

A System for Intellectual Property Protection

Christoph Busch, Frank Graf, Stephen Wolthusen, and Armin Zeidler
Department Security Technology, Fraunhofer Institute for Computer Graphics
Darmstadt 64283, Germany

ABSTRACT

An integrated system for the protection of data both on computers and in analog representation is presented. Based on the automatic and mandatory encryption of all data on storage media, authenticated encrypted communication channels, and digital watermarking technology, the system protects data from misappropriation while working as an extension to the operating system, making the security mechanisms fully transparent for legitimate users. Even analog representations of the data objects are still protected through the use of digital watermarking and can be traced back. An outline of the system architecture along with information on the prototype is given.

Keywords: Use Control, Digital Watermarking, Encrypted File System, Mandatory Access Control, Virtual Private Network

1 INTRODUCTION

The protection of intellectual property, whether industrial trade secrets, classified governmental material, or something as mundane as illegal copies of software or entertainment products¹ is certainly not a new problem, but has been exacerbated by the availability of inexpensive high capacity storage media and seamless worldwide internetworking systems.

The algorithmic building blocks addressing the various facets of access and use control, confidentiality, and authenticity are well established and have been the subject of extensive research. What has been mostly lacking, though, is an integrated mechanism for data protection in a networked environment that integrates seamlessly into the existing IT fabric while not unnecessarily burdening the user[16, 4]. The case of closed systems is well researched[2, 10, 6, 12, 15] and can be handled by existing operating system implementations. Networked systems are more difficult to secure, since, among other problems, one cannot presume the network connection to be trusted and must always be prepared to presume networked components compromised; other problems include trust relationships across administrative domain boundaries and the extent of the required trusted code base in such situations[8].

Another important issue is auditing in such distributed systems. There are a number of applications in which the availability of an audit trail for each individual access to given data is desirable, even when coupled with mandatory access control. Again, this is easily accomplished within a closed system but becomes difficult to enforce in a distributed system (a typical example of this is the creation of copies behind the back of the original supplier of the data object). Even if the legitimate user of a client system is trusted, though, the risk of an attacker accessing the data when unobserved or simply stealing the system in question remains.

Even if the problems of system security were all solved, one is continually confronted with the issue of “low technology” security breaches. Going to great lengths to protect data on a computer system and its leakage is a rather futile effort if an inside threat can simply transfer the desired data to analog media (printouts, audio recordings) and walk out of the building unimpeded. Unless one is able to counter this threat effectively, investing in computer security alone is not going to improve the overall situation. As will be discussed later, we have found an integrated solution which – at least partially – addresses this problem as well since we do not subscribe to the view stated in [4] that it is negligible.

¹One should keep in mind that between 1987 and 1994 the core copyright industries in the U.S. grew twice as fast as the rest of the economy[14], while the total copyright industries accounted for 5.72% of the U.S. GDP, so these concerns are quite substantial

In our proposed approach, an integrated system addressing the most severe problems identified was designed and subsequently implemented as a prototype on a reference platform under the internal code name CIPRESS (Cryptographic Intellectual Property Rights Enforcement System). The following is an outline of the design and the prototype implementation. The latter is important since the basic design must be translated into the functionality provided by the host operating system.

Protection is achieved of both data on computer systems – through mandatorily encrypting all data on storage media and going through communication channels, rigorous access and use control – and after transferring the data to analog media through digital watermarking.

The system described is rather complex and has many interdependencies. As a consequence, it is rather difficult to create a narrative thread in the description, so it may be necessary to follow the cross-references occasionally.

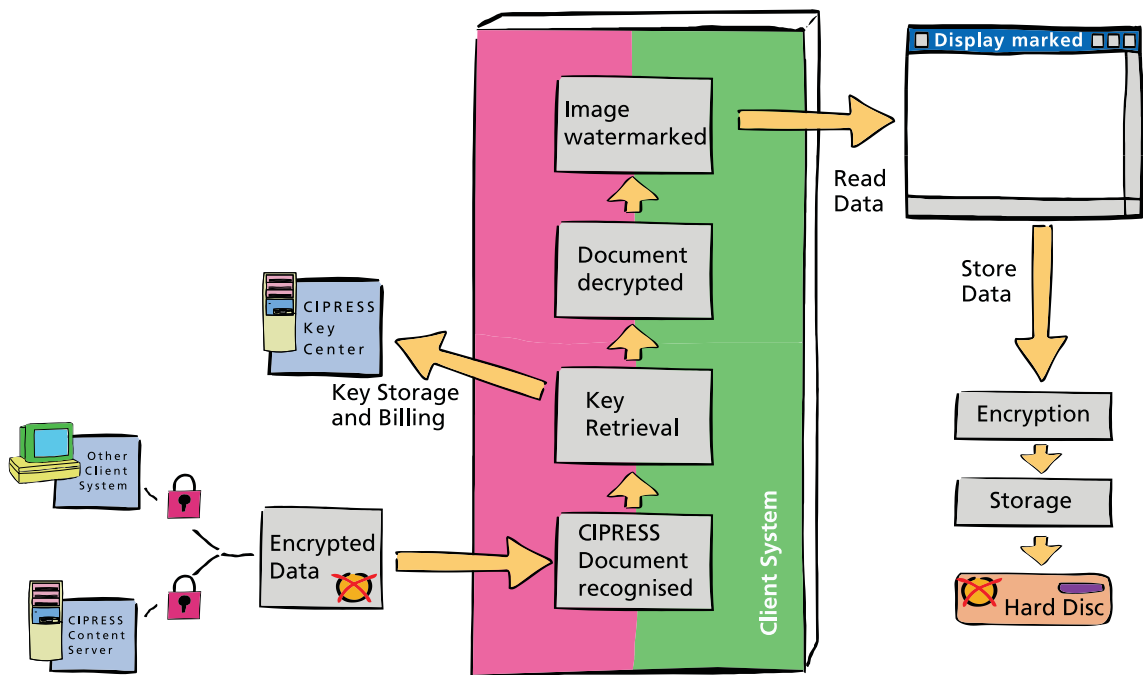


Figure 1: Data flow in CIPRESS

2 ARCHITECTURAL OVERVIEW

2.1 Client Systems

CIPRESS consists of three distinct types of systems. Those in the first group are referred to as “client” systems and are workstations for ordinary users, although administrative tasks on servers can also be performed if a user has sufficient privileges. All storage media of client systems are transparently and mandatorily encrypted at all times without the possibility for user intervention; this occurs at the operating system level.

For most files (since they are not shared between clients) this is done using a secret “master key” specific to the given client machine outside the control of the user. Files are decrypted only for loading into volatile memory and are encrypted again upon storing the files. Since the master key is unique to its client system, it is not possible to transfer master key protected files through storage media or other operations on file systems (sharing file systems over the network) to unprotected systems or even other CIPRESS clients. Other files, referred to as “registered” files are known to the central authority in CIPRESS, the so-called Key Center (see section 2.3). These

files originate from users on a client system² and are subsequently registered with the Key Center and stored at a storage facility, the Content Server (see section 2.2). Within the confines of proper access rights being available to users, groups of users or the general public, the document subsequently becomes accessible to the outside world. Registered documents are also always encrypted, but using keys retrieved from the network (see section 3) which are discarded immediately after use and are specific to user and system. In any case, data is unencrypted only when in volatile memory. Note that importing unencrypted files (e.g. from unprotected application servers) can be read with no problems.

Another component of CIPRESS acts as a choke or filter on network connections, again outside the control of the user. It is possible to restrict connections to certain addresses or address ranges³ (although the latter should best be handled by a firewall to reduce administrative overhead); the more important feature is to establish secure channels between hosts (again, possibly specified for entire subnets). The establishment of the secure channel is preceded by identification and authentication of both parties involved and obviously requires a matching implementation on the far end. Trusted communication is particularly important when it comes to two special types of systems, Content Servers and Key Centers (see sections 2.2 and 2.3). As noted above, local data is restricted to the given client machine which can receive data from the outside world unimpeded – provided that the given addresses are not blocked – but not transfer data to the outside. This is possible only by making the given data object known to the Key Center, i.e. registering the document.

In any case (file system or network access) files into which a digital watermark can be embedded are marked with an invisible (inaudible etc.) fingerprint of the current user before the data is released to the applications, ensuring that the documents are marked with the identity of the last user even after an analog representation has been created (e.g. by photographing the screen, creating printouts). Together with the markings embedded by the Content Server a seamless protection for both digital and analog representations is thus achieved.

2.2 Content Servers

The second type of system is the “Content Server”. As the name implies they act primarily as data management facilities for data within the CIPRESS system. Several access methods have been defined, among them WWW access, a proprietary CIPRESS protocol, and the matching platform file server mechanism. While not strictly necessary for content distribution (or as a file server if only downloads are needed), the Content Server is needed for client systems to export data to others; all exports (referred to as “registration”) are routed through a Content Server which forwards the required information to the Key Center (see section 2.3) which does not know about the data itself, marks the data with special digital watermarks (see section 3.2) and then acts as the first level storage facility for the data objects.

The general architecture assumes that releasing data is allowed to occur under user control only between systems of the same security level, i.e. from one CIPRESS client to another. Releasing data to the outside must occur through an administrative process put into place by the Content Server administrator; this must ultimately be dealt with by a human since automatic processes can be fooled.

Content Server systems are to be grouped into domains for administrative purposes if an organization maintains several systems (e.g. for load balancing purposes). Within one Content Server domain⁴ the administrator has full control over the contents (even in unencrypted form; this can be disabled if the need should arise and the responsibility can instead be relegated to the Key Center administrator) and is also the authority to turn to for security clearances for releasing material.

In the case of client system users sharing data with other authorized users, an application on the client system contacts the Content Server over a secure, authenticated channel, and up-

²in this terminology, an administrator of a server uploading data for general consumption is also considered a user

³We are assuming TCP/IP as the network protocol, the terminology used here reflects this.

⁴Obviously, users in a given group such as a research group or corporate organization will use their own domain for sharing sensitive documents; crossing domain boundaries is also possible but is useful only for publishing and similar specialized applications

loads data to the Content Server. This data is then registered with the Key Center which issues a re-encryption key and the necessary keys for digital watermarking. The Content Server then embeds the watermarks if possible and deposits the encrypted documents; these may now be downloaded by all authorized systems and can also be stored on other file servers (e.g. CD juke-box archives, backup systems) without any impediment to security.

2.3 Key Center

The Key Center is a central facility in the CIPRESS system. It has knowledge of all users authorized to use client systems, as well as of all Content Server systems which are attached to the Key Center. Most of this information kept in a database on the Key Center can be derived from an external directory server (e.g. X.500, LDAP), the prototype design chose to forego this and make the system self-contained.

The most important functionality provided by the Key Center, though, is access control (multilevel access control with permissions for different operations on the data objects, comparable to security labels), which regulates access to registered data objects (remember, each data access to a registered document requires the possession of the specific key, hence access verification occurs on each document access). The Key Center – as implied by its naming – also houses the keys required for each registered document and user and generates new keys for new data objects (in the event of a object being registered) or new legitimate users accessing data objects. Since the Key Center is queried for each access, it is able to monitor each access to a data object, thus implementing a use control mechanism. Such mechanisms obviously have the advantage of always enforcing consistent rights and maintaining seamless access logs for purposes such as auditing and possibly billing while on the downside the network traffic generated by such security servers can be quite considerable and is mandatory.

We have aimed to strike a balance by distinguishing between local data objects and those to be exchanged over system boundaries; the working set for each user should, at any given time, be rather small (as opposed to e.g. files needed for system operation) and, since this working set is usually created by manual interaction, delays of at most a few seconds before documents are displayed should be tolerable.

The Key Center is, by its very nature, the most sensitive unit in the entire CIPRESS system. Compromising a single client has only very modest effects on security; compromising a Content Server is, depending on whether the operator has elected to retain the uploaded plaintext, either a very limited problem (since only encrypted data is lost, along with whatever is uploaded to the Content Server during operation as a “rogue” server) or compromises all backups of unencrypted data on the server. A loss of the Key Center, though, has disastrous consequences since any document ever injected into CIPRESS can be decrypted and information as to which user accessed which document when also becomes the property of the attacker. Consequently, extreme care must be given to both the software integrity and physical security of the Key Center. At the same time, the availability of the Key Center must be given at all times.

3 DESIGN AND IMPLEMENTATION

The overall design objective was to create a system that, during normal, legitimate operations, did not differ – from the user’s perspective – in terms of facilities provided and handling from a standard operating environment. This implies that users should be able to continue working with the application programs in existence before deployment of the security system without any modification to the application software. The motivation here is that one must keep re-training efforts down and minimize inconvenience to users if one is to hope for a sufficient user acceptance of a security system that is to be deployed on a large scale. These requirements imply that the security mechanisms be embedded at the operating system level while at the same time not modifying the visible behavior of the operating system.

3.1 Re-Encryption

One of the central notions within the CIPRESS system is *Re-Encryption*. Each tuple (U_i, O_j) consisting of an user and a data object is assigned a key K_{U_i, O_j} . The key assignment operation occurs centrally at the Key Center and depends on unique identifiers for both users and data objects in a system. The identification is performed through the use of X.501[9] *distinguished names*, while data objects are identified by means of an unique hash value⁵ digest code (SHA-1[13] was specified for this purpose). This mechanism also allows the implementation of an elaborate access control mechanism at the Key Center. Since each access to a document requires the provision of a key (in the basic design the keys are discarded after one time use, regardless of the fact that subsequent read operations by the same user U_i on the same document O_j will use the same key K_{U_i, O_j} . This is to minimize the duration during which the key might be observable to an attacker who has already penetrated the system. However, in some situations it might be helpful to provide a key cache, thus trading off the visibility of and control over each data object access for lowering the communication overhead. In such cases the key cache should be located in tamper-proof hardware).

The Re-Encryption must take place upon load and store operations. In the case of file systems, this can be accomplished by inserting an additional layer of logic into the file system wherever supported by the host operating system; a similar mechanism is required for network transmissions. Most modern operating systems provide such facilities.

3.2 Digital Watermarking

Digital watermarking[11, 1, 3] can address the problem that analog representations are essentially unprotected since they are required by the legitimate workflow of the user. This technique, closely related to steganography, imperceptibly[5] embeds a second signal (a hidden message) within a carrier signal such as image or audio signals. The main distinction between steganographic messages and digital watermarks lies in the robustness of the latter. This is required, since unlike in the first case, where one hopes that any attacker is unaware of the hidden message and the carrier signal is usually not disturbed, the attacker is possibly aware of the marking in the signal and wants to destroy the signal; it is also highly desirable for the watermark to be highly resilient to manipulations of the carrier signal, even if the attacker is unaware of the fact that the misappropriated material contains a marking (e.g. image editing).

To that end, CIPRESS specifies that, during the registration process of a data object, a digital watermark containing a secret identification of the data origin shall be embedded in the data object; additionally, the identity of the last users accessing a document are also embedded. This allows the tracing of unauthorized copies to the likely source of the leak even across a chain of users based solely on an analog representation.

3.3 Implementation

The system was implemented in 1998 and 1999 using the Microsoft Windows NT[®] 4.0 operating system as a foundation and has been deployed in a field trial with the German Dermatological Association with great success; the key issue to point out is that the physicians were able to use the secure system after one hour of training without further incidents and used the system to exchange medical information during an extended field trial that lasted one year and included several dermatologist's offices.

4 PRIVACY CONSIDERATIONS

A system as presented above – while challenging in its design and implementation – presents even more delicate problems when it comes to the problem area of privacy. The European Community recently passed stringent regulations in this regard[7]; there are even harsher regulations requiring employee consent and notification for monitoring the working habits and other information regarded to be under the privacy protection in some of the EC member states. In such a

⁵note that this excludes open streams as data objects

situation the deployment of the system as well as any use of the information gathered from the access logs will have to be done in careful consideration of the legal framework in force at the given location.

It should be noted that such a system inherently provides the possibility for key escrow and detailed usage profiling. This is true for internally deployed systems as well as scenarios in which a single entity (trust center) handles the key material from several others. In the case of internal Key Centers, mechanisms must be put in place which prevent the use (and manipulation) of data for curiosity, personal gain, and similar motives. In the latter case the relationship between the trust center and the trusting party must be governed by some agreement clarifying liabilities and disclosure policies towards law enforcement⁶. In any case, encrypting the source document with a different system from the one used to provide use and access control (which itself is then subject to local regulations) is always an option.

Beyond the legal aspects one should also be aware of the inherent dangers of such a use control system since unless a compelling reasoning exists, users may strongly resent the notion that their use of every document and implicitly their work patterns are constantly monitored.

5 ACKNOWLEDGMENTS

The authors would like to thank the Mitsubishi Corporation, Tokyo, Japan for its generous support and close cooperation in supporting this research.

References

- [1] ANDERSON, R., Ed. *Information hiding: first international workshop, Cambridge, U.K., May 30-June 1, 1996: proceedings* (New York, NY, USA, 1996), vol. 1174 of *Lecture Notes in Computer Science*, Springer-Verlag Inc.
- [2] BELL, D. E., AND LAPADULA, L. J. Secure computer systems: Mathematical foundations and model. Tech. Rep. M74-244, The MITRE Corp., Bedford MA, May 1973.
- [3] BONEY, L., TEWFIK, A. H., AND HAMDY, K. N. Digital watermarks for audio signals. In *1996 IEEE Int. Conf. on Multimedia Computing and Systems* (Hiroshima, Japan, 1996), pp. 473-480.
- [4] CHOUDHURY, A. K., MAXEMCHUK, N. F., PAUL, S., AND SCHULZRINNE, H. G. Copyright protection for electronic publishing over computer networks. *IEEE Network Magazine* 9, 3 (May/June 1995), 12-21.
- [5] COX, I. J., AND MILLER, M. L. A review of watermarking and the importance of perceptual modeling. In *Proc. of Electronic Imaging '97* (February 1997).
- [6] DION, L. C. A complete protection model. In *Proc. IEEE Symp. on Security and Privacy, Oakland, CA* (Apr. 1981).
- [7] EUROPEAN COMMISSION. Directive 95/46/EC of the European Parliament and of the Council of 24 October 1995 on the protection of individuals with regard to the processing of personal data and on the free movement of such data. *Journal of the European Communities* L281 (November 1995), p31ff.
- [8] GARY GROSSMAN. Immediacy in Distributed Trusted Systems. In *Proceedings of the 11th Annual Computer Security Applications Conference, New Orleans, Louisiana, December 11-15, 1995 (CSAC '95)* (1109 Spring Street, Suite 300, Silver Spring, MD 20910, USA, December 1995), IEEE, Ed., IEEE Computer Society Press.

⁶While a Key Center can be physically located anywhere with Internet connectivity to avoid the problem of key escrow, the seemingly attractive use of offshore facilities does have its drawbacks when it comes to the possibility liability litigation and the possibility of a tacit assumption that some of the material may not be legal in the originating jurisdiction

- [9] INTERNATIONAL TELECOMMUNICATION UNION. *ITU-T Recommendation X.501: Information technology - Open systems interconnection - The directory: Models*. Place des Nations, CH-1211 Geneva 20, Switzerland, November 1993.
- [10] K. BIBA. Integrity Consideration for Secure Computer Systems. Tech. Rep. MTR-3153, MITRE Corporation, 1975. Also published as ESD-TR-76-372, USAF Electronic Systems Division (1977).
- [11] KOCH, E., AND ZHAO, J. Towards robust and hidden image copyright labeling. In *Proc. of 1995 IEEE Workshop on Nonlinear Signal and Image Processing* (Halkidiki, Greece, June 1995), pp. 452–455.
- [12] MCLEAN, J. A comment on the ‘basic security theorem’ of Bell and LaPadula. *Information Processing Letters* 20, 2 (Feb. 1985), 67–70.
- [13] NATIONAL INSTITUTE FOR STANDARDS AND TECHNOLOGY. *Secure Hash Standard (SHA)*. Gaithersburg, MD, USA, Apr. 1995.
- [14] STEPHEN E. SIWEK AND GALE MOSTELLER. Copyright Industries in the U.S. Economy: The 1996 Report. Tech. rep., Economists Incorporated, 1996.
- [15] UNITED STATES DEPARTMENT OF DEFENSE. Trusted computer system evaluation criteria. Tech. Rep. 5200.28, US Department of Defense, 1985.
- [16] VAN FABER, E., HAMMELRATH, R., AND HEIDER, F. P. The secure distribution of digital contents. In *13th Annual Computer Security Applications Conference, San Diego, California, December 8–12, 1997: proceedings (ACSAC’97)* (1109 Spring Street, Suite 300, Silver Spring, MD 20910, USA, 1997), IEEE, Ed., IEEE Computer Society Press, pp. 16–22.