The following paper was originally published in the
Proceedings of the First USENIX Workshop on Electronic Commerce
New York, New York, July 1995.

DigiBox: A Self-Protecting Container for Information Commerce

Olin Sibert, David Bernstein, and David Van Wie
Electronic Publishing Resources, Inc.
Sunnyvale, California

For more information about USENIX Association contact:

1. Phone: 510 528-8649
2. FAX: 510 548-5738
3. Email: office@usenix.org
4. WWW URL: http://www.usenix.org

The DigiBox:
A Self-Protecting Container for Information Commerce

Olin Sibert
David Bernstein
David Van Wie

Electronic Publishing Resources, Inc.
460 Oakmead Parkway
Sunnyvale, California
1 408 774 6100
info@epr.com

Abstract

Information Commerce is a business activity carried out among several parties in which information carries value and is treated as a product. The information may be content, it may be returned usage and marketing data, and it may be representative of financial transactions.

In each of these cases the information is valuable and must be kept secure and private. Traditional approaches secure the transmission of that information from one point to another; there are no persistent protections. Protection of all of these components of information commerce for all parties in a transaction value chain is necessary for a robust electronic infrastructure.

A prerequisite to such an environment is a cryptographically protected container for packaging information and controls that enforce information rights. This paper describes such a container, called the DigiBox™. EPR has submitted initial specifications for the DigiBox container to the ANSI IISP Electronic Publishing Task Force (EPUB) within the User/Content Provider Standards Working Group (WG4).

1 Introduction

As services and products in modern commerce increasingly take electronic form, traditional commerce is evolving into electronic commerce. This includes both creation and enforcement of various agreements between parties in an electronic commercial relationship. It also includes enforcing the rights of these parties with respect to the secure management of electronic content or services usage, billing, payment, and related activities.

To save money, to be competitive, and to be efficient [1,2], members of modern society will shortly be using new information technology tools that truly support electronic commerce. These tools provide for the flow of products and services through creators', providers', and users' hands. They enable the creation, negotiation, and enforcement of electronic agreements, including the evolution of controls that manage both the use and consequences of use of electronic content or services. In addition, these tools support "evolving" agreements that progressively reflect the requirements of further participants in a commercial model.

Participants in electronic commerce [3,4] will need rules and mechanisms such that:

1. Information providers can be assured that their content is used only in authorized ways;

2. Privacy rights of users of content are preserved; and

3. Diverse business models related to content can be electronically implemented.

The Internet and other information commerce infrastructures will require a management component that enforces such rules, ensuring a safe, coherent, fair, and productive community. This management component will be critical to the electronic highway's acceptance. Without rules to protect the rights of content providers and other electronic community members, the electronic highway will comprise nothing more than a collection of limited, disconnected applications.

Analysts have concluded that content will constitute the largest revenue-generating component of the information superhighway [5]. It is also clear that unfettered access to content requires that content providers be able to maintain control over literary or copyrighted assets. Many analysts conclude that this will be one of the key bottlenecks in the implementation and deployment of New Media.

2 Information Commerce and Digital Value Chains

Information commerce is often considered a wholly new concept, made possible only through the use of networks and computers. In fact, a robust information economy has existed for centuries, involving trafficking in physical *representations* of information such as books, newspapers, and so on. Because such commerce involves physical goods, there is a non-negligible floor to the cost of handling information goods. The new aspects of the electronic information economy are that the information itself is the entire product and that the product can be distributed at negligible marginal cost.

The traditional information economy in physical goods is publisher-centric, because creation of information goods—particularly low-cost goods—requires a substantial manufacturing investment. Figure 1 illustrates a simplified traditional information economy: physical goods flow from a publisher (manufacturer) to a customer, in response to orders and followed by payments. The author's relationship with the publisher may be more lightweight, but the author is nonetheless dependent on the publisher to report sales and make royalty payments in accordance with the author's contract. In addition, a financial institution provides payment processing and clearing services for all parties.

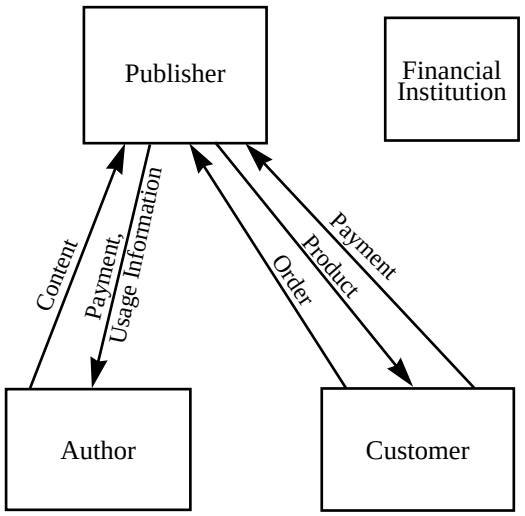

Figure 1. Traditional information economy.

Because of the flexibility afforded by electronic mechanisms, information commerce is evolving from indirect, advertiser-supported, mass-audience media to a new, niche-audience-oriented business model. In this system, members of the electronic community, with or without the economic support of advertising, pay providers directly for what they want to receive. Business-to-business purchasing is steadily evolving into a direct electronic ordering model.

Figure 2 illustrates the flexibility possible in new electronic information commerce models. Although there is still a role for publishers, this role no longer involves physical goods. Rather, the publisher is responsible for packaging and aggregating information goods and control information,

then making them available to customers. Similar to a manufacturing/distribution/retail chain for physical goods, the electronic model permits information retailers, and even end customers, to repackage and redistribute different aggregations of information while ensuring that the appropriate control rules are maintained. A clearinghouse ensures that usage information and payments are provided directly to authors and publishers; the payments themselves are made through traditional financial institutions. Because control rules are associated with information, a variety of payment and other business models can be associated with the same content (e.g., *purchase* versus *pay-per-use*).

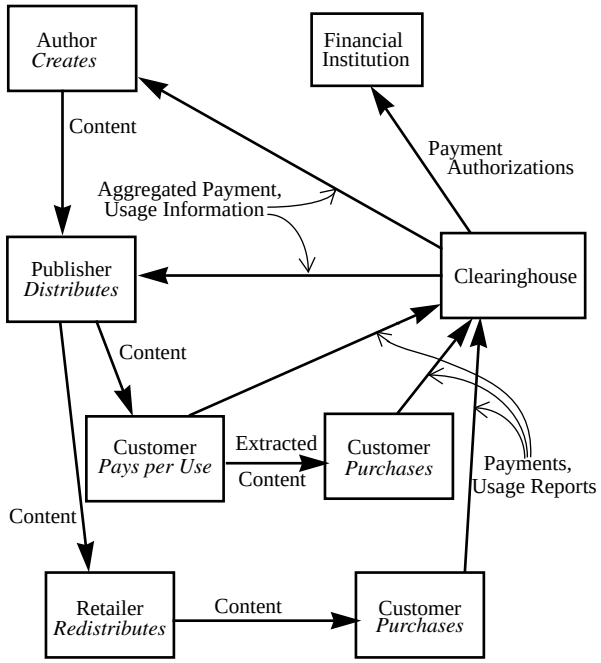

Figure 2. Electronic information economy.

The conversion from traditional commercial distribution channels requires key foundation technologies and results in a fundamental shift in existing infrastructures. This channel transformation will create a new electronic digital distribution industry. Digital distribution employing the DigiBox container architecture and its associated support environment, InterTrust™, can play a critical role in this transformation of the communication, media, and information technology markets.

2.1 Protecting All the Information in Information Commerce

The very properties that make "the net" attractive as a distribution medium—ease of manipulating information in electronic form—also appear to make these protections intractable. Addressing this dichotomy requires a paradigm shift in computer architecture to introduce the concept of a "secure processing" environment in which protected information can be manipulated without being subject to external tampering or disclosure. A prerequisite to such an environment is a cryptographically protected "container" for seamlessly packaging information and controls that enforce information use rights.

The DigiBox described by this paper is such a container.

The need for various information commerce computers and appliances to interoperate requires that this container format and its access methods be standardized. EPR has submitted initial specifications for the DigiBox container to the American National Standards Institute (ANSI) Information Infrastructure Standards Panel (IISP) through the Electronic Publishing Task Force (EPUB) in the User/Content Provider Standards Working Group (WG4).

The primary goal of information protection is to permit proprietors of digital information (i.e., the artists, writers, distributors, packagers, market researchers, etc.) to have the same type and degree of control present in the "paper world." Because digital information is intangible and easily duplicated, those rights are difficult to enforce with conventional information processing technology. Many types of rights (compensation, distribution, modification, etc.) are associated with the various elements of information commerce, and these information property rights take many forms. At a high level, there is the legal definition of "copyright," codified in U.S. law [6–9] and the Berne Convention. This gives copyright holders a legal right to control how copyrighted information is handled. In addition, various high-level rights are conferred by contractual arrangements between primary rightsholders and other parties.

For example, the protections needed for content elements incorporate the licensing provisions for the intellectual property rights of the content rightsholders. In a broader sense, these rights include control over several activities: the right to be compensated for use of the property; the right to control how content is distributed; the right to prevent modification of content by a distributor; "fair use" rights; the rights to the usage data, privacy rights of individuals, and so on.

In the realm of physical goods, these rights are enforced by a combination of legal and technical means. However, the technical means can be (and are) unsophisticated because the technology for violating rights is relatively expensive and time-consuming—in comparison to equivalent activities with respect to digital information. Photocopying a book or copying a video cassette is inherently more labor intensive and costly than copying a file. So, while defeating technical means of enforcement is (relatively) expensive, it can be done—and often the legal means to deter this are inadequate.

2.2 Information *Commerce*—Not Just Payment

Rights protection is also a fundamental aspect of commerce. Commerce is not just a way for two parties to pay each other for something. Rather, it is an extraordinarily rich web of relationships among parties that concerns payment, negotiation, control, advertising, reporting, auditing, and a variety of other activities. These activities are important aspects of the transaction relationships. Often the information carried in these reports, audits, and the like is highly valuable and highly confidential, perhaps even more valuable than the content that is the subject of the information commerce at hand. These activities too are performed and controlled in the "paper world" by legal and technical means, but there are no widely used models for their electronic equivalents.

Figure 3 shows some of the operations that could occur in true electronic commerce, using the Internet World-Wide Web [10] mechanisms as an example. Creators originate content and apply rules (e.g., "pay author $1.00/use") for its use. Distributors repackage content, applying additional rules (e.g., "pay $5.00 for the collection, then pay the creator," "report use of each item"). Users receive content and operate on it, generating billing reports and usage reports that are delivered to a clearinghouse and paid or summarized back for the originating parties. This structure is very rich and is capable of supporting many business models. There are multiple flows of information in many different directions amongst the parties involved in the transactions.

Another example is that of an advertiser (acting as distributor, or with a distributor). The advertiser might have a rule that offers a discount, or no charge at all, but only if the user views the advertisement and agrees to have that fact reported to the advertiser.

It is relatively simple to devise schemes for parties to pay each other electronically (for example, DigiCash [11], NetBill [12], Open Market [13], SNPP [14], NetCheque [15], First Virtual [16], etc.). Payment, however, constitutes only one—and perhaps the simplest one—of the means in which parties in commerce interact. All the other information commerce components must be accomplished with the same needs for security, privacy, and integrity. In fact, these aspects of electronic commerce, including rights protection, are strongly intertwined in the digital economy, because much digital commerce concerns information and innovative business models for information commerce.

3 Existing Approaches to Information Commerce

Information proprietors employ a variety of technological protection approaches today. These approaches are generally "point solutions," in that they protect a specific type of property in a specific context and enforce only specifically defined rights—typically only the right to compensation for use. Because the technologies are limited, the market is fragmented, and there are no general protection solutions.

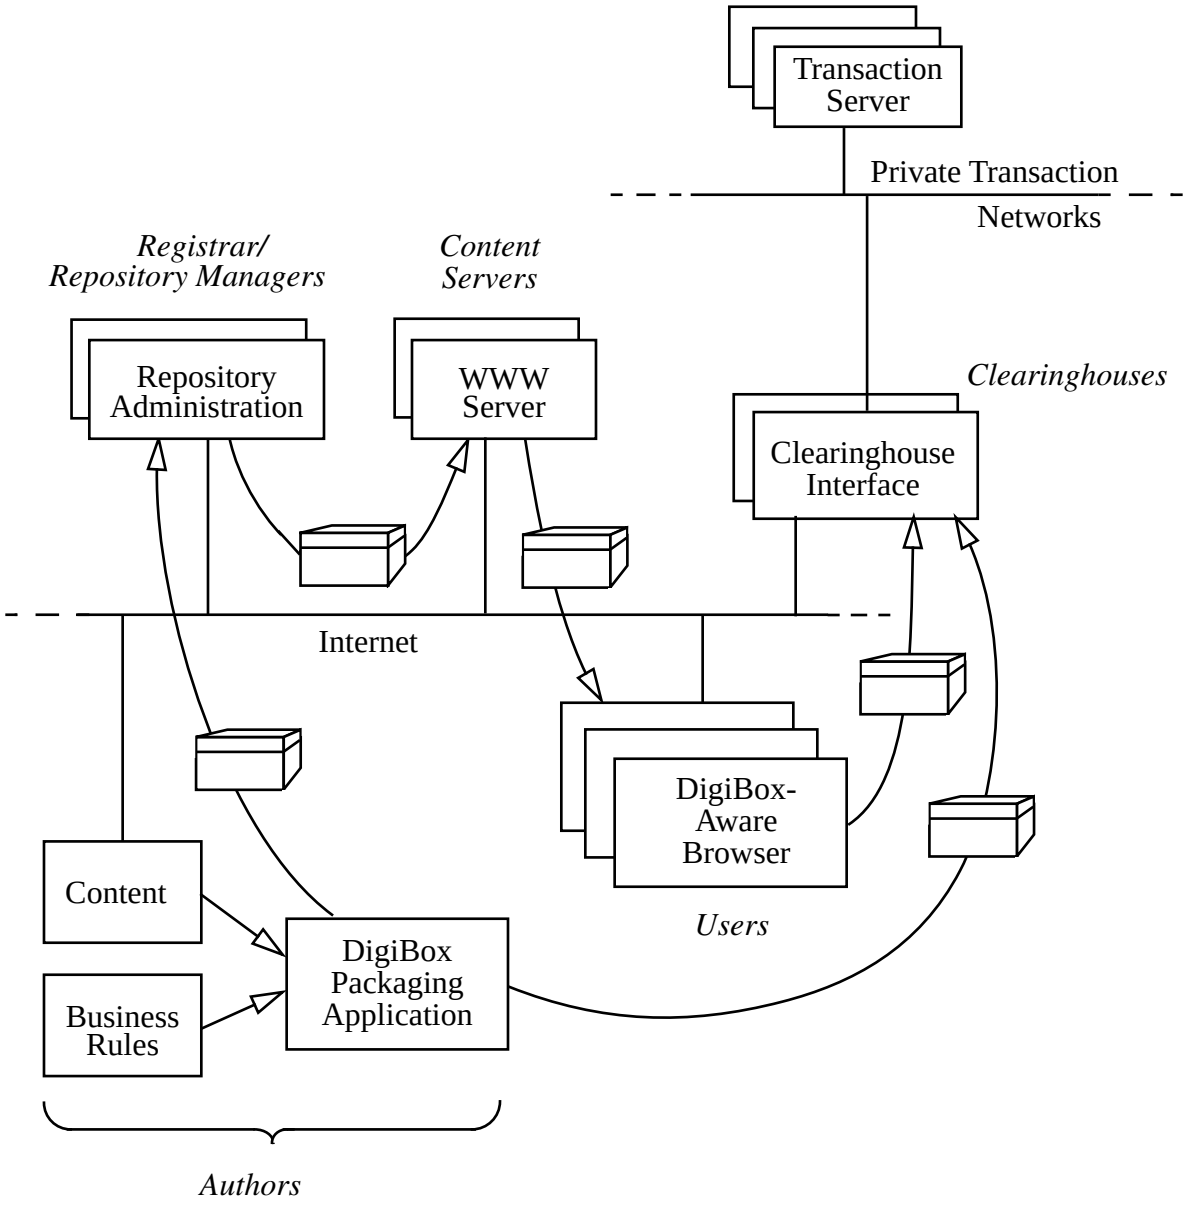

Figure 3. Multi-party Internet information commerce.

3.1 No Protection

Much digital property is distributed without any technological enforcement for property rights, on the assumption that legal means suffice. This approach works well enough for many low-value properties, but it has the disadvantage of raising the price to legitimate users who must pay for both their own and illegitimate use. In many cases, however, this cost is negligible, and no protection is an economically sound choice. Even for content that is free, however, a creator may wish to impose some rules for reporting or some access control. Of course, privacy rights of users will be a concern to many.

3.2 License Managers

For some valuable software properties, license managers are used. Because a software property is dynamic (executable), it is feasible to restrict it so that it functions properly only through interaction with a license manager process. In general, there is no protection of usage data in these schemes. In some cases this technique has been applied to content protection, but only with limited success [17, 18].

3.3 Cryptographic Unlock

Some static properties (fonts, for example; also some installable software) are protected by a simple "unlock" scheme: a purchaser makes a purchase, for example by telephone with a credit card, and receives a cryptographic key in return. This key can then be used to "unlock" one property from some widely distributed medium (e.g., CD-ROM or network download). This mechanism is relatively inflexible, and its inherently manual nature makes it expensive.

3.4 Billing Schemes

Various billing schemes (as mentioned above) permit purchase of information following what is essentially an electronic check or electronic credit draft model. These methods are suitable for conventional transactions, but not for the enormous volumes of (individually) very low-value transactions that would be generated using a complex digital property.

3.5 Secured Delivery

Various secured delivery systems (e.g., SSL [19], SHTTP [20]) share the same problems as cryptographic unlock, but in a network context. They are only point-to-point solutions, with the information (content, usage data, etc.) at each site being left unprotected once the delivery has occurred. Furthermore, they are inherently online systems: it is not practical to decouple the delivery of information from payment for its use.

4 Information Protection Architecture: InterTrust and DigiBox

EPR has produced the InterTrust Virtual Distribution Architecture to solve unmet, critical needs of electronic commerce. Almost any imaginable information transaction can be supported by InterTrust. A few examples include distribution of content (e.g., text, video, audio) over networks, selective release of data from a database, controlled release of sensitive information, and so on. InterTrust can also support the secure communication of private information such as EDI and electronic financial transactions, as well as delivery of the "back channel" marketing and usage data resulting from transactions.

DigiBox is a foundation technology within InterTrust. It provides a secure container to package information so that the information cannot be used except as provided by the rules and controls associated with the content. InterTrust rules and controls specify what types of content usage are permitted, as well as the consequences of usage such as reporting and payment.

Within InterTrust, DigiBox containers can enforce a "distributed electronic contract" for value-chain activities functioning within an electronic distribution environment. This unique approach underlies EPR's information metering and digital rights protection technology. Electronic commerce infrastructure participants can use InterTrust to substantially enhance their network, security, or payment method solutions.

The DigiBox is a container for both digital property (content) and controls. It is used in conjunction with a locally secured rights protection application (discussed further below) to make content available as governed by arbitrarily flexible controls.

The DigiBox container mechanism is implemented in a set of platform-independent class libraries that provide access to objects in the container and extensions to OpenDoc and OLE object technologies. DigiBox allows rights management components to be integrated with content in highly flexible and configurable control structures. Digi-

Box rights management components can be integrated with content in a single deliverable, or some or all of the components can be delivered independently. DigiBox rights management components enable true superdistribution [21] and can support virtually any network topology and any number of participants, including distributors, redistributors, information retailers, corporate content users, and consumers.

4.1 Content

The digital information in a DigiBox (one or more "properties") is information in any form. It may be mapped to a specific compound object format (e.g., OpenDoc, OLE, PDF), or may be application specific.

Further, it may be delivered in stream or other communication-oriented forms, not just in a file-like container.

4.2 Controls

Controls specify rules and consequences for operations on content. Controls are also delivered in a DigiBox, and the controls for a property may be delivered either with the property or independently. Controls are tied to properties by cryptographic means.

Because controls can be delivered with properties in a container, the DigiBox supports superdistribution.

4.3 Commerce

Commerce takes place governed by controls. This may involve metering, billing for use, reporting of usage, and so on. These operations take place locally in a secure environment, and they generate audit trails and reports that must be reported periodically to clearinghouses.

5 DigiBox Implementation

The DigiBox is a structure that can hold, in a protected manner, information commerce elements of all kinds: content, usage information, representation of financial transactions (e.g., electronic cash), and other digital elements of information commerce.

5.1 Container Logical Structure

Figure 4 shows the logical structure of properties and control sets in two containers. Container C_1 holds two properties, P_1 and P_2, and one control set, CS_1, that applies to property P_1; container C_2 contains two control sets and no properties. As shown in the example, each of these elements has a title attribute to provide a human-readable description of the element and, for control sets, an attribute indicating to what other elements the control set applies.

A control set specifies rules and consequences, such as pricing, reporting, and so on, for the properties to which it applies. A user holding just this container could use (e.g., view, print) content from P_1—though only as specified by CS_1. Because there is no control set applying to P_2 in that container, P_2 would not be usable in any way.

A user holding both containers could use property P_2, as specified by CS_2, and in addition has the choice of whether to designate CS_1 or CS_3 when using P_1. CS_3, which describes itself as "discount," is likely to be the user's preferred choice.

The DigiBox includes several elements: organizational structures, properties, controls, and supporting data items. Almost all the information in a DigiBox is encrypted, as described below, and access to the encrypted form is provided through a storage manager as appropriate, depending on how the DigiBox is delivered (e.g., as a file or as a data stream).

5.2 Container Physical Structure

Figure 5 is a schematic picture illustrating the physical structure of a DigiBox container. (Some elements have been omitted for clarity.) It begins with a *container header* structure containing descriptive and organizational information about the container. Part of the container header is encrypted (both for secrecy and for integrity protection); the rest is public organizational informa-

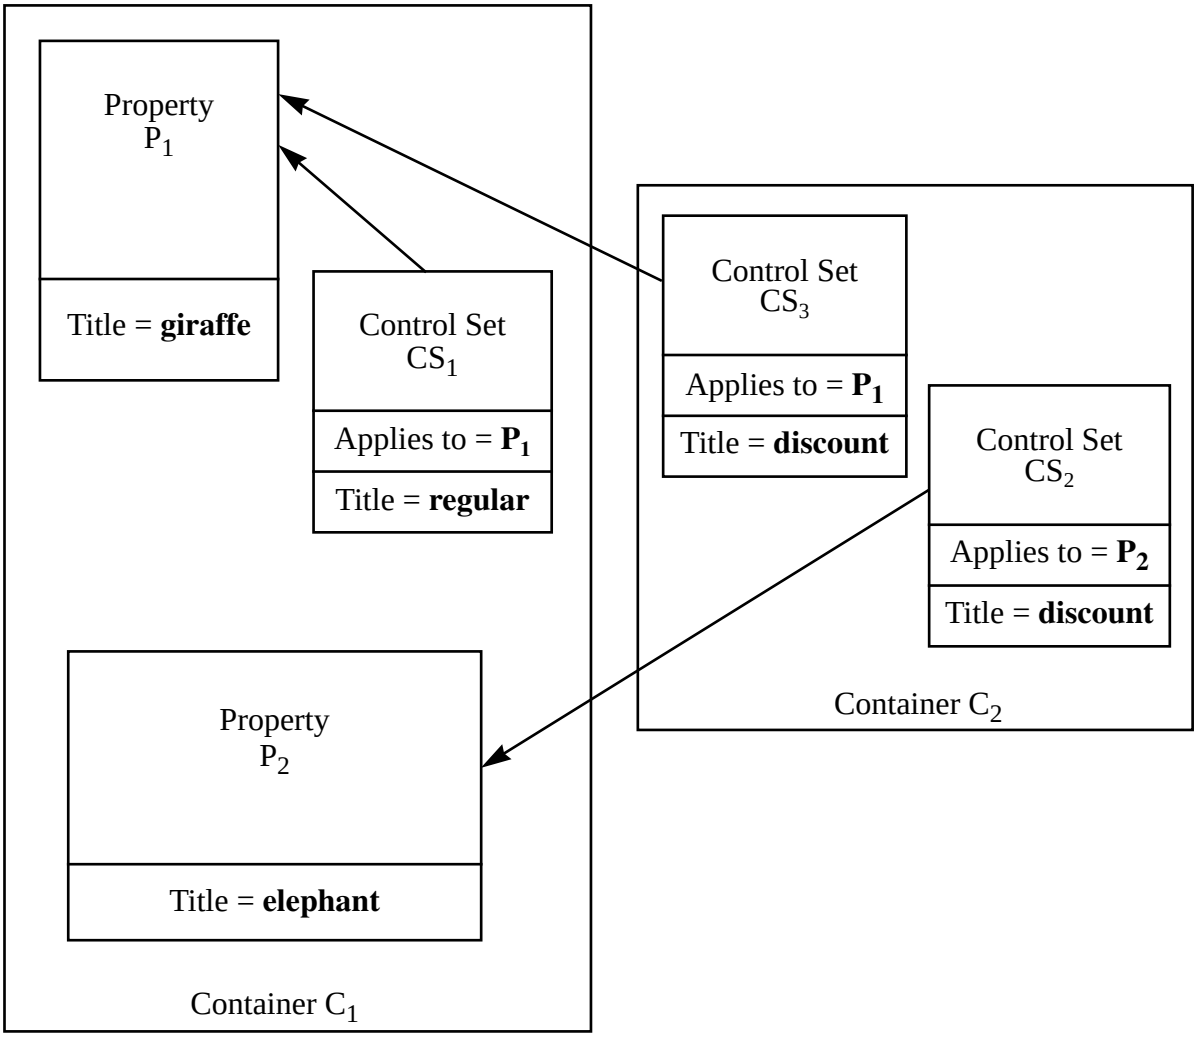

Figure 4. Container logical structure.

tion. The header is followed by additional container-wide structures such as the *transport key block (TKB)* and the *container table of contents (TOC),* some of which are encrypted and others not.

These organizational elements are followed by the structures defining the container's content (e.g., *properties* and *control sets*). As shown in the figure, a property is represented by a *property header*, *property attributes,* and data blocks composing the property. As shown, the header is encrypted and the attributes are not; the data blocks may be wholly or partly encrypted, or not at all, depending on security requirements.

The figure shows an example property consisting of a multimedia property formed from a pair of synchronized data streams for audio and video. In this example, each video block is mostly unencrypted so that access can be rapid while still maintaining reasonable security—encrypting even 10 percent of an MPEG stream renders it effectively useless for illicit copying. On the other hand, the audio is entirely encrypted, and each audio block

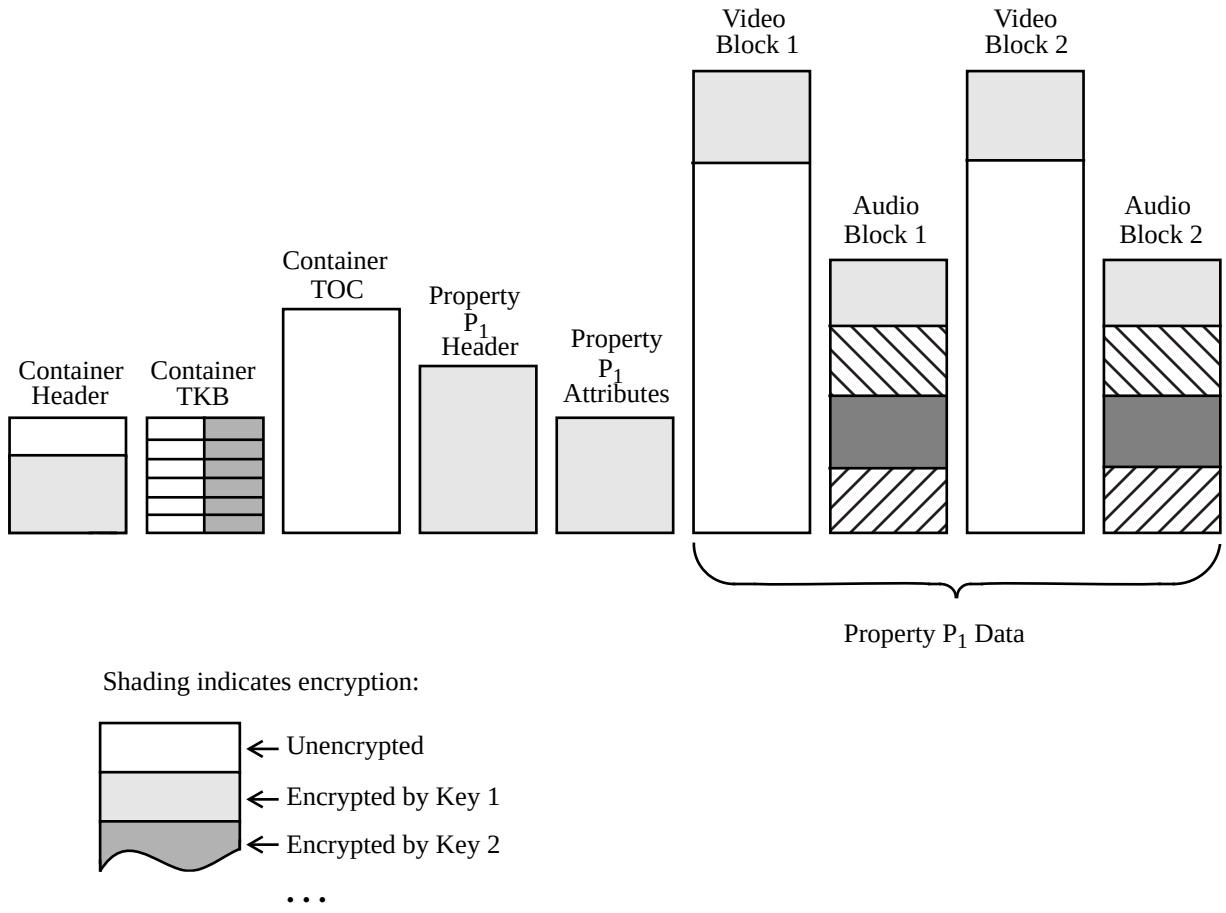

Figure 5. Container physical format.

uses four distinct keys, because the content proprietor requires much stronger security for audio than for video.

A property is represented as one or more property sections, each of which is independently associated with control information, and which may also be stored and accessed independently. A property, for example, might be a collection of clip-art images, and each image might be a property "chunk," with its own control specifying how that image's creator is compensated.

Controls can map to property chunks at arbitrary granularity and can enforce arbitrary organizational structures within the property (such as a file hierarchy). Controls can apply to individual bytes, frames of a movie, segments of a musical piece, and so on, because the mapping is performed by a control process specified by the control structure, not simply via a table-driven data structure.

5.3 Cryptographic Techniques

The high-level elements in a DigiBox are encrypted with a *transport key* that is normally derived (by exclusive OR) from two parts: one that is delivered in the DigiBox itself, encrypted with a public key algorithm, and the other that is stored in protected storage locally. The locally stored part is shared among all the local nodes capable of processing that DigiBox, but the part in the DigiBox is unique. This separation provides protection against accidental or malicious disclosure of either part.

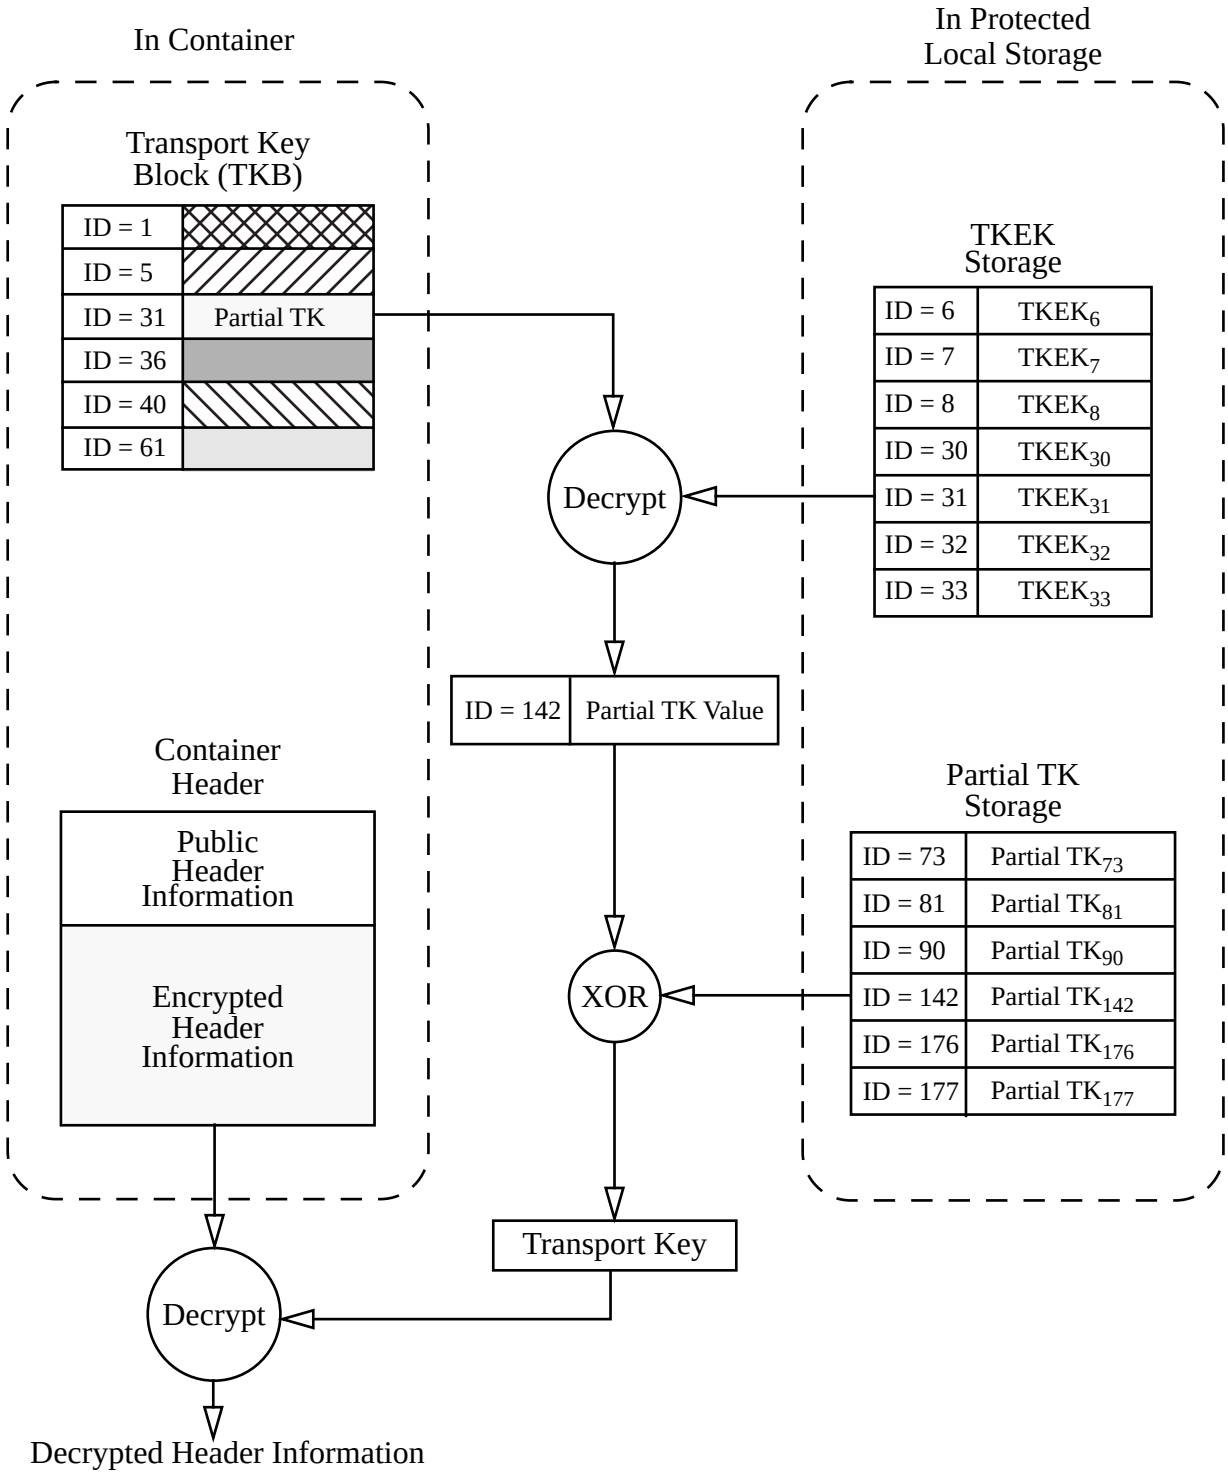

Figure 6. Container transport security.

Figure 6 illustrates how the transport key (TK) is derived. The transport key block (TKB) contains one or more slots, each of which contains a partial transport key encrypted under a different transport key encrypting key (TKEK). Each TKB slot identifies the TKEK used, and a matching TKEK is

selected from local protected storage. Decrypting the slot yields a partial TK, which is combined with its corresponding partial TK again from protected local storage to yield the actual TK for decrypting the container header.

The data for the property itself is encrypted with other keys ("content keys") that are themselves delivered in encrypted high-level structures; this approach permits the keys for a property to be delivered entirely separately from the property or its controls. Multiple keys, in a wide variety of key-mapping schemes, are used to encrypt the data, limiting the loss that would occur from disclosure of any one key.

All DigiBox control structures are both encrypted and verified for integrity with a cryptographic hash function. Several cryptographic algorithms are supported for these control structures (principally for export control reasons), and arbitrary algorithms are supported for encryption of the data.

5.4 Security Characteristics

The DigiBox cryptographic structures are designed to be secure even in the face of loss of individual key components, and to minimize the damage in case a key or processing environment is compromised. The system is designed to provide commercially acceptable risks and losses for a variety of business models.

The basic algorithms are strong: Triple DES [22] and RSA [23] are preferred. This security is, of course, only as strong as the tamper-resistance of the local processing environment. The preferred implementation of DigiBox processing relies on a "secure processing unit" (SPU) that contains a CPU, memory, program storage, and key storage in a single tamper-resistant hardware package. Although these are not widely available today, the variety of applications they might support makes it likely that such SPUs will become widely integrated into common computing platforms. When running in an SPU, the DigiBox processing and control mechanisms are sufficiently well protected to support most commerce applications.

In the absence of an SPU, other approaches are useful for many business models. In fact, a software-only implementation is sufficient for many applications, because much content is of relatively low value and is used in a context (business to business) where a modest level of fraud is both less likely and more tolerable. As long as the software is moderately difficult to defeat and tools to defeat it have no legitimate purpose, business models can be supported where some risk of loss is acceptable. In the world of electronic commerce, just as for traditional commerce, security is not absolute: it is just a factor to balance against the cost of loss and fraud.

6 Conclusions

The DigiBox is one component of a general-purpose electronic commerce solution that rests on three basic principles: rights protection, interoperability, and strong security.

Electronic commerce, and information commerce in particular, needs a robust information protection mechanism, including rights protection and controls, not just payment systems. As the electronic world evolves, however, and moves forward from simply emulating traditional transactions into entirely new business models, rights protection and control will become the predominant issues.

Protection of intellectual property rights in information requires strong cryptography as well as a flexible infrastructure for controlling use of the information. A standard protected container for information is necessary to support interoperability—most existing schemes tightly bind the creator of protected information and the software that processes it. A standard container can rationalize information commerce and reduce costs for all participants.

In the long term, general-purpose secure electronic commerce will need pervasive deployment of tamper-resistant hardware devices to perform secure processing of protected content. However, as these solutions are developed, many business models can be accommodated with weaker or less complete solutions because the risk and expected losses are commercially acceptable.

Business-to-business purchasing is steadily evolving into a direct electronic ordering model. Future communications and media markets will become increasingly segmented and specialized in response to customer preferences and needs and involve increasing, and more sophisticated, direct interaction between consumers and providers. These markets and their value chains (with or without intermediary distributors) will require secure metering and control tools that enable a user to efficiently and economically tailor resources to his or her own desires.

During the next decade, digital delivery of traditional electronic products, such as information databases and software, will be joined by a rapidly growing array of both New Media and electronically distributed traditional content. The conversion from traditional models requires key foundation technologies and will result in a fundamental shift in current infrastructure. This transformation will create a new distribution industry. Digital distribution employing a universal content and commerce container can play a critical role in this broad economic transformation.

7 References

[1] A. Chandler and H. Daems, "Administrative Coordination, Allocation, and Monitoring: A Comparative Analysis of Accounting and Organization in the U.S.A. and Europe," *Accounting, Organizations and Society,* 1979: 3–20.

[2] O. Williamson, "The Modern Corporation: Origin, Evolution, Attributes," *Journal of Economic Literature* XIX (1981): 1537–1568.

[3] Office of Technology Assessment, *Accessibility and Integrity of Networked Information Collections.* Washington, D.C.: U.S. Government Printing Office, July, 1993.

[4] E. Hollings, *Communications Competitiveness and Infrastructure Modernization Act of 1990.* Washington, D.C.: U.S. Government Printing Office, report of the Senate Committee on Commerce, Science, and Transportation, 12 September 1990.

[5] R. Benjamin and R Wigand, "Electronic Markets and Virtual Value Chains on the Information Superhighway," *Sloan Management Review,* Vol. 36 No. 2 (1995).

[6] U.S. Constitution, Article 1, Section 8, Clause 8 (1787).

[7] U.S. Copyright Act of 1978

[8] 17 U.S.C. s107

[9] 17 U.S.C s102(a)

[10] T. Berners-Lee, R Caillian, and J.-F. Groff, "The World Wide Web," *Computer Networks and ISDN Systems,* Vol. 25 (Dec. 1992), pp 454–459.

[11] D. Chaum, "Achieving Electronic Privacy," *Scientific American,* August 1992, pp 96–101.

[12] M. Sirbu and J. D. Tygar, "NetBill: An Internet Commerce System," *IEEE CompCon Proceedings*, March, 1995, pp 20–25.

[13] D. Gifford et al., "Payment Switches for Open Networks," *IEEE CompCon Proceedings,* March, 1995, pp 26–31.

[14] S. Dukach, "SNPP: A Simple Network Payment Protocol," MIT Laboratory for Computer Science, Cambridge, MA, 1993.

[15] B. C. Neuman and G. Medvinsky., "Requirements for Network Payment," *IEEE CompCon Proceedings,* March, 1995, pp 32–36.

[16] First Virtual, Inc. "Introducing the First Virtual Internet Payment System," 1994.

[17] A. K. Choudhury, et al., "Copyright Protection for Electronic Publishing over Computer Networks," June 1994, *IEEE Network Magazine*.

[18] J. Erickson, "A Copyright Management System for Networked Interactive Multimedia," *Proceedings of the 1995 Dartmouth Institute for Advanced Graduate Studies,* 1995.

[19] K. Hickman, “SSL Reference Manual,” Netscape Corporation World Wide Web Site, `http://www.netscape.com/newsref/std/sslref.html`, 1994.

[20] E. Rescorla and A. Schiffman, “The Secure HyperText Transfer Protocol,” Internet Draft draft-resorla-shttp-0.txt, 1994.

[21] B. Cox, “Superdistribution,” *Wired*, Sept. 1994, pp 89-92.

[22] U.S. National Bureau of Standards, “Data Encryption Standard,” *Federal Information Processing Standards Publication,* FIPS PUB 46-1, Jan. 1988.

[23] R. Rivest, A. Shamir, and L. Adleman, “On Digital Signatures and Public-key Cryptosystems,” *Communications of the ACM*, Vol. 21 (Feb. 1978), pp 120–126.