim of ensuring that data packets that belong together are not stored next to each other on the CD (interleaving). When the CD is read, the interleaving is reversed, so that an error such as a scratch is divided up into smaller individual errors that can be rectified by using a suitable number of parity bits. This allows data to be read even if the scratches are several millimetres in length (but not too deep). However, even with this method extreme damage to the surface also leads to data loss. If the optical drives are damaged it normally has not affect the media ensuring that the data is not impacted.

Dirt on the lens and natural ageing caused by negative environmental influences can make the reflected laser beam increasingly asymmetrical, leading to read errors on the drive. However, these so-called tracking errors on the drives can be fully rectified by electronic means. It is important to note that data loss does not occur as a result of these read errors since the media itself is not affected by changes in the optics of the drive system.



Brief overview of media life span

The remarks made regarding the mechanical wear and tear in hard disk systems generally hold true for the **life span of tapes**.

Here too it is justifiable to talk of a life span of 5 to years, as they are subject to extreme mechanical strain, even if this is only the case during access operations and despite the fact that tapes are not in continual operations.

Yet with tape-based solutions it is often pointless to philosophize about the life span of the media, as general security considerations entail that data migration is carried out at much shorter intervals for this medium than its theoretical lifespan might indicate.

Medium	Lifespan / years
	g=guaranteed p=presumed
Acid-free paper	several 100 (g)
PET black and white film	≤ 1000 (p)
Acid paper	70 - 100 (g)
Celluloid film ^{*)}	50 - 70 ^{*)} (g)
Optical Disc BD, UDO	> 50 (p)
Optical Disc DVD	≥ 30 (p)
Magnetic tape	5 - 10 (g)
Special magnetic tape	≤ 30 (p)
Hard disks	5 - 7 (statistical)

^{*)}Manufacturer gives life span estimate of several hundred years

Tab.: 5.5.2.1: Overview of media lifespan

Apart from film, the only digital media that can claim to guarantee life spans at the upper levels of legally required retention periods is optical. The manufacturers of Blu-ray and UDO, for example, guarantee a media life spans of at least 50 years.

Of course it is not a surprise that no one has yet tested the possible limits of the life span for optical discs. However, all new generations of media are subjected to artificial ageing processes in order to obtain better estimates. It is already possible today to make highly confident estimates that BD and UDO have life spans of more than 50 years. Yet on the market it is best to stay on the safe side, which is why for the time being the figure is set at 50 years given the sensitivity and value of the data being archived.

Currently some well-known manufacturers of hard disks have begun a debate within their own industry regarding the limited lifespan of hard disk-based systems. This has led to some confusion on the market. Yet at least this debate has once again made it clear how important it is to replace media and subsystems in good time to ensure data security. Indeed, life spans that are considerably lower than the limits given in Table 5.5.2.1 are under discussion. In view of this, the reader is referred to the quotation in Section 4.4.1 (Remarks on the security of hard disk systems in practice), which has a similar message.

As far as optical media are concerned, recent years have seen enormous leaps in the confidence with which estimates of the lifespan of such systems can be made. Yet it is also worth mentioning that even optical storage solutions make use of drives and libraries (jukeboxes) that involves mechanics and electronics. However, the read/write heads of the drives do not come into contact with the media, thus not subjecting them to wear and tear. Jukeboxes are also designed so robustly that storage solutions based on optical discs outlive all other digital solutions for long-term archiving.

Considerations of particular importance for optical media

In terms of reliability, any storage solution is only as good as the media. Cheap products in the CD and DVD sector have long since flooded the market and even many professional users have not been able to resist the allure of “bargain” prices. Recently there have been an increased number of complaints by these very users, who have found that data stored on this consumer media has either been lost or damaged. Particularly in the case of discs and tapes it is important to repeatedly emphasize that real quality is only really ensured by the highest-quality products. Cheap CDs can give up



the ghost after only one or two years. Unfortunately this has led some analysts in the industry to come to the conclusion that the life span of optical media has previously been over estimated. Yet for professional quality optical technology it can be demonstrated that media life is very long.

6. SUMMARY AND CONCLUDING REMARKS

When faced with the necessity of finding solutions for the storage or archiving of extensive collections of data, it is essential to take a close look at the range of technologies used. Considerations include economic, technical and corporate strategy.

If this analysis of the available technologies and solutions is coupled with precise knowledge of the documents that are to be stored and how “active” the data will be throughout the retention period (i.e. in terms of the expected frequency of access throughout the entire lifecycle of the documents), then the following key insights can be summarized:

- ✓ Hard disk-based solutions have their strengths, particularly in the first stage of document lifecycles, during which access frequency is very high.
- ✓ Tape-based solutions have their strengths, particularly for backup purposes (keyword: hybrid solutions in combination with disk arrays).
- ✓ Optical solutions have their strengths, particularly in the field of long-term archiving (and in special cases can also be used for backup purposes as part of hybrid solutions, such as in applications for small and medium-sized companies).
- ✓ All technologies have their strengths and weakness. A combined solution can bundle strengths and neutralize weaknesses.
- ✓ In consideration of the entire lifecycle of many documents, it is increasingly the case that finding storage solutions for storage management tasks is not a question of deciding on one solution over another, but of finding the most effective combination of solutions.



7. Glossary

Blu-ray Disc	<p>A Blu-ray disc – BD – is an optical storage medium the size of a CD, and with several layers similar to a DVD, yet with considerably higher storage capacity. The Blu-ray disc takes its name from the blue laser light it uses, which has very short wavelengths. This allows the data tracks to be placed more closely together than on a DVD; the pits and lands are also smaller, which allows for a higher storage density. It's most important specifications are the laser wavelength of 405 nm, the numerical aperture of 0.85 and the thickness of the data layer, which is only 0.1 mm.</p> <p>BD, which was developed by a consortium of over 130 members and was supported by the Blu-ray Disc Association, is designed as a one-sided medium with several storage layers. In the one-layer format its storage capacity is 25 GByte, and with two layers 50 GByte. The simple data transfer rate is 9 MByte/s, although drives have been developed that can considerably improve the data transfer rate through higher rotation speeds. Currently 4x drives are available on the market.</p> <p>Companies that have opted for Blu-ray technology include Apple, Dell, Hewlett Packard, Hitachi, LG, Panasonic, Pioneer, Sony, Samsung, TDK, Thomson, and Yamaha.</p>
BD-ROM BD-RE BD-R	<p>BD-ROM: Distribution medium BD-RE: Rewritable storage medium BD-R: write-once storage medium</p>
DAS	Direct Attached Storage: DAS refers to a mass storage medium that is directly connected to a server (without a storage network). They may be individual internal or external hard disks or a disc array consisting of several hard disks.
DDS magnetic tape drive	<p>Digital Data Storage (DDS) is a follow-up format of Digital Audio Tape (DAT) for use in data backup.</p> <p>The ongoing development of DDS has led to several different versions being launched onto the market, not all of which can be played on older DDS drives (incompatibility). The individual DDS versions (DDS-1, DDS-2, DDS-3 and DDS-4) differ in tape length and thus in their storage capacity and data transfer rate.</p> <p>DDS-4 has a storage capacity of 20 GByte in uncompressed form and 40 GByte in compressed form. Its tape length is 125 m, and the data transfer rate is 4 MByte/s.</p>
Disc	An optical or magneto-optical data storage device such as MO / UDO / Blu-ray disc.
Disk	A storage device based on electromagnetism, such as a hard disk, in which the mechanical components are physically connected to the storage device in a casing.
DMS	Document Management System
DVD	Digital Versatile Disc. An optical storage medium with a capacity of 4.7 GByte per layer/side; a double-layer/double-sided disc has a capacity of 9.4 and 18.8 GByte respectively.
ECM	Enterprise Content Management
FC	Fixed Content – data that can no longer be altered.
FCS	Fixed Content Server
HDD	Hard Disk Drive
HD-DVD	High Density DVD , formerly Advanced Optical Disc (AOD) , an optical storage medium that is read and written to using blue laser light. The HD-DVD has a storage capacity of 15 GByte with one layer, and 30 GByte with two layers. The data transfer rate is 4 MByte/s.
HD-TV	High Definition Television (HD-TV). This technology has primarily been supported by Acer, Fuji Photo Film Co., Fujitsu Ltd., Hewlett Packard, Hitachi Maxell, Intel, Kenwood, Lenovo, NEC, Onkyo, Ricoh, Sanyo and Toshiba.
HSM	Hierarchical Storage Management
ILM	Information Lifecycle Management
iSCSI	iSCSI (Internet Small Computer Systems Interface) is a protocol for describing SCSI data transport via TCP/IP networks.
MO / MOD	Magneto Optical Disc. Magneto-optical storage media consist of a thin layer of vertically magnetized material sandwiched between a transparent polycarbonate protective layer and a reflective layer. Storage is carried out magnetically and can only be conducted after the material has been heated by a laser. The data is read by making use of the magneto-optic Kerr effect with a laser beam, which is reflected in different ways by the magnetised and non-magnetized regions. MO is currently being replaced by its purely optical successor UDO.
MTTF	Mean Time To Failure: The average number of operating hours before a hard disk irreparably fails.



MTBF	Mean Time Between Failures: The average number of operating hours before a hard disk fails but can still be repaired.
NAS	Network Attached Storage: Mass storage units connected to the local area network to increase storage capacity. Normally NAS is used to avoid the effort and expenditure required to install and administrate a dedicated file server.
NDAS	Network Direct Attached Storage: NDAS refers to storage media (mainly hard disks) that can be connected directly to a network without the need for a PC or server, and which appear on the target system as a local hard drive.
Optical memory	<p>The strengths of optical storage media lie in their high storage density and storage capacity, long lifespan, high reliability and stability, and low manufacturing costs.</p> <p>Data storage on optical storage media is based on the thermal effect of a laser beam on the plastic substrate on the surface of the storage disc. Some optical storage media are read-only, others can only be written to once, and others are rewritable many times over.</p> <p>The key criterion for optical storage media is storage density. DVDs, for example, which use shorter laser wavelengths and several storage layers, can attain storage capacities up to 17 GByte, which is 27 times that of a CD.</p> <p>Media based on blue laser light can attain storage capacities of 100 GByte and higher.</p>
PDD	The Professional Disc for Data (PDD) is an optical medium developed by Sony that is only used in Sony video appliances.
RAID	Redundant Array of Independent Disks: An arrangement of several hard disks to form a volume (of a logical unit / drive). To compensate for the shortcomings of a single hard disk, where a mechanical defect inevitably leads to data loss, additional parity information is calculated from the data to allow lost data to be reconstructed in the event of the failure of one particular hard disk.
SAN	Storage Area Networks: A separate network in which computers (servers) and storage devices are connected independently of each other – normally via glass fibre optical cables and switches.
SAS	Serial Attached SCSI
UDO	<p>Ultra Density Optical (UDO) is an optical storage technology that is in the process of replacing magneto-optical storage media. The UDO1 technology developed by Plasmon offers a storage capacity of 30 GByte and a data transfer rate of 8 MByte/s on a 130-mm disc format. The UDO2 format that is currently available on the market offers a capacity of 60 GB with a data transfer rate of 12 MByte/s.</p> <p>Similarly to the Blu-ray disc, the UDO technology makes use of short-wavelength blue-violet light (405 nm) and has a numerical aperture (NA) of 0.70 (UDO1) or 0.85 (UDO2). The use of short-wavelength light allows a considerable increase in the density of pits and lands compared to light with longer wavelengths. Its storage density ranges between 1.15 GByte/cm² and 4.65 GByte/cm². UDO discs are available as write-once (WORM) and rewritable variants.</p>
WORM	Write Once Read Many: a medium that can be written to only once but read many times over.
True WORM	Physically determined write-once properties (non-deletable)
Soft WORM	Write-once properties that are determined by software (not determined physically).

Tab. 7.1: Glossary of terms used in this White Paper (selection only)

Original title: „**Daten sicher aufbewahren – Speichertechnologien im Überblick**“


English translation: Döring Sprachendienst GmbH
(Graphik charts: Fachagentur Dr. K. Engelhardt)





Storage technologies – every with individual strength:

- Optical:** Secure long-term data storage
- Tape:** Backup against data loss
- Hard Disk:** Active data and backup



Original title: „Daten sicher aufbewahren – Speichertechnologien im Überblick“
 English translation: Döring Sprachendienst GmbH

© Fachagentur Dr. K. Engelhardt (2006, 2007, 2008)
 Tel.: +49-6373-89.11.33 • Fax: +49-6373-89.11.43
 eMail: ke@dr-k-engelhardt.de
www.dr-k-engelhardt.de