RJ10332 (A0410-018) October 18, 2004
Computer Science

IBM Research Report

Content Immutable Storage: Truly Trustworthy and Cost-Effective Storage for Electronic Records

Windsor W. Hsu, Lan Huang, Shauchi Ong
IBM Research Division
Almaden Research Center
650 Harry Road
San Jose, CA 95120-6099

Research Division
Almaden - Austin - Beijing - Haifa - India - T. J. Watson - Tokyo - Zurich

Content Immutable Storage: Truly Trustworthy and Cost-Effective Storage for Electronic Records

Windsor W. Hsu, Lan Huang and Shauchi Ong

IBM Almaden Research Center San Jose, CA 95120
{windsor, lanhuang, ong}@us.ibm.com

July 27, 2004

Executive Summary

Maintaining trustworthy records is essential for an organization to function effectively. In this white paper, we analyze what is required of the storage system to affordably enable trustworthy electronic record keeping. We note that the traditional approach of storing electronic records on WORM optical storage is no longer practical given the large and growing volume of records that have to be stored and managed today. Furthermore, the products that have recently been introduced to overcome the deficiencies of WORM optical storage fail to offer adequate protection against record tampering. Therefore, in this paper we introduce the concept of *Content Immutable Storage (CIS)*, storage designed specifically to satisfy all of today's electronic record-keeping needs. CIS provides secure data immutability to robustly protect records from any modification, online or nearline accessibility to ensure that records can be stored and retrieved in a timely fashion, and a low total cost of ownership.

1 Introduction

Proper record keeping is essential for an organization to function effectively. The fundamental purpose of record keeping is to establish solid proof and details of events that have occurred. Therefore, at a minimum, the records must be trustworthy--meaning that they can be relied upon to provide irrefutable evidence of all of the events that have been logged. As described in [2], the record management system must: reliably store all of the records for an extended period of time, securely protect the records from any modification, quickly locate every record pertinent to an enquiry, and faithfully deliver the records intact to an agent seeking proof or details of events. Such requirements are increasingly mandated by regulations, especially for electronic records, which could potentially be deleted and modified without leaving so much as a trace. With the increasing volume of records and legally mandated long retention periods, trustworthy yet cost-effective record keeping systems are increasingly necessary.

To ensure their trustworthiness, electronic records have traditionally been stored by making irreversible changes to an optical storage medium, such as a Write-Once-Read-Many (WORM) optical disc, CD-R and DVD+-R. Market forces as well as physical and/or technological constraints, however, prevent WORM optical storage from improving much in performance and storage density. Organizations thus must store their rapidly growing volume of critical records on an increasingly large number of WORM optical discs. Managing such a large number of discs is a massive, time-consuming, error-prone and expensive process. More importantly, records pertinent to an enquiry cannot be easily located within such a system and delivered in a timely fashion.

Attempting to both improve the accessibility to electronic records and decrease the cost of their storage, several vendors have recently introduced systems that use rewritable magnetic disks with function-rich software to protect against overwriting. Such an approach, however, compromises the trustworthiness of the records because it offers inadequate protection against record tampering, especially in view of the high stakes involved and the likelihood of malicious and inside attacks. Moreover, some of the systems require a new storage interface and extensive changes to applications and system software so that the total cost of deploying the system is high. Several WORM storage products based on magnetic tapes have also been launched recently. These offer higher volumetric efficiency than WORM optical storage, but do not support direct access, which is required for the timely retrieval of relevant records.

In this white paper, we establish the basic requirements of a storage system that satisfies all of today's electronic record keeping needs. We refer to such storage as *Content Immutable Storage (CIS)*. CIS provides secure immutability of data to robustly protect records from any modification, online or nearline accessibility to ensure that records can be stored and retrieved promptly, and a low total cost of ownership. We also discuss principles and practices for properly implementing CIS.

2 Content Immutable Storage

The key requirement for Content Immutable Storage (CIS) is to enable the primary objective of record keeping, namely to establish solid proof and details of events that have occurred. This means that CIS must be able to reliably store all of the records for an extended period of time and securely protect the stored records from any modification. In addition, it must allow every record that is pertinent to an enquiry to be discovered quickly--within days and sometimes even within hours [1]. Given the large volume of electronic records today, CIS must, therefore, efficiently support some form of direct access mechanism such as an index for the records. The large volume of records, coupled with the long retention period, further requires that CIS be very cost effective. This is the case especially because there is a tendency in the short term to view the records largely as an overhead needed just for satisfying the current intense regulatory scrutiny.

The requirement for reliably storing records over an extended period of time is similar to the basic requirement of any storage system and includes preventing any loss of data due to disasters, system failures, equipment obsolescence, *etc.* Loss of data could also result from intentional destruction of the storage system to remove damaging evidence. However, one considering such an action would be better off by safeguarding the records because large scale destruction of records would be evident and could result in the presumption of guilt. Also, deliberate destruction of records can be effectively mitigated by imposing physical security and remotely mirroring the data.

The more pressing need is to protect against clandestine modification, including destruction, of selected records during their storage. Modification of the records could result from software bugs and from user errors such as issuing the wrong commands and replacing the wrong disks during service actions. Given our increasing reliance on electronic records, the potential gain from intentionally manipulating and altering the records is huge. Thus, CIS would also have to protect the records from intentional attacks. These malicious attempts to compromise the records could come in the form of hacking and viruses. The more menacing threat is that they could be inside jobs that are launched by disgruntled employees, company insiders, or even conspiring technology experts. CIS must, therefore, offer extremely secure protection against record modification.

In some environments, the requirement is that a record once created, can never be altered or deleted within its retention period. After its retention period expires, the record can be shredded, *i.e.,* completely removed from the system. In addition to such term-based retention functionality, some regulations further require the capability to hold and preserve a record indefinitely or for some specified duration after a triggering event. Such advanced features, when required, must again be implemented very securely to protect against malicious attacks, even those from the "inside," and from conspiring technology experts.

3 Principles and Practices

In view of the high stakes involved in record tampering and the likelihood of malicious and inside attacks, the enforcement of record immutability in CIS must be extremely secure.

In the following list, we present principles and practices adapted from computer security that are essential for implementing CIS properly:

- *Increase barrier to attack.* Increase the cost and conspicuity of any attack against the system. This can be achieved by reducing the pool of people and the organizations who possess the relevant expertise, for example, by using non-universal components and custom hardware for key parts.

- *Focus on end-to-end trust.* Take a holistic approach to increase the trustworthiness of the overall system rather than focusing exclusively on any one component. Applying this principle, there is little point in making the immutability enforcement really secure if it can be bypassed. We must ensure that the requests always go through the enforcement mechanism. Such complete mediation of requests can be achieved, for example, by securely binding the storage media with the enforcement mechanism through physical locks.

- *Limit what has to be trusted.* If a system is large and complex, we cannot be sure whether it is correct, let alone secure. In a complex system, we should isolate the trust-critical modules and make them simple, verifiable and correct. In computer security terminology, this principle states that we must have a small trusted computing base. It implies that the enforcement of record immutability in CIS should be isolated from higher level functionalities such as storage virtualization and object management.

- *Use a simple, well-defined interface between trusted and untrusted components.* Restrict traffic into the trusted components to only legitimate requests to reduce the possibility of compromising the trusted components and to compartmentalize the components and, thereby, limit any error propagation. For example, if the trusted component that enforces record immutability and the untrusted component that provides object management share the same address space, the untrusted object management component could easily access the internal structures of the trusted immutability enforcement component, and thereby, compromise it.

- *Trust, but verify.* Every component must be verified to ensure that it works as intended and that it interacts as designed with other components. Preferably, the verification should be performed on each and every operation so that any fault can be quickly discovered, isolated and corrected.

Preserving records securely from any modification or tampering is critical, but to enable truly trustworthy record keeping, CIS needs to further ensure that every record that is pertinent to an enquiry can be promptly located and retrieved. This means that CIS must offer a high retrieval rate so that all of the relevant records can be retrieved quickly. Having good raw performance alone, however, is not sufficient because with the large volume of records today, scanning all of the records stored to discover those that are relevant to an enquiry is no longer practical. Instead, the data must be organized in such a way that all of the records relevant to a query can be found directly. In other words, to facilitate timely record discovery, it is essential for CIS to support direct access mechanisms, such as an index, efficiently. CIS should, for example, support fast random access and fine write granularity so that such access mechanisms can be updated efficiently.

The large volume of records today and their long retention periods also make a low-cost storage solution imperative. To achieve low cost, CIS should ride market forces and leverage existing components that are in volume production. Using commodity parts, for example, brings not only economies of scale and market competition, but also leverages the large investment in development that is possible and worthwhile only with volume production. Using such parts, however, seems to contradict the above principle of using customized components for key parts, but it really does not. We need to use non-standard parts only where they are necessary to securely enforce record immutability. The remaining parts should be standard. For example, CIS could rely on magnetic disks for storage if there is a secure mechanism for ensuring that the stored records are immutable. Similarly, CIS does not require a new interface to comply with the principle of using a simple and well-defined interface between trusted and untrusted components. Instead, it should leverage an existing standard interface that is simple and well-defined so that it can achieve highly trustworthy record storage with a low cost of integration and deployment.

4 Summary

Proper record keeping is essential for the effective functioning of an organization. With our growing reliance on electronic data processing, records are increasingly generated in large volumes and in electronic form, which makes them vulnerable to undetected deletion and modification. In this white paper, we clearly establish the basic requirements of a storage system that satisfies today's electronic record keeping needs. We note that, with the large volume of records today, traditional approaches to storing electronic records, such as writing them to WORM optical storage, is too slow and expensive. Furthermore, recent attempts to overcome the inadequacies of WORM optical storage do so at the far greater expense of compromising on the requirement to prevent any record modification, including malicious and inside tampering. Therefore, we introduce *Content Immutable Storage (CIS)*, storage that is specifically designed to satisfy all of today's record keeping requirements. CIS provides secure immutability of data to robustly protect records from any modification, online or nearline accessibility to ensure that records can be stored and retrieved in a timely fashion, and a low total cost of ownership. We also discuss principles and practices for properly implementing CIS.

References

[1] COHASSET ASSOCIATES, INC. The role of optical storage technology. White Paper, April 2003.

[2] WINDSOR W. HSU, AND SHAUCHI ONG. Fossilization: A process for establishing truly trustworthy records. White Paper, IBM Research, July 2004.