FAST LTA
We secure Petabytes.

Background

Sustainable IT - Long-Term & Energy-Efficient Archiving of Data

Data archiving is exposed to special requirements as „long-term storage“. Compared to the usual investment periods of three years, energy consumption and maintenance take on a higher priority in the archive. IT components and providers must also withstand critical selection. We take a close look at the individual aspects and compare technologies.

February 2022

Long-term data storage

Data archiving is the displacement of data that is no longer or only rarely needed in the productive environment. Although the trend is towards the „active archive“, which can provide data quickly and by random access compared to classic, tape-based offline archives, the main purpose is still: to store data securely and reliably for years and decades. The cost factor also plays an important role. Since archives generally have longer lifetimes than the three to five years generally applicable in IT, factors such as energy consumption, maintenance effort and service level agreements also play a greater role. Investments, on the other hand, are mostly incremental (expansion of capacity); replacement of storage systems is avoided as far as possible due to the constantly growing volumes of data and the necessary data migration.

The topic of sustainability in the archive thus includes the following aspects:

Hardware: Durability and availability, carbon footprint, energy efficiency, expandability, protection against failure

Software: Reliability, certification, compatibility with standards, data migration, security against attacks

Service: Maintenance effort, calculable costs, spare parts availability

Sustainable hardware

The selection of the hardware used is generally strongly geared to the intended use. Since long-term storage is by definition designed for a service life of several years or decades, special requirements must be met here.

Storage security

Since the original data is eliminated during archiving - in contrast to backup, where data is copied and deleted again after a certain time - data loss in the archive is final. Data must therefore be stored in such a way that data loss can be effectively ruled out over the projected period of use. These measures include protection against the failure of individual data carriers (for tape: copies, for hard disks: redundancy), protection against the failure of complete systems (geo-redundancy, outsourcing of data carriers if possible), and protection against unintentional deletion or overwriting (WORM).

These criteria cannot usually be met by standard IT and require special archive storage.An important aspect, especially when using hard disks, is the prevention of correlated failures („epidemic failure“). In most storage arrays, a large number of identical hard disks are used, often even with manufacturer-specific firmware.

Since hard disks from the same production also age at the same rate, there is a risk of a kind of „epidemic failure“ when the first specimens fail, which makes it necessary to replace the entire storage system in good time.

Scalability

Since data is not usually deleted in the archive, the required capacity is constantly increasing. One exception is archives that only serve to fulfill legal requirements without the user having any interest in retaining the actual data. In this case, data is to be deleted or destroyed as soon as possible after expiry of the legally prescribed retention periods, which is why optical storage or even paper files that can be easily destroyed are still frequently used in this area.

All other archives serve the purpose of data backup. Since data is often also referred to as the „new gold," it is not deleted - the likelihood of later use of big data is constantly increasing with the proliferation of Big Data analytics and AI applications.

Archive storage must therefore be massively scalable. Since parallel data throughput does not usually need to grow with capacity, this is referred to as scale up storage.The underlying infrastructure should be as lean as possible, but allow for maximum growth in storage capacity. Scenarios of several petabytes are quite realistic.

Conserving resources

The longer a system is in use, the higher the energy costs are reflected in the total bill. Since access to archive data usually decreases with storage time, the storage of older data should also consume less or even no energy. This can be done by tiering, i.e., offloading very old data to offline-only storage. However, this requires additional data management and further storage systems, which again increases complexity and expense. Better are systems that are modular and have efficient energy management.

Another aspect is the carbon footprint of the components used. The selection of the materials used and the compensation of the CO2 emissions generated by production, delivery and disposal contribute to the climate balance.

Longevity & availability

Although old systems may not be as energy efficient as the latest generation, continued operation is in many cases more resource-efficient than early disposal. The longevity and availability of any necessary spare parts avoids problematic waste and thus contributes directly to the environmental balance. The choice of data media in particular makes a significant contribution to longevity here. Rapidly changing generations with limited downward compatibility not only force costly and error-prone data migration, but also create additional hazardous waste.

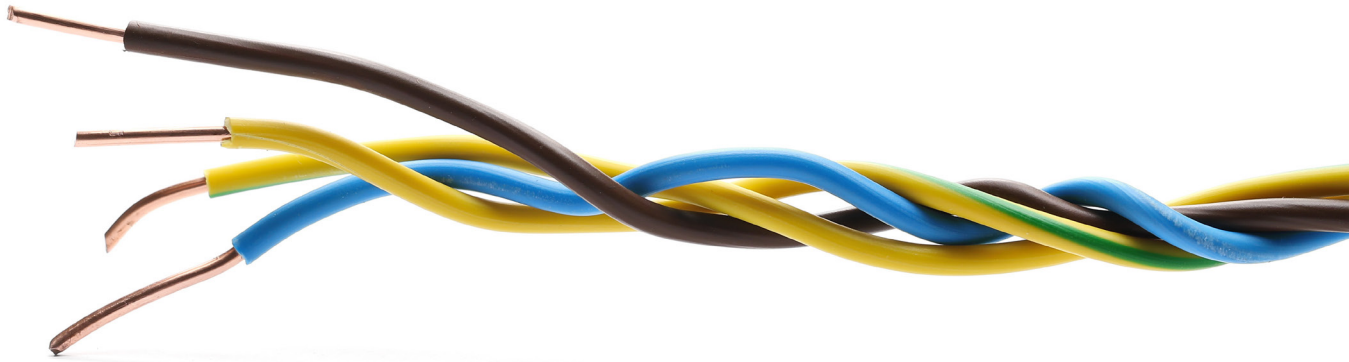

Sustainable software

The right software makes the difference between usable hardware and useless electronic waste. If the software for a hardware solution is no longer developed or essential infrastructure is no longer supported, the overall system loses its usefulness.

Standards, Certifications & Data Migration

An important criterion for a long-lasting software concept is the support of standards. This is the only way to ensure that a storage system can be replaced relatively easily if this should become necessary. Compatibility with industry standards ensures interchangeability while preserving data - including that of the connected software solutions. Data migration - if necessary - is also only feasible at all if the applicable standards are adhered to.

In many areas of archiving, a number of certifications are also mandatory in order to meet legal requirements. Without certification, the technically best storage system cannot be used for compliant archiving without the threat of far-reaching consequences in the event of an audit. Since such certifications involve considerable financial and organizational effort for manufacturers, there is a vested interest in the long-term validity of the solution offered - and thus the long-term further development of the corresponding software.

Reliability & Security

The classification as reliable software is not subject to any generally valid criteria. However, since today‘s storage systems are not self-sufficient, completely self-developed systems, but are made up of countless software components, regular provision of updates and security patches is absolutely essential, even after decades. The support of industry standards requires the integration of (open source) libraries, which, the more widespread, serve as gateways for malware in the event of security vulnerabilities. Prompt reaction on the part of the manufacturer is just as necessary as open communication.

But the functionality of the administration software can also contribute to longevity. If far-reaching backup measures can be configured via the administration, there is no danger of data loss in the event of a disk failure, for example. Special automatic mechanisms such as continuous snapshots provide additional protection against the consequences of a ransomware attack and also protect the data outside WORM archives.

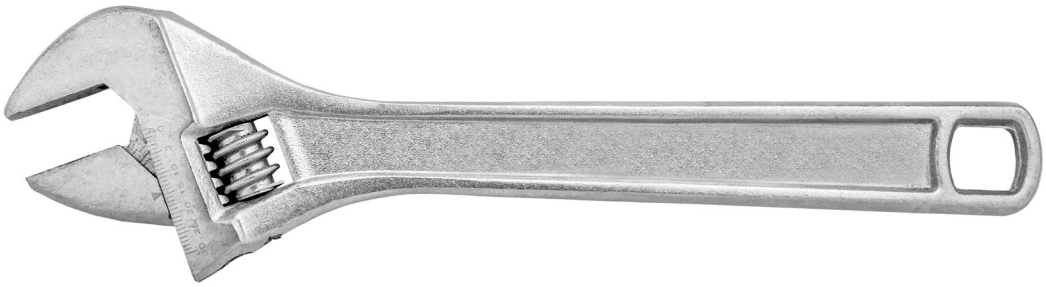

Sustainable Service

The more complex IT becomes, the more important the associated service level agreements are. Since these are usually mandatory for professional systems, the long-term predictability of maintenance costs contributes to the overall balance.

Maintenance effort

Generally speaking, storage systems with a low mechanical content are less susceptible. However, since flash memories that completely dispense with mechanics are likely to be too expensive for data archiving in most cases, it is important to weigh things up. Hard disk storage contains rotating magnetic disks that are susceptible to shocks. Fast server disks must also be in constant operation to prevent the read/write heads from gumming up.

Tape storage is based on fragile magnetic tapes that are used in mechanically complex drives. Tape libraries, which use additional robotics to exchange the media, are used above a certain storage requirement.

Costs & Spare Parts

Service calls cost money, generate travel and often require a partial or complete shutdown of the system. So doing without is desirable, and not just for environmental reasons. Systems that can be serviced by customers themselves, for example by replacing individual data carriers, are simpler. The availability of components even after decades plays an important role here. If systems are designed in such a way that successor models can also be used, the maintenance effort is considerably lower in case of doubt.

Especially when it comes to maintenance after the usual service life of three or five years, many manufacturers often increase the costs - for good reason. The longer systems are in use, the more likely they are to require maintenance. In addition, many companies are more interested in selling new solutions than in the long-term continued operation of existing systems. As a result, systems are often decommissioned and scrapped that basically work relatively well but have become disproportionately expensive to maintain.

Technologies compared

For archiving, different technologies perform differently depending on requirements. Tape, hard disk storage and cloud offerings compete with different promises.

Tape

Due to the low media costs, tape storage has been used most frequently for data archiving in the past. While the costs per TB of storage are indeed unrivaled, tape is also becoming increasingly less important in the archiving sector because data on tape storage is not available or only available to a limited extent due to its strict linearity. For use, therefore, it is always necessary to restore to accessible media - hard disks or flash - which increases effort, costs and complexity. In individual cases, the supposed cost advantage quickly turns into the opposite, as additional systems and management software are required.

Due to the high mechanical content of tapes and the associated components, the maintenance effort is relatively high - and can continue to increase considerably over the service life. Long-term maintenance contracts are therefore usually associated with considerable surcharges after a few years and are difficult to calculate.

The longevity of tape is also often advertised as being due to backward compatibility. Since LTO-8, however, this only applies to one generation, which was previously the case for twice as long. So if you want to benefit from the higher capacity of modern LTO media, there is a risk of regular data migration - in addition to the costly and less resource-conserving replacement of the drives.

As far as storage security is concerned, tapes are relatively good as individual media - but always require copies to safeguard against failure, which multiplies capacity and takes a considerable amount of time.

However, even older generation tapes tend to remain available for many years and are very inexpensive. Tapes are compact, and many IT managers have automated their handling over the years. If later use of the archived data is not planned, or only in individual cases, tape storage can continue to serve as archival storage.

FAST LTA
We secure Petabytes.

Cloud archive

Increasingly, there are also cloud offers for data archiving. The name already indicates the main application purpose: Amazon calls its long-term storage offering „Glacier“. The focus is on long-term storage; data retrieval - depending on the configured policy - is associated with considerable waiting time. Additional local storage systems are necessary for the availability of at least part of the data.

The cost structure is also different from that of in-house systems. Instead of paying for the storage system, the customer pays relatively little for storage - and relatively much for retrieval. If individual data is to be restored, this results in manageable costs. However, these storage systems are less suitable for use as „active archives“.

It should not be forgotten that cloud archives are also based on corresponding storage hardware. Depending on the availability level selected, data is stored on huge hard disk arrays or via tape libraries. Transparency about CO2 balance and compensation, energy efficiency and service life of the hardware used is often not available to the customer, especially with international providers.

Even more difficult to answer is the question of security and any necessary certifications. Often, storage outside the state or at least Europe is an obstacle to use. In general, you have to rely on the promises made by the providers. Extremely high availability promises are supposed to convey security, yet a problem in accessibility often results in the failure of a number of services - even if the actual storage service is not affected at all.

The advantages of cloud storage are easy scalability and low maintenance.

Silent Cubes and Silent Bricks: hard disk based specialist storage

With Silent Cubes and Silent Bricks, FAST LTA offers special systems for archiving. Silent Cubes and Silent Brick WORM are certified as storage with hardware WORM sealing for compliant and GDPR-compliant archiving. And Silent Bricks represent a technology alternative for tape libraries as large VTL archives in the petabyte range.

The software has proven itself for decades and is constantly being further developed. Since the first Silent Cubes were installed in 2008, the storage products have been performing their service for thousands of customers - in most cases largely maintenance-free in the background in accordance with the „Silent“ designation. The products are not affected by current security vulnerabilities such as Log4j due to strict quality control.

Hard disk systems are random access and very fast in the data carrier network. In addition, each storage module uses disks from three different production batches to prevent data loss due to correlated failures. Replication to a second site further increases security through geo-redundancy.Silent Cubes and Silent Bricks are designed as modular storage and scale into the petabyte range. Replication to a second site further increases security through geo-redundancy.

Silent Cubes and Silent Bricks are designed as modular storage and scale into the petabyte range. Individual storage modules can be switched off when not in use, Silent Bricks can even be physically removed from the system (Air Gap). This enables massive energy savings, since only the active storage modules actually consume energy. The modules regularly „wake up“ themselves for self-checking and self-repair.

FAST LTA
We secure Petabytes.

The systems also stand up well in terms of their CO2 footprint. With the introduction of the Silent Cube DS, FAST LTA is starting to completely offset the CO2 emissions during production, delivery and disposal. For future developments, the use of recycled materials, e.g. for housings, is increasingly planned. FAST LTA encourages its customers to use climate-neutral and renewable energy for operation.

Storage systems from FAST LTA are designed for ten or more years of use. The use of standard hard drives (SATA) allows for extensive compatibility across multiple generations. In the event of a hard disk failure, the affected data carrier can be easily replaced. A special model is not required; the in-house developed controller automatically adjusts the available capacity to the existing configuration. This ensures the availability of spare parts even after years.

Maintenance contracts are also adapted to these terms and can be renewed at the same conditions. The well-developed service network is available 24/7/365 on demand.

Bottom line

The benefit of a storage system is oriented to the intended application. With the increasing replacement of completely „cold“ archives by active storage systems with high data availability, the focus on sustainability is also changing. However, hard disk storage must overcome a number of challenges that can only be met by modular storage systems developed specifically for archiving.

Clever energy management, a balanced carbon footprint, durable hardware and software, and long-term service at predictable costs form the basis on which even „online“ storage systems can be described as sustainable.

With Silent Cubes and Silent Bricks, FAST LTA offers two solutions that are precisely tailored to the needs of long-term data storage, whether for compliant and GDPR-compliant archiving or as a large VTL archive.

FAST LTA
We secure Petabytes.

PROULDY MADE IN EUROPE

FAST LTA | Ruedesheimer Str. 11 | 80686 Munich, Germany | info@fast-lta.com | **www.fast-lta.com**