A Credential Based Approach to Managing Exceptions in Digital Rights Management Systems

Jean-Henry Morin and Michel Pawlak

Abstract. While Digital Rights and Policy Management (DRM / DPM) technologies have matured to become mainstream they are now poised to thrive in all aspects of our daily electronic lives and digital assets. A key issue remains however unsolved with respect to managing exceptions in the context of DRM enabled systems and information assets. This paper addresses this issue defining the problem and proposing a model based on credentials allowing to dynamically account for lawful unanticipated usage situations while still maintaining a given level of persistent protection, governed usage and audit trails.

1 Introduction

The DRM industry is currently struggling with interoperability issues. Several recent advances and initiatives in this field are encouraging and should ultimately set the ground for an initial basic level of interoperability among the currently incompatible systems. Let's just mention here the recently published MPEG-REL [1] and RDD [2] ISO standards based on work done within MPEG-21 [3] [4]. Or work done by InterTrust [5] [6] [7] also very active and involved in several initiatives focusing on interoperability issues such as the recently formed *Coral Consortium* [8]. It is a cross-industry consortium involving some 30 key players from entertainment, consumer electronics, telecommunications, and technology industries.

While this is a critical issue and an enabling factor for the broad endorsement and deployment of DRM based systems, whether in the entertainment or enterprise sector, there still remains a hard problem to be addressed. How do DRM enabled systems manage or are able to deal with so called *exceptions*? In order to further emphasize this critical issue, let us cite the *Copyright Balance* principles that should underline public policy regarding DRM as recently outlined by E. Felten in a column of CACM [9]: "*Since lawful use, including fair use, of copyrighted works is in the public interest, a user wishing to make lawful use of copyrighted material should not be prevented from doing so by any DRM system.*". This sound principle is exactly at the forefront of this work making the case for such "Exception Provisioning" in DRM enabled systems.

This paper is structured as follows. After further describing the issues and objectives of expressing exceptions as credentials, section 2 presents a usage scenario illustrating the limits of traditional DRM solutions, and how an exception based approach could help. A revisited taxonomy and scenario is presented in section 3. Section 4 presents the proposed model and a possible architecture is described in section 5. Section 6 outlines related work positioning the research issue and our approach. Concluding remarks and future work are presented in section 7.

Copyright 2006 University of Geneva - http://cui.unige.ch/ASG

1.1 Issues and objectives

Let us briefly describe the issue. In a global DRM enabled information market, and provided there is a need for governed content usage (not all content requires governed usage), we assume all digital assets to be persistently protected. We also assume that the content follows the super distribution model [10] [11][11][12][11][12][13] where the rules governing its usage are cryptographically attached to the content either directly and or can be dynamically acquired on-line. In both cases, it is reasonable to postulate that rights holders cannot anticipate all usage situations within the set of rules, and hence are definitely not in a position to anticipate most exception situations where some rights should be waived while still maintaining a given level of persistent protection and governed usage. This is especially true considering a global world wide market still having complex, often contradictory national and international regulations and legal frameworks. Even if these issues were solved from a legal standpoint, there would remain a tremendous technical overhead in accounting for exceptions and waivers beforehand. Imagine a picture of a 100 kilobytes requiring a 1 Mb policy. This also becomes critical when considering mobile devices such as PDAs, cell phones, sensors, etc. with limited resources. In addition, traditional DRM based approaches rarely address such issues since it is basically considered to defeat the purpose of enforcement for rights managed content. In this context we argue that there is a legitimate need and reason to introduce such a *controlled gray zone* as most users (e.g. faculty, home users, etc.) shouldn't be considered as criminals while in legitimate usage situations such as space shifting, personal use, fair use, etc.

1.2 Expressing Exceptions as Credentials

We propose an interesting alternative approach. A credential based approach by which a DRM module would provide a "hook" to evaluate locally held credentials which could have precedence over the attached rules and be traceable (i.e. auditable). The process could be rather straightforward as it would be comparable to the existing verification of locally held licenses in the users license-store. For example, let's imagine that blind and visually impaired users are provided with such a credential due to their disability. Or an academic holds a credential, delivered by the university, showing his affiliation and status. Such credentials would be stored on the users computer (e.g. in a credential store) and made available to the DRM module (enforcement point) when evaluating rights at runtime.

Such an approach would be rather elegant and would allow to accommodate many situations where explicit rule specification would simply be too cumbersome or simply impossible to anticipate and formalize. In the case of fair use. It is commonly agreed that non commercial use of copyrighted material in academic environments is free. Being an academic staff member or a student would allow to have an academic credential delivered by the university. In a general way, this approach would allow to capture generic rights management in the form of groups or communities. Being a member of a group provides a generic right with respect to content when accessed by its members. Further refinement could consider a hierarchy of credentials for example

Copyright 2006 University of Geneva - http://cui.unige.ch/ASG

within a company where management would be provided credentials with broader rights than those of staff members.

2 The Traditional Approach without Exception Management

We first need to taxonomize the traditional approach to DRM, without exception management, in terms of resources and roles in order to introduce the general scenario used to analyze and revisit the approach in the light of exception management.

2.1 Taxonomy

2.1.1 Resources

- **Content**: A digital content is a resource requiring its use to be managed. Content, also referred to as digital assets, holds value in a patrimonial sense and as such is a strategic resource for its right holder. This value can evolve over time. This applies to any kind of content. Even if the term *content* is used in this document, it can be replaced at anytime by *asset*. Digital Assets have to be protected by policies governing the way they are used. Such protected contents are called *rights enabled contents*.
- **Policy**: A policy defines at a strategic level how a content can be used, when, by whom, in what context, etc. Policies are defined through rights expression languages (REL) and cryptographically associated to contents by content and rights owners. Policies may involve the existence of multiple contents and respective policies.
- **License**: A license is the counterpart of a policy. Licenses represent the DRM implementation of a set of rights granted to a user defining the usage of the content. Licenses are acquired by users in transactions and may involve fees. Licenses may apply to a single identified content or to a set of contents. The combination of usage rights granted in a license is evaluated against the policies governing content usage. This evaluation process defines the terms and conditions of use when applicable.
- **License store**: The license store holds the licenses granted to and held by users for subsequent use.

2.1.2 Roles

- **Owner**: Owners are content or rights owners or their representatives. The right owner may not always be the same as the content owner or its creator, but for clarity, we group these roles as it has no impact on the proposed model. Owners define policies governing content usage and cryptographically associate these policies to the content.
- **License Manager**: The license manager is the entity providing Licenses. Licenses can be created and distributed in advance or can be created on the fly when a user

Copyright 2006 University of Geneva - http://cui.unige.ch/ASG

need them to use a rights enabled content. Depending on the business the owner may or may not be the license Manager.

- **Content User**: The content user holds or acquires licenses from the license manager in order to access rights enabled content.
- **Enforcement Point**: The enforcement point is a trusted process residing close to the user in charge of evaluating policies governing content usage based on a transaction involving the license manager or locally held licenses. Based on this evaluation content rendering and usage is either granted, denied or further negotiated.

2.2 The Traditional Scenario

The following scenario presents the traditional approach to rights and policy managed content. Let's consider the case of a pharmaceutical research company with its pre-commercial drug data. Decision is taken to release under embargo some data to customers. This company being aware of the strategic value of this content, decides its usage should be managed since the content will leave corporate security and administrative boundaries (i.e. outside company intranet and perimeter). The content is thus rights protected with DRM using a set of policies corresponding to its intended purpose defining the terms and conditions governing its use.

More specifically, the chosen policy defines a set of usage conditions such as for example:

- The content expires on a given date
- Contributors to the content have full access to it
- Unknown users cannot access the content
- Selected customers get a read only access to the content based on their status
- The content cannot be extracted from its persistently protected form.

This set of policies is intended to ensure proper protection of the content, and that no unauthorized access will be granted. Selected users receive licenses matching those policies, providing them the right to access the content. Other users trying to access the content need to get a license from the license manager.

Over time, the enterprise starts working with other partners such as research partners. These partners may need to access some content like the content we just presented. If so they probably signed contracts giving them the right to access some content for specific usage over a given time period. Thus they receive licenses corresponding to the rights they have. Figure 1 illustrates this scenario.

Copyright 2006 University of Geneva - http://cui.unige.ch/ASG

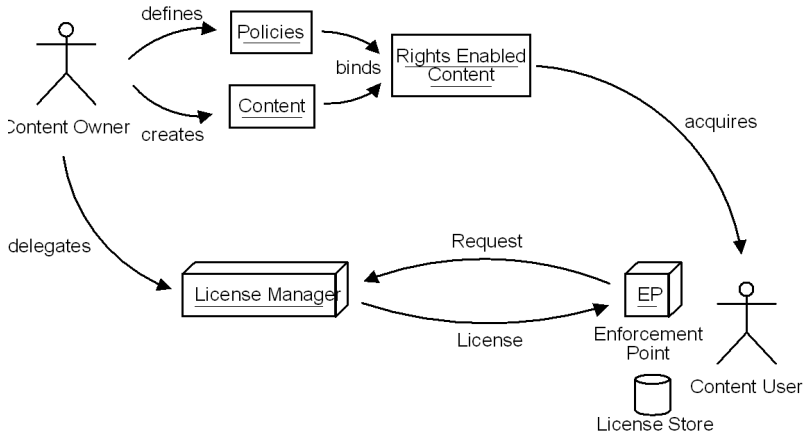

Figure 1. Traditional scenario illustration

Further, these partners may be located in countries having mandatory legal requirements to provide access to some government authority (e.g. ethics committee or drug agency) for auditing or control purposes.

Since such policies cannot be anticipated by the content owner nor could all such exceptional situations be reasonably encoded due to their number and variety, a mechanism is required for legitimate use in such exceptional circumstances. Obviously, a key requirement for such a mechanism is the ability to log and thus be able to track and monitor such exceptions.

As a result, and considering the initial guiding principle whereby "*lawful use of copyrighted material should not be prevented by any DRM system",* it is legitimate to argue that there is a need for exception management in DRM systems. Furthermore, there are many examples of such situations where exceptions could be required. One of the most cited one being *fair use*, which is still a major issue in the DRM space. We intentionally introduced the notion of monitoring of exceptions as a requirement as it is legitimate for a content owner to be able to monitor possible unjustifiable and repeated abuse. This issue will be further discussed in the proposed model.

2.3 Analysis

The previous scenario showed a case where unanticipated, although legitimate, access is forbidden by existing policies. Possible solutions of adding new policies to the content are not realistic in the long run and would not scale for large distribution. Exception situations will be numerous, depending on many criteria that cannot be foreseen. Embedding policies capturing these exceptional situations directly in the content would lead to tremendous overhead compared to content size and / or value.

Copyright 2006 University of Geneva - http://cui.unige.ch/ASG

One has to strike a reasonable balance between security, commercially viable risk level and usability.

As these policies handle particular situations and do not represent usual cases, they certainly do not apply to most content users. Distributed content should be protected using only *core* policies - policies satisfying most cases or at least the ones considered critical to the right owner. Exceptions and other unanticipated policies should be efficiently handled separately.

Enterprises can define core policies as they know best how they usually want their assets to be protected and managed. Even if they are able to deal with some exceptions, such as partnership exceptions, etc., they are probably not aware or qualified to cope with all possible situations that exist or might arise, be they legal, regional, etc.

This mandates for a specific exceptions management model for DRM systems and leads to the following revisited scenario.

3 The Credential Based Approach for managing Exceptions

The traditional DRM approach described above highlights the problems and limitations of an approach where unanticipated exceptions cannot be managed. This section revisits this scenario based on a credential approach for managing exceptions.

3.1 The Revisited Taxonomy

This revisited taxonomy introduces the main terms used in the context of this work which are specific to the credential based exception management approach. It extends the taxonomy of the previous section.

3.1.1 Resources

- **Exception**: special situation indicating that even if policies in place do not allow access to a piece of content for a given user, the latter may have legitimate reasons to access it.
- **Credential**: A Credential should be understood as an administrative token, certifying that an given user has been granted a special status by a locally administered Credential Manager (e.g. faculty staff, blind user, student, etc.). It is authenticated by the granting Administrative Credential Manager and bound to a specific user. The validity of the credential may be limited in time, renewable and can also be revoked.
- **Credential Store**: Physical space containing all credentials bound to a user.
- **Short Lived License**: Short Lived Licenses are generated by Exception Managers and provided to users when an exception has been detected and verified. They are meant to give an exceptional access to a content, and their validity is thus limited in time. Short Lived Licenses can give more or less rights depending on the type of detected exception.

Copyright 2006 University of Geneva - http://cui.unige.ch/ASG

3.1.2 Roles

- **Administrative Credential Manager**: The entity (local) that emits, revokes and manages credentials. Can be any structure, such as an enterprise, an academic entity, or a national entity. It does not have to be known by the Content Owner neither at credential generation time, nor at content creation time; but it has to be able to prove its legitimate existence as well as the motivation leading to the credential emission.
- **Enforcement Point:** Extension of the previous *enforcement point* definition. If access to the rights enabled content is denied, the enforcement point now searches for credentials available in the user's credential store as one of them may qualify and thus raise an exception. If any credential is found, available credentials are sent to the Exception Manager along with content identification for further verification. If not, access is denied or a license is requested from the License Manager, as it was done in the traditional scenario. This however requires to be provisioned by the specific DRM providers.
- **Exception Manager**: It is an extension of the traditional License Manager role it may or may not be the same. In the later case it would be delegated to an external entity in charge of that role. It verifies if a credential can entitle access to a piece of rights enabled content. The Exception Manager checks if the credential is valid, if it has not been revoked and its applicability for the content. Thus it verifies if the Administrative Credential Manager has legal existence and for what reason the credential has been emitted. If the credential passes all checks, a short lived license may be granted providing access to the content for a limited time. Moreover, the operation is logged as possible further proof of legitimate activity.

3.2 The Revisited Scenario

The revisited scenario extends the traditional scenario, adding an exception management dimension to the traditional DRM based content protection.

In this context, the enterprise owning the content defines the core policies needed to effectively protect its assets. This set of policies do not contain exception policies. Instead the enterprise delegates to two other kinds of entities the management of possible exceptions.

The first entity, the Administrative Credential Manager, is intended to certify that a user legitimately holds some status based on a specific situation, affiliation, etc. This entity may be unknown to the content owner but must be local to the user (i.e. the user must have some form of relation to the administrative credential manager). This credential granting process is independent of existing policies. Its only purpose is to certify a specific status of one of its affiliates in the form of an authenticated and verifiable credential (i.e. token).

The Exception Manager's role is to verify if a credential could entitle its holder for an exception when existing policies might not be applicable. If the entity detects that one of the credentials is applicable, it generates a short lived authorization that grants access to the content and send it to the user as a short lived license.

Copyright 2006 University of Geneva - http://cui.unige.ch/ASG

Back to the traditional scenario, once the content is successfully protected, the enterprise starts distributing it. When a user holds legitimate rights to an exception such as in the previous example of an ethics committee or drug agency, a request together with the potentially applicable credentials is sent to the Exception Manager for verification. The verification is manifold. First, it verifies if the Administrative Credential Manager really exists and is allowed to generate credentials. Then it asks for the motivations that led to granting the credential. After having completed these verifications, if at least one of provided credentials can apply, the event is logged with all required information, a corresponding short lived license is generated and returned to the enforcement point, allowing the user to access the content. This scenario is illustrated in Figure 2.

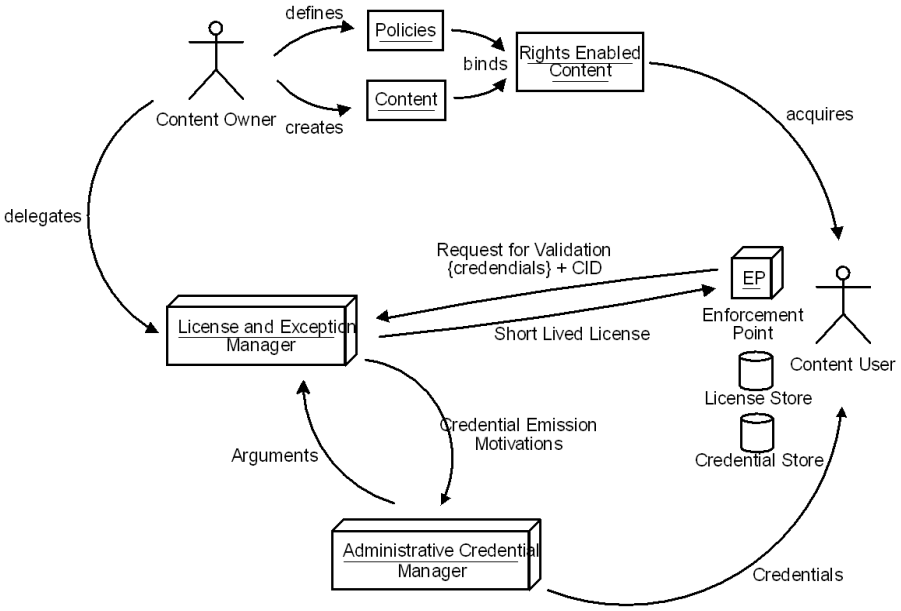

Figure 2. Revisited scenario illustration

3.3 Analysis

This revisited scenario depicts a situation where content protection, exception credential creation, exception verification and corresponding authorization are decoupled. The proposed approach provides greater flexibility than the traditional scenario allowing unknown Administrative Credential Managers to inform Enforcement Points that an exceptional situation may be taken into consideration if the user has no explicit rights to access the content.

Copyright 2006 University of Geneva - http://cui.unige.ch/ASG

While providing flexibility this approach still gives final control to the Exception Manager by allowing it to verify the legitimacy of the exception. Content Owners only have to care about the way they wish to protect their assets, ad hoc decisions being taken by the Exception Manager in case of exceptional situations. A model illustrating this scenario is proposed in next section.

4 Model

This section presents the credential based model for managing exceptions in DRM systems. We first present the specifics of the content protection process when using exceptions before describing the exception management itself.

4.1 Content protection

Content protection in the context of an exception based model differs from its traditional representation. This section explore the main differences introducing or refining the concepts of core policies, certification delegation, exception handling delegation and rights distribution.

4.1.1 Core Policies

At the very beginning of the content protection process the definition of policies is driven by the need to protect a content asset. But this process follows a path leading from this simple content protection to the need of having flexibility in any situation. Following this path results in producing heavy content overloaded with policies needed to be able to deal with all particular situations that may be encountered.

In the exception based model we propose, only *core* Policies should be associated to contents. Core Policies are the set of policies needed to efficiently protect the content. These policies have to reflect enterprise strategy, the most important requirements concerning the content and all *usual* situations that may occur. Thus policies embedded into the rights enabled content should not include other considerations, such as policies dealing with extremely rare situations, that could be considered as exceptions.

In this context all policies added to provide further flexibility not in the scope of usual policies are considered as potentials exceptions and should thus be handled using the credentials based exception handling model.

4.1.2 Credential Properties

Credentials have the following set of properties:

- **Known Source**: Credentials must contain information about the Administrative Credential Manager who generated them, in order to be able to verify its legal existence as well as the motivations that led to credential generation.

Copyright 2006 University of Geneva - http://cui.unige.ch/ASG

- **User Bound**: Each credential is bound to a single user or role, affiliated to the Administrative Credential Manager, able to prove that he is the legitimate owner of the credential.
- **Limited validity**: Credentials are limited in time, their validity period is included in the credential.
- **Revokable**: The Administrative Credential Manager has the ability to revoke a credential it has generated at any time.
- Note that information about the nature of the credential, the reasons explaining why it has been created are not embedded into the credential. This approach allows to modify the scope of credentials generated by an Administrative Credential Manager for a single user, by widening the set of motivations, narrowing it or refining it, without having to revoke the credential and having to generate new ones. Thus providing additional flexibility, while being able to retain control over the number of credentials.

4.1.3 Credential Generation

The model delegates the generation of credentials that may lead to an exception to Administrative Credential Managers. They indicate that credential owners can legitimately ask for the rights to access a piece of content in a given context.

Resulting credentials do not provide any direct access grant to a piece or type of content, but only indicates that even if their owner does not have the rights - in the form of a license - to access a piece of content, if the credential is recognized, he may be entitled to the right to access the content due to an exceptional situation.

4.1.4 Exception Handling Delegation

As stated before, the goal of the credential based model is manifold. First, it provides a way to reduce the size and complexity of rights and policy managed contents. Second it provides more flexibility in handling special or unanticipated situations as content needn't be modified to deal with new exceptions. Finally, it simplifies the role of content owners allowing them to produce contents and protect them with the most important and representative policies, not having to deal with all possible situations.

Thus the business is provided the ability to delegate the particular situations that may lead to exceptions. In this model, exceptions are detected, verified and handled by an Exception Manager. Its role is to detect possible exception situations and to handle them, without having to inform directly the producer of the content, nor having to modify the content in order to adapt to occurring exceptional situations. Activity logging should be done for possible audit reasons.

4.2 Exception management

In this section we explore in further details the process of rights verification, exception detection and short lived license acquisition.

Copyright 2006 University of Geneva - http://cui.unige.ch/ASG

4.2.1 Rights Verification

A central role in the exception based model we propose is the rights verification process. As stated before, the way the enforcement point manages rights verification in our model differs from the usual way. Figure 3 depicts the underlying sequence of actions that have to be completed.

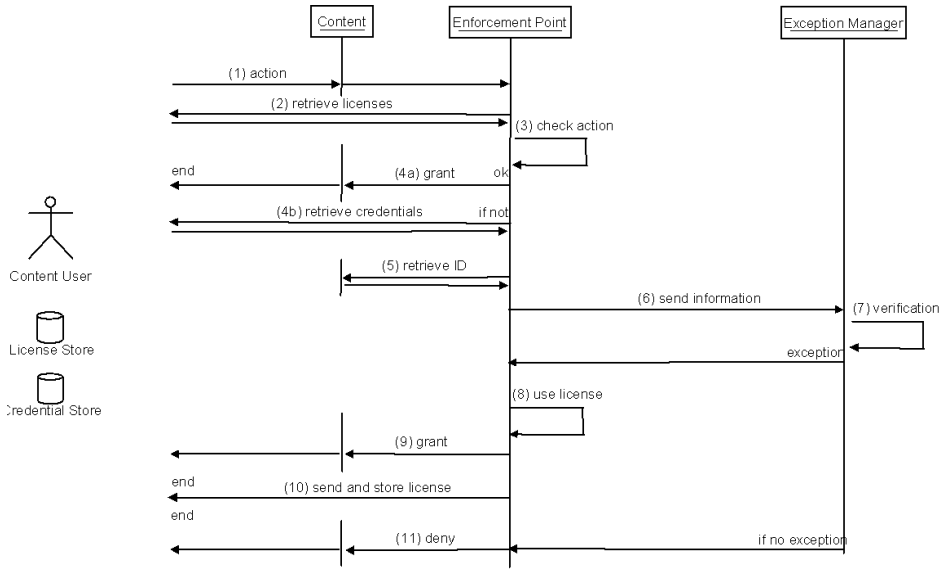

Figure 3. Rights Verification Sequence Diagram

When a user wants to access content (1), the licenses it possesses are taken from its license store (2) and the enforcement point tries to use them for the specified action (3). This part of the process is exactly the same as done traditionally. If existing licenses match content policies, access is granted (4a). If none of the licenses are applicable to the content, available credentials are taken from the local credential store (4b), content identification is extracted (5) then these information are signed and sent (6) with the information about the way the content is being accessed, to the Exception Manager for further verification (7). This next step tries to detect possible exceptions instead of simply denying access to the content. The enforcement point then awaits for an answer which can eventually be a short lived license, if an exception is considered, and uses it (8) to then grant access to the content (9) and store the license (10) or a deny if not (11).

Copyright 2006 University of Geneva - http://cui.unige.ch/ASG

4.2.2 Exception Detection

When the exception manager receives the credentials, as well as content identification and the usage context, it tries to detect if their combination are applicable for an exception. For each credential multiple steps are involved. These are illustrated in Figure 4.

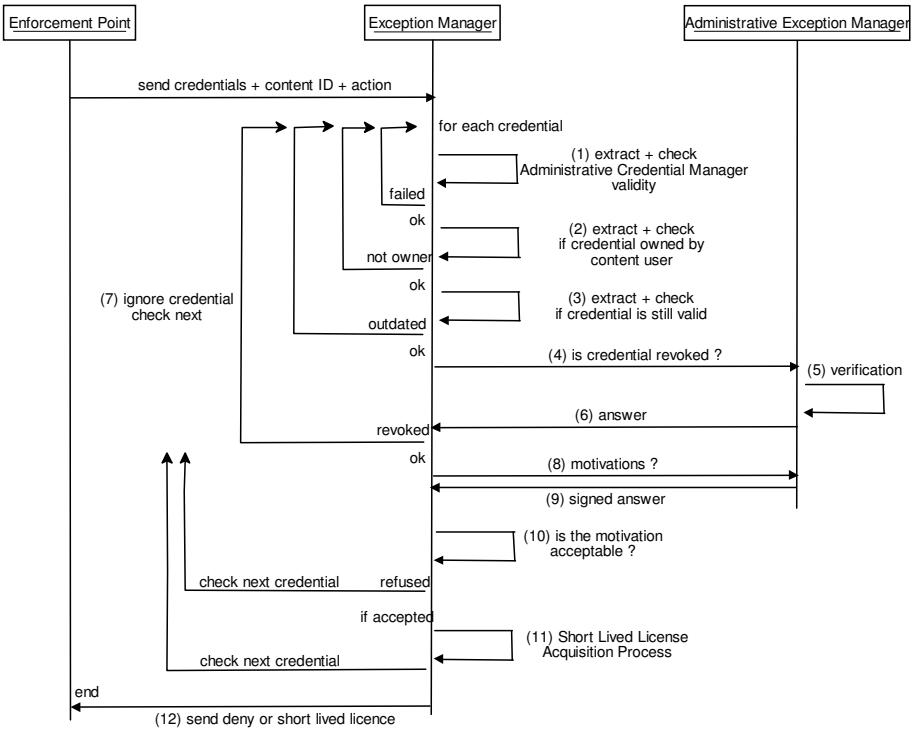

Figure 4. Exception Detection Sequence Diagram

First, the exception manager has to verify if the credential has been generated by an existing and valid Administrative Credential Manager (1). To achieve this task, the credentials have to be examined in order to retrieve information about their creator, then verify their legal existence. The next step is to verify if the credential really belongs to the user trying to access the content (2). If it is the case, the exception manager checks if the credential is still valid (3) and asks the credential manager if it has not revoked it (4). Administrative Credential Manager verifies it (5), then sends an answer (6). Credentials not complying to any of these rules are ignored (7). Last step is then to check if the credential can be applied to the content in the context in which the content is being used, to do so the Exception Manager asks the Administrative Credential Manager for the motivations that have led to a credential generation (8) and the Manager sends back its signed answer (9). This answer may include textual information that can be analyzed, parsed, it may also contain any other

Copyright 2006 University of Geneva - http://cui.unige.ch/ASG

kind of information such as a certificate emitted by a content owner indicating that a contract has been signed by both parties, or even another credential emitted by another recognized Administrative Credential Manager. If this last verification succeeds - i.e. if any of the retrieved information is accepted (10) - an exception is applicable and the short lived license acquisition process can start (11). When all credentials have been verified, a short lived license or a deny is sent back to the enforcement point depending on the result of the process (12).

4.2.3 Short Lived License Generation

The short lived license generation process is started when an exception has been detected and is applicable. This is a recursive process creating a license based on all exceptions that have been detected as applicable for a single access to a rights enabled content.

At this stage, the Exception Manager knows that it has to deal with an exception situation and knows what credentials have raised what kind of exception. The short lived license is built incrementally analyzing all exceptions. In order to emit such a short lived license some precautions have to be taken in order to manage issues of precedence and potential conflicting exceptions.

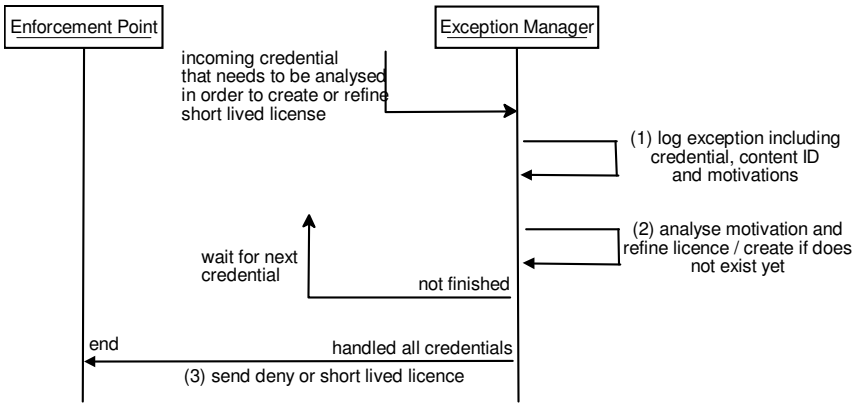

Figure 5. Short Lived License Acquisition Sequence Diagram

Figure 5 presents the different steps of this process. First, each exception has to be logged for traceability purpose (1). The log has to keep all required information to justify the exception. This includes the identification of the content, the credentials that led to an exception, the motivations signed by the Administrative Exception Manager and the context of use, i.e. the type of access that is being foreseen on the content. Once all required information have been logged, the rights the specific exception may grant to the user are compared to the rights granted by previous exceptions, and the license is refined (2). Differences may occur due to the provided

Copyright 2006 University of Geneva - http://cui.unige.ch/ASG

reasons. For instance, a first credential may raise an exception with motivation "academic use", and a second credential may indicate that there is a "research agreement with the content owner". First credential would allow limited use, but second one would allow access to additional features, or a more detailed output. Once all exceptions have been handled, the short lived license can be generated (3).

The log of all exceptions is needed in order to be able to detect Administrative Credential Managers abusing the system - and eventually blacklist them -, and keep a global trace of content usage.

The validity of the license will be usually short (from a single access to a few days validity) as each credential can be revoked at any time. But the effective validity is a matter of specific policies bound the content owner. Further the nature of the grant may also be limited, providing only limited access to the content. The final decision is thus left to the Exception Manager responsible for this task.

5 Architecture Overview: Attribute Certificates

The basic idea behind the proposed approach is to make use of a credential based scheme. This raises however the issue of who and how these credentials are managed. To this end, we propose the use of PKI infrastructures which are already well established techniques. Moreover, certification authorities are accustomed to handling similarly sensitive aspects of security. The model would also perfectly fit the operation of such services with registration authorities, issuing services, revocation lists, etc.

Instead of using X.509 public key certificates (PKCs), we propose to use Attribute Certificates (ACs), RFC3281 [14], having a similar structure to PKCs without the public key. ACs can hold attributes specifying relevant information such as roles, affiliations or whatever is needed to evaluate exceptions.

Such credentials would be delivered to the user, together with other administrative tokens, passwords, etc., by the institution / organization to which the user is affiliated. A credential would hold several information such as a known lifetime (expiry date), a unique ID (social security number, employee or student number, etc.) within the domain of the institution delivering the credential, and any other relevant information that should be used when evaluating whether or not an exception or waiver is applicable.

From thereon, the DRM system, upon deciding whether or not to render the content, could be required by the user to first check for locally held credentials. Then based on these credentials, further actions could be undertaken in order to acquire the corresponding license and thus grant the user access based on his situation. The important point to note here is that basically the content remains persistently protected. It is processed just as if it were in a situation without exception request. The rendering is done within the usual trusted renderer and basic rules, identified as mandatory for example can still be enforced.

Copyright 2006 University of Geneva - http://cui.unige.ch/ASG

6 Related Work

Other research works aim at proposing solutions to protect the copyright in a balanced way for copyright holders and users. The problem of managing exceptions is considered a hard problem and has been mainly explored in the context of fair use and rights expression languages.

In [15], authors explore how rights management systems can be designed and implemented in a way that preserves the traditional copyright balance, especially with copyright's concern for the public domain and for the legitimate fair use. The authors are against leaving the determination of fair use in the rights holder's hands. Indeed they emphasize the fact that collective public interest may run contrary to the rights holder's individual interest and thus there may be a strong incentive for the rights holder to deny access.

As ourselves the authors doubt that system designers will be able to anticipate the range of access privileges that may be appropriate to be made of a particular work. Their trusted third party approach to handle fair use access to content can be compared to our Exception Manager. The main difference with our approach is that the authors see little prospect for development of private escrow agents, while we insist on the need for them. Indeed, the authors argue that content owners are "*unlikely to pay voluntarily for an institution that facilitates low cost or free access to their works"*, but in our context content owners may *need* to provide *controlled access* to the content in some *exceptional* cases. As examples we may quote again the case of financial auditors or ethical commission needing to legitimately control sensitive documents whose policies would not allow access. In this context, the situation has to be detected, then the exception has to be controlled by the content owner's trusted party in order to provide access if the request is legitimate and avoid all other unauthorized accesses. Our model exactly targets such situations.

In [16], the authors suggest certain accommodations that DRM architectures, and especially their rights expression language components, should make to adequately express certain core principles of copyright law. They focus on two recommendations. The first recommendation proposes changes to the XrML REL vocabulary to be able to highlight limitations on copyright exclusivity in cases such as fair use or first sale and rights transfer situations. The second one, goes toward the need for the creation of an Open Rights Messaging Layer.

Indeed, they highlight current lack of rights messaging or transaction protocol that would provide standardized means for retrieving and disseminating rights information and policies, and issuing rights grants or permissions. They propose an overview of an architecture involving a rights proxy server listening to verbose rights requests, verifying rights of users and able to generate DRM licenses then send them based upon its decision criteria. The resulting licenses would enable certain uses under conditions that the commercial transactions would not generally allow.

While this architecture can be compared to our model, the main difference is that our approach adds an external authority certifying through Attribute Certificates that the user trying to access the content is doing it in a specific context. In our context we consider that content owners protect their content according to a set of policies. The rights issuing is done through the transmission of credentials possessed by users to the exception manager who then verify the legitimacy of incoming requests. If incoming

Copyright 2006 University of Geneva - http://cui.unige.ch/ASG

credentials are not recognized no access is granted. In the case of [16] the proposed REL extension is used to provide alternative rights retrieval capabilities. A combination of these two approaches could be used to provide hints to the enforcement point indicating which credentials can potentially be recognized by the Exception Manager. This could lower the number of credentials sent to the latter and thus lower bandwidth usage. Then, if none of these credentials grants access to the content, the other ones may be sent for verification. Nevertheless, this might cause further issues in the case of an evolution of policies and recognized exceptions at the Exception Manager side, since even if some of the provided credentials may grant access, this access may be more restrictive access than some of the unsent ones.

7 Conclusion and Future Work

This paper proposes an original model for managing exceptions in DRM enabled systems. An exception is any unanticipated situations that cannot be handled of foreseen by the initial set of policies associated to a piece of content.

In order to provide a flexible scheme to manage exceptional situations and ultimately to reduce the size and complexity of rights and policy enabled content, the proposed approach decouples persistent content protection from the management of exceptions. The model delegates to a dedicated role the detection and dynamic management of exceptions as well as the generation of a short lived license granting access to contents. Based on this model, this paper proposes a novel approach and architecture using X.509 Attribute Certificates to implement Credentials.

The model presented in this paper implies a multiplication of credentials. As a single user may possess dozens of credentials requiring to be verified for exceptions, optimizations will thus be needed. The model already provides a way to minimize the number of credentials emitted for a single user by a given credential manager. Other optimizations could include bandwidth usage reduction. This issue may be tackled by only sending credentials potentially relevant to the content being used. The underlying optimization may be a structured classification of credentials allowing a smart credential filtering at enforcement point.

Another issue to be considered is the case of distributed exception verification. In the proposed model, the Exception Manager is presented as an unique entity. Nevertheless, such centralization should be avoided in real case situations, where thousands of users may need to access simultaneously content managed by a single exception manager. Thus new issues due to the decentralization of this role would have to be tackled.

Finally, the role of Exception Manager may be a good argument mandating for the existence of Content Clearing Houses. The latter were presented in the Nineties as aggregators of License Managers and rights clearing centers. This aggregation was not providing any additional added value for their clients. By aggregating Exception Managers, Content Clearing Houses would be able to capitalize knowledge acquired about Administrative Credential Managers, their behavior and their validity. This knowledge is mandatory to be able to easier detect frauds, and cannot be available if exception management is parceled out. As previously stated, centralizing roles may

Copyright 2006 University of Geneva - http://cui.unige.ch/ASG

not be a good idea, thus a way to build distributed Content Clearing Houses should be explored.

References

[1] MPEG-REL, Retrieved Sept 2004, from http://www.contentguard.com/MPEGREL_home.asp

[2] MPEG-RDD, Retrieved Sept 2004, http://www.rightscom.com/default.aspx?tabid=1172

[3] MPEG-21, Retrieved Sept. 2004, from http://www.itscj.ipsj.or.jp/sc29/29w42911.htm#MPEG-21

[4] MPEG-21, Retrieved Sept. 2004, from http://www.chiariglione.org/mpeg/working_documents.htm

[5] W.B. Bradley and D.P. Maher, "The NEMO P2P Service Orchestration Framework", in *Proceedings of the 37th Hawaii International Conference on System Sciences*, January 2004

[6] R. Koenen, J. Lacy, M. MacKay, and S. Mitchell, "The Long March to Interoperable Digital Rights Management", January 2004, http://www.intertrust.com/main/research/papers.html

[7] CNET News.com, "Tech powers seek antipiracy accord", Retrieved Oct. 2004, from http://news.com.com/Tech+powers+seek+antipiracy+accord/2100-1025_3-394347.html

[8] The Coral Consortium, visited June 2005, http://www.coral-interop.org/

[9] E. Felten, "DRM and Public Policy", in *Communications of the ACM*, V. 48, No. 7, July 2005, p. 112.

[10] R. Mori and M. Kawahara, "Superdistribution: The Concept and the Architecture", Transaction of the IEICE, Vol. E 73, no. 7, July 1990, pp. 1133-1146.

[11] R. Mori and S. Tashiro, "The Concept of Software Service System (SSS)", Transaction of the IEICE, J70-D.1, Jan 1987, pp. 70-81

[12] Brad Cox, "Superdistribution", Wired Magazine, September 1994, pp 89.92.

[13] Brad Cox, "Superdistribution Objects as Property on the Electronic Frontier", Addison-Wesley, 1996.

[14] Farrell, S., Houslez, R, "An Internet Attribute Certificate Profile for Authorization", RFC 3281, The Internet Society, April 2002.

[15] Dan Burk & Julie Cohen, Fair Use Infrastructure for Copyright Management Systems, 11 Harv. J. Law & Tech., 2002

[16] Mulligan, D., Burstein, A., and Erickson, J. "Supporting Limits on Copyright Exclusivity in a Rights Expression Language Standard. A requirements submission to the OASIS Rights Language Technical Committee.", Samuelson Law, Technology & Public Policy Clinic, and The Electronic Privacy Information Center, August 2002

Copyright 2006 University of Geneva - http://cui.unige.ch/ASG