

Towards Meeting the Privacy Challenge: Adapting DRM¹

Larry Korba
National Research Council of Canada
Ottawa, Canada
Larry.Korba@nrc.ca

Steve Kenny
Independent Consultant
Den Haag, The Netherlands
Stephen_MH_Kenny@yahoo.com

Abstract

There are many requirements for achieving the privacy needs as expressed in law. Currently there is no commonly accepted technical approach for meeting these privacy requirements. An often-fruitful way for uncovering solutions to new challenges is to examine how current technologies used in quite different applications may be adopted to meet the specific challenges. In this paper, we examine the prospect of adapting systems designed for Digital Rights Management for the purpose of Privacy Rights Management for European Community application. We begin by outlining the legal requirements for privacy under the European Union Data Directive. After an overview of digital rights management systems, we describe adaptations for transforming a DRM system into a privacy rights management system. In the conclusions we detail the strengths and weaknesses of the approach.

1. Introduction

Privacy issues facing developed societies today are made complex by incompatible ideologies and policies between the different countries, the Internet and the growth of new technologies in general. In this context, privacy issues are complex, from a technological perspective, due to Directive 95/46/EC of the European Parliament and the Council of 24 October 1995. This legislation, referred to as The Directive [1], describes the protection of individuals regarding the processing and free movement of their personal data. Many of the provisions of this Directive have the potential to become global de facto standards for e-business.

In this paper we investigate the potential of adapting Digital Rights Management (DRM) systems for the purpose of managing personal data held and controlled by organizations. For the purposes of this work, we define Privacy Rights Management (PRM) as the management of personal information according to the requirements of The Directive. The purpose of this work is to develop a framework for a broader integration of privacy services that would mitigate certain privacy-concerning characteristics of e-commerce systems in general. As well, through this exploration we uncover pertinent research issues that must be solved in order to develop robust, new, privacy-enhancing technologies.

1.1 Legislative Imperative

The right to privacy in the EU is defined as a human right under Article 8 of the 1950 European Convention of Human Rights and Fundamental Freedoms (ECHR). The implementation of this

¹ NRC paper number: NRC 44956

Article can be traced to The Directive. Similar legislation and enforcement structures to the European model exist in Canada, Australia, Norway, Switzerland and Hong Kong.

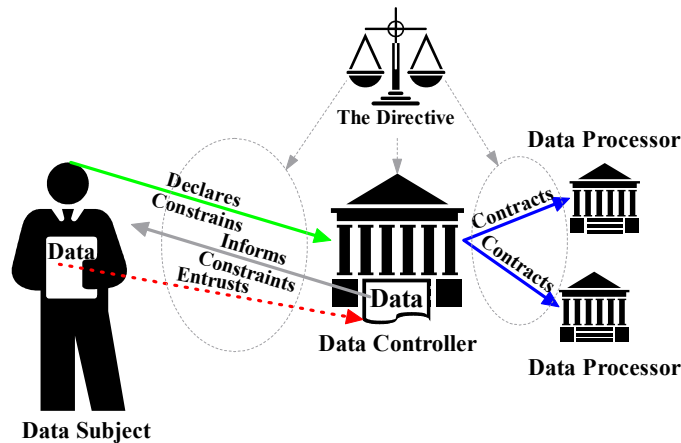


Figure 1. A schematic representation of the roles of the three entities defined in the Directive.

The Directive applies to all sectors of public life, with some exceptions. It specifies the data protection rights afforded to “data subjects”, plus the requirements and responsibilities obligated for “data controllers” and by association “data processors” [2]. This triad structure of entities balances data subject fundamental rights against the legitimate interests of data controllers (see Figure 1). The Directive places an obligation on member states to ratify national laws implementing its requirements. The implicit principles and constructs of The Directive define the enforcement and the representation of data protection. The terms privacy and data protection are often used interchangeably, though we are aware that in other contexts the two terms are not necessarily equivalent.

Table 1. European Union Privacy principles.

<i>Principle</i>	<i>Description</i>
1. Reporting the processing	All non-exempt processing must be reported in advance to the National Data Protection Authority.
2. Transparent processing	The data subject must be able to see who is processing his personal data and for what purpose. The data controller must keep track of all processing it performs and the data processors and make it available to the user.
3. Finality & Purpose Limitation	Personal data may only be collected for specific, explicit, legitimate purposes and not further processed in a way that is incompatible with those purposes.
4. Lawful basis for data processing	Personal data processing must be based on what is legally specified for the type of data involved, which varies depending on the type of personal data.
5. Data quality	Personal data must be as correct and as accurate as possible. The data controller must allow the citizen to examine and modify all data attributable to that person.
6. Rights	The data subject has the right to improve their data as well as the right to raise certain objections regarding the data controller's execution of these principles.
7. Data traffic outside EU	Exchange of personal data to a country outside the EU is permitted only if that country offers adequate protection. The data controller assures appropriate measures are take place in that locality if possible.
8. Data processor processing	If data processing is outsourced from data controller to processor, controllability must be arranged.
9. Security	Measures are taken to assure secure processing of personal data.

The data subject is a natural person who can be identified by reference to one or more pieces of data related to his physical, physiological, mental, economic, cultural or social identities. Even data associated with an individual in ambiguous ways may be deemed reasonably identifiable. Following Article 1 of the ECHR, the fundamental right to data protection falls not to the nationality of the data subject, but as an obligation to a relying party of the data subject [3]. The relying parties are the data controller and, by association, the data processor.

The data controller is an entity that determines the purpose and means of processing personal data, and is defined as the holder of ultimate accountability as it relates to the correct processing and handling of the information from the data subject. The data processor is an entity that processes personal data on behalf of the data controller.

Privacy principles (Table 1) abstracted from the complexities of legislation have been developed to simplify compliance with privacy regulations. Analyzing an approach using the principles as a guide, offers a fruitful means for determining the effectiveness and pitfalls of the approach. We use these principles to consider the appropriateness of adapting systems and ideas currently used in DRM for PRM.

1.2 Business Imperative

DRM systems are not without controversy regarding privacy. Since DRM systems track what users purchase, how often they access material, when they use it, it is clear that these systems may be used to track detailed activity of subscribers [4]. Currently, divisions are opening between content providers and technology developers regarding intellectual property protection, versus privacy protection. Technology providers are faced with attempting to please their corporate customers, i.e. content providers, who are being subjected to revenue atrophy from copyright abuses, versus the potential alienation tracking solutions generate within their customers.

DRM system mechanisms for capturing and tracking of personal data have incited concern from strong data protection bodies. It is clear that design assumptions such as extensive notification of organisational privacy policies coupled with controllability via external privacy auditing from reputable firms will be insufficient to quell concern.

Our position is that, notwithstanding the privacy issues with DRM systems, aspects of DRM architecture have features that would allow the development of a system-based approach to data protection compliance, i.e. Privacy Rights Management. Privacy Rights Management offers a solution for the paradox in which content deliverers find themselves. It embeds The Directive into a technology framework for protecting data subject information. Such an architecture may be applied to the management of personal data for many types of e-commerce applications. By implementing European style data protection rights ubiquitously through PRM, individuals are able to engage personalised content-provision business models, such as pay per play, in confidence that all their personal data is being processed legitimately.

The rest of this paper is organized as follows. Section 2 we state the problem we are addressing in this work, followed by a description of a basic DRM system. In Section 3 we describe the architecture of a PRM system, drawing parallels between its components and their counterparts in a DRM system. We show what the changes required to transform a DRM system into a PRM system. Section 4 describes mechanisms to express privacy using ODRL. Section 5 proffers a discussion on this analysis.

2. Problem Statement

Under The Directive, the data controller has a major data protection compliance responsibility. There are currently no technical solutions that would meet all aspects of The Directive. The problem focus of this paper is the development and analysis of a PRM architecture that meets the requirements indicated by the privacy principles of The Directive. Interestingly, Digital Rights Management technology, developed for protecting intellectual property rights, appears to offer the potential as a foundation for meeting these requirements. The next three sections of this paper describe how a generic DRM system may be adapted to offer Privacy Rights Management. We first start with an overview of digital rights systems.

2.1 DRM overview

Originally conceived to facilitate controlled distribution of digital content and to combat breaches of copyright law, digital rights management (DRM) involves all aspects of content distribution, ranging from content locking mechanisms, through content metering, to payment processes, and record keeping. DRM architectures support description, trading, protection, monitoring and tracking of the use of digital content. DRM technologies can control file access (number of views, lengths of views), altering, sharing, copying, printing, and saving. These technologies may be contained within operating systems, program software (e.g. specialized viewers), or in the hardware of a device. Figure 1 illustrates a simplified DRM system. In order to present the concept of privacy rights management we adopt the client-server rather than the peer-to-peer architecture for reasons of simplicity.

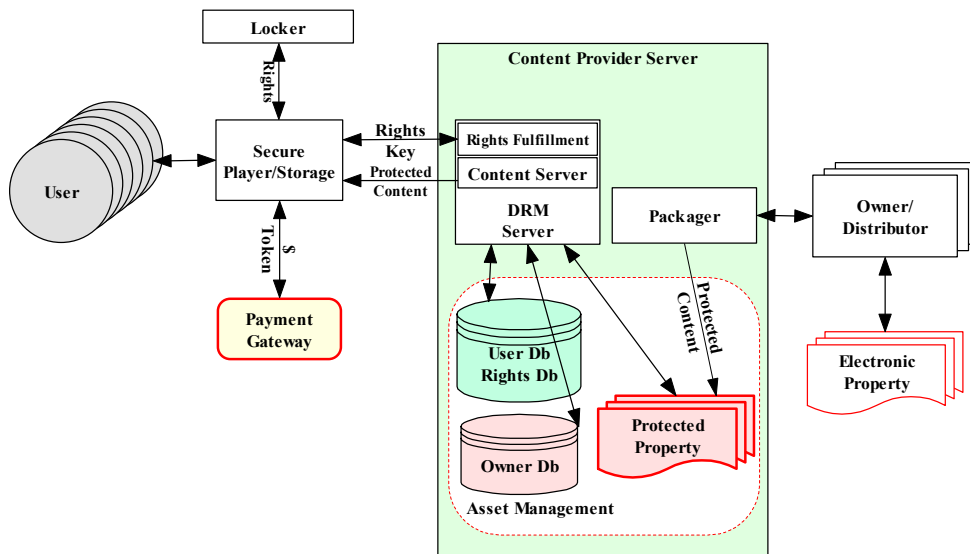


Figure 2. A simplified DRM system focussing on client-server architecture.

As illustrated in Figure 2, a DRM system operates in the following fashion. An owner or distributor submits its electronic property to a packager that encodes the property into an appropriate format for eventual end use. The packager encrypts the content to guard against unauthorized use, and adds metadata concerning the content. The metadata not only specifies the content, but also may hold information regarding how a user may gain access to the content. The DRM Server, sometime known as the Rights Fulfilment Server, manages assets stored within various databases. An important concept that forms a foundation for DRM is the separation

between the content and the rights for access to the content. Rights describe precisely what a user is allowed to do with the content. Typically some sort of language is used to express those rights (for example: XrML [5], and ODRL [6]). The Rights Management Language implements the business model for the commercial distribution of the content, providing details concerning different types of purchase models, use models, etc.

In order to view or play DRM managed content, the user must deploy client software on his computer. This client software handles user authentication and provides secured access to the content. The intention here is to ensure that only those entitled to a file will be able to access it. A challenging element of DRM is to ensure that the content may not be saved for other uses when it is available in the clear for legitimate processing (e.g. viewing an e-book, or video file, or listening to an audio file).

3. Privacy Rights Management

As is clear from the description of privacy principles in Table 1, there are many demanding requirements placed upon the data controller. In order to examine the possibility of meeting these requirements, we propose the system shown in Figure 3.

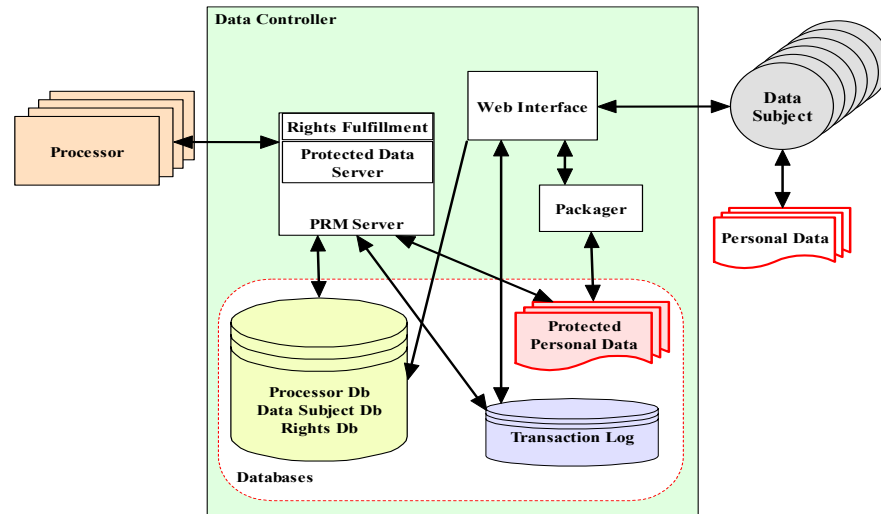


Figure 3. A simplified privacy rights management system showing the three participants: data subject, data controller, and data processor.

Key participants for the system include the data subject, the data controller (in this case a single data controller) and one or more data processors. PRM manages personal data from the data subject, the originator and the owner of the personal data. The Directive defines the authorities and boundaries of the relationships between each of the participants.

The data controller manages gathering, storage and processing of data subject data. The responsibility is enforceable through both national data protection authorities and the importance of preserving data controller reputation. There may be many data processors associated with a PRM system. A data processor may be an element operating under direct management of the data controller, or it may operate as a separate entity, under a contractual arrangement with the data controller. Since the data controller enlisted the data processor to render services, liability ultimately falls to the data controller for correct data processor operations. In this PRM system there may be many data processors dealing with data from many different data controllers. This

situation is very similar to the DRM case where a user may interact with many different content providers.

When comparing Figure 2 and Figure 3, definite parallels may be drawn between the PRM and DRM system components. Table 2 outlines these comparisons.

Table 2. DRM and PRM system component parallels.

<i>PRM Component</i>	<i>PRM Comments</i>	<i>DRM Component</i>	<i>DRM Comments</i>
Data Subject	The Data Subject entrusts a relatively small amount of data for management by the data controller. The data controller manages the data, including its distribution to data processors.	Owner	The Owner entrusts its electronic property to the DRM server for distribution. In contrast to its PRM counterpart, the Data Subject, there are relatively few Owners as compared to Data Subjects.
Data controller Web Server	The data controller acts as the enforcer of usage requirements associated with personal data, with accountability provided through detailed logging. The web server provides an interface providing data subjects with several views such as the 'objection view' where they can access, rectify, revoke and maintain their personal data. Data controllers are provided with management views of the PRM system operation.	Content Provider Web Server	The Content Provider Web Server provides an interface allowing owners to maintain personal data and for management of the system. The Owners may track usage and other information regarding their data.
PRM Server(s)	Privacy rules implementing triad entity rights, preferences and requirements are handled here.	DRM Server	The DRM server contains rules implementing the way in which owners' property and subscribers' interactions are managed.
Personal Data	Data provided by the data subject traceable to them in some way.	Electronic Property or Asset	The electronic property (content) entrusted to the DRM system for controlled distribution by the owner.
Protected Personal Data	PRM protected property is personal information entrusted to the data controller, held and distributed using data protection. The number of entities among which the property is shared (data controller and data processors) is smaller than in the DRM protected property scenario.	Protected Property	Protected property is held and delivered using data protection. Access to the property is controlled via a rights usage policy. There may be a very large number of people gaining access to the protected property.

The PRM server block provides base PRM services. Personal data in a PRM system plays a similar role to that of Protected Property in a DRM system (see Table 2). The Data Subject owns the data, and entrusts it to the PRM server wherein it is protected and managed by the data controller. Data subject profiles are treated as electronic property assets in DRM. In order to perform its functions, the server block must maintain and use different sets of data. As well, it

will manage data exchanges with processors to meet potentially widely varied processing objectives.

The PRM server maintains several databases. A rights database provides information regarding how personal data is managed within the system. There are also databases containing processor and data subject reference information, as well as activity logs for collecting information regarding system operation. Interestingly, while there is the potential for unbridled user tracking in a DRM system, when adapting DRM to PRM the tables are turned where the activities of both data controller and data processor are monitored. PRM data subject tracking would be strictly in accordance with the stipulations of The Directive's Article 7.

At the organisational level, there are also important distinctions between DRM & PRM. System elements within DRM models may well be operated by different legal entities. Thus the partner selection criteria for a privacy-conscious firm will naturally consider the degree of trust a potential partner presents regarding its privacy practices. One foresees several ways to achieve that credential, with the most obvious being extensive notification of organisational privacy practice, augmented with strong controllability via external privacy auditing. One aspect of interpretation of the security stipulations from The Directive from the perspective of Dutch data protection law is that contracts between data controller and data processor must provide assurance that data processors will enforce a security policy as rigorous as the one to which the data controller is subject. Service Level Agreements inclusive of bi-lateral audit rights are appropriate here.

3.1 DRM evolution to PRM

The three aspects of DRM functionality of interest to PRM are Asset Creation, Asset Management and Asset Usage.

Asset Creation (as illustrated in Figure 4) supports rights creation and validation. Rights validation ensures content may only be created from *existing* content if the rights exist to do so. Rights creation allows rights to be assigned to content. Below we schematically illustrate asset creation.

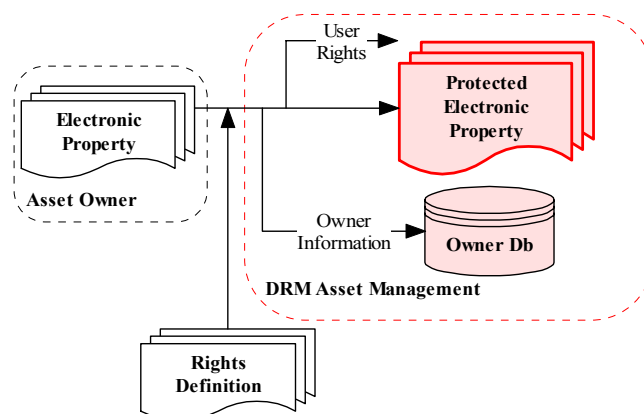


Figure 4. DRM asset creation.

The driving purpose behind DRM - content distribution management - relates easily to data protection constructs constraining the exchange of personal data. Article 6 (d) of The Directive builds arguments related to the responsibility of data quality on the part of the data controller and, by association, the data processor. Similarly, Article 12 (b) will require the data controller to

provide the data subject with opportunity to amend his or her personal data. In addition, the data must be of consistent quality in all its instances, and a retention period of personal data that is either based upon legitimate grounds or consented to must be upheld.

Asset Management supports the access and retrieval of both content and metadata in distributed databases. Asset Management also provides logging functionality. Article 6 (c) requires data controllers to process volumes of personal data that are minimised for the task at hand. More centrally, PRM asset management maintains data subject's rights over their data, which would be managed by a PRM asset management rules engine. The rules engine also codifies data controller interests so that, for instance, a data subject objection to a processing request, may not be complied with by the data controller, if the data subject has not explicitly consented to the processing.

The PRM system must implement a high degree of monitoring of subject data usage. As well, the monitoring process itself must be protected. With multiple data processors operating on personal data, the data controller is in a high-risk position if any data processor engages in illegitimate processing. A data controller requirement will therefore be for each contractually engaged data processor to maintain cryptographically protected log files [7] relating to the operations on every individual's personal data. In addition to meeting the requirements of The Directive, it would offer the data controller a means for data controller monitoring of privacy performance via log analysis.

Asset Usage supports permission management and (depending upon definition) audit trail functionality that permits the usage environment to honour rights associated with content. This offers a means for monitoring and tracking content use. Below we illustrate some functional elements of a DRM system especially required for content usage monitoring.

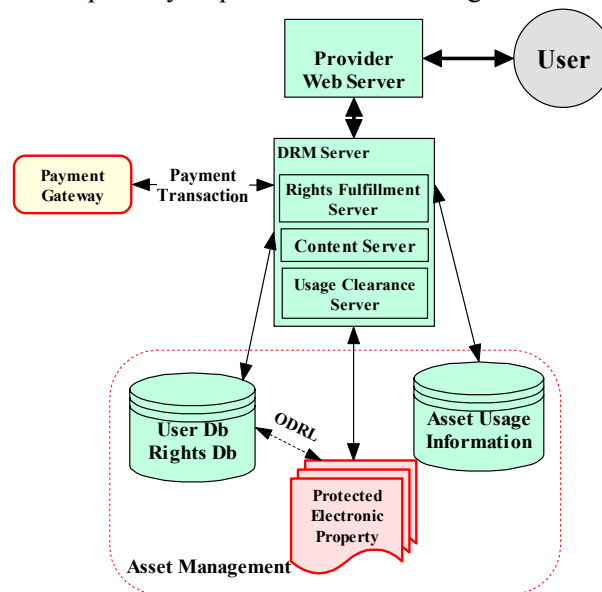


Figure 5. DRM asset usage.

The PRM server must extend the core logic associated with DRM server asset usage so as to support the PRM operational context: a large number of different owners of electronic property, many distributed data processors, as well as major responsibilities under The Directive (see Figure 5). Three key entities contained within a DRM server are the Content Server (CS), the

Rights Fulfilment Server (RFS) and the Usage Clearing Server (UCS). These entities are present, but reconfigured, in a PRM server. In DRM, the CS standard task is to distribute cryptographically packaged content, accessible by retrieving content and rights keys. In PRM, this is similar: the CS manages the controlled distribution of personal data assets. A significant difference in a PRM secure container however is that it may have a varied granularity level of asset protection and auditing requirements based upon role-based rules dictating and auditing access on the grounds of consent or other permissions as specified under Article 7 of the Directive.

The functionality provided by the RFS in DRM ranges from providing payment receipts to recording asset accesses and device sets. In PRM, the RFS enables the tracking of processor use of subject data. The Asset tracking databases must be tamperproof, to prevent unauthorised changes to the tracking records. Article 6 (b) of The Directive may be implemented by appending a retention period to personal data. This retention period is transfer-independent. If the period agreed to is 30 days, and the data controller passes this data to a data processor after 15 days, the data processor must conform to the remaining 15-day retention period. Once the retention period is exhausted, all instances of the personal data must be erased. Given this requirement, RFS functionality may be extended to coordinating asset usage information databases. This extension is required to meet Article 6 (b) of the Directive. To support temporal semantics, a secure timing mechanism linked to the database is required.

Basic UCS functions include recording and analysing transaction data. From Figure 5, the PRM server is advised by the rights database of the degree to which personal data may be disclosed to other parties, according to original data controller data capture conditions. There must be sufficient granularity in the operation of the usage clearance server to link different purpose specifications to different parts of data subject data. This permits implementation of both Article 6 (b) through the ability to identify which (element of) data is needed for each purpose, and Article 6 (d) via the retroactive and proactive updates necessary to assure data accuracy (plus audit trails) in the relationship between data protection concepts such as purpose specification and the personal data itself.

Underlying these PRM requirements is a concept of data subjects controlling their personal data in much the same way as content owners or distributors control and monitor access and use of their digital content using DRM approaches. Interestingly, DRM systems gain maximum leverage from personal data through tracking consumer activity and subjecting that output to data mining at a clearing agency. In PRM, data subjects are able to visualise and influence the amount, quality and granularity of tracking information generated from their data.

4. Expressing Privacy

In this section, we describe entity modelling for a PRM system based upon the Open Digital Rights Language (ODRL). ODRL is a standard vocabulary for expressing the terms and conditions for the use of assets [6]. Our approach may also be applied to extensible rights mark-up language (XrML) [5]. The XrML approach is described in a future paper.

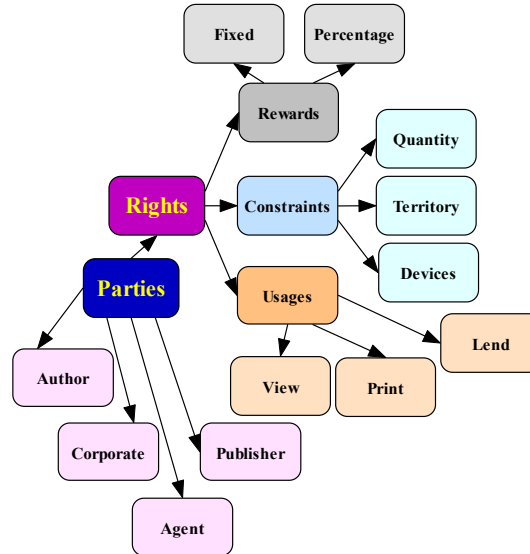


Figure 6. ODRL rights and parties model.

Modelling content is necessary in a PRM system because personal data is a non-homogeneous asset, in terms of its sensitivity, post processor download control, and also in terms of the data subject's ability to control some part of the asset. Since both personal data and usage tracking data are personally identifiable, they are, in the sense of data protection, one and the same thing. Since one can think of granularity as being descriptive metadata about a data subject herself in addition to usage information available at different levels of granularity to the asset viewer, then a standard vocabulary for the degree of granularity regarding both content and tracking information is a need which vendors and indeed standards bodies would do well to consider, though we do not consider further here.

DRM rights describe permissions, constraints and obligations between users and contents. DRM business models such as pay per play rely on client software receiving rights, formatted in rights languages, expressing the number of times a song can be played for instance. Rights metadata defines control over that content. For instance, a client may *view* but not *edit* a document. In PRM, these rights are configured to allow the data subject to exert control over personal data as permitted by data protection legislation. The data controller when dealing with content interprets and enacts those rights. It follows that the rights entity in a PRM system is a relevant target for privacy expression.

ODRL Specification 1.0 proposes a base set of semantics useable for PRM proposes including rights holders and the expression of permissible usages of assets. Consider below a DRM entity model for ODRL.

In terms of parties' expression, PRM is primarily interested in multiple processors all of whom must enforce and be advised of the processing preferences and requirements for assets. These preferences and requirements may be denoted by for instance jurisdictional origin and self determination metadata constructs appended to those assets. In a PRM system, the jurisdictional requirement regarding rights implies that a bi-directional operator would replace the unidirectional attribution between parties and rights in Figure 6 as rights for any personal data must match the legal requirements of the country of origin for the data subject.

In terms of rights expression, there is a need to consider a vocabulary translated from The Directive to describe ODRL access rights for profiles, data subject metadata profiles in terms of granularity and tracking extensiveness, and also the contingent responsibilities passed to interacting processors. The current forms of ODRL *agreement*, *permissions* and *constraints* abstract elements, as they relate to the rights entity in the specification's data dictionary possess syntax that may be applied to a PRM system. The *agreement* element represents a concatenation of the entity's asset, context, party and permissions so as to express agreements between processors for specific rights over the assets of personal data. Specifying expression containers and linkages may be effectively used to generate data protection service level agreements (SLAs) between different legal entities operating under a PRM system.

Within the permissions abstract element two abstract entities have particular value to PRM systems: reuse permissions and transfer prepermissions.

<Permissions> Abstract Element

<i>Reuse</i>	<i>Transfer</i>
<Modify>	<Sell>
<Copy>	<Lend>
<Annotate>	<Give>
	<Lease>

These metadata definitions give the data subject an unprecedented level of control over the processing of their data by disparate processors within a PRM system. The reuse abstract entity offers syntax applicable to reuse of some part of personal data, while the transfer abstract entity implies temporal constraints applied to personal data actions. This can be instantiated in expression fragments through ODRL-defined expression linking. In this aspect we find the semantic basis to realise in part our earlier description of finality as required by The Directive. Once a retention period is exhausted, a processor has an obligation to delete or to make anonymous the personal data related to the asset. Further, the *modify and lend* abstractions would be key to instantiating versioning accountability for the data controller, for managing revisions of the personal information and for enforcing the responsibility to maintain an accurate version of personal to every implicated processor. This latter aspect is clearly related to notions of quality and data subject rights.

<Constraints> Abstract Element

<i>Bounds</i>	<i>Temporal</i>	<i>Spatial</i>	<i>Aspect</i>	<i>Target</i>
<Count>	<DateTime>	<Country>	<Quality>	
<Range>	<Accumulated>		<WaterMark>	
	<Interval>			

The *bounds* abstraction may be applied for the benefit of the data controller, to model control of onward transfers of personal data. Indeed, the data subject herself may also make use of this, and in doing so, would be provided with a new level of control. In fact, the level of control could exceed that prescribed by The Directive.

The *temporal* abstraction represents important definitions for a PRM system, in view of the need for a timestamp tag. Ideally maintenance of this tag should be independent of any processor's

infrastructure. The retention period functionality discussed earlier would be a timestamp abstraction.

The *spatial* abstract entity is an important tag for designating the country of origin of the data subject in PRM. As an indication of EU nationality, personal data must conform to the control and processing restrictions related to EU Community law – as we have illustrated through the simplification of *principles* (Table 1). For instance, if this entity indicates US citizenship, then in effect since there is currently no legal requirement to execute PRM system functionality for that user.

The *aspect* abstract entity appears to be the ideal focus of The Directive’s data quality requirements and the requirements of data controllers to enforce these in processing. The *target* abstract entity is particularly interesting because of Article 8 of The Directive, regarding national implementation. Such an entity would limit the transfer of assets, even within the same legal entity, to uses similar to the original purpose of processing. In Figure 7 we summarise the key changes needed to DRM metadata for PRM.

In addition to the expression technology for access rights to an asset, the P3P protocol also offers a ready-made data transfer platform that, in terms of data subject privacy preference expression, is generally judged to be sufficiently rich.

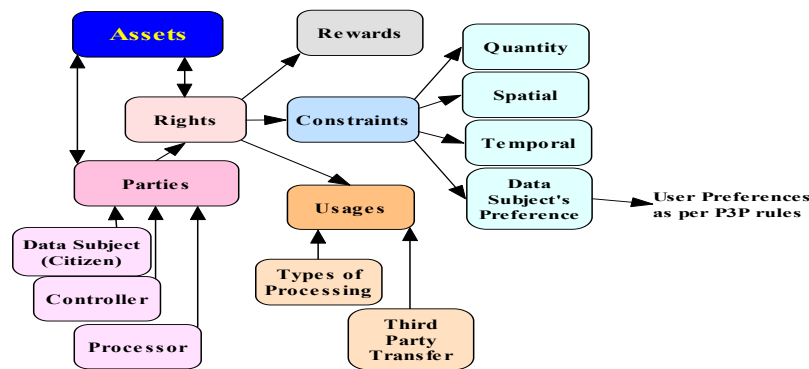


Figure 7. PRM entity model.

Because the rights attribute encapsulates semantic expression over assets, and because the Data controller, when dealing with assets, interprets and enacts rights, it follows that the rights entity in a PRM system is a key component of the PRM server’s data controller and data processor management activities. Clearly this system element holds the assurance responsibility of enforcing legitimate processing, which may be realised through periodic examination of processor log files. A computational map of the appearance of legitimate processing for a given scenario is the main prerequisite for automated analysis of this crucial controllability parameter.

5. Discussion and Conclusions

To clearly understand the potential of how effective adapting DRM to meet the demanding requirements for PRM would be, we analyse PRM requirements and implementation challenges with respect to the facilitation principles (Table 3).

Table 3. PRM requirements and implementation challenges in meeting the privacy facilitation principles.

<i>Principle</i>	<i>PRM Requirements</i>	<i>Implementation Challenges</i>
1. Reporting the Processing	The PRM server tracks the data processors with which the data controller web server has processing arrangements. It tracks: data processor identity, processing type, and the assets upon which processing is applied.	There are many possible data subjects (many millions) and data processors (many hundreds both local and remote). This contrasts dramatically with DRM wherein the number of content owners (Data Subject counterpart) is limited and there are many millions of subscribers (Processor counterpart). Only a limited number of content owners would be active at a time. In the case of PRM, all of the Data Subjects may expect reports, and there may be many hundreds of data processors active at any time. Therefore despite individual asset size (Data subject information) being small, highly scalable approaches are required to manage the logging and reporting processes required in PRM in order to meet scalability requirements for the PRM server.
2. Transparent Processing	The PRM server provides data subjects the ability to view data controller / processor operations on subject data on an “on demand” basis.	DRM systems are designed to meter usage of content. In a PRM system, adaptation of usage tracking through secure distributed logging techniques is required. Centralized management via asset usage monitoring is the common approach in DRM for asset metering. However, this approach may not scale well. A major challenge for implementing this principle is that it is not currently possible to determine exactly what a processor is doing or what it has done with the personal data it has received.
3. Finality & Purpose Limitation	When adapted for PRM, DRM means for specifying and enforcing requirements on processing of tangible assets may similarly protect personal data such as data subject consented retention periods.	There are many data subjects, each with potentially different rights specifications especially with regards to rights specifications for processing. Scalability challenges present themselves here if the data controller holds the rights, and there are many distributed data processors that must access each data subject’s rights for processing. One way to mitigate this challenge would be to distribute rights as well as the personal data. Functionally (if not legally), this distributes responsibility for data protection enforcement from the PRM server to the data processor. A means for maintaining data controller-linkable responsibility is facilitated by the rights granting model specified by ODRL. For instance, personal data can travel separately with only information on where to get permission to process. At the time of processing, a request would be made to the Data controller, which would in turn return an ODRL “License”. The “License” would contain the permissions and the conditions (time, territory, tracking state, etc) for processing. DRM systems support similar functionality for e-media distribution.

4. Lawful basis for data processing	The data controller may only process personal data on certain grounds. These grounds which must be replicated in all data processors of personal data.	The central problem of processing enforcement has yet to be solved. Considering DRM, It is difficult to ensure that once a user receives a license to use electronic material, that it is not processed in a manner that was not intended (for instance, in the case of music, making copies, or converting to other formats). The first step, however, is to standardize a data protection definition language as a starting point so as to control parsing. For instance, there may be an exchange of credentials between the data processor must possess in order to grant permissions. In this case, the effectiveness of this approach depends upon the trust between the data controller and data processor. Since once the data is in the clear at the data controller, any sort of processing is possible.
5. Data Quality	Quality relates to specified attributes the asset must maintain. Effectively it must be as correct and accurate as possible for all who deal with the data.	If the data controller maintains a central repository of subject data and controls access for each data processor request, there is a reasonable likelihood data quality may be maintained. However, this approach is not scalable. On the other hand, if the personal data is distributed to provide scalability, the data must first of all be protected, and secondly, it must be possible to assure the data is consistent throughout should the data subject requires amendments.
6. Rights	The data subject has the right to determine and maintain the correctness of the relevant personal data held by the data controller. The data subject also has the right to raise objections as to the behaviour of the data controller and processors. For a PRM system, this requires editing provisions and a communication channel for raising objections.	In DRM asset management, owners may transfer content to the server for distribution. Content editing is performed by the owner offline, on a master copy of the content. To support an on-line editing function, some sort of online editing tools would be required. As well, access to the editing function must be authorized. Also, a communication channel for raising objections is required. An effective tool for raising objections would also include data mining tools of processing transactions. The objective would be to provide evidence of contract, or privacy breaches.
7. Data traffic outside the EU	The PRM server block should enforce grounds for data transfer on the basis data adequacy and exceptions – for EU nationals of different member states, as well as say American nationals who express a self-determination for EU data protection applied by a PRM system.	This requires the ability to identify the nationality of the data controller, and data processors, and the enforcement of suitable logic appropriate to origin. Unfortunately, there are currently no foolproof technologies to determine geographic location of users (although Quova Inc. [8] purports to have a solution). As well, rules systems to support multiple countries would be extraordinarily complicated.
8. Processor processing.	The PRM server must decide when it will outsource data processing, and on the correct grounds	It is clearly challenging for the data controller to enforce processing among widely distributed data processors, apart from recourse to third party auditing. Negotiation techniques between data

	in a dynamic arrangement. Key for this operation is the enforcement of data controller rules.	controller and data processor could determine a likelihood of compliance, but not enforcement.
9. Security	The data controller is responsible for ensuring data processors apply uniform security standards across all data controllers.	DRM secures content for distribution – PRM builds on this in an adaptive way as data protection prescribes – such as relating authentication to data sensitivity.

It appears from Table 3 that adapting DRM systems holus bolus would accommodate PRM functionality with relative ease, forming a technical implementation of the privacy principles. However, there are areas that require further research and development. For instance, protection against unlawful processing and data traffic outside of the EU are two key areas potential technology development.

In the former case, a means for tracking the actual processing performed by a data processor is needed. A DRM system is well suited to track the time a data processor requests and receives data for processing, however it is not designed to restrict, track and record the actual processing performed. Once a data processor receives the data in the clear for an expressed purpose, the data processor may simply do what it wishes with the data. This information “leakage” by data controllers or data processors would be difficult to detect. To remedy this situation at least two approaches may be taken. One involves development of a means for determining the actual processing done by the data processor. Another involves deploying a reputation management and reporting system to assess over time which data processors may be most trusted to deal with personal data. Another possibility for accomplishing processing management might be a specialized type of sandbox wherein the personal data would be entrusted to the processor only if the processing to be performed by the subject may be verified before and after processing operations.

With respect to the issue of data traffic outside the EU, one aspect of this issue is the ability to determine geographical location of data controller and processor representations. There have been techniques and at least one service [8] developed to determine geographical location. Unfortunately, these approaches are far from foolproof. One means of circumvention involves the deployment of dynamic proxies. A further complication to dealing with data traffic outside the EU is that privacy laws do not have consistent electronic implementations that would facilitate any sort of automatic negotiation or decision-making around how to deal with subject data.

Other challenges exist regarding adoption of DRM architecture for PRM: third party tracking, scalability, and DRM purpose. Regarding third party tracking and scalability, DRM was developed to support delivery and protection of potentially vast amounts of electronic property from typically just a few owners or distributors. In PRM, relatively small amounts of data are collected from a very large number of citizens, where the citizen entrusts information to a data controller. All Data subject data must be tracked for use. This data must be managed, kept confidential and must be editable by the data subject to assure accuracy. A PRM system is designed to keep data protected as well as track the sharing of personal data. It is clear that conventional DRM systems potentially require extensive redesign to support this demanding application. Incorporating a Trusted Third Party approach wherein, a data controller or processor must “check out” information each time it is used may appear to offer a solution to this issue. Unfortunately, this approach adds considerable overhead to data controller and processor activities as well as being a single point of failure. An alternate approach might be the delegation

of the use-tracking function to the data controller. While this would distribute the tracking function load, it would also require a high degree of *trust* between the data subject and the data controller.

Given that DRM systems may be used to profile individuals, one may question the value of considering such systems to implement privacy rights management to uphold data subjects' privacy. It is important to understand that in the PRM architecture we describe, the tables are turned; the digital material of value is user data. It is treated like the immaterial goods controlled in DRM. Rather than tracking purchasers of immaterial goods (music video, etc.) our system tracks the use of personal data by data controllers.

As we have illustrated, simply protecting data in storage and transit is no longer enough when considering The Directive. In our approach, we propose an adaptation of DRM functionality to provide privacy rights management for individuals. Given the embryonic commercial status of the privacy market and its projected economics in a commercial environment placing increased value in integrity, a PRM investment appears extremely worthwhile both in terms of what is necessary to come close to achieving compliance with current legislative requirements, and what is required to meet corporate privacy policies towards building a stronger trust relationship with clients. On the other hand, while the application of digital rights management appears to offer promise for privacy rights management, a fully scalable implementation to support The Directive would be challenging.

References

- [1] Official Journal L 281, 23/11/1995 p. 0031 – 0050.
- [2] L. Deitz, Privacy and Security – EC's privacy directive: protecting personal data and ensuring its free movement, *Computers and Security Journal*, V. 17, N. 4, pp. 25-46.
- [3] Council of Europe Convention 108: Convention for the Protection of Individuals with regard to Automatic Processing of Personal Data. Available at: <http://conventions.coe.int/treaty/EN/Treaties/Html/108.htm>
- [4] J. Feigenbaum, M. Freedman, T. Sander, A. Shostack, Privacy Engineering for Digital Rights Management Systems, Oric, of the ACM Workshop on Security and Privacy in Digital Rights Management 2001. Available at: <http://citeseer.nj.nec.com/feigenbaum01privacy.html>
- [5] XrML is being contributed to the standards body OASIS Rights Language Technical Committee as its foundation technology. More information can be found at <http://www.oasis-open.org/committees/rights/> or at <http://www.xrml.org>
- [6] Open Digital Rights Language (ODRL) Available at: <http://www.odrl.net>
- [7] B. Schneier, J. Kelsey, *Secure audit logs to support computer forensics*, *ACM Trans. on Information and System Security*, Vol. 2, No. 2, 1999, pp. 159-176.
- [8] Quova, Inc. at: <http://www.quova.com/>