

## *Full Length Research*

# LONG-TERM DOCUMENT PRESERVATION SYSTEM

Evans Nyanyu Makwae

Archivist iii Judicial Service commission (JSC), Kenya.E-mail: nyanyuevans@gmail.com/nyanyu2004@yahoo.com

Accepted 12 October 2017

---

**This paper analyzes the requirements and describes a system designed for retaining records and ensuring their legibility, interpretability, availability, and provable authenticity over long periods of time. In general, information preservation is accomplished not by any one single technique, but by avoiding all of the many possible events that might cause loss. The focus of this system is on preservation in the 10 to 100 year time span—a long enough period such that many difficult problems are known and can be addressed, but not unimaginable in terms of the longevity of computer systems and technology. The general approach focuses on eliminating single points of failure - single elements whose failure would cause information loss - combined with active detection and repair in the event of failure. Techniques employed include secret sharing, aggressive "preemptive" format conversion, metadata acquisition, active monitoring, and using standard Internet storage services in a novel way.**

**Key words:** long-term preservation, archiving, archives, format conversion, metadata acquisition, Internet storage services.

---

**Cite This Article As:** Makwae EN (2017). LONG-TERM DOCUMENT PRESERVATION SYSTEM. *Inter. J. Acad. Lib. Info. Sci.* 5(8): 239-246

## INTRODUCTION

Recorded history is a history of documents and records; individuals and groups keep information in recorded form for later retrieval. The documents and records are descriptions of events, financial transactions, government activities, stories, opinions and so forth. Often a single document is used both for managing a transaction as well as recording the fact of the transaction for later review. In the last 50 years, computer systems and information automation have moved work processes and records online. Electronic records, however, do not have the same longevity properties as physical documents, and many unsolved problems remain.

It is important, then, to be able to retain such records, and to ensure their legibility, interpretability, availability, and provable authenticity over very long periods of time. In general, information preservation is accomplished not by any one single technique, but by avoiding all of the

many possible unfortunate events that might interfere with preservation or cause loss. A system for long-term storage and archiving must guard against all of the problems which might cause inadvertent or unplanned loss.

The desired lifetime of personal and business records ranges from weeks to months to years and, in some cases, to decades. Thus, it is important to protect records against events that are foreseeable in those time frames. However, the technical means for information preservation for centuries or millennia are hard to imagine. The researcher therefore defines "long-term" as a 10 to 100 year span.

The researcher begins description of a system for long-term preservation with a survey of the variety of events which may affect documents over time, and he describes the approach he takes to protect documents against such

events. The second part of the paper gives an overview of a system designed and prototyped which combines these approaches to provide for guaranteed long-term archives of documents and information. Finally, it's concluded with an analysis of the system and a comparison to related work.

## RATIONALE

The purpose of this paper is to analyze the requirements and describe a system designed for retaining records and ensuring their legibility, interpretability, availability, and provable authenticity over long periods of time. The paper assesses causes of information loss and how to avoid them, media failure, site failure, organizational failure, software failure, malicious modification or destruction, loss of interpretability, loss of context, loss of guarantee of authenticity, authorization failure and loss of privacy. As no extensive research has been conducted in this area in Kenya, the study also fills the gap in the long-term document preservation system literature related to this area.

## RESEARCH OBJECTIVES

1. To assess causes of information loss and how to avoid them.
2. To analyze media failure
3. To assess site failure
4. To analyze organizational failure
5. To ascertain software failure
6. To establish malicious modification or destruction
7. To assess loss of interpretability
8. To analyze loss of context
9. To assess loss of guarantee of authenticity
10. To establish authorization failure and loss of privacy

## RESEARCH METHODS

The researcher used descriptive survey, of a system designed for retaining records and ensuring their legibility, interpretability, availability, and provable authenticity over long periods of time. Description, design and implementation of a system which provides for long-term archive and data preservation, based on the principles outlined is described in details by reference to the functional steps used to perform the major operations. However, the author contacted the storage departments in case need were felt to clarify some integrity and also to fill the gap.

## LITERATURE REVIEW

The design and implementation of a system which provides for long-term archive and data preservation, based on the objectives.

### Causes of information Loss and How to Avoid Them

The general approach taken for reducing the risk of loss of information is to acknowledge that loss is inevitable, but that overall reliability can be improved using redundancy and active correction. That is, one can attempt to design a system in which there are no "single points of failure"—single elements whose failure would cause information loss—and combine this with active detection and repair of failure. Avoiding single points of failure is a fundamental principle applied to the design of reliable systems, and is essential in a system designed to provide long-term access to archived documents. Some single points of failure are easily recognized, such as using a single physical device or single location. Other risks, such as file formats, protocols, software implementations, or encryption algorithms, are less obvious, but no less important.

### Media Failure

Any practical digital media for recording data—disk, tape, CD-ROM, and the like—will eventually fail. Failure of recording media comes sooner than many imagine. Currently the technique used for avoiding media failure is periodic media refresh—read in the digital data, check for errors using error correction techniques, and rewrite on new media. There are many variations on this theme, including RAID disks and backup strategies. One cannot rely on any single media copy for data safety (Patterson et al., 1988).

### Site Failure

A more serious cause of data loss can come when the physical location for holding the media is somehow threatened. For example, through earthquake, flood, electrical storm, or other kinds of physical disasters, entire sites holding data may be in jeopardy. Frequently the method employed for protecting against site failure is to backup the data off-site. However, this approach has limitations. Over time, backup media may fail undetectably. The offsite backup storage location itself may be at risk. The author's approach is to use distributed multi-site redundant storage, spreading information among multiple physical locations in different geographic areas, and to include periodic checks of storage integrity.

### **Organizational Failure**

Even with multi-site redundant storage, if multiple sites are controlled by a single institution, storage is at risk if the organization fails—loses funding, goes bankrupt, or decides to discontinue its business for data storage. In some cases – when an organization is managing its own multi-site storage system – organization failure isn't a risk that needs to be guarded against, because the need for long-term archives might also go away at the same time as the storage. Some systems, such as LOCKSS employ multiple storage sites, each run by independent organizations with separate financial backing (Rosenthal and Reich, 2000). The author's approach is to use multiple independent storage providers.

### **Software Failure**

Another kind of difficulty occurs when there are bugs, errors or security flaws in software. If each storage site is running the same software, there is the possibility that a single flaw in the software will put the integrity of multiple sites at risk. For this reason, the author propose using standards-based protocols for access to data storage, where different storage sites are running different implementations of the storage software, so that the integrity and reliability of stored data does not depend on the integrity and reliability of any single implementation.

### **Malicious Modification or Destruction**

For any of a number of reasons, long-term archives may be subject to attack. It is insufficient to rely on multiple copies without mechanisms for ensuring that a concerted attack on storage sites will not cause information to be destroyed or changed. The author focuses on increasing the number of sites which must be simultaneously attacked, and obscuring the identity of those sites for a particular document as a way of reducing the risk of malicious modification or destruction (Dingledine et al., 2000).

### **Loss of Interpretability**

According to Lorie, (2001) because the goal is to keep retrievable records of human activities, and not just sequences of bits, it is important to ensure that the data stored can be interpreted in the future. Thus, a long term archive needs to be concerned with the data formats used and not just bit sequences. The relatively brief history of computer based document processing is full of document formats for which interpreters are difficult to find or are unreliable. While others have imagined using "virtual machines" to store interpreters for individual file formats, these themselves may be considered "single points of failure." Rather, the author proposes storing

multiple representations of the same document at the time of creation of the archive. In addition, he also propose that as many of those renditions as possible be "archival quality." What is an "archival quality" file format? It is one for which there are multiple interpreters, written in different languages, running on different operating systems, and for which there is a specification precise enough that it is possible to build a credible interpreter based solely on reading the specification. Planning to migrate documents to new formats in the future, when the stored file formats are "almost obsolete," is not a practical solution. Delayed migration risks the possibility that, if information about a format is lost, it will be too late to migrate! Thus we emphasize early conversion, at time of archive, to one or more formats of archival quality.

### **Loss of Context**

In a large archive of data, each document must be specifically identified and its context supplied; the context includes information such as the source of the document, its intended distribution, whether it is a draft or a final report, and other kinds of contextual information that, in conjunction with the document content, turn a document into a record. Often, in day-to-day use of documents and content, this kind of contextual information is implicit—not explicitly represented and therefore at risk of being later forgotten. Thus, the author includes explicit acquisition of contextual (descriptive) metadata as an integral part of an archival record. For example, descriptive metadata for a text document might include the author, topic, creation date, distribution, and version of the work. Descriptive metadata for a photograph might include information about the subject, the date the photograph was taken, and the photographer. Metadata also may suffer from "loss of interpretability"; thus it is important that metadata be stored in a format of archival quality.

### **Loss of Guarantee of Authenticity**

Records of transactions may be subject to manipulation. For physical documents, while the technology for forgery may improve, the technology for detecting forgery also improves, to the point where it is frequently possible to be assured of the authenticity of a physical document. However, the situation is more complicated for electronic records; simple maintenance processes or media refresh may cause loss of clues about document origins or dates and interfere with processes to manage physical custody of records which might otherwise be used to determine authenticity. The author's approach to long term guarantees of authenticity is to request that the storage services that are holding individual pieces of data also act as "notary" services, maintaining timestamps of receipt, possibly validating the authenticity of the document at the time of receipt, and individually guaranteeing the storage

service's records of such facts. The author proposes that each original record have multiple notary services which individually attest to data authenticity, avoiding a single point of failure.

### **Authorization Failure**

There are situations where multiple data storage locations might still be subject to centralized regulation, legal intervention, or other events affecting archived data even though the effects (such as data change or destruction) might be contrary to the original wishes of the principal creating the archive. For this reason, it is useful to consider the possibility of ensuring that the storage repositories are held under multiple legal jurisdictions, such that a legal intervention in a small subset of locations will still not cause the archived information to be lost, modified, destroyed, or revealed in ways inconsistent with the archiver's intent.

### **Loss of Privacy**

While many individuals and organizations may have records or information that are not particularly sensitive or private, almost every organization has some records that are private. There are many different approaches to trying to guarantee privacy in storage systems, through both technical and operational means. Most of these systems, however, have a "single point of failure" because they rely on operational integrity of the storage system. In some systems, the data is stored encrypted, using a symmetric-key based encryption scheme, such as DES or AES, which uses random (or pseudorandom) encryption keys. The security of such key-based encryption relies on the inability of a polynomial time adversary to successfully recover the key and on the assumption that a particular underlying mathematical problem, such as factoring large prime numbers or determining quadratic residues, is "hard" and will remain so. Recent history has shown that algorithm security is relatively fragile (Wang et al., 2005).

Key-based encryption in a long-term archive suffers from two major drawbacks. First, the privacy of the document is based on the inability of an adversary to break the encryption within the time period during which the document must remain private. Clearly, if a document that must remain private indefinitely is archived, there is a problem, because an adversary would have an indefinite amount of time to break the key. Second, key-based encryption requires remembering an additional piece of information: the key. If the key is lost, so is the ability to decrypt the document. If the keys are stored online, they must be kept securely for the same duration as the document. Merely encrypting the key with another key does not reduce the complexity of the problem. Using key-based encryption may result in the use of non-

random keys, such as keys based on a typed password. The minimal amount of entropy possible in passwords that can also be easily remembered leads to easily successful dictionary attacks. To avoid the difficulties inherent in key-based encryption, the author turns to "secret sharing" which provides for threshold storage. An adversary who obtains fewer than "shares" is unable to learn any information about the document, no matter how much compute power is applied. Threshold secret sharing does not require an encryption key, resulting in one less piece of data that must be remembered over the lifetime of the document. (Shamir, 1979)

## **RESULTS AND DISCUSSION**

### **System for Long-Term Archives**

The author now discusses the design and implementation of a system which provides for long-term archive and data preservation, based on the principles outlined above. The system design is discussed by reference to the functional steps used to perform the major operations.

#### **Select Content to be archived**

Any system for long-term archives needs a mechanism for selecting the content to be archived. There are a wide range of possibilities, because archiving might be an automatic part of any kind of record transaction system, or might be initiated by a manual process. In our initial implementation, selection is a manual process, but other kinds of automatic selection are easy to imagine.

#### **Prepare for Archive**

There are several steps involved in preparing material for storage in an archival repository, independent of the manner in which storage is guaranteed.

#### **Convert Content to Multiple Archival Formats**

To protect against format obsolescence, we propose converting data from its original form into one or more "archival formats", such as PDF/A, XHTML, or DNG. By archiving the same document in multiple formats, each document becomes its own virtual "Rosetta stone," aiding future interpretability and mitigating the chances of loss due to format obsolescence. (ISO 19005-1, 2005)

#### **Package Together the Original, Conversions, and Metadata**

There is some risk that documents, conversions and metadata will become separated, if stored separately.

While there are many ways of linking multiple components, the author has taken the straightforward approach of creating a single file which contains the original document, the metadata, and the converted (archival format) documents. Once metadata is gathered and the document is converted into appropriate formats, the author maintains the relationship between them by bundling them together and storing them as a single “archival package.”

### **Distribute to Remote Sites**

Having prepared the material in a form suitable for archiving, the resulting package must now be stored in a way that will guarantee its availability and integrity over a long period of time.

This process involves several components. The basic principle that links these steps is the use of “secret sharing” in which an (efficient) implementation is used for creating  $N$  “shares” of data (where each “share” is (nearly) the same size as the original bundle), for distribution to  $N$  sites, any  $K$  of which are necessary and sufficient for reconstructing the original package.

### **Select Sites for Distribution**

Secret sharing involves distributing packages of data to  $N$  sites, where any  $K$  of the sites’ data are necessary to reconstruct the original. The larger  $K$  is, the more impervious the system is to malicious attacks on data privacy. The larger  $N-K$  is, the more impervious the system is to inadvertent or intentional data loss. The larger  $N$  is, the more costly data storage is; the larger  $K$  is, the more costly data retrieval will be. If privacy is not a requirement, the author can use  $K=1$ , meaning every “share” is really just a copy of the original data.

Many of the threats facing long term storage, such as site failure, organizational failure, authorization failure, and malicious loss, are addressed by distributing copies of the archival data to numerous distinct remote locations. one must select remote sites such that the probability of loss due to geographic, organizational, economic or political failure is mitigated. This process involves some amount of evaluation of the qualities of the sites: Are they in the same geographic location? Do they share the same risks for political instability? Are they controlled by the same financial institution?

In the current implementation, the author manually configures the selection of remote sites, but selection might be automated. Selection should take into consideration the geographical location and control of the site, in addition to the perceived trustworthiness, reliability and availability of the providers. Reliability ratings may be measured based on past observations (Cooper, 2002). Sites with a long history of availability or that are operated by well established organizations may be

considered more likely to survive in the long term. The study also considers other characteristics of the organization that operates a site. To reduce the risks of malicious loss, selecting sites operated by industry competitors might be a good choice, as they might be less likely to collude.

According to Cooper, (2002) sites could also be chosen based on existing contractual relationships. Parties interested in archiving their documents may enter into a contract, in which each party agrees to store data from the other. Techniques such as auctions and bidding for storage space are also possible. Based on the selection of sites for distribution and the evaluation of the reliability and risks associated with those sites, one can choose the variables  $N$  (number of shares) and  $K$  (threshold of shares to reconstruct). Determining the number of shares to generate and the threshold for recovery involves a tradeoff between preserving the document and protecting privacy (Wylie, 2001). The system under study provides flexibility by allowing the user to configure their desire for privacy versus document preservation. Support for public documents is provided by using a threshold of one, in which case each “share” is simply a copy of the original archival package.

### **Use Secret Sharing to Create Distribution**

Given  $N$  and  $K$  established above, the study uses a novel and efficient algorithm for secret sharing to produce shares of the archival package. Shamir-style secret sharing requires generating random numbers (to be used as coefficients of polynomials) and doing arithmetic (multiplication, addition and division) on integers modulo a prime. The key insight that leads to improved performance was the use of 257 as the prime number, on a 32-bit processor. Each byte of the original document is treated as an individual “secret.” This allows for efficient generation of multiple random numbers at a time (as a string of  $K$  random bytes), the use of a pre-computed lookup table to eliminate modulo division operations and 32-bit arithmetic with CPU-supported modulo operators to compute the polynomial value for each of the  $N$  shares. The results pack into vectors of 8-bit numbers, with a separate “overflow” bit-vector for those cases where the polynomial value equals 256 exactly. Each “share” is thus 12.5% larger than the original data.

### **Send Shares to Remote Sites**

Once shares of the data have been generated, they must be sent to remote sites for storage. One of the evaluation criteria for site selection can be whether the sites are based on different operating systems and different storage system implementations, to protect against simultaneous loss of information. In order to increase the likelihood of finding such different, independent services,

the author has chosen to focus on storage sites where storage and access can be obtained using standard Internet services. The study's implementation uses email. The author select storage providers among Internet service providers (ISPs) around the world. The author used SMTP to transmit shares to the ISPs, and POP or IMAP to retrieve, enumerate, or test the shares provided. the shares were indexed at a storage provider using additional headers and/or message subject. This system can easily be extended to use web hosting (with WebDAV for storage and HTTP for access), FTP-based services, or other standards-based storage services.

### **Manage Archives**

At this stage, the documents originally archived have each been shared and stored in multiple repositories. Then it must be ensured that the data will persist unaltered for the desired lifetime of the document. Guaranteeing that data remains available and consistent is a difficult problem.

The lifetime of traditional archiving media, such as paper, may likely be maximized using a hands-off approach. Because physical paper media may be damaged by excessive handling, the content may be inspected for authenticity, and then placed in a physically secured location, such as a bank vault, and guarded from access. The archiver is generally confident that media placed in the safe will remain intact for a long period of time.

Digital media is susceptible to a wide range of situations that may jeopardize its integrity. Physical media for digital data, such as magnetic or optical disks, are fragile and subject to loss, ranging from random bit errors to entire disk crashes. Unlike paper media, the lifetime of digital media is much less predictable. The study therefore takes the position that digital media requires a form of active management: a periodic process that verifies the integrity of the digital data and performs any necessary recovery steps as the result of a detected failure. The following steps are necessary for active management of the archived material.

- **Organize Records of Content**

The study envision the management of records about archived data, searching the archives, and organizing the index of archived data to be a separate function from the archival storage itself. To support this, "proxy" files are created for each archived file, where the proxy itself contains a record of the identity of the services where the individual shares are stored, as well as a copy of its metadata. A simple XML representation is chosen for this information. Such records can be searched, shared, displayed, and used as the material for indexing records on the local disk.

Part of the reason for separating archiving from archive management is that the technology and requirements for search and retrieval, indexing, and so forth are likely to change more rapidly than the archives themselves.

- **Store Data for a Long Time**

Given that commercial Internet Service are being used Providers with email storage banks, this part of the system is not specifically a part of the prototype implementation. However, commercial storage services tend to use standard industry techniques for ensuring long-term data storage and integrity. Most utilize offsite backup; many use RAID storage and other storage network techniques for improving the reliability of their service. Routine hardware maintenance and upgrades add additional storage space into the system and provide media refresh to further reduce the chances of media failure.

- **Monitor Stored Material for Loss or Modification**

Simply distributing shares of a document to various locations is not enough to guarantee that the data will remain available for a long period of time. It is inevitable that data stored on digital storage media will eventually suffer from loss or corruption. Shares may be intentionally deleted or modified by a malicious site or user. Corruption may not be realized until the data is read as part of an attempted retrieval operation, when it may already be too late to recover from the damage. The researcher therefore takes a proactive approach, and periodically monitors the availability and integrity of archived shares. Accidental loss may result from media or hardware failure, site or organizational failure, or operational errors such as accidental deletion. We also must consider malicious sites, which may intentionally destroy or alter a share while claiming it is available or unmodified. The monitoring process will detect any loss or modification, regardless of the cause. To protect against the threat of loss, we periodically query storage sites for the availability of shares. The data may be retrieved to guarantee that it is, in fact, accessible. Protecting against corruption or integrity failure requires more than verifying availability: it must also be ensured that the contents are unaltered. Verifying the integrity of a share does not require remembering the share data. Instead, one must only know small facts about the data, which are significantly smaller in size. Question and answer style of verification is used, where the question is such that only an unmodified copy of the data could answer the question correctly. The questions must be unknown to the storage provider prior to being asked. Monitoring as a self-serve process has been implemented, where users are responsible for monitoring

their own archived data. Any party given the share location and question and answer sets, however, may perform the monitoring process. As monitoring does not require knowledge of the data, the envision supplementing self-monitoring with the use of independent third party monitors to perform periodic checks and notify the owners upon detection of loss.

- **Repair When Data Loss Occurs**

When a share is determined to be corrupt, one can take corrective action to repair the damage. Given at least the threshold number of the shares, one can recover the polynomials used to produce the shares, and an individual share may be reconstructed. This regenerated share is then resubmitted to the original site or distributed to a new site. If users themselves detect loss, they may retrieve and regenerate shares. Independent monitors may be authorized to initiate a repair process by using protocols for dynamically generating additional shares or changing the threshold of an encryption without requiring the reconstruction of the original data (Wang, 2005 ... et al).

### **Recover Archived Content Records**

Shares of archived documents are dispersed amongst a wide range of Internet accessible services. Location information is stored in a simple "proxy" record on the user's local system. Knowing the share locations is required to retrieve a document from the archive, the proxy records is guaranteed to survive as long as the archived data. a simple process of recreating the proxy records of archived content from the archives themselves as a mechanism for disaster recovery is implemented. This is a fairly time consuming process (since it requires enumerating all of the archived content), but the ability to reconstruct the index using only the archived content is demonstrated. It is possible to periodically archive the index of archived content itself as a way of speeding this recovery. It is worth noting that as long as a sufficient number of shares of the data remain available in the network, their corresponding proxy record can be reconstructed, and so the monitoring and recovery of share data protects the proxy records as well.

### **Retrieve Data from Archive**

Archived data must remain accessible regardless of the state of the document submission system. The retrieval process is independent, requiring only the ability to communicate with the remote sites holding shares of a document, using standard Internet protocols such as POP, IMAP, FTP, or HTTP. At archive time, authorized document retrievers are given proxy records containing all of the information necessary to locate and obtain

shares to recover the original archival package, including the identity of each remote site, their last known location, and the necessary handles to retrieve a share. Access control to archived documents is controlled by knowledge of the share locations and handles. Given the vast potential space of share locations and handles, only users who have been given the proxy record can successfully locate enough shares of a given document.

To recover a document from the archive, a user issues a retrieval request with a particular proxy record. At least the threshold  $K$  number of the identified shares are retrieved and verified for authenticity. Then Lagrangian polynomial interpolation is used to reconstruct the contents of the original package, which contains the original document and metadata. Retrieval is implemented as an automatic process given a proxy record. In the absence of such a process, these steps may be performed manually.

### **Delete Archived Data**

To delete a document from the archive, it is necessary to convince a sufficient number of shareholders to delete their shares. Of course, the ability to delete a document introduces the risk that a document will be deleted inadvertently or maliciously. Thus, the access control for deletion must be controlled carefully.

### **Evaluation**

In this section, the prototype system is evaluated, particularly in the context of design decisions made during implementation and deployment. The challenges and practical problems faced are discussed which were previously overlooked and only realized by actually attempting to build such a system.

### **Recommendations and opportunities for further research**

Studies have been conducted Long-Term Document Preservation System. It is unfortunate that not a single known study wholly focuses on a system designed for retaining records and ensuring their legibility, interpretability, availability, and provable authenticity over long periods of time. It is apparent that Long-Term Document Preservation Systems in Kenya should be done to preserve our ephemeral documents for posterity.

### **CONCLUSION**

The objectives of this paper were to assess Long-Term Document Preservation System. As per the assessment, it is clear that preservation systems in Kenya have dramatically changed and they have emerged along with

changing social needs. Preservation of document heritage in our organizations has become one of the key areas in Kenya where modern ICTs are fully utilized. However, none of these preservations has been able to challenge the basis of long term preservation systems.

## REFERENCES

- Adobe Systems Incorporated, The Digital Negative (DNG), <http://www.adobe.com/products/dng/> G.R. Blakley, Safeguarding Cryptographic Keys, Proceedings of the National Computer Conference. (1979).
- B. Cooper, H. Garcia-Molina, "Peer-to-peer data trading to preserve information" *ACM Transactions on Information Systems* 20, 2. (2002).
- B. Cooper, H. Garcia-Molina, Bidding for storage space in a peer-to peer data preservation system, Proceedings of the 22nd International Conference on Distributed Computing Systems. (2002).
- R. Dingledine, M. Freedman, D. Molnar, The Free Haven Project: Distributed Anonymous Storage Service, Proceedings of the Workshop on Design Issues in Anonymity and Unobservability. (2000).
- H.M. Gladney, R.A. Lorie, "Trustworthy 100-Year Digital Objects" *ACM Transactions on Information Systems* 22, 3. (2004).
- A. Goldberg, P. Yianilos, Towards an Archival Intermemory, Proceedings of IEEE Advances in Digital Libraries. (1998).
- A. Goldberg, P. Ylanos, Y. Chen, J. Edler, A. Gottlieb, S. Sobti, A Prototype Implementation of Archival Intermemory, Proceedings of the Fourth ACM International Conference on Digital Libraries. (1999).
- ISO 19005-1, Document management – Electronic document file format for long-term preservation. (2005).
- R. Lorie, Long Term Preservation of Digital Information, Proceedings of the First ACM/IEEE Joint Conference on Digital Libraries. (2001).
- D. Patterson, G. Gibson, R. Katz, A Case for Redundant Arrays of Inexpensive Disks (RAID), International Conference on Management of Data. (1988).
- G. W. Paynter, Developing Practical Automatic Metadata Assignment and Evaluation Tools for Internet Resources. Proceedings of the Fifth ACM/IEEE Joint Conference on Digital Libraries. (2005).
- D. Rosenthal, V. Reich, Permanent Web Publishing, Proceedings of FRENIX Track: 2000 USENIX Annual Technical Conference. (2000).
- A. Shamir, "How to Share a Secret" *Communications of the ACM* 22, 11. (1979).
- X. Wang, Y.L. Yin, H. Yu, Finding Collisions in the Full SHA-1, The 25th Annual International Cryptology Conference. (2005).
- J. Wylie, G. Ganger, P. Khosla, M. Bakkaloglu, M. Bigrigg, G. Goodson, S. Oguz, V. Pandurangan, C. A. N. Soules, J. Strunk, "Survivable Information Storage Systems" DARPA Information Survivability Conference and Exposition II, 2. (2001).