CEOS

Working Group on Information Systems and Services
Data Stewardship Interest Group

Archive Technology Evolution

CEOS
Data Stewardship Interest Group

Doc. Ref.: CEOS.WGISS.DSIG.ATE
Date: June 2023
Issue: Version 1.0

Contents

INTRODUCTION AND SCOPE

The Data Stewardship Interest Group part of the Working Group on Information Systems and Services (WGISS) provides a forum to exchange information and lessons learned related to data archiving, preservation and curation.

ESA ensures the preservation of a copy of all its Earth Explorer, Heritage Missions and Third Party Missions data through a state-of-the-art technology robotic archive located at ESA/ESRIN premises in Frascati (Italy).

The archive is part of a wider ESA EO Space Data Preservation System which has the objective of providing the required infrastructure and services to ensure archiving, preservation and retrieval of EO data records and technical information. In addition, the ESA EO Space Data Preservation System provides support to the cooperation activities with national and international organizations in the data preservation domain.

The main challenge of the archive is the sheer volume of data, the copious number of files and the diversity of formats. To keep up with the always growing amount and variety of data produced by different missions and sensors, the archive is constantly enhanced to facilitate the archival process by using technological innovation and practices. To save energy, thus making the archive and its hardware environmentally friendly, green computing technologies are favourite.

This white paper summarises the results of the activities and the study performed at ESA, to evolve the ESRIN Archive to the next technology step in order to assure seamless preservation of the always increasing volume and variety of data records. It describes the functionalities and the architecture of the archive, the data archived and the different challenges driving its evolution. The paper also provides (not exhaustive) information about archiving and preservation solutions that are available on the market or implemented by other agencies.

The content of this white paper is based on online surveys results and is not meant to be exhaustive. Suitability for any purpose of use is not guaranteed and left at risk of the user.

ARCHIVE SOLUTIONS IN EARTH OBSERVATION

Space data are generated in a variety of formats depending on its origin and its level of processing. Data is usually stored in the native format generated by the receiving stations, together with auxiliary data, documentation and other types of binary files.

The Open Archival Information System (OAIS) Submission Information Package (SIP), Archival information package (AIP) and Dissemination Information Package (DIP) are widely adopted by producers, consumers and archive managers to preserve data and make it available to its designated community.

Archives based on Hierarchical Storage Managers allow the management of preservation actions in different storage levels in a transparent way for users, both in terms of ingestion and access. Such an archive usually has three levels of storage with two different technologies, SSD disks and magnetic tapes although the largest volume is stored exclusively on tapes since disk storage is used only in the processes of ingestion and access, deleting the information when it has already been transferred to the final file in the tape library or to the user.

The system, since the data always ends up on the tapes, does not have a theoretical size limit, although this depends on the manageability of the tape technology used. Users view the archive as a single access point where they get all files without being immediately aware if they are online or need to be retrieved from tapes. Access to files located on tapes can take a few minutes depending on the workload of the library or a few hours if the tape is not in the library and an operator must retrieve it from the off-line archive.

ESA EO Space Data Preservation System Solution

The EO Space Data Preservation System (SDPS) is composed of several (sub-) systems and services and has the main objective to ensure ESA Heritage and Earth Explorer Missions, and ESA managed Third Party Missions (TPM) EO Space Data Packages holdings (i.e. Space data and associated technical information) preservation, and to support the cooperation activities across ESA and with national and international organizations in the data preservation domain.

The focus of the SDPS is the Space Data Packages which consist of:

- Space-borne data held or owned, and managed by ESA
- Payload instrument data from ESA and/or Third Party satellites
- Spacecraft telemetry data, auxiliary and ancillary data about satellites and instruments
- Higher-level data processed from the above data types
- Technical and Scientific Information related to Space-borne Data (from ESA, Industries, and Scientists) necessary or useful for intelligibility, usability and exploitation
- Software and Data Analytics Tools related to Space-borne Data necessary or useful for usability and exploitation

The Space Data Preservation Archive, together with a second archive outside ESA premises, provides the data archiving and preservation function for all ESA managed Earth Observation data (excluding Copernicus Sentinels). The Space Data Preservation Archive also provides a back-up function for space data from other ESA directorates.

Two copies of all the data are archived at ESRIN premises in two separate buildings using different technology to comply with internationally recognized preservation guidelines. Two tape libraries are the final tier of the multi-tier storage environment.

Main Issues and challenges of the ESA Space Data Preservation System Archive

Over the course of its operative life, many challenges have arisen and a consequent upgrade of the hardware, software and network environment have been performed to enhance and increase the overall performances of the ESA EO Archive.
Some outstanding challenges required an overhaul of the environment or data flow, and additional upgrades are being taken into consideration to handle different operational scenarios.

The current main challenges of the archive are the following:

- *Data produced is always increasing*. At the time of writing, the system is storing about 15.6 PB of EO data in approximately 1.2 billion files. The system is prepared to receive high daily volumes of data, currently reaching 20TB on specific days and normally handling around 13TB. Annual growth of about 5PB is projected due to future missions planned to be launched;
- *Performance degradation when dealing with small files*. The file-system accessed by data client is partly based on a physical storage with Hard Drives. The filesystem metadata shared the physical storage and the I/O of the daily operations affected performances. Reading, moving, copying and deleting lots of small files required constant updates of the metadata and affected performances;
- The *current system has a limitation on the number of files it can store* since the addresses that allow access to the files on the tapes are stored on a 1TB LUN that cannot be extended or are very complicated to extend and requires a significant downtime;
- *The Archive Manager ceased being maintained by its vendor.* A migration towards an alternative is therefore required and it has proven to be a complicated matter. Alternative Archive Managers are often not compatible with the current metadata format meaning that all data (17PB) needs to be read from tapes to be migrated to the new solution. Some of these alternatives offer licensing schemes where you pay for what you read or archive. This is not applicable in the case of ESA's archive due to constant data migration activities.

A serious of minor challenges affects daily operations and are a consequence of the hardware solution used in the Archive:

- Media failures recovery takes time and uses resources;
- Tape libraries support requires lots of data collection even for obvious failures;
- Tape library support often result in requests to update drives firmware;
- Longevity of technology has side-effects (maintenance costs, obsolescence of hardware);
- Tape prices are not predictable;
- Extracting data from tapes is slow and complicated.

ARCHIVING TECHNOLOGY TRENDS

A non-exhaustive description of hardware and software solutions available on the market is provided below.

Hardware

Various storage systems currently in use or under development could be used and be suitable to manage the volumes of information required by the ESA EO archives.

LTO technology	Linear Tape Open (LTO) Ultrium is a high capacity, single reel tape storage solution developed and continually enhanced by Hewlett Packard Enterprise, IBM and Quantum and promoted by the LTO Program. It's a powerful, scalable and adaptable tape format that helps address the growing demands of data protection. It's also an open format, licensed by some of the most prominent names in the storage industry to ensure a broad range of compatible tape drives and cartridges. (www.lto.org) This technology is delivered to the customer in versions that increase the capabilities of previous versions. Currently the current version is the LTO-9 which has an uncompressed storage capacity of 18TB and 45TB compressed per cartridge. The writing time of the full tape is 12 hours 30 minutes. Since version 8, the LTO consortium has introduced the WORM (Write Once Read Many) tape functionality that allows to avoid overwriting the data already stored on the tape but allows to add data at the end until its capacity is completed. In the roadmap of the product are currently planned versions up to 14 with estimated capacities of 576TB uncompressed and 1.44PB compressed. No writing times are announced. The next version LTO-10 doubles the capacity of LTO-9 in storage and the write time is slightly less, 12 hours and 7 minutes. LTO through 7 versions are compatible with hardware from two previous versions. Versions 8 and 9 are only compatible with the previous version of read/write hardware. There is no information about the expected compatibility in future releases. Another interesting feature that has been introduced in the LTO family since version 5 is the LTFS or Linear Tape File System. The Linear Tape File System (LTFS) is a file system that allows files stored on magnetic tape to be accessed in a similar fashion to those on disk or removable flash drives. It requires both a specific format of data on the tape media and software to provide a file system interface to the data. Several software vendors have implemented LTFS in conjunction with LTO equipment. In the long term, LTO requires constant migrations to keep the data alive since the new hardware is not compatible with previous versions (only the previous one in LTO-9) and therefore when you need to change the hardware due to its obsolescence, a complete migration of the file to the new media will be required. As a technology based on magnetism, it is sensitive to magnetic fields or radiations that can totally or partially destroy the contents. According to

	the sources consulted, in 2022 there are four manufacturers of LTO cartridges, Sony, Fujifilm, IBM and HPE.
Folio Photonics	Folio Photonics is developing a new optical disk technology based on polymers and advanced nanophotonic materials that are embedded into multi-layer films via an efficient co-extrusion manufacturing process. This film is then attached in DVD-form disks to be used by end-users and is readable by BluRay lasers. Folio disks are still on a prototype phase. In this phase they have already proven that the materials are correct to write data permanently on them and that the reading and writing of data on the disk at commercial speeds is feasible. They have also confirmed that the manufacturing process is good in a commercial perspective. The technology has 6 patents (2 pending) but they are claiming that if is open on reading, as their disks can be used in currently in use BluRay drives. By now only a Folio proprietary prototype of drive is available. They expect the first generation to be available in early 2024 with a capacity between 0.8 and 1 TB per disk. In the roadmap of the product they have planned capacities of up to 10TB per disk. They also plan to make automated libraries available to handle large numbers of disks simultaneously in the way tape libraries currently work. This type of technology is random access so data access times are reduced compared to linear access technologies. It is a WORM technology so it is not possible to modify or delete the data after writing them, it does not require energy for its conservation and it is not alterable by radiation. Folio indicates that its laboratory tests have defined a longevity of the storage medium of +120 years although its resistance to the problems that other optical storage discs had such as scratching or the separation of the writing medium from the plastic that supports it is not indicated. 
IBM Storage	International Business Machines Corporation (IBM) is a technology company. IBM is offering various types of storage for different uses such as SAP, NVIDIA, mainframes, cloud, virtual disk and others. As solutions for archives it has tape libraries based on LTO-9 currently being the TS4500 and TS4500 R8 the most suitable for the uses that ESA / ESRIN has to perform. The newest TS4500 R8 is a storage solution that is designed to help midsize and large enterprises respond to storage challenges such as high data volume, growth in data centers, increasing cost of data center storage footprints, difficulty migrating data across vendor platforms, and increased complexity of IT training and management as staff resources shrink. The TS4500 tape library combines reliable, automated tape handling and storage with high performance in an open systems and enterprise environment. Incorporating the IBM Linear Tape-Open (LTO) Ultrium tape and 3592 drives, the TS4500 tape

	library offers outstanding retrieval performance with typical cartridge move times of less than 3 seconds. The TS4500, installed with the High Availability (HA) option, provides dual active accessors for redundancy and can double the robot performance during tape move operations. The TS4500 tape library can be deployed as a single frame library and  upgraded to a maximum of 18 frames, with a combination of LTO and 3592 frames. It supports both LTO-9 and other proprietary IBM tapes such as Type 3592. It is modular and new modules can be incorporated that can expand the active capacity of the library up to 17 frames and 23,000 cartridges. This configuration allows you to store an active volume of 417 PB. Obviously if you mix it with an offline file you would have an unlimited capacity as currently. In second place, IBM is offering its two services on cloud for Block and Object Storage. These are public cloud storage services with three tiers of storage called Standard, Vault and Cold Vault for different purposes and with different costs. As an interesting service they offer what they call Smart Tier that classifies the data in the three previous tiers and stores them automatically depending on this classification. The costs per GB in this last option are a little higher than if the operation is done manually in all three tiers. It has distribution in different geographical areas and variable costs of handling information. You only have recovery costs in the Cold Vault.
Microsoft Glass Drive	Project Silica is developing the world's first storage technology designed and built from the media up to address needs for a long-term, sustainable storage technology. Microsoft stores data in quartz glass: a low-cost, durable WORM media that is EMF-proof, and offers lifetimes of tens to hundreds of thousands of years. Microsoft use different technologies to write and read the data: ultrafast femtosecond lasers to write, and polarization-sensitive microscopy using regular light to read. As a consequence, the Silica storage system guarantees true airgap by design for the storage media; it's physically impossible to accidentally overwrite data during reading, as there is simply not enough power to modify the glass material. The mechanical design of the media library also makes it impossible for media to find its way back into a writer, further guaranteeing the security of archived data for its entire lifetime.

	As a storage technology, Silica offers volumetric data densities higher than current magnetic tapes (raw capacity upwards of 7TB in a square glass platter the size of a DVD), and using beam steering of the laser 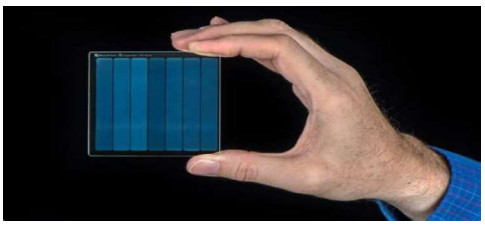 beam, we're able to achieve system-level aggregate write throughputs comparable to current archival systems. It is still a Microsoft research project and the information available on how Microsoft intends to market the product is scarce. It seems that they will integrate it into their data centers as a long-term archive storage, offering it as a service of their clouds. In 2019 they published a news about a proof of concept with Warner Bros to validate the technology in real cases, apparently successfully. However, only a few more cases of possible uses appear and it seems that, despite the interesting and promising technology, there is not much interest in Microsoft for it. It is interesting to note that GitHub, a company owned by Microsoft, chose another technology for the very long-term storage of free software hosted on its service. There are other examples of research in this type of glass storage but there is no information on its marketing status, such as that of the University of Southampton's Optoelectronics Research Centre
Hitachi Vantara	Hitachi Vantara is the company Hitachi has dedicated to all data-related services and products, from storage to cybersecurity, analytics and cloud services. Vantara has different offerings for the following types of storage: • Block Storage. • Object Storage. • Software-Defined Storage. • File Storage. • Cloud Storage. • Hyper-converged Infrastructure. • Mainframe Storage.

	The range of storage solutions is very wide and their analysis requires knowledge outside the scope of this report. Object Storage is provided with two main components, Content Platform which is an object storage management software and Content Platform S Series Storage which is the hardware platform that physically stores objects. This combined solution is aimed at allowing users to build hybrid clouds easily. It allows you to combine public and private clouds using Content Platform Gateway and other technologies. In the case of hardware, it has two HCP S11 and S31 disk arrays that cover a maximum of 288PB in the S11 up to 1.3EB in the S31. They allow a modular growth since you can start with the number of racks needed and incorporate new ones as needed. You can mix SSDs and HDDs on the same computer, and access and management is done with two built-in servers active at the same time. 
NetApp	The NetApp portfolio includes intelligent cloud services, data services, and storage infrastructure that helps organizations manage applications and data everywhere across hybrid multi-cloud environments. They are providing different solutions for the storage market and more and more they are focused on cloud services. Therefore, some products are fitted to an on-premise installation. Among them it is worth to mention the FAS family of storage appliances.  From FAS2750, that stores 1.4PB per unit to the FAS9500 that keeps 14.7 PB. They rely on ONTAP operating system which is providing key features for a preservation archive like ESA/ESRIN. The WORM feature for disks, mirroring, data compaction and others would help to guarantee the access to the data.
XenData	XenData is a global provider of data storage solutions optimized for creative video, medical imaging, video surveillance and other applications that require very high capacity file storage. XenData archive systems provide secure long-term storage on LTO data tape, optical disc cartridges and/or the cloud object storage.

	 XenData X100 archive systems manage LTO robotic libraries that scale to 100+ PB. They support leading enterprise libraries, including from HPE, IBM, Oracle, Qualstar, Quantum and Spectra Logic. X100 systems use the power of modern servers to simplify the archive architecture. The X100 is available as a clustered system with no single point of failure and as a single high-performance server. It is powered by XenData Archive Series software which is trusted by many of the world's largest media companies. Adding an object storage S3 Server Interface to an X100 system allows files to be written to and read from the archive from anywhere worldwide using fast and secure HTTPS. Using the S3 Server interface allows an organization to keep the attractive aspects of LTO on-premises archives which include cost effective scalability and adds the ease of distributed access which has traditionally been associated with public cloud storage.
Qstar Kaleidos	The QStar Kaleidos is an S3-compliant object storage platform that enables enterprises and service providers to build reliable, private, hybrid or public cloud storage environments that deliver reliability, security and unlimited scalability. 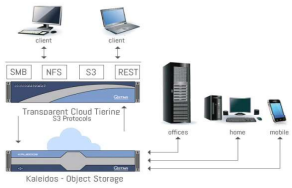
Overland Tandberg Libraries	Overland Tandberg mission is to distribute sustainable and scalable worldwide solutions to ensure business continuity.  They provide an entire family of products, from servers to disk and tape-based storage systems.
Spectra Logic	Spectra Logic defines, designs, and delivers innovative data protection and backup solutions through tape backup, deduplication, and disk-based backup, recovery, and archive storage solutions. Spectra provides storage solutions, either software, object storage and tapes to fulfil large volume needs.

	With tools like StorCycle and BlackPearl they offer automated tier-management and hybrid cloud platform to allow integration of single system storages in the cloud or on-premise. Different models Tape Libraries of different capacities are offered with the smaller Spectra Stack offers up to 10PB/rack in LTO tapes and can scale by adding new racks that will be managed independently with StorCycle software. The bigger Tfinity Exascale that can reach 1EB on LTO tapes.
Oracle Storage	Oracle markets three products for storage purposes. Oracle ZFS Storage Appliance is a high performance, enterprise storage system that is optimized for high workloads and cloud integration. A unified storage system that enables consolidation of block, file, and object storage, it is available in flash and disk configuration. In the description, Oracle focus on their high performance characteristics but also some integrated storage management is shown. It mentions up to 8PB on SSD disks and up to 21 PB on HDD. For ESA/ESRIN archive this could be a solution that substitutes the current HSM solution adding more online capabilities. In second place, we can mention Oracle StorageTek tape libraries. They provide immutable, offline storage to protect customer data while providing easy, automated access for compliance, governance, and historical-preservation purposes. Organizations use StorageTek tape libraries to protect against cyberattacks and to archive data for years with lower energy consumption and less cost than disk-only alternatives. In the StorageTek family there are several products that go from the SL150 with max 8.1PB of data to the SL8500 that can store up to 57.6 EB. We could not find any references to standard access protocols, such as S3 or LTFS, compatibility with this library. There could be third party software that could address this issue translating them into StorageTek's language.
Piql Film	Piql provides a digital storage based on photosensitive 35mm film. It could mix also digital information with visual contents to help on the interpretation of the data or represent some of the data (like documents) in a human-readable way. Its main advantage is that is an offline self-contained and auto described media. It stores not only the data but the description of it and the software to recover it. It has a capacity of 120GB per reel but the archive can grow without limitation as the reels are offline. The tests done by Piql claims that the data can be accessed in more than 750 years in standard conditions. In the roadmap of the product, it is expected to have a capacity of 200GB within 2023. Piql also provides vaulting services for the reels in the Arctic World Archive that guarantees a +1000 years of live of the data without any intervention in a secure site.
DNA storage	There is a research project at Harvard University for data storage using a 4-letter code in nucleotides. The interest in this type of storage is that it has a higher density and durability than other silicon-based ones being

	1000 times denser than SSDs and 300 times more than magnetic tape. The storage allows you to save digital files as in any other medium. However, currently the cost of storage is prohibitive due to the cost of the chemical synthesis of DNA, being currently $ 3500 / MB. They continue to investigate the system to fundamentally reduce this cost but its possible availability in the short term is not indicated. Being an organic storage, longevity and mutation resistance validations would be necessary like other "living" elements that could corrupt the stored data.
M-Disk	Millenniata, Inc. is a global data storage company that has created the M-DISC™ DVD, which they claim is the world's first archival disc to last up to 1,000 years. The M-DISC engraves data into a patented rock-like layer that is resistant to extreme conditions of light, temperature and humidity – outlasting all other archival optical discs on the market. Its technology is different from previous CD or DVD disks because it is not based on organic dyes. 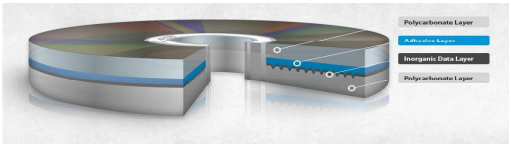 The biggest disk commercialized with this technology has a 100GB capacity and it is usable in currently available BluRay writers and readers. On the website of the company is not showing any library capable of handling multiple disks at a time. Such libraries used to exist in the past years and we found that some are still working like the LB-DH7 Series Data Archiver. The maximum capacity of this particular library is only 1.9 PB but several racks could be placed together. Maybe other types of "BlueRay-like disks" could be used on it to increase the capacity but it is uncertain now.
Quantum	Quantum technology and services help customers capture, create and share digital content – and preserve and protect it for decades. Quantum offers four products on tape storage. Of these, two could be used to cover the needs of the ESA / ESRIN file that are Scalar i6 that reaches up to 14PB and the most suitable Scalar i6000 that stores up to 253.8PB. Both are based on LTO technology and have WORM capabilities. They also have a product called HyperScale Archive Storage that mixes the aforementioned tape libraries with two software products, StorNext for file file management and ActiveScale Object Storage for object-based file systems. Through StoreNext, the user can access the data with standard protocols such as NFS, SMB or S3 in addition to Quantum's own interfaces. ActiveScale software allows you to manage hybrid clouds by integrating different technologies and manufacturers.

	Finally, Quantum also offers different disk appliances both SSD and HDD that can be used in line with the libraries to build a complete archive.
Hewlett-Packard	In the case of HPE, it has file and object archiving solutions based on SSD and HDD disks. It also has some tape libraries but its capacity is limited for the requirements of the ESA/ESRIN file. HPE is shifting all of its storage and archiving services to the cloud and offers HPE GreenLake with a combination of block, file, and object storage in the same cloud service.
Active Archiving Concept	In an active archive, data is intelligently and dynamically managed across a high performance, optimized storage architecture comprising high-end flash devices, intermediate flash storage, cloud, disk and tape. The objectives are too free up primary space, optimize storage costs, reduce environmental impact, deliver on compliance legislation and provide maximum protection in the event of data loss, compromise and theft. Intelligent software continuously tracks, manages and moves data up and down the storage tiers using rich metadata, indexes, directories, tags, and global namespaces to unlock archives and deliver high-performance search and retrieval across the storage spectrum.

Software

Archive management cannot reside exclusively in adequate storage systems but requires applications that support the different activities that preservation requires. In addition, if you want to have an OAIS-based preservation file, you must meet all the requirements that are defined in the standards and that are not covered by hardware access software. Below it is provided a description of software elements that are related to the management of the archive in all its areas, from communication with storage, movement of data between tiers for optimization to preservation management.

Alluxio	Alluxio is an open-source virtual distributed file system (VDFS) and was created at the University of California, Berkeley. Alluxio provides a data abstraction layer for computation frameworks, enabling applications to connect to numerous storage systems through a common interface. Alluxio enables data orchestration for compute in any cloud. It unifies data silos on-premise and across any cloud to give you the data locality, accessibility, and elasticity needed to reduce the complexities associated with orchestrating data. In summary, it allows to access multiple storage media in a common interface. It is widely used in multiple organizations and provides a free open-source license but also an Enterprise version. It does not mention specific support to any tape technology but as it supports S3 and other access protocols, once the tape system supports these interfaces it is likely that they can be used in Alluxio.

Atempo	Atempo has data management solutions to preserve data ecosystems for mid-sized and corporate organizations. Atempo protects endpoints, applications, servers and large storages including the cloud. Their product Miria for Archiving is a file archive solution for long-term data retention and data retrieval. It allows multiple storage destinations with an agnostic approach. Storage media include disk, tape, object, cloud, NAS or piqlFilm all managed in the same way. It also supports automatic data movement between storage tiers relying on a configuration for each dataset type. The product supports storage from many different providers such as Oracle, AWS, Azure, Spectra and many more. No information about the licensing or purchasing options is found on their website and neither OAIS support is mentioned.
Versity	Versity Storage Manager (VSM) is a large scale archive data management tool for cost-effectively protecting extensive data collections on any mix of storage devices and cloud services. VSM interfaces with existing enterprise business applications and backup solutions through a standard POSIX file system interface, making it easy to expose low-cost archival storage resources to end-users. Rich and flexible archive policies optimize the data flow both in and out of cloud storage services and archival storage devices such as tape libraries, disk arrays, and on-premise object storage systems. Versity claims to be a solution to migrate from Oracle HSM because it supports a zero data migration The pricing is done per site in a subscription model.
MinIO	MinIO offers a high-performance S3 compatible object storage. It is the only object storage suite available on every public cloud, every Kubernetes distribution and the private cloud. MinIO is software-defined and is 100% open source under GNU AGPL v3. It has functionalities to manage the data lifecycle, automated tier movement, bucketing, encryption or multi-cloud storage. As mentioned, it has an open-source license but also provided paid services in two bundles that have some different functionalities.
LABDRIVE	LABDRIVE is a digital research data management and digital preservation solution developed by LIBNOVA. It is commercialized as a SaaS product, available in EOSC cloud, and it is a result of Archiver project. They are claiming to be scalable to high volumes and throughputs reaching 15.87 PB with 600 million of files in 30 days. As a pure cloud solution, a pay per service model is established. AWS is their cloud provider. It is worth mentioning that, as developed for Archiver project, the validation of the system has been made with CERN, EMBL_EBI, DESY and PIC datasets. No information about real users found.
nageruHive	This is a data management, archiving and preservation tool compliant with OAIS developed by nageru Solutions. It is a service platform to

	acquire, interpret, process, transform, store and retrieve all kinds of data. Its chains of services can be adapted to different input, output or processing needs and it has been designed with high-volume, high-throughput in mind and using hybrid approach to the storage. It also provides AI functionalities to be used during or after the data ingestion. Among them it can be mentioned automatic classification, metadata extraction, anomaly detection or correlation of data. Nageru has passed ESA tests to provide a Long-Term Archive for  Copernicus program and it is waiting to become an ESA provider for it. It can be purchased as cloud services with different pricing plans depending on the needs or in a license-basis for on-premise deployments.
Arkivum	Arkivum provides archiving and preservation solutions to different verticals like pharma, research, education, heritage and corporations in general. Arkivum supports a broad range of data types, formats and sources, and provides integrations to a range of different systems. On their website they mention that the solution has an automated ingestion module to make the contents to enter the archive, a safety module to ensure the integrity and security of the data and a preservation module to provide tools for the access in the long-term. They also talk about search, access and reporting tools availability.  In their information, they show performance of 25TB of data ingested (440000 files) in 24h hours what fulfils current ESA/ESRIN archive needs. As part of the solution, they are using Archivematica for the packaging of the contents. No pricing or licensing information is found in their website. For the research domain, it is known that it can be purchased through SaaS model.
Preservica	Preservica have both cloud-hosted SaaS and on-premise software to support different kinds of customers and needs. Preservica's standards-based (OAIS ISO 14721) active preservation software combines all the critical capabilities of successful long-term digital preservation into a single integrated platform. It keeps content safely stored, makes sure it

	can be found and trusted, provides secure immediate access, and automatically updates files to future-friendly formats. They have developed some open-source tools used within the preservation industry like DROID but their solution is not open-source. On the cloud service, they sell it through 3 packages hosted in AWS or Azure. For the size of the ESA/ESRIN archive the Enterprise package is the one to use as it is the only one with the storage limits good for ESA/ESRIN archive. It includes all functionalities needed for a proper management of the preservation.
Seagate Lyve	Lyve™ is an edge-to-cloud mass storage platform from Seagate®. It allows to build private cloud solutions with a claimed "best value per PB". It includes an open-source object storage called CORTX. CORTX is an open-source software-defined object store backed by Seagate and designed, built, and maintained by a growing community of data scientists and big data and enterprise storage experts. CORTX enables the data to communicate directly with the storage drives without an intervening file system. This direct-to-drive architectural feature improves performance and reliability while giving finer control over storage drives.
IBM	HPSS is disk and tape storage software designed to manage and access exabytes of data at high data rates. Files can be on disk, disk protected by tape, disk space-managed by tape, or stored directly to tape. For files on tape, an option determines if files are recalled to disk or accessed directly from tape. Standalone, HPSS presents its own file system directly to the user. It is a HSM software that support different disk and tape vendors. Among them is Oracle. French National Space Center (CNES) and several NASA services are using HPSS on its storage needs. 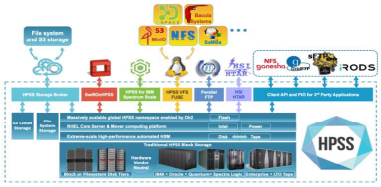 There is another solution that IBM markets for Archiving and storage management, IBM Spectrum Archive. It gives organizations an easy way to use cost-effective tape drives and libraries within a tiered storage infrastructure. IBM Spectrum Archive simplifies data movement between flash/disk and tape, helping to lower costs and provide air gap capabilities without the need for proprietary tape applications. It supports LTFS in three versions, single drive, Library and Enterprise editions. No support for Oracle tape libraries is directly available.
Piql	Based on Piql's "no-migration" strategy for the preservation, Piql iVM is a product to ensure the future access to data independently of the

	software that is available in the moment of the preservation. It relies on an abstract virtual machine that do not depend on any specific hardware and that it can be executed in future computers like now. By now it supports the access to PDF, TIF, JPG and WAV files and also SIARD databases. It can be easily extended to other types of formats. It is licensed in packages depending on the number of files to be preserved.
Qstar Archive Storage	QStar Technologies provides the software, and infrastructure needed to implement active archive storage solutions for enterprises. With the use of LTO (Linear Tape Open) and either LTFS (Linear Tape File System) or TDO (Tape Disk Object) technologies, we can provide scalable, cost-effective solutions for any archive storage need. TDO is a Qstar proprietary technology that can combine physical magnetic tape storage and the Cloud (local, hybrid or remote) to create a cohesive data storage system. LTFS provides an industry standard tape format for sharing archive data. Tape as NAS or Tape as Cloud solutions can be obtained with Qstar software.

Cloud Providers

Cloud offerings are becoming an active trend within archives and preservation market. There are plenty of cloud providers with similar services, storage on several tiers, processing power and backup. Some of them also provide other services like database hosting, AI services, etc. that could be used together with the archiving ones. Among these providers is worth to mention Amazon Web Services AWS, Google Cloud, Microsoft Azure, OVH, Oracle, Safespring or Cloudferro.

Although the storage prices are possibly lower than in on-premise solutions, it is important to note that many of these providers charge for the recovery of the data. In some of them it would be almost impossible to budget the costs needed for the services needed by a space agency or organization as ingestion or access volumes are not really known. A transparent approach to budgeting is key when applying cloud services to archiving and preservation. Cloud archiving is a trade-off between accessibility and cost. Lower cost means limited accessibility. Exit strategies need to be implemented as it is easy to get locked onto an specific cloud-vendor if the customer relies too much on specific services and future changes on the conditions would lead to unexpected costs. Not to forget there are also legal issues regarding the placement where the data is placed or processed. In some of these providers, regioning is available but the data owner needs to ensure that its data is following its legal rules. Service-levels need to be clearly defined to be able to ensure high-quality and sustained end-to-end service to users.

Standards

In addition to Agencies' internal standards and OAIS as an international standard for the development of digital preservation systems, other standards and good practices that are in use in other preservation repositories are indicated below.

- CoreTrustSeal is an international, community based, non-governmental, and non-profit organization promoting sustainable and trustworthy data infrastructures. It offers to any interested data repository a core level certification based on the Core Trustworthy Data Repositories Requirements. For repositories that store information as valuable this certification provides a framework to guarantee access to data. It provides additional guidelines to those given by OAIS.
- For cloud solutions, SWIPO is a framework to be taken into account. It is a group of multiple stakeholders in the cloud world facilitated by the European Commission with the aim of developing procedures that guarantee the correct application of European directives on the free flow of non-personal data. In this area, the ISO/IEC 19086-1:2016 standard seeks to establish a set of common cloud SLA building blocks (concepts, terms, definitions, contexts) that can be used to create cloud Service Level Agreements (SLAs) is also interesting.
- Other possible rules that apply to data repositories and information systems in general would be ISO/IEC 27001 which is the world's best-known standard for information security management systems (ISMS) and their requirements. Additional best practice in data protection and cyber resilience are covered by more than a dozen standards in the ISO/IEC 27000 family. And also ISO/IEC 27040:2015 which provides detailed technical guidance on how organizations can define an appropriate level of risk mitigation by employing a well-proven and consistent approach to the planning, design, documentation, and implementation of data storage security.
- Finally, the FAIR principles have been developed with the intent that these may act as a guideline for those wishing to enhance the reusability of their data holdings. Distinct from peer initiatives that focus on the human scholar, the FAIR Principles put specific emphasis on enhancing the ability of machines to automatically find and use the data, in addition to supporting its reuse by individuals. Following these FAIR principles would improve the reusability of the data, and the importance of the archive itself consequently.
- All these standards, procedures and good practices are being used by ESA-like organisations around the world and market interest in them is growing every day.

Existing tools (excerpt)

LTA Copernicus	The ESA Long-Term Archive is a place where Earth observation data is stored under the European Copernicus Earth observation program. The platform collects data from satellites, describing various parameters related to phenomena on our planet. LTA provides access to archival data in a cloud that is used by scientists, researchers and public administration officials e.g. for the implementation of projects in scope of marine and land ecosystems analysis, environmental protection, atmosphere, climate change, crisis management and security. The Long Term Archive Service has the objective to preserve the Copernicus mission products for the long term (i.e. a longer period than the contractual coverage of the activity). It is essential that there should be no loss from the archive.

	Copernicus LTA, as stated before, is managed as cloud services provided 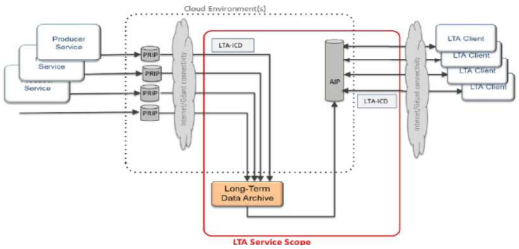 by different suppliers with a clear Service Level Agreement. Although the long-term archiving of the data is a goal for LTA, it also serves for the purpose of give access to the Copernicus data to the research community. For this reason, the LTA has high-performance requirements either for the ingestion of data and for the access to it. The LTA has defined an open interface that all supplier must comply with in order to receive access requests from Copernicus SciHub website. This website is responsible of managing to which provider should the request be redirected in order to get the best response times to customers.
Planetary Defense LTA	In 2022, ESA's Planetary Defense Office has awarded a tender to provide the design of a new Long-term Archive solution for Space Safety Program data. No hardware or software are part of the products of the project but a full design of the system to be developed. No further public info is available.
Planetary Science Archive	The Planetary Science Archive (PSA) is the European Space Agency's (ESA) repository of science data from all planetary science and exploration missions. The PSA is being developed and operated within the ESAC Science Data Centre (ESDC) in collaboration with ESA's planetary Science Operations Centres, also located at the European Space Astronomy Centre near Madrid, Spain and with the missions' instrument teams mostly in Europe. The ESDC also hosts other science archives for 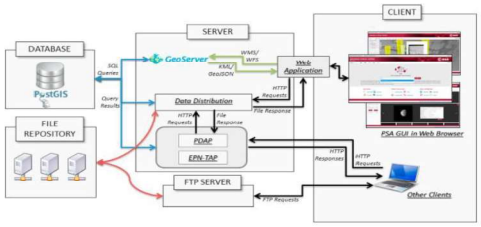 ESA space science astronomy and heliophysics missions. Its primary goal is to provide access to researchers to the science data generated on the different planetary missions. The access is made through a web interface in which the user can search for metadata and download the required products. A command line interface is also available. They are trying to increase interoperability by adding recognised science specific protocols. The PSA is based on a modular

	and flexible 3-tier architecture. The storage layer hosts a data repository of around 45TB (in 2018) and a database.
ESA Historical Archives	The Historic Archives of the European Union are hosting historical information from ESA as they reached an agreement to receive the data in 1989. Since then, more than 50000 original papers have been preserved in the HAEU. In 2018, ESA embarked in a program to open these archives to the public through a digital transformation program. As stated the Historical Archives are focused in documents, photographs, audio and video recordings, rather than data.
ESAC Science Data Centre	The ESAC Science Data Centre (ESDC), located at the European Space Astronomy Centre (ESAC), Madrid, Spain, provides services and tools to access and retrieve observations and data from ESA's space science missions (astronomy, planetary science and heliophysics). The majority of ESA's space science mission archives are developed and maintained by the ESDC, in coordination with the science operations centres, the instrument teams and the consortia of the various missions. An example is shown previously with the PSA. This data centre hosts the archives from different missions and programs like XMM-Newton, GAIA, ESA Hubble or EXOSAT to mention just some. In 2012 they defined the ESAC Science Archives based on three pillars: • Enable maximum science exploitation; • Enable efficient long-term preservation; • Enable cost-effective archive production by integration of projects. All ESAC Science Archives are based on a 3-tier concept: data layer for the data storage and metadata modelling in the database, application layer to separate the data from its presentation and interface layer either from GUI or script based.
NASA	The Planetary Data System (PDS) is a long-term archive of digital data products returned from NASA's planetary missions, and from other kinds of flight and ground-based data acquisitions, including laboratory experiments. The PDS is a federation of nine teams geographically distributed. Each of them are in charge of six science discipline nodes, two of support activities and the last one of the project management. They have published a roadmap about the strategy they are following on building the archive and supported services to researchers. However no technical information about it has been found. Another example is the HEASARC Archive Research Center. It is the primary archive for high-energy astronomy orbital missions observing at extreme ultraviolet, X-ray and gamma-ray wavelengths. Several of these missions, including Chandra, Fermi and XMM-Newton, fall under NASA's **Physics of the Cosmos (PCOS)** theme. They started using

	CDROM jukeboxes to store the data but they moved to RAID disks to enhance the accessibility. It is a much smaller that ESA/ESRIN archive, storing 86TB in 2017 but having higher access ratios (12TB/month). EOSDIS is the NASA's Earth Observing System Data and Information System. It is designed as a distributed system, with major facilities at NASA's Distributed Active Archive Centers (DAACs) located throughout the United States. These institutions are custodians of EOS mission data and ensure that data will be easily accessible to users. EOSDIS DAACs process, archive, document, and distribute data from NASA's past and current Earth-observing satellites and field measurement programs. It is more than a preservation archive but a full access to the data generated in many different ways. It is kind of similar to Copernicus ecosystem. As such, it is managed with a Content Management System (CMS). One of these DAACs is the LP DAAC (Land Processes DAAC). The LP DAAC processes, archives, and distributes land data products to hundreds of thousands of users in the earth science community. **Process**: Raw data collected from specific satellite sensors, such as ASTER aboard NASA's Terra satellite, are received and processed into a readable and interpretable format at the LP DAAC, while other data undergo processing in other facilities around the country before arriving to the LP DAAC to be archived and distributed to the public. **Archive**: The LP DAAC continually archives a wide variety of land remote sensing data products collected by sensors onboard satellites, aircraft, and the International Space Station (ISS). The archive currently totals more than 3.5 petabytes of data, the equivalent of listening to 800 million songs, and distributes data to over 200,000 global users. **Distribute**: All data products in the archive are distributed free of charge through NASA Earthdata Search and USGS EarthExplorer search and download clients. The LP DAAC supports tools and services, like the Application for Extracting and Exploring Analysis Ready Samples (AppEEARS), which allows users to transform and visualize data before download while offering enhanced subsetting and reprojecting capabilities. At the time of writing this report no additional information is found but we requested some from LPDAAC directly. NASA is showing off big interest on providing high-value services through the use of AI/ML.
CERN	CERN was one of the organizations leading the Archiver project for developing archiving and preservation cloud services for the research community. The solutions selected on this project were validated with CERN data but there is no information about CERN adoption of any of them. CERN mentions that they preserve 420PB of scientific data, photos, minutes, memoranda, etc. They are founder member of Data Preservation in High Energy Physics collaboration group. Each group of

	experiments used to preserve their data with their resources but with the knowledge gained in Archiver, they claim to have unified preservation services but also higher added-value services. BaBar experiment uses T10000 tapes to keep their 800 TB of raw data plus 6PB of processed data. In addition, in 2012 they copied the entire archive to a different site. The copy took 1 year to finish.
NOAA NCEI -	NCEI (National Centers for Environmental Information) is USA's leading authority for environmental data, and manage one of the largest archives of atmospheric, coastal, geophysical, and oceanic research in the world. It also contributes to the NESDIS mission (NATIONAL ENVIRONMENTAL SATELLITE, DATA AND INFORMATION SERVICE) by developing new products and services that span the science disciplines and enable better data discovery. 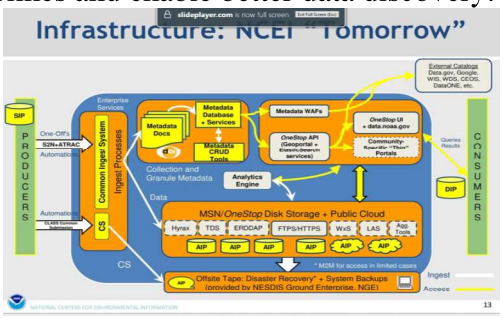 NCEI provides archive services for much of the data collected by NOAA scientists, observing systems, and research initiatives. The NCEI archive and backup copy contains more than 37 petabytes of data. Data stewarded by NCEI must include metadata about provenance, authenticity, the technical environment necessary to use the data object, what preservation actions have been taken, and what intellectual property rights apply to the data object. To provide long term preservation for digital and analog data holdings, NCEI follows guidance from many best practice communities to ensure that data are accurate and authentic for the long term. To support and ensure data preservation, NCEI routinely performs regular media migration, data integrity checks (e.g., comparing cryptographic hash values prior to and following file movement between storage media), daily and weekly media content backups, and other practices to ensure access to data objects. NCEI follows guidance from NARA for backup processes, storage media, etc. In 2017 they published their plans and designs for the evolution of their preservation infrastructure. Below you can find a mapping of OAIS structure to the real entities that perform the jobs.

Eumetsat	EUMETSAT Data Centre provides a long-term archive of data and generated products from EUMETSAT. Set up in 1995, the data centre has developed to become a state-of-the-art archive serving all EUMETSAT satellite programmes. Currently around 5PB of data volume are stored in the data centre and made available to the data centre users. 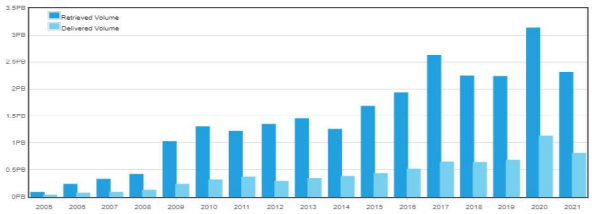 EUMETSAT manages the archive with UMARF, Unified Meteorological Archive and Retrieval Facility. UMARF is a multi-mission, archiving and retrieval facility, including on-line web users access capabilities and was developed after EUMETSAT decision in 1997.
EOSC	The European Open Science Cloud (EOSC) is an environment for hosting and processing research data to support EU science. The ambition of the European Open Science Cloud (EOSC) is to provide European researchers, innovators, companies and citizens with a federated and open multi-disciplinary environment where they can publish, find and re-use data, tools and services for research, innovation and educational purposes. This environment will operate under well-defined conditions to ensure trust and safeguard the public interest. The EOSC enables a step change across scientific communities and research infrastructures towards • Seamless access; • FAIR (Findability, Accessibility, Interoperability and Reusability) management; • Reliable reuse of research data and all other digital objects produced along the research life cycle (e.g. methods, software and publications).

RECOMMENDATIONS AND WAY FORWARD

Key Findings

- There are solutions to keep using current storage technologies, mainly tape libraries, with new software to handle the limitations of the current solution in place. No need to change the entire infrastructure because of HSM end of life.
- Hardware providers are developing hot, warm and cold storage solutions that can complement the current infrastructure and, defining a roadmap, could substitute it in the long-term.

- Some WORM storage options are already available or being developed that are based on non-magnetic media, like optical disks or crystal based. There are some that could help on the long-term perspective without needing to migrate as the media lasts for a longer time.
- New disruptive storage technologies are on research but, for instance, DNA or crystal storage are not ready yet to be commercialized and it is not seen a short-term perspective to have them available.
- More and more, organizations are moving to cloud-based services because of their flexibility and cost effectiveness in the short-term. However, the step to cloud needs to be planned because long-term costs and vendor locking could be an issue in the future.
- Complementary to pure cloud solutions, the hybrid-cloud approach is in many hardware and software providers. For organizations with already existing infrastructure is interesting to see how a mix between public access on the cloud plus cold on-premise storage would benefit in the long-term to address pure-cloud flaws.
- The preservation software market is very active and growing. In addition to the classic providers in this area, some new players are showing interesting solutions also with a very open-minded approach to avoid extra costs or cloud locking. SaaS is a trend on this market but private cloud solutions are also available.
- Security and sustainability are also worries in this market. Many providers are showing some environmentally friendly messages but it is not clear how much are they real or just green-washing. In terms of security, mainly if moving to cloud services, the security standards mush be enforced. ISO27001 and others from the family must be available.
- OAIS is the standard for preservation which is recommended to be followed. In addition, other standards could improve archive operations in different aspects.
- Big cloud providers are charging for requests or operations what make almost impossible to make a budget in advance.
- Some European cloud providers are not charging for the recovery or the model is more manageable.

Sustainability

IT systems and service providers are showing an increasing interest in offering sustainable alternatives to perform business as usual. In most of the examples shown in this report, their websites indicate that they are aware of this issue.

For the particular case of the ESA/ESRIN Archive and since its main function is the digital preservation of data in the long term, the use of storage media is a key element to be taken into account for the reduction of energy consumption and therefore of emissions. Tapes, optical disks or digital films are the best solutions on this issue. The figures shown in the following table are based in publicly available information.

Product	Idle Energy Consumption	Max Energy Consumption	Cooling requirements
Oracle StorageTek SL8500	958 w	1358 w	3,271 BTU/h
Hitachi G350	No data	226 w	743 BTU/h
Sun Server X3-2 / X6-2	No data	229 w (*)	781 BTU/h (*)
Spectra T950	190 w (**)	1250 w	4,347 BTU/h
IBM TS4500 L25 7 TS1160 + 3 LTO8 drives	580 w	903 w	2328 BTU/h
SpectraLogic TFinity	No data	No data	No data
Sun Server X7-2	No data	229 w	781 BTU/h
Hitachi Vantara HCP S31 32PB	No data	198720 w	678062 BTU/h

(*) Based on the available data from Sun X7-2
(**) Based on the information shown in the library itself. Average consumption

These data are based on public information that will need to be checked thoroughly at the time of purchase, if decided so, of new elements. If it is finally decided to change the tape library for another model, it would be appropriate to request this data in the proposal and evaluate them accordingly.

In case of cloud providers, most of them explain in their web their sustainable approach to their services.

- Google Cloud claim its carbon-neutrality since 2007 and shows a strategy to use carbon-free energy in 2030. It also has tools to estimate emissions of their cloud services.
- Amazon Web Services is also focused on efficiency to become a 100% renewable energy service in 2025. As part of Amazon, AWS is committed to achieve net-zero carbon by 2040. It has a Customer Carbon Footprint Tool to calculate carbon emissions from the usage made on its cloud.
- Microsoft Azure services has a clear message around sustainability. They are trying to use 100% renewable energy in 2025, water positive in 2030 and zero-waste by 2030. In 2018 they published a white paper to demonstrate the carbon benefits of cloud computing against on-premises installations.
- OVH is claiming itself as the pioneer of sustainable cloud services. OVH shows in the website good values of PUE (Power usage effectiveness), WUE (Water usage effectiveness) and CUE (Carbon usage effectiveness) against the rest of the industry.

It is interesting to mention the Uptime Institute that is an unbiased advisory organization focused on improving the performance, efficiency, and reliability of business critical

infrastructure through innovation, collaboration, and independent performance certifications. They have published reports about creating a sustainable strategy, reduce energy footprint and tackle greenhouse gases among others.

Costs

Another important factor when defining the strategy for the evolution of the archive in the long term are the costs associated with the operation of the systems. Since the technologies analyzed in this report cover both on-premise and cloud installations, an estimate of the cost models available in the different solutions are presented below.

Product	**Costs model**
Versity	Subscription pricing model fixed and taking into account the estimated growth. They claimed that they don't use capacity-based pricing.
Alluxio	It has a free Community Edition and a paid Enterprise Edition with much more features. Subscription pricing for Alluxio Enterprise Edition is based on the number of nodes in the environment.
Atempo	No public info about the pricing model.
minIO	Open-source license and two levels of enterprise licensing based on the storage managed with the software.
nageruHive	Cloud SaaS model with packages of services including processing, storage and support oriented to Archive. Also offers on-premise subscription models based in the processing and support levels.
Arkivum	No public info about the model.
Google Cloud	Basic computing power based on the type of servers. Cloud Storage: 4 tiers paying per GB/month with minimum storage duration depending on the tier. Data Operations: changes or retrievals of information. 3 classes of operations with different pricing depending on the tier. There are tasks even inside Google Cloud that consume operations (more than one in some cases).
AWS	Basic computing power based on the type of servers. Cloud Storage: 2 tiers paying per GB/month; Data Transfer: Cost per GB; Requests: depending on the operations, costs per requests are charged.
Microsoft Azure	Basic computing power based on the type of servers. Cloud Storage: 4 tiers paying per GB/month depending on the redundancy. Data Transfer and transactions: depending on the tier, charges per operation exists.
Cloudferro	Basic computing power based on the type of servers. Cloud Storage: 1 tier paying per GB/month.
Safespring	Basic computing power based on the type of servers. Cloud Storage: 2 tier paying per GB/month; Data egress: per GB/month from 1TB free.
OVH	Basic computing power based on the type of servers. Cloud Storage: 2 tier paying per GB/month; Data ingress: free or per GB/month depending on the tier; Data egress: per GB/month.
Oracle Cloud	Basic computing power based on the type of servers.

	Cloud Storage: 3 tier paying per GB/month; Requests: depending on the operations, costs per requests are charged.
Qstar	They are pricing the software based on the capacity of the storage they manage through it. They mentioned 1PB storage license is around 35,000€ with 5 years maintenance included. They have special pricing models when no tender is needed.

Recommendations for ESA archive evolution (as example)

After analyzing the market, the current situation of the ESA/ESRIN Archive and the future needs, the best solution for updating the Archive is a hybrid cloud model, which integrates the current tape storage infrastructure with a flexible model that facilitates access and improves its online capabilities. The work-life of the StorageTek SL8500 tape library that ESA/ESRIN currently has is sufficient to cover the needs of the coming years and therefore, the recommendation is to continue using it for some years now. Secondly, and since Oracle has stopped supporting the HSM solution that currently uses the Archive and also counting the limitations it has, it is clear that this solution must be replaced, either by a new HSM software or similar or by an integration into cloud services.

Moving certain services to cloud providers oriented to Digital Archive services will give power and flexibility to archive operations and improve efficiency aspects. In addition, there are solutions that allow integrating these cloud services with on-premise infrastructure and therefore reusing the existing library.

Given that OAIS is the applied standard for the preservation services offered by the ESA/ESRIN Archive, and together with the cloud model, ingest and access services could be moved to the cloud provider. In addition, control and administration services could also move to the cloud and complement them with automatic collection, artificial intelligence and operational control services.

Conclusions

This white paper summarised the results of the activities performed at ESA, to evolve the ESRIN Archive to the next technology step. It described the functionalities and the architecture of the archive, the data archived and the different challenges driving its evolution. The paper provided not exhaustive information about archiving and preservation solutions that are available on the market or implemented by other agencies, and which can be used for potential evolution of own archives.