

# The Need for Preservation Aware Storage \*

## A Position Paper

Michael Factor  
IBM Haifa Research Lab  
Haifa University Campus  
Mt Carmel, Haifa 31905, Israel  
factor@il.ibm.com

Dalit Naor  
IBM Haifa Research Lab  
Haifa University Campus  
Mt Carmel, Haifa 31905, Israel  
dalit@il.ibm.com

Simona  
Rabinovici-Cohen  
IBM Haifa Research Lab  
Haifa University Campus  
Mt Carmel, Haifa 31905, Israel  
simona@il.ibm.com

Leeat Ramati  
IBM Haifa Research Lab  
Haifa University Campus  
Mt Carmel, Haifa 31905, Israel  
leeat@il.ibm.com

Petra Reshef  
IBM Haifa Research Lab  
Haifa University Campus  
Mt Carmel, Haifa 31905, Israel  
petra@il.ibm.com

Julian Satran  
IBM Haifa Research Lab  
Haifa University Campus  
Mt Carmel, Haifa 31905, Israel  
satran@il.ibm.com

### ABSTRACT

*Digital Preservation* deals with ensuring that digital data stored today can be read *and interpreted* tens or hundreds of years from now. At the heart of any solution to the preservation problem lies a storage component. This paper characterizes the requirements for such a component, defines its desirable properties and presents the need for preservation-aware storage systems. Our research is conducted as part of *CASPAR*, a new European Union (EU) integrated project on the preservation of data for very long periods of time. The position presented was developed while designing the storage foundation for the *CASPAR* software framework.

### 1. INTRODUCTION

One of the most challenging problems faced by technologists today is the *Digital Preservation problem*, namely, how to ensure that digital data being stored today can be read and interpreted many years (tens or hundreds years) from now. The challenge is twofold: the *bit preservation* aspect deals with the ability to read the bits from the media, whereas *logical or information preservation* targets the ability to interpret and understand the information stored in those bits.

There are numerous examples of applications and industries that require long-term preservation, partially driven by com-

---

\*Work partially supported by European Community under the Information Society Technologies (IST) programme of the 6th FP for RTD - project *CASPAR* contract IST-033572. The authors are solely responsible for the content of this paper. It does not represent the opinion of the European Community, and the European Community is not responsible for any use that might be made of data appearing therein.

pliance and regulation. These include: medical retention regulations in the healthcare industry; pharmaceutical companies that need to preserve their research, development and filing application records for decades; aircraft design records in the aerospace industry; satellite data accumulated by space missions (e.g., of NASA and ESA); cultural and heritage records, and many more.

Due to the constant growth in the amount of long-lived data, as well as new compliance legislations that are emerging worldwide, this problem has been gaining increasing attention from the research community, government and semi-government agencies. The Open Archival Information System (OAIS) ISO standard [7, 8], developed by NASA, is one example. It targets the preservation of knowledge rather than the preservation of bits, and provides a set of concepts and guidelines for building preservation systems. At the heart of every OAIS preservation system lies a *preservation-storage component*. This is the portion of the system that manages the long-term storage and maintenance of digital material entrusted to the OAIS.

A storage system is deemed *preservation-aware* if it supports preservation applications. *OAIS-based storage* is a specific type of preservation-aware storage, which is based on OAIS notions, functions and information model. This paper summarizes findings that were developed in the initial phase of a research-related project in the area of Storage Systems for Data Preservation. The project is being conducted as part of *CASPAR*, a new European Union (EU) integrated project on the preservation of data for very long periods of time [3]. The objective of *CASPAR* is to build a framework to support the end-to-end preservation "lifecycle" for scientific, artistic and cultural information based on existing and emerging standards, most notably the OAIS model. As part of *CASPAR*, we plan to implement a new storage concept, called *Preservation DataStores*, for OAIS-based storage supporting new functionalities and extensions specific for preservation. Preservation DataStores are object-based and utilize open standards (e.g., XAM and OSD [17, 16]).

In the rest of the paper we review the main OAIS concepts

that have implications on the storage system, and provide the characteristics of data and data lifecycle in a preservation system. We then analyze the aspects that differentiate a preservation-aware storage system from other (non-aware) archival systems. In particular, we investigate what constitutes a 'good' preservation-aware storage system. We argue that in order to better preserve data and understandability for long periods, a new type of storage must emerge that will take preservation considerations into account. We call this new type of storage *Preservation Aware Storage*.

## 1.1 Related Work

The storage aspect of digital preservation has been attracting more attention lately (for example [1, 2, 10, 18]), and this trend is likely to increase. However, so far these studies have all concentrated primarily on the bit preservation aspect of the problem, and suggested how traditional storage systems can be used to address the new challenges posed by long-term preservation.

The papers [1, 2] compare traditional (shorter-term) storage environments to long-term preservation environments, focusing on the *bit preservation* aspect of the problem. In [1], Baker et. al contrast a transaction processing enterprise system, which is designed to meet customer-specific or application-specific needs for performance, with a large-scale library archival system which requires long term availability and higher reliability; the differences in cost considerations are also compared. Specific threats to long-term systems are listed, including storage failure, disk errors, outdated media, software and hardware. The paper recognizes the difficulty in preserving the 'context' and structure of the data, but does not talk about its implications on the storage system. Challenges in building long-term archival storage systems are outlined, such as supporting high level of replication and scalability. The paper is based on a case study at the British Library, thus focusing on *paper archives*. In [2] the authors further concentrate on the *reliability* aspect of long-term storage systems, taking into account faults due to humans and organizations in addition to hardware and software. Similarly to [1], it focuses on the reliability of the data bits.

In [18], Storer et. al study the new *security threats* that endanger the secrecy, availability and integrity of a long-term archive. The paper considers security problems such as secrecy and user authenticity, as well as slow attacks (on keys and data) which could not be possible otherwise. The paper provides a comparative analysis of a number of existing systems, and shows how today's systems deal or do not deal with these new security threats. POTSHARDS [5] is a prototype system that addresses some of these security issues.

Over the past years, data grid technology [10, 9] has been used as a foundation technology for providing a persistent archive in the context of preservation projects. Many of these projects were commissioned by NARA, the National Archives and Records Administration, and supported by the Library of Congress and NSF. As a persistent archive, it manages the retention of a digital record as well as the context that describes the origin, relevance and authenticity of the record in a distributed and saleable manner. Data grid is a middleware software mechanism which builds on today's

underlying storage systems.

The white papers [15, 14] provide further motivation to the general data preservation problem as well as requirements on the overall preservation system. The National Digital Information Infrastructure and Preservation Program [13](NDIIPP) is a collaborative initiative run by the Library of Congress. Its goal is to develop a national strategy for digital preservation. NDIIPP develops a rich set of formats appropriate for preservation and provides a high level architecture. The architecture emphasizes the need to support federation of archives; it does not refer to specific storage aspects or functionalities of the system, nor it is compared with short term archives.

## 2. THE OAIS MODEL

The *Open Archival Information System (OAIS)* [7, 8], an ISO standard since 2003 (ISO 14721:2003 OAIS), specifies how digital assets should be preserved for a community of users – from the moment digital material is ingested into the digital storage area, through subsequent preservation strategies, to the creation of a dissemination package for the end user. OAIS is a high-level reference model and as such is flexible enough to use in a wide variety of environments; it expects more detailed steps and workflows to be developed by the implementing institution. Below is a very brief summary of the concepts that are most relevant to the discussions in this paper.

Of particular interest is the OAIS *Information Model*. This provides a high-level description of the information objects managed by the archive. An **AIP - Archival Information Package** is the information package that is stored and preserved by the OAIS. It consists of the preserved information, called the *content information*, accompanied by a complete set of metadata, as depicted in Figure 1.

Note that the accompanied 'metadata' is well compartmentalized. It consists of the *representation information* that is required to render and interpret the object intelligible to its designated community. This might include information regarding the hardware and software environment needed to view the content data object or a specification of the data format. The other metadata, called the *Preservation Description Information (PDI)* is broken down by OAIS into four well-defined sections:

**Reference information** - a unique and persistent identifier of the content information both within and outside the OAIS (e.g., UUID).

**Provenance information** - the history and origin of the archived object.

**Context information** - the relationship to other objects (e.g., the hierarchical structure of a digital archive). For example, this may be a set of PDF documents, each representing a chapter of a book; or a collection of objects representing digitized images of an art collection.

**Fixity information** - a demonstration of authenticity, such as checksums and cryptographic hashes, digital signatures and watermarks.

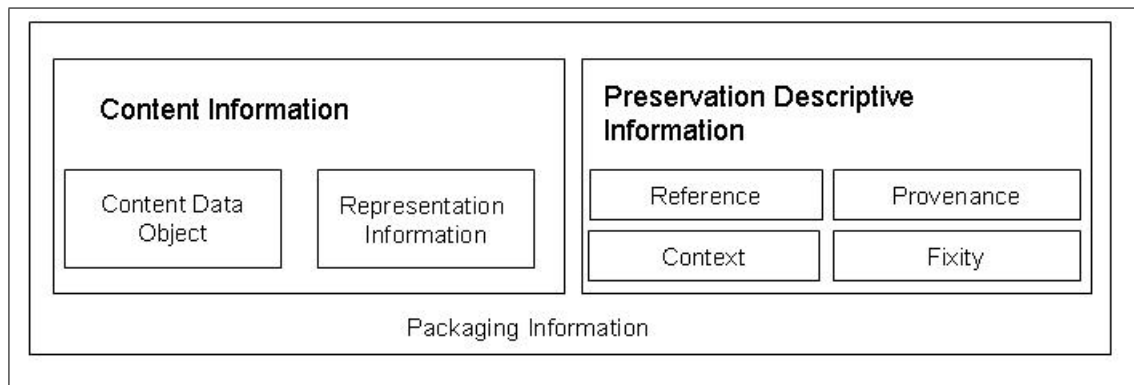


Figure 1: AIP - Archival Information Package as defined by the OAIS Information Model

**Data Migration** is another aspect of the OAIS functional model that has bearing on the storage system. Migration is the act of moving data from one system to another due to a change. This may be triggered by a variety of reasons, for example the decay of the storage media, obsolescence of hardware and software, or a change in the copyright or external environment (e.g., organization). The OAIS reference model identifies four primary digital migration types:

1. Refreshment - bit to bit copy of the entire media onto newer media of the same type.
2. Replication - copying data onto newer media which is not necessarily of the same type.
3. Repackaging - copying data while changing the placement of the components within a data object.
4. Transformation - copying data while performing format change on the data.

## 2.1 OAIS-based Storage

Today, more and more storage systems offload advanced functionality and structure-awareness to the storage layer. Functions that were traditionally carried out by the application or the operating system are gradually becoming integral parts of an 'intelligent storage system'. Object-store devices (OSDs)[4], for example, offload space allocation and security to the storage device. The eXtensible Access Method (XAM) storage system [17], represents a newly emerging storage initiative carried out by the Storage Networking Industry Association (SNIA). XAM bundles together multiple pieces of data and metadata for access under a common globally unique external name (XUID). Functions such as bit-to-bit data migration, block-level data integrity, and even encryption, are carried out by advanced, intelligent disks and tapes. Some systems (e.g. provenance-aware storage system (PASS) [12]) already track the provenance of data at the storage level rather than storing it in a standalone database<sup>1</sup>.

Today, the OAIS model assumes traditional storage as its underlying archival component and relies on other components of the data management system to provide higher-level

<sup>1</sup>The paper [12] discusses the advantages of managing provenance at the storage system; PASS is prototyped at the file system level.

functions such as packaging of the data or the computation of provenance and fixity. Under this assumption, the storage system merely stores the bits. As more intelligent archival solutions emerge, it will become possible to incorporate some of the concepts and processes specifically defined by OAIS into the intelligent storage component, resulting in *OAIS-based* storage. OAIS-based storage can provide a more robust infrastructure for preservation systems. The next section provides specific examples for such extensions.

## 3. PRESERVATION VS. ARCHIVING: ARE THEY DIFFERENT?

Digital archiving refers to the ability to safely store and access digital data. Digital preservation is a special case of digital archiving, where the lifetime of the stored data exceeds the lifetime of the program/format used to interpret the data as well as the lifetime of the media that stores the bits (see the results reported in [14] on the question "What does Long-Term Mean?"). It is therefore natural to ask: What are the aspects that differentiate a data preservation system from a traditional archival system that stores and accesses data for relatively shorter periods? To address this question, we first characterize the data that is stored in a preservation system and the unique processes it undergoes during the long preservation life-cycle.

**Characterization** Digital data in preservation systems has the following characteristics:

- The data is read-only.
- Most content data is *cold* and is rarely accessed during its lifetime. New metadata can be added later on.
- The data is heterogeneous in type, size and value (some data are more important than others).
- The data quantity is too large for on-line storage.

**Processes** The following preservation-related processes may affect the data:

- The process of packaging the content into an AIP requires the metadata to be stored (sometimes this will be a substantial amount of additional data) together with the data in order to interpret

it. Some of it is referential data, representing either relationships within various data objects in the system or references to representation information.

- The additional information needed for interpretation (the object's representation information) may reside elsewhere, outside the system that stores the actual data, and may also evolve over time.
- The data must be migrated and possibly transformed to support the obsolescence of formats, hardware or software.
- The data is likely to be accessed by a system that is not the same as the system which originally ingested and stored it.
- The likelihood that some of the data will be lost or be corrupted over time is high, especially when data is migrated and transformed repeatedly.

Given the characterizations above, what are the special requirements of a preservation-aware system? What should a preservation-aware system support beyond the services supported today by a short-term archive system? The following are some specific examples.

1. Extend the data migration process to handle media degradation as well as replication, repackaging and transformation (due to format conversion and software migration), and to possibly add new metadata. This requires the storage to react to events that are external to the storage, for example events triggered by a preservation planning component. It also extends the notion of provenance to include the event of migration to another data format.
2. Maintain an extended notion of object integrity that goes beyond bit integrity, e.g., referential integrity.
  - Ensure that links point to existing/valid locations. This requires an awareness of certain metadata fields that represent links (both internally to the system and externally).
  - Ensure linkage updates. Representation information links into registries may have to be updated due to migration of the registries.
  - Support partial integrity of the information object, as well as information on parts that are intact and those that are corrupted.
  - Package data at the lower media level so it can be self-contained and ensure that all the object's components are co-located physically.
3. Support long term aspects of *data integrity* and *encryption*.
  - The algorithms employed to compute data integrity and encryption may need to change (due to loss of security<sup>2</sup>, or when better algorithms are developed), along with the keys. If these changes are anticipated, the updates can be done during the migration process.

<sup>2</sup>For example, the recent findings that require replacing SHA-1 with SHA-256

- Export the cryptographic mechanisms that were used to compute integrity and encryption, to support access to the data by a different system. This requires building a representation information for the fixity.
4. Ensure *readability* of the data by a different system in the future by developing and supporting global self-described formats for disks and tapes.
  5. Support a *graceful loss* of data. Some portions of the data are likely to be lost or corrupted over time. A good preservation system must ensure that if some data is lost, the other data in the system that is still intact can be read or interpreted.

## 4. DISCUSSION

We laid the groundwork for an advanced type of storage that is *OAIS-based and preservation-aware*. The storage must be aware of a core set of metadata in order to guarantee referential integrity and packaging of the data, both at the logical and physical layer. The system must support data integrity over time and should offer a rich set of functions when data is migrated. We further argue that the implementation of this concept is better optimized if it is architected as part of the storage device.

### 4.1 Answers to Some FAQs

*Why is data preservation a storage problem?* One may claim that all of the above can be addressed by an adequate preservation application that is built on top and uses an archival system that is agnostic to the preservation needs and only stores the bits. Although a valid approach, this is very much against the current trend of offloading functionality (such as encryption, mining, analytics and policies) to the storage itself. Moreover, there are some examples of specific functionality that will clearly benefit if executed close to the data and not above it. The examples are: (i) Integrity/fixity. Instead of moving data to the application, the functionality can be performed near the data and only a short answer (either Boolean 'valid/not valid' or short numeric hash value) is transferred to the application. (ii) Provenance. Since some of the events that affect provenance (e.g., chain of ownership and the transformations) are already performed and managed by the storage system itself, it is beneficial to have the storage layer maintain provenance. (iii) Functions coupled with migration. These include data transformation, re-computation of integrity and encryption, updates of external links to registries, management of copies, and more. Since migration will be performed by the storage layer, integrating these functionalities into the storage systems and not the application above it will result in substantial optimizations.

*Why can't preservation be treated as yet another metadata?* After all, it's all about metadata, which is already treated by archiving systems that are not necessarily preservation aware. The simple answer is that a preservation system must be aware of the semantics of some important core sets of metadata. (i) The representation information link must be collocated (logically and physically) with the data as long as the data is preserved. This link is likely to change with time. This metadata is at the core of preservation and its integrity

and persistency must be treated accordingly. (ii) Metadata fields that are maintained by the archival system and of which only the storage system is aware, must be treated with appropriate preservation mechanisms. The best example is integrity/fixity - if calculated by the device, only the device can update it.

*Isn't it only a question of universal formats?* Universal formats for data on media that have packaging capabilities and are self-contained play an important role in an overall preservation solution, but do not provide a complete answer. (i) To support future migrations and be self-contained, data must be packaged at both the physical and logical levels; self contained formats and packaging formats guarantee this only at the logical level. (ii) Other functions mentioned above (e.g., bit integrity and referential integrity, provenance) are not addressed by universal formats.

## 4.2 Future Work and Acknowledgments

A Preservation DataStore is a new OAIS-based preservation-aware storage that is being implemented as a research-prototype by IBM. Our initial realization of a Preservation DataStore will primarily utilize OAIS and eXtensible Access Method (XAM). It will also build upon the Object Store Devices (OSD) standard, as well as the filesystem interface. The Preservation DataStore is being designed as an infrastructure component of the CASPAR framework, which will demonstrate its validity in preserving cultural, artistic and scientific knowledge. To bring global dimension to its work, CASPAR also plans active collaboration and teaming with relevant digital preservation initiatives outside the EU, with national and international projects such as Chronopolis, INTERPARES [11, 6] and with NARA (US National Archives and Records Administration).

*Acknowledgments:*. Foremost, we wish to thank the general data preservation community, our CASPAR partners, and in particular David Giarretta (CASPAR's coordinator) who exposed us to this problem and made us think about the unique storage aspect of it. We had valuable discussions on this subject with Reagan Moore and the SRB team in San Diego SC. Finally, thanks to our colleagues from the storage department at the IBM Almaden Research Center for asking the right questions.

## 5. REFERENCES

- [1] M. Baker, K. Keeton, and S. Martin. Why traditional storage systems don't help us save stuff forever. Technical Report 2005-120, HP Laboratories Palo Alto, June 2005.
- [2] M. Baker, M. Shah, D. Rosenthal, M. Roussopoulos, P. Maniatis, TJ Giuli, and P. Bungale. A fresh look at the reliability of long-term digital storage. In *Proc. European Systems Conference (EuroSys)*, April 2006.
- [3] CASPAR - Cultural, Artistic and Scientific knowledge for Preservation, Access and Retrieval. <http://www.casparpreserves.eu/>.
- [4] Michael Factor, Kalman Meth, Dalit Naor, Ohad Rodeh, and Julian Satran. Object storage: The future building block for storage systems. a position paper. In *Local to Global Data Interoperability - Challenges and Technologies, Sardinia Italy*, pages 119–123, June 2005.
- [5] Kevin Greenan, Mark Storer, Ethan L. Miller, and Carlos Maltzahn. Potshards: Storing data for the long-term without encryption. In *SISW '05: Proceedings of the Third IEEE International Security in Storage Workshop (SISW'05)*, pages 12–20, Washington, DC, USA, 2005. IEEE Computer Society.
- [6] InterPARES - The International Research on Permanent Authentic Records in Electronic Systems. <http://www.interpares.org/>.
- [7] ISO 14721:2003, Blue Book. Issue 1. *CCSDS 650.0-R-2: Reference Model for an Open Archival Information System (OAIS)*, 2002.
- [8] Brian F. Lavoie. *The Open Archival Information System Reference Model: Introductory Guide*. DPC Technology Watch Report 04-01, 2004.
- [9] Reagan W. Moore, Joseph F. JaJa, and Robert Chadduck. Mitigating risk of data loss in preservation environments. In *MSST '05: 22nd IEEE / 13th NASA Goddard Conference on Mass Storage Systems and Technologies (MSST'05)*, pages 39–48, Washington, DC, USA, 2005. IEEE Computer Society.
- [10] Reagan W. Moore and Richard Marciano. Building preservation environments. In *JCDL '05: Proceedings of the 5th ACM/IEEE-CS joint conference on Digital libraries*, pages 424–424, New York, NY, USA, 2005. ACM Press.
- [11] R.W. Moore, F. Berman, D. Middleton, B. Schottlaender, J. JaJa, and A. Rajasekar. Chronopolis - federated digital preservation across time and space. In *Local to Global Data Interoperability - Challenges and Technologies, Sardinia, Italy*, pages 171 – 176, June 2005.
- [12] Kiran-Kumar Muniswamy-Reddy, David A. Holland, Uri Braun, and Margo Seltzer. Provenance-aware storage systems. In *Proceedings of the 2006 USENIX Annual Technical Conference*, June 2006.
- [13] The National Digital Information Infrastructure and Preservation Program - The Library of Congress. <http://www.digitalpreservation.gov/>.
- [14] Michael Peterson. *The Coming Archive Crisis, a white paper*. SNIA Data Management Forum, Nov. 2006.
- [15] Andres Rodriguez. *Preserving the Last Copy: Building a Long-Term Digital Archive, A white paper*. Archivas, Inc., 2004.
- [16] SNIA - Storage Networking Industry Association. *OSD: Object Based Storage Devices Technical Work Group*. [http://www.snia.org/tech\\_activities/workgroups/osd/](http://www.snia.org/tech_activities/workgroups/osd/).
- [17] SNIA - Storage Networking Industry Association, Data Management Group. *XAM (Extensible Access Method)*. <http://www.snia-dmf.org/xam/>.
- [18] Mark W. Storer, Kevin M. Greenan, and Ethan L. Miller. Long-term threats to secure archives. In *Proceedings of the 2nd International Workshop on Storage Security and Survivability ( StorageSS 2006)*, Alexandria, VA, pages 9–16, October 2006.