The Long March to Interoperable Digital Rights Management

ROB H. KOENEN, SENIOR MEMBER, IEEE, JACK LACY, MICHAEL MACKAY, AND STEVE MITCHELL, MEMBER, IEEE

This paper discusses interoperability of digital rights management (DRM) systems. We start by describing a basic reference model for DRM. The cause of interoperability is served by understanding and circumscribing what DRM is "in the whole." Then we outline and contrast three different approaches to achieving interoperability. One approach relies on flexible network services to provide functionality where it is needed, perhaps by bridging different systems. We describe an experimental service orchestration system (NEMO) that enables such an approach.

Keywords*—*Digital media distribution, digital rights management (DRM), standards, trusted computing, Web services.

I. INTRODUCTION

Digital rights management (DRM) is a collection of technologies that enable technically enforced licensing of digital information. DRM makes it possible for commercial publishers to distribute valuable content electronically, without destroying the copyright holder's revenue stream. DRM can also be used in other settings to enable safe distribution of digital content including, for example, document management within and between corporations, protected e-mail, medical patient records handling, and government service access.

At a minimum, a well-designed DRM system provides the following.

Governance: DRM is different from classical security and protection technologies [1]. Conventional media distribution systems based on conditional access techniques protect media during transmission using a control model based on direct cryptographic key exchange. DRM systems, on the other hand, implement control, or governance, via the use of programming language methods executed in a secure environment.

Secure Association of Usage Rules With Information: DRM systems securely associate rules with content. These rules determine usage of the content throughout its life cycle. Rules can be attached to content, embedded within content (e.g., via watermarking), or rules can be delivered independently of content.

Persistent Protection: DRM systems are designed to protect and govern information on a persistent basis throughout the content's commercial life cycle. Protection is frequently provided using cryptographic techniques. Encrypted content is protected even as it travels outside of protected distribution channels.

The use of DRM in commercial end-consumer media distribution is controversial for several reasons. DRM allows content providers to create licenses that are different from, and more rigidly enforceable than, the *de facto* generally understood licenses that have accompanied traditional media (CDs, VHS tapes, and DVDs). Conversely, the nature of today's DRM technology makes it difficult to automate accurately some existing usage conventions, such as the United States' fair use traditions or European privacy expectations.

DRM license enforcement requires security safeguards on home equipment to protect the interests of content vendors. Although it is common for basic utility vendors to install security systems around home metering systems (e.g., cable television, water, electricity and natural gas), some consumers are wary of DRM systems operating on their family PC, which is used for many personal tasks besides presenting media.

Traditional media distribution (before the mid-1990s) has been tied to physical media, such as music CDs and video tapes. Making and distributing high-quality copies of music and video was difficult for the average consumer. Successful business models have been well established around the processes of manufacturing, distributing, merchandising, and charging consumers for individual copies of a work. Early electronic distribution systems have likewise been built around the notion of digital copies of works ("copy control systems"), but this paradigm is becoming less relevant as it becomes easier for consumers to manage content as disk files on their home network, in their cars, at work, and in school.

It is easy today to find consumers who would think it appropriate to pay full price for a second factory-pressed

Manuscript received September 12, 2003; revised December 22, 2003.

The authors are with InterTrust Technologies Corp., Santa Clara, CA 95054 USA (e-mail: mmackay@intertrust.com; rob@intertrust.com; lacy@intertrust.com; mitchell@intertrust.com).

Digital Object Identifier 10.1109/JPROC.2004.827357

copy of a favorite music CD, but who have few misgivings about downloading free (unauthorized) digital compressed copies of music for which they (or someone in their family) already own a commercial CD. Consequently, consumers are developing their own ideas of what the right business models should be for commercial music licensing. Commercial publishers are scrambling to work through the business and technical hurdles to deploying business models that protect their interests and are acceptable to consumers, device manufacturers, and service providers.

The result is the emergence of DRM-enabled digital music services, such as Roxio's Napster service (originally known as pressplay), Apple's iTunes Music Store, Musicmatch Downloads, and others. Apple's music service has so far been the most popular with consumers, but we have not yet heard the last word in legal online music distribution [2]–[4]. BuyMusic, Musicmatch, MusicNow, Napster, and numerous others use Microsoft's Windows Media Audio format, which bundles DRM capability with an audio codec and a file format. Apple's iTunes uses an open standard audio codec [MPEG Advanced Audio Coding (AAC)] and a proprietary DRM system. The Microsoft and Apple formats are not compatible. Microsoft's format is supported on the largest variety of portable music players, while Apple's format is currently supported on only one—its own iPod. (Reportedly this is the current top-selling music player [3].) At the time of writing, no portable music player supports both formats.

This paper focuses on the issue of DRM interoperability. There are several reasons why DRM interoperability is desirable. The content industry desperately needs to deploy legitimate content services that compete favorably (based on features, not on price) with unauthorized free services. A simple and seamless user experience must be part of that goal, and DRM interoperability is necessary to achieve it.

Content providers and e-commerce service providers would like to see a healthy business climate from which they can multisource essential technologies like DRM, especially when these technologies must adapt rapidly to evolving industry needs and consumer expectations. The DRM market is strongly influenced by network effects: a DRM technology becomes more valuable as it becomes more widely adopted. Thus, there are strong forces pushing DRM technology providers toward interoperability, even as vendors attempt to differentiate their products based upon features.

While many people have articulated a goal for media distribution where any content is available to anyone, anytime, anywhere on any useful device using viable business models, significant barriers exist to the goal of an interoperable and secure world of media-related services.

- Overlapping *de facto* and formal standards.
- Implementation technologies are not interoperable.
- Consumer devices cannot locate and connect to needed services.
- Web services standards do not bridge services spanning Web distribution and personal area network protocols.
- Impedance mismatches between different trust and protection models.
- No unified notion of content governance useful in peer-to-peer (P2P) distribution models.

We outline some of the possible approaches to achieving interoperability and discuss related issues. We start in the next section by describing a basic reference model (RM) for DRM. The cause of interoperability is served by understanding and circumscribing what DRM is "in the whole." We then outline and contrast three different approaches to achieving interoperability. One approach relies on flexible network services to provide functionality where it is needed. Finally, we describe an experimental service orchestration system (NEMO) that enables such an approach.

II. Toward a DRM Basic RM

Commercial practice across a variety of DRM systems has matured to a point where robust technical patterns can be identified as a basis for establishing a DRM basic RM.[1] In this section, we consider the architecture of current DRM systems in order to identify common technical elements and the requirements they try to address. Proceeding from this analysis, we then outline an RM that may serve as a basis for coordinating evolution and interoperability of next-generation DRM systems. Establishing a general vocabulary and a set of reference concepts is the first step in building a framework for interoperability of heterogeneous systems.

A. *Current DRM Architectures and Industry Practice*

Fig. 1 illustrates an abstract system architecture based on DRM application and service elements representative of a variety of contemporary commercial DRM systems. Key concepts in this diagram are as follows.

- Content and associated usage rights enter the system through a packaging process, typically under the authority of the content licensor.
- Packaging services produce protected content and either full licenses, or rules and metadata as input to a licensing and reference service. Licenses can usually be personalized based on the particular parameters of the license-requesting party [5].
- Consumers use a local consuming application to transact with the licensing and reference services for licenses, and interact with streaming or download services for acquisition of the protected content. Often, the licensing service provides the reference to the correct content and associated distribution source.
- The consumer may be licensed to transfer protected content to another peer system (e.g., other "full-featured hosts"), or to a portable device with DRM capabilities. Portable or "tethered" devices interact with the DRM system by proxy via a more capable upstream

[1] The CEN/ISSS Digital Rights Management Final Report [16] provides an overview of evolving DRM technical architectures with the goal of "identifying the current status of DRM usage and possible means to ensure effective implementation of DRM in the marketplace."

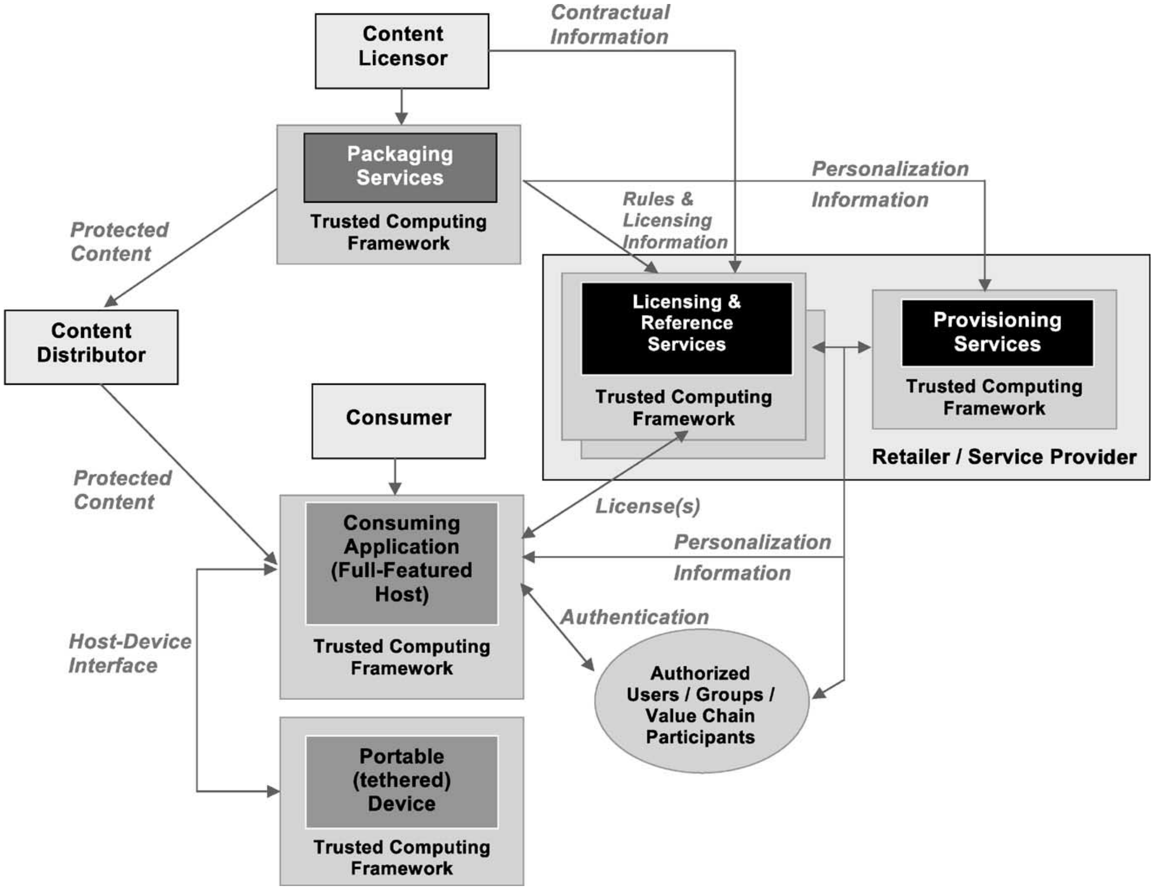
Content Licensor
Contractual Information
Packaging Services
Trusted Computing Framework
Protected Content
Personalization Information
Rules & Licensing Information
Content Distributor
Licensing & Reference Services
Trusted Computing Framework
Provisioning Services
Trusted Computing Framework
Retailer / Service Provider
Consumer
Protected Content
Consuming Application (Full-Featured Host)
Trusted Computing Framework
License(s)
Personalization Information
Authentication
Host-Device Interface
Authorized Users / Groups / Value Chain Participants
Portable (tethered) Device
Trusted Computing Framework

Fig. 1. Abstract DRM systems elements.

system (e.g., the "full-featured host"). The host may for example create a restricted form of the original license better suited to the capabilities the device, or may buffer or cache certain usage information on behalf of the less capable device.

Each of the elements in Fig. 1 may consist of multiple systems in a real-world implementation. For example, licensing services may embody an entire distribution value chain consisting of retail, subscription or download services.

Each element may be hosted by different business entities, acting in cooperation with other parties' systems based on contractual business relationships. Current deployment scenarios for DRM systems involve mutually well-known business partners, carefully architected technical responsibilities, and negotiated business relationships. However, increased business automation and more dynamic business relationships create the need for flexible provisioning and management of DRM infrastructure.

DRM applications and services (consumption, packaging, license services, provisioning services, etc.) are all built on elements of the trusted computing framework, which includes secure software distribution and execution environments, trusted identity management, secure policy and rule processing and enforcement, supporting cryptographic functions and key management, and tamper resistance. Provisioning services support adding new participants and services, and supplying DRM systems with supporting software, certificates, etc.

The ability to programmatically configure and manage trusted and secure relationships between the participants and the underlying DRM technology is paramount [6]. All of the parties in the value chain must trust that distributed content or information and its source are authentic, is accessible only by intended or contracted receivers, and is used by those receivers consistently with the contracted rights. Devices and services must be qualified as trustworthy and then maintained as such.

B. *Value Chains and DRM Systems*

Understanding roles in the commerce value chain and how these interact with DRM services is essential.

A detailed model of roles involved in electronic copyright management systems was developed by the European Commission-funded Imprimatur project. Completed in 1998, the goal of Imprimatur was to "understand and analyze the context in which Electronic Copyright Management Systems are to be developed," and which "reflect[s] current business practices for trading and licensing multimedia documents [by identifying] relevant roles, their relationships and corresponding transactions" [5]. Roles and responsibilities addressed by the Imprimatur model include the following.

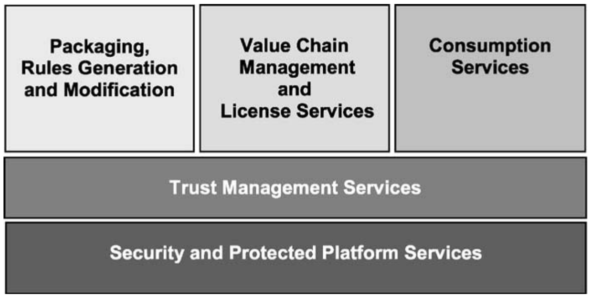
Packaging, Rules Generation and Modification
Value Chain Management and License Services
Consumption Services
Trust Management Services
Security and Protected Platform Services

Fig. 2. DRM basic RM.

- The creator—the party responsible for delivering their creation to the creation provider.
- Creators may assign exploitation rights to a rights holder (e.g., a collection or licensing agency).
- The relationship between creators and rights holders and associated contracts are maintained in an IPR database.
- The media distributor is expected to pass appropriate royalties to the rights holder according to the current payment details stored in the IPR database.
- The purchaser (consumer) may use the creation, and if they generate a new composite document based on it then they also become a creator. In order for the purchaser to perform functions associated with the creator role, they must have obtained the required permission from the corresponding rights holder of the original creation. Rights holders of original creations automatically have rights on composite creations—the flow of royalties is determined according to the IPR database.

Few DRM systems take all of these types of roles, relationships, and activities directly into account as part of their intrinsic design, leaving contract management and auditing and accounting issues to a diverse array of largely unintegrated back office systems. With increased end-to-end systems automation and sophisticated digital content manipulation and aggregation services, models like Imprimatur will likely receive increased attention in new architectures. Possibly the most thorough attempt to date in a single DRM system was undertaken by InterTrust in its Commerce system [7].

C. DRM Systems Functionality

The proposed basic DRM RM is illustrated in Fig. 2. We now frame the functional characteristics of the five main domains of our proposed basic RM.

1) Packaging, Rules Generation, and Modification: The point of entry to the DRM-managed content and governance life cycle includes technologies supporting content packaging, specification of rights and associated data, and generation and modification of digital items.

a) Content Packaging: Content packaging is the process of preparing content for DRM protection—placing content into a secure container, usually by encrypting it, associating the necessary identifiers and metadata, and logging and cataloging the content, its identifiers and metadata, and its cryptographic material. Consumers and associated consumption processes may also be enabled to package their own content.[2]

Content packaging can be closely associated with rules and license generation or may be completely independent from it. Content identifiers couple the protected content with rules and content protection keys. Therefore, rules, packaged content, and content keys may be generated together or separately, at the same time or at different times. They may be delivered together, through the same channels, or separately, at different times, through different channels. In a production environment, content may be packaged initially without rules. Alternatively, content may be packaged on-demand and immediately associated with rules.

The content may contain directions as to where licenses or offers associated with the content can be acquired or other offer metadata that can be used to automate downstream distribution processes.

Content protection is typically accomplished using cryptographic processing, where content protection keys are made available to one value chain participant or consumer, and are not exposed in the clear to other value chain participants or consumers. Key management procedures can bind or associate a content package to any security principal, including individual consumers, devices, certain types of secure media, or content-sharing networks (e.g., a network of home media devices). Associating content with a consumer allows the protected content and license to be transported to other systems on which the consumer is also authorized.

b) Rules Generation and Modification: Any authorized member of the value chain from packager to consumer may create rules to be associated with a content package. Rules may be used to govern consumer access to content as well as to govern the actions of other value chain members on the content or information associated with the content. For example, usage rules may require authentication on access or usage, or require license updates to be obtained before operating on the content.[3]

Rules may specify consequences such as generation of audit records based on content usage actions or attempts at usage, such that the audit records are securely delivered to a designated authority prior to execution of the action governed by the rule.

Rules are often associated with the whole piece of content, but may also be managed at the granularity of a content subelement (e.g., stream, component, etc.). Rules can also be associated with a class of content (e.g., all content belonging to a particular owner, all audio content, all low-bitrate content, etc.) rather than a specific content instance.

Rules can be delivered as separate files (e.g., a license), or combined with the protected content (integrated with the content data format itself), or both. Alternatively, the rules

[2] The term "consumer" typically refers to retail end users but may also apply to other value chain participants—regardless, consumers are participants of the managed value chain and may participate in a broader class of functions than strictly consumption and rendering.

[3] For example, expired rights might require license updates to enable access or usage.

can be provided as input to value chain management and licensing services or applied in conjunction with processes for resolving references to the content.

Rules, terms and conditions, and consequences can be represented in a variety of different ways. For example, one approach is to use a standardized rights expression language such as the MPEG-21 Rights Expression Language (REL) [8] or the Open Digital Rights Language (ODRL) [9]. Alternatively, rules may also be encoded in formatted text (such as XML or named key-value pairs), or possibly via compiled or interpretive code as part of an application.

In some systems, it is possible to modify or extend rules after their initial creation. For example, value chain management and licensing services may support the ability to select and apply rules that have been updated to reflect up-to-the-minute changes in business offers, regardless of when the content was packaged and placed into the system.

In the final stage of rules generation, rules are embedded into data structures that can be linked to the content. There are a variety of mechanisms available for packaging rules. For example, sets of rules may be organized into "offers" that describe the content and the associated license for presentation to a consumer or other value chain member. Offers may be delivered to a content distributor, who may choose to present some or all of the offers to other participants further down the value chain. Associated collateral information and promotional content can be included in a separate package for use in retail promotion and downstream distribution.

2) Value Chain Management and License Services: A common characteristic of systems that support nontrivial operational models (such as subscriptions, superdistribution, push-distribution, etc.) is the ability to produce, modify, assemble, and aggregate rules and negotiate conflicts involving rules from one or more sources.

Consumer licenses are sometimes the result of a collaboration of multiple value chain participants. Authorized value chain members may insert new rules into the licensing structures, using processes that are themselves governed. The rights of various services to interact with the content's distribution process may be encoded in rules delivered directly to the service or that are referenced using the same identifiers or references that are associated with the content.

Value chain management services may include posttransaction processing (e.g., allocation of the value exchanged such as financial payment, usage data, etc.) per contractual obligations [5]. Such posttransaction processing rules can be included in the license associated with the content (whether packaged together with the license or separately), or created as an electronic contract covering specific offers or content and delivered separately.

Historically, the terms by which value chain participants are allowed to interact with the content and rights to its use are expressed via contractual relationships between creators (or creation providers) and other value chain participants. We anticipate that contractual relationships may be automated using similar mechanisms (e.g., electronic contracts) as those used to control access to content by consumer applications. Contracts may be encoded using a contract expression language [10], similar to RELs used for encoding content usage rights. Electronic contracts are then delivered to participating entities and used by trusted applications to manage content distribution rights. The ways in which these terms are delivered and managed are discussed in greater detail in the next section.

Frequently, rights and contractual obligations associated with a piece of content already exist as a result of prior interactions with the content (e.g., as part of prior distribution arrangements). Rights discovery refers to a set of functions provided either by technically automated or other means, such as conventional business processes, for referencing these existing rights and obligations.

a) Value Chain Management: Value chain management refers to those system facilities that track, serve, and govern value chain participants. Value chain participants have interests in the distribution of products and provide decision-making, reporting, and other processing services affecting the digital content under their control. Just as rules govern the use of protected content, rules and policy govern the ways in which value chain participants interact with one another and with their associated content.

Static value chain management refers to approaches where offer and consumption rules are computed at content packaging time. An expression of rules can be distributed with content packages for examination or modification by other participants in the value chain.

In the static model, content packages are created for a particular set of distribution participants. The value chain management process is parameterized at packaging time with information about the known and identified participants, and the packager output conveys the necessary information in advance of actual participation. Once packaged, modification to the value chain information is governed by the associated rule set. The upshot of this early-binding approach is that unanticipated business model changes might necessitate content and/or rules repackaging from an original source.

The **dynamic value chain management** model is late binding. In the dynamic model, rules governing the use of value chain information are accessed on demand through network services, rather than being carried as they were encoded at packaging in an early-bound and immutable configuration. Rather than copying packaged files to each value chain participant, content may be distributed by reference [10]. The rights to the content are distributed based on these references and the references may be incorporated in or used by other structures, such as licenses. Reference services fulfill requests for content consumption by consulting their current rule sets [10].

Dynamic value chain management allows for modification of the value chain information as references to the content move through the distribution channel. The dynamic model allows content to be packaged without advance knowledge of distribution configurations. Distribution configurations can change in response to new contracts, law, or business models. In addition to enabling greater adaptability and responsive-

Explore Litigation Insights

Docket Alarm provides insights to develop a more informed litigation strategy and the peace of mind of knowing you're on top of things.

Real-Time Litigation Alerts

Keep your litigation team up-to-date with **real-time alerts** and advanced team management tools built for the enterprise, all while greatly reducing PACER spend.

Our comprehensive service means we can handle Federal, State, and Administrative courts across the country.

Advanced Docket Research

With over 230 million records, Docket Alarm's cloud-native docket research platform finds what other services can't. Coverage includes Federal, State, plus PTAB, TTAB, ITC and NLRB decisions, all in one place.

Identify arguments that have been successful in the past with full text, pinpoint searching. Link to case law cited within any court document via Fastcase.

Analytics At Your Fingertips

Learn what happened the last time a particular judge, opposing counsel or company faced cases similar to yours.

Advanced out-of-the-box PTAB and TTAB analytics are always at your fingertips.

API

Docket Alarm offers a powerful API (application programming interface) to developers that want to integrate case filings into their apps.

LAW FIRMS

Build custom dashboards for your attorneys and clients with live data direct from the court.

Automate many repetitive legal tasks like conflict checks, document management, and marketing.

FINANCIAL INSTITUTIONS

Litigation and bankruptcy checks for companies and debtors.

E-DISCOVERY AND LEGAL VENDORS

Sync your system to PACER to automate legal marketing.

WHAT WILL YOU BUILD? | **sales@docketalarm.com** | **1-866-77-FASTCASE**

Smarter legal research.