



The MPEG-21 Rights Expression Language

A White Paper

Rightscom Ltd
10 Leake Street
London SE1 7NN

Tel: 020 7620 4433
Fax: 020 7620 0760

Contents

1	Executive Summary	3
2	The Business Case	3
3	The DRM landscape	5
4	The Role of Standards	6
5	Introducing MPEG-21	6
5.1	<i>Background to the MPEG Initiative</i>	6
5.2	<i>The Seven Architectural elements</i>	7
5.3	<i>The specifications so far</i>	8
5.4	<i>Using the MPEG Standards</i>	8
6	Why a Rights Expression Language is key	9
6.1	<i>What goes into a rights expression language</i>	9
6.2	<i>Developing the MPEG REL</i>	10
6.3	<i>The Architecture Rough Guide</i>	11
6.4	<i>Defining the Rights Semantics</i>	12
7	Business Model Support	12
7.1	<i>REL and DRM Technologies</i>	12
7.2	<i>Some implementations</i>	13
7.3	<i>New Business Models</i>	14
8	MPEG REL in a wider context	14
9	The future	16

1 Executive Summary

Analysts agree the trade in intellectual property rights will drive much of the revenue through digital networks. Whether content is consumed by an end user connecting a device to the Internet or by a TV viewer using a set top box, both activities will depend on a reliable process for protecting rights. And unless the rights to the content can be packaged within machine-readable licences, guaranteed to be ubiquitous, unambiguous and secure, which can then be processed consistently and reliably, it is unlikely that content owners will trust consign their content to networks.

The MPEG Rights Expression Language (REL) is designed to provide the functionality required by content owners in order to create reliable, secure licences for content which can be used throughout the value chain, from content creator to content consumer. The MPEG experts who created the REL were guided by a set of forty eight requirements, agreed by content owners, technology vendors and intermediaries.

The MPEG REL is constructed on a modular basis with the following components:

- REL Core Schema
- REL Standard Extension
- REL Content Extension

This modular design provides the inherent extensibility of the language and is designed to be:

- Flexible – enabling the creation of licences to support any kind of business model
- Scalable – enabling the creation of profiles to support a wide variety of consumer devices
- Extensible – enabling the creation of specific, autonomous extensions for use in vertical markets, both open and closed
- Technology agnostic – enabling support for any kind of proprietary or standardised enforcement technology

The MPEG REL is one of the critical parts of the DRM tool set. For implementers it provides a single means to support the trade in rights and other types of access controlled digital goods and services across platforms. For content owners it provides a means to create licences that can be exchanged and interpreted by different channel partners across the value chain. And for content consumers it provides the ability to access content through different clients, interpreting a common rights expression language.

2 The Business Case

Content owners and rights-holders worldwide are engaged in preparing strategies and implementations for secure on-line digital content delivery. From the largest media companies to individual creators, the opportunities and urgency of on-line delivery are increasingly well understood. Secure delivery of digital content will open global markets, reduce distribution costs, create more intimate contact with consumers and enable innovative business models and services that will allow content providers to better serve their customers. However, on-line delivery also brings with it some specific challenges, such as

technical complexity, implementation costs and fundamental business change. And of course, there is the fundamental issue of protecting content in its digital form.

Preparations for secure on-line delivery are made more urgent and arguably more challenging because they are taking place against a background of widespread Internet piracy directed at physical content. Assets that were formerly comparatively safe between hard covers or on fixed media are being illegally extracted to new digital formats and exchanged without fee or recompense to the content owners. This process seems to be following an exponential curve in all sectors of the content industry. For instance, figures recently released by the British Phonographic Industry show that in Great Britain, sales of audio CDs were down by 15% in Q2 of 2002, the fall attributed to the state of the economy and to piracy¹. The publishing industry is also suffering. Scanning of books and journals is fast becoming a way of life in several areas of the world, with the resultant loss of revenue from potential sales. While this is not yet a major threat to literary fiction, it presents a major challenge to professional, Scientific Technical and Medical (STM) and text book publishers.

The imperative to develop secure systems for on-line distribution is therefore increasingly urgent. However, the technology barriers to achieving a secure delivery channel under the control of the content owners often seem insuperable to content owners. This effectively prevents them from reaping the huge potential benefits from on-line secure distribution, namely lowered costs and bigger markets. The upshot is that they are left in the unenviable position of having to rely on increasingly piracy prone analogue modalities or face fundamental changes in business models, for instance to sponsored publication.

From a management perspective, one of the major perceived barriers to the development of extensive secure online markets for content is that the security applications available from vendors appear to be non-interoperable. Whereas in the analogue space, content owners are used to being able to choose technology and be relatively sure that the way they package the content will not prove an impediment to its consumption (a CD produced on one pressing system is interoperable with all CD players), this does not seem to be the case in the digital environment. Worse still, different encryption and player systems often dictate that content owners have to develop multiple work flows for content packaging, dramatically increasing production and distribution costs.

However, there is a solution, part of which has been delivered by the content industry itself. For many years, content owners have been developing consistent content identification and description systems. Initially driven by the needs of the physical supply chain, such systems are now increasingly pervasive. With consistent and industry standard identification and metadata processes of this nature, content owners have prepared the way for the major change digital brings, namely the conversion of the market from the sale of content to the sale of access based on intellectual property rights. With its ability to identify and describe its assets in a logical, consistent manner, the content industry has provided a platform for the automation of licensing, both in the business to business and business to consumer contexts.

¹ Financial Times, 13 August 2002

Machine to machine, automated licensing, whereby licences can be exchanged and interpreted by software, will bring the full benefit of on-line delivery to maturity. However, while it is possible now to identify and describe assets securely with industry standard systems, what has been missing up to now is an industry standard method for creating and managing licences. By integrating its identification and descriptions systems with a standard licensing language, it would be possible for content owners to identify, describe and package content with a licence that could be read by any DRM system. This would have huge implications for interoperability, in that content owners would be able to deliver their content through any DRM distribution channel, irrespective of the technology.

This is precisely what the MPEG REL offers and why MPEG has produced the standard. The rest of this paper provides more detail on the background to the MPEG REL and some insights into its functionality and future.

3 The DRM landscape

The past ten years saw rapid development in technologies to protect and trade content on line. A highly competitive market in software for encryption, watermarking, authentication, tracking and tracing emerged. Content owners, wishing to select a technology, were faced with an almost impossible task. Caught between hesitation and fear of selecting a technology that might become obsolescent, many opted to do nothing.

Nevertheless they continued to work on supply chain issues, many of which have latterly taken into account both off-line and on-line modalities. New identifier systems, such as the Digital Object Identifier, the International Standard Audio Visual Number and the International Standard Work Number, were developed. These are complemented by innovative metadata schemes, such as ONIX (for the international book trade), the SMPTE data dictionary (for broadcasting) and the data dictionary for the recording industry associated with the GRID initiative.

The shake-out at the end of the dot.com era was helpful for content owners. Many of the start up companies that had been leading the field in technologies for digital rights management failed or were acquired by larger companies. This Darwinian process led to a greater clarity in the market place and recognition of the need for some standards in the digital rights management technology field.

These two strands of work, the Management of Digital Rights (identification and metadata, including the expression of rights) and the Digital Management of Rights (encryption, watermarking, authentication etc), have been gradually pulled together in a single standards forum, the Moving Picture Experts Group, MPEG (formally Working Group 11 of the ISO/IEC Joint Technical Committee, Sub-committee 29). MPEG, now in its fifteenth year of operation, has enabled IT companies, Consumer Electronics manufacturers and content owners to agree a growing set of standards to support the secure delivery of content over networks. Its success can be gauged by the development of the digital TV industry, which is entirely based on MPEG standards.

While these two strands of work might seem to be very different, they both have one thing in common – they are driven by a need to create an efficient market for content subject to intellectual property rights which will satisfy the needs of consumers of content as well as its creators, owners and distributors. Without such a market, investment in digital content and the mechanisms by which it can be traded and consumed will never be justified.

4 The Role of Standards

It is significant that the earliest standards were made to enforce honesty among people. At the dawn of trade, unfair dealings and incorrect measurements were resented and systems were soon developed to enforce fairness and transparency. The earliest written standards were for weights, probably from measures developed in the Euphrates basin. In 4000BC the Egyptians followed with a standard for linear measurement, the cubit, which formed the basic measurement standard for the pyramids and other great monuments. Conflict too, has driven standards - war drove Napoleon to impose the metric system to ensure that ordnance used in battle was interchangeable between armies.

Today, commercial dealings increasingly rely on standards. In the global economy, where a common infrastructure of understanding is essential, standards are used both to create markets and to maintain them in an efficient state. This is true in both business and technical senses. A good example of the former is exemplified by the Bolero standard for the shipping industry, which used to suffer exceptionally high documentation costs because of a complete lack of standards for management of bills of lading. With the introduction of the Bolero paperless messaging system, the shipping industry has seen documentation costs dramatically fall and efficiency massively increase. A good example of a successful technical standard for the media industry is the MPEG-2 specification, which was fundamental to the creation of digital television.

Standards are not new to the information and content industries. The ISBN standard was developed to meet the needs of the book industry supply chain. It has been astonishingly successful and now it is inconceivable that the publishing industry could operate without it. The ONIX standard will no doubt prove to be just as successful. Both have proved their worth by reducing costs and enhancing interoperability.

In brief, standards are valuable because they improve the conditions under which business is conducted, by promoting efficiency and communication which together can lead to a reduction in costs, growth of markets and increased profitability.

The business case for the MPEG Rights Expression Language, a technical specification for the creation of machine readable messages to enable rights management, succeeds or fails on precisely the same basis. Its primary purpose, which it shares with that first, primitive standard for weights, is to ensure the precise description of the agreements between participants in the value chain. By providing this "secure mediation" the MPEG REL will support the development of the on-line market in content, bringing buyers and sellers together, in an authenticated, transparent fashion, across all those digital rights management systems and domains that implement the extensive range of MPEG standards.

5 Introducing MPEG-21

5.1 Background to the MPEG Initiative

Established in 1988, the Moving Picture Expert Group is formally Working Group 11 of ISO/IEC Joint Technical Committee 1, Sub-Committee 29. Since its inception, it has produced a range of hugely important standards that have created new industrial sectors. For instance, MPEG-1 enabled such products as Video CD and MP3 while MPEG-2 enabled such products as Digital Television set top boxes and DVD. Subsequent standards such as MPEG-4, the standard for multimedia for the fixed and mobile web, and MPEG-7, the standard for

description and search of audio and visual content, are only now being incorporated into consumer products but will doubtless prove to be as influential as their predecessors.

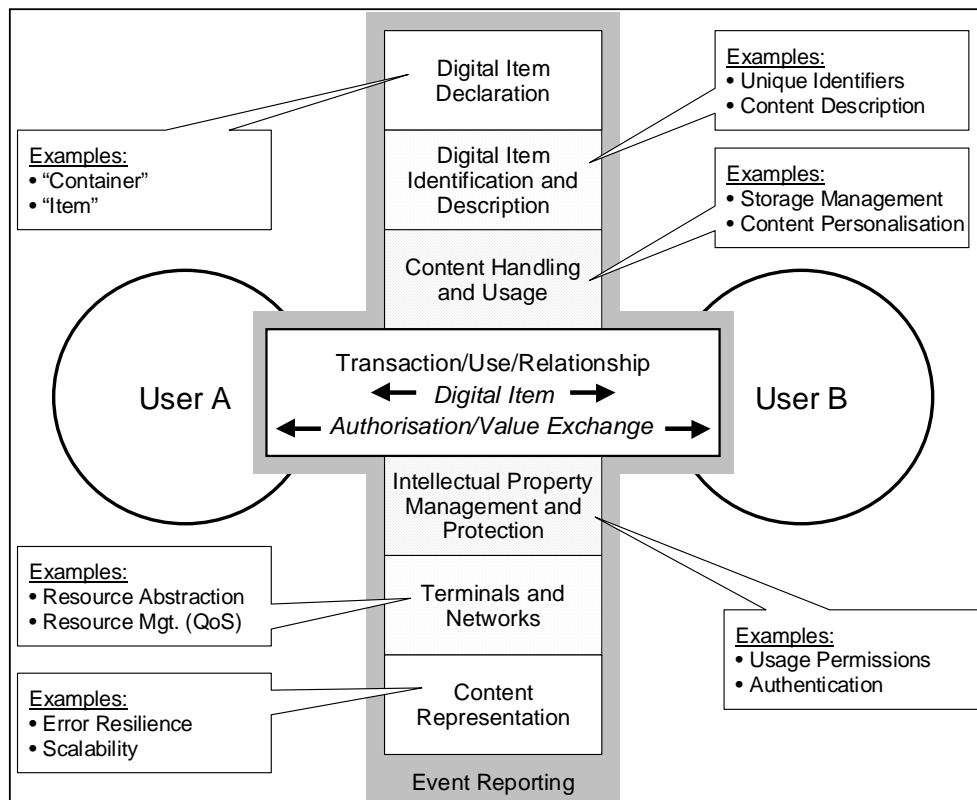
While the previous standards created by MPEG focused on audio and audio-visual, the latest standard, MPEG-21, is formally content agnostic. Dubbed, the "Multimedia Framework", work on MPEG-21 started in June 2000. Being content agnostic, the purpose of MPEG-21 is to develop a set of specifications which can support the delivery of rich multimedia content in a secure environment.

One of the important aspects of MPEG-21 is that it should be seen as a baseline specification, as its "framework" description suggests. For MPEG-21 is intended to provide the underlying specification infrastructure on which content verticals can build further specifications, that will be interoperable between verticals by virtue of their common foundations. More will be said about this later in this paper.

5.2 The Seven Architectural elements

As part of its early work, the MPEG-21 experts produced a Technical Report, which laid out the proposed architectural elements for the future specifications. The elements were identified through an analysis of the functional requirements for secure online delivery of rich multimedia content.

Seven elements were identified, which will work together to provide a secure, auditable environment, capable of delivering content to consumers on demand, in a format appropriate to the end user device, while satisfying value chain participants' need for transparency and security. A diagram shows the seven elements.



While this vision of MPEG-21 is being gradually superseded by the actual MPEG-21 specifications in development (which do not map directly onto the seven elements), the diagram is very familiar and remains useful. So, for instance, the digital item identification and description element has been addressed in the Digital Item Identification Specification which is already finalised. The Rights Expression Language is part of the Intellectual Property Management and Protection element, shown in the diagram as “Usage Permissions”.

5.3 The specifications so far

To date two specifications have been finalised. They are:

- MPEG-21000, Part 2 – Digital Item Identification
- MPEG-21000, Part 3 - Digital Item Declaration

The Digital Item Identification specification defines how different identification schemes can be managed in the MPEG environment, while the Digital Item Declaration describes the contents of a multimedia package.

In addition to the two completed specifications, six others are currently in development. Numbers given refer to parts of MPEG-21.

4. Intellectual Property Management and Protection (dealing with security processes)
5. Rights Expression Language (the subject of this White paper)
6. Rights Data Dictionary (supportive of the REL)
7. Digital Item Adaptation (dealing with reconfiguring content to suit device capabilities)
8. Reference Software (to guide implementations of the specifications)
9. File Format (dealing with issues arising from the need to have a compressed file format for all kinds of content)
10. Digital Item Processing (dealing with methods for processing digital items)
11. Evaluation Tools for Persistent Association (dealing with tools for evaluating watermarking technologies)
12. Test Bed for MPEG-21 Resource Delivery

This list, though long, is not the end of the MPEG-21 story, for it is anticipated that several further specifications in the series will be necessary. As always, these will be only produced in response to articulated industry requirements.

5.4 Using the MPEG Standards

All MPEG standards are “big” standards, each with several parts. For instance, MPEG-4 has ten parts and MPEG-21 already has nine. However, MPEG has always worked on the assumption that parts of its standards can be used individually and choosing some elements while not using others is fairly common. For example, the MPEG-2 standard comprises 10 parts, including Systems (Part 1), Video (Part 2), Audio (Part 3), Advanced Audio (Part 7) and Delivery (Part 10). It is very unusual for a system to use all these parts; for example parts 3 and 7 are almost never used in conjunction. Similarly an MPEG-2 based video surveillance system would only use Part 2 and disregard all other 9 parts. It is anticipated that implementers will adopt exactly the same approach

to MPEG-21. This process enables implementers to select those parts of the standard they require, while leaving the rest. The MPEG REL is a prime example of this process and has been designed to be capable of implementation on its own.

To understand a bit more of the thinking behind the MPEG-21 specifications and why individual parts can be used on their own, it is necessary to recognise the difference between the infrastructural and the application domains. What MPEG is attempting is to ensure that different *proprietary* applications can work together by using a common, standardised *infrastructure*. Much of this has to do with messaging, ensuring that the various vendor technologies and products available in the market place are designed to understand a common lingua franca defined through the different parts of the MPEG Standard.

So, for instance, MPEG has specified the Digital Item Declaration as a way of describing the contents of a Digital Item (MPEG's term for a digital object). While different proprietary products and technologies will certainly have different competitive characteristics, such as their ability to render content, they should all be able to understand the Declaration (really, the message) that is made about a Digital Item in a format called the "Digital Item Declaration" (or DID). This common format will be an advantage to device manufacturers, who will only have to engineer their technology to understand a single message structure and it will also benefit content owners who will have a single template (the DID) to follow when packaging content for the MPEG compliant environment.

This concept of consistent, standard messages lies behind much of the MPEG-21 development. And given that one of the most important aspects of MPEG-21 concerns the secure delivery of content, one of the most important areas of messaging is concerned with Intellectual Property Rights. And this is why MPEG has developed a Rights Expression Language to complement its other messaging specifications.

6 Why a Rights Expression Language is key

The Rights Expression Language is increasingly seen as one of the cornerstones of MPEG-21, for a number of significant reasons.

Given that the trade in rights is likely to be one of the drivers of network commerce, it will be essential to have a process by which the rights to content can be packaged within machine-readable licences, guaranteed to be unambiguous and secure, which can then be processed consistently and reliably by MPEG compliant devices. Without such a system, it is unlikely that content owners will trust the other MPEG technologies.

Second, given that the MPEG specifications are about defining messages that enable proprietary technology to be designed in a technically interoperable way, the Rights Expression Language is *the* key to technical interoperability between proprietary DRM systems.

This is why MPEG has devoted considerable resources, both in time and expert input, to the development of its rights expression language specification.

6.1 What goes into a rights expression language

The basic construct of a Rights Expression Language is the rights expression, which describes a permission granted to a user or consumer of protected content. The expression can be very simple, such as, "this content may be

printed/played”, or very complex, such as, “this content can be played on Thursday 7th March at 6.00 a.m., provided that the device meets the following criteria, e.g. “the device must be equipped with a secure processor” and the user has contacted www.anyplace.com first in order to provide the details, e.g. name, address, age. Such rights expressions can be generated by any person authorised to grant permissions over protected content, and can be passed along a value chain, sequentially obligating parties in return for grants.

In order for a rights expression language to be machine-readable, it must be based on a syntax that is recognised, such as XML (eXtensible Mark-up Language), the well established W3C standard, widely acknowledged as a fundamental web technology. And in order for the language be able to convey instructions, it must contain enough unambiguous semantics (units of meaning that can be understood by a computer) to be able to create the expressions.

Finally, so that the expressions created can be protected, there must be some measures built in, such as the ability to digitally sign rights expressions (i.e. technologically verify that the expression is genuine and can be trusted), so that their authenticity and tamper-resistance can be verified.

6.2 Developing the MPEG REL

The process of developing the MPEG-21 Rights Expression Language began with the definition of a set of 48 requirements, agreed by experts (drawn from IT, CE and the content industry) to be critical to the success of the language. These requirements ranged from the expressiveness of the language (to ensure that it should be able to frame any kind of permission a rights owner might wish to grant) to the security of the language (to ensure that it should be possible to protect a rights expression from unauthorised tampering). The requirements also covered functionality such as the conditions which must be met in order to exercise a grant of permission.

Having defined and agreed the requirements, MPEG then called for technology proposals. This is normal practice in many technology standards organisations. Several submissions were made, but the eventual choice was reduced to two submissions. The process of selection took ten days, during which MPEG experts, both those with an interest in the outcome (because they were to be either users or implementers of the language) and those with no interest (such as MPEG XML specialists), exhaustively tested both submissions. The eventual choice was made in favour of XrML, based on original work conducted at Xerox Parc in the early 90s.

While the structure and basic semantics of the language were already comprehensively established in the XrML submission (i.e. the language could create structure expressions to define actions that must happen before other actions can take place), most of the semantics for the actual permissions themselves have been developed by MPEG, on the basis that semantics for rights do not exist elsewhere and thus could not be sourced from another standard, (a regular practice in MPEG). They have also been the subject of extensive discussion and agreement.

Work on the REL specification began in December 2001. Since that time, the work has progressed smoothly and the text of the standard has been extensively revised. At present, the specification is under ballot by the National Body members of MPEG, who will present their final comments in July 2003. Once these have been dealt with, the text will be promoted to a Draft International


Standard, prior to its final publication as a full International Standard in the late fall of 2003.

The publication of the International Standard will not be the end of the story of course. Like all standards, amendments and revisions will be made as the years go by. This is particularly important in a standard such as the REL, where content owners will continually call for added functionality to support the richer and more complex on-line business models they develop in the future.

6.3 The Architecture Rough Guide

The MPEG REL adopts a simple and extensible core data model for its key concepts and elements. The data model for a rights expression consists of four basic entities and the relationships between these entities. This basic relationship is defined by the REL assertion "grant". Structurally, an MPEG REL grant consists of the following:

- The principal to whom the grant is issued
- The right that the grant specifies
- The resource to which the right in the grant applies
- The condition that must be met before the right can be exercised



Principal

A Principal encapsulates the identification of principals to whom rights are granted. Each Principal identifies exactly one party, which may be a person, a terminal (end user device) or a network node. The party is identified by information unique to that individual or machine. The Principal type supports a full range of secure identification technologies.

Right

A Right is the "verb" that a principal can be granted to exercise against some resource under some condition. Typically, a Right specifies an action (or activity) or a class of actions that a principal may perform on or using the associated resource.

Resource

A Resource is the "object" to which a principal can be granted a right. A Resource can be a digital work (such as an e-book, an audio or video file, or an image), a service (such as an email service, or B2B transaction service), or even a piece of information that can be owned by a principal (such as a name or an email address).

Condition

A Condition specifies the terms, conditions and obligations under which rights can be exercised. A simple Condition is a time interval within which a right can be exercised. A slightly complicated Condition is to require the existence of a valid, prerequisite right that has been issued to some principal. Using this mechanism, the eligibility to exercise one right can become dependent on the eligibility to exercise other rights.

The Core data model will be enhanced by a number of so-called "Extensions" which will add both functionality and applicability. For instance, the Standard Extension greatly enhances the range of Conditions that can be applied to grants, thus providing extra functionality. And of crucial importance to content owners is the Content Extension, which defines a specific set of rights, such as Play and Print, which are applicable to content. It is anticipated that the REL will quickly have many Extensions of this type. For while MPEG has defined a basic Content Extension it is quite likely that distinctive content verticals will require rights not defined in the Standard.

6.4 Defining the Rights Semantics

In order to support the REL, MPEG-21 is also defining a specification for a Rights Data Dictionary, with a methodology for defining and cataloguing terms necessary for the description of rights. The MPEG-21 RDD specification, though independent, is closely linked to the REL specification in order to ensure the interoperability of rights semantics across different content types. This will greatly enhance the value of the REL to content owners.

7 Business Model Support

7.1 REL and DRM Technologies

While the REL will become a vital part of the DRM landscape, it is solely concerned with setting business rules. It does not provide any encryption functionality for content, though it does link to processes for ensuring the rights expressions themselves are tamper proof and capable of authentication.

In many ways, the relationship of the MPEG REL and digital rights management technologies is similar to that between identification and metadata schemes and digital rights management technologies. While it would be possible to trade content through digital rights management systems without standard identification and metadata schemas, the content would not be interoperable across vendors or within the content vertical, making commercial relationships with distributors, consumers and users infinitely more difficult. Every vendor would have its own system, forcing content owners to identify and describe content for each vendor. In much the same way, and with most of the same

adverse consequences, it would be possible to employ a non-standard rights language. Most damagingly, this would tie content owners to a particular digital rights management system, which would be commercially very unwise for any content owner.

While the MPEG REL is designed to be implemented between DRM systems and services proprietary or otherwise in the MPEG context, it is also capable of enabling content to be moved between MPEG and other environments. This is possible because it is written in XML, the standard language of the web, which can easily be translated, just like any other language, into other systems. This DRM agnosticism means that a rights expression, once written, can be passed from one DRM system to another, providing a very substantial measure of interoperability.

7.2 Some implementations

To date the MPEG REL has not yet been implemented. However, the technology on which it is based, XrML, has been implemented in a number of major vendor applications.

One of the first was the Microsoft eBook Digital Asset Server system. While this used an early version of XrML, the implementation demonstrated the basic capabilities of the language, thus serving as a signpost for the way forwards.

A more fully functioned implementation, based on XrML 2.0 (the baseline technology for the MPEG REL) is that created by the Dutch software company DMDSecure. The company produces a platform for the delivery of secure digital media, including streaming and downloads based on a variety of business models, such as subscription and promotion. The purpose is to provide consumers with the widest choice for the consumption of secure content while at the same time providing publishers and other content Owners with the widest reach. A vital part of this is setting the rights attached to the content so that it can be driven through several different enforcement engines. Using XrML 2.0 enables DMDSecure to use a single language for each of the enforcement systems fed by the platform. This is significant because it enables content owners to trade content through a variety of different digital rights management technologies, thus producing an important measure of interoperability the owners and consumers alike.

Another significant implementation has been created by Integrated Management Concepts (IMC), which provides software applications for Financial Project Management. The company primarily works in the Aerospace, Engineering and Computing (A/E/C) market spaces.

IMC's product, Microfusion Millennium, integrates with project management tools such as Primavera and MS Project. In earlier versions simple access control was enough but with larger workgroups it was decided that some DRM-like capabilities embedded in the security/access policy framework would be advantageous. Using XrML, upon which the MPEG REL is based, the new version of Microfusion Millennium has enabled project managers or project leads to assign rights (such as View, Print and Delete) on an individual basis to workgroup members, constrained by a specific timeframe. This has enabled the company to meet its customers request for more granular access policies and has proved to be a huge product differentiator.

7.3 New Business Models

The limitations of current digital rights management technology is substantially caused by the lack of a means to express in an unambiguous, precise, machine readable way the complex permissions content owners may wish to grant to distributors and consumers. While it is currently possible to grant simple print, play and read permissions to consumers, anything more complex, such as conditional time licences, is not possible outside the context of highly proprietary solutions. Nor is it possible to create business relationships with distributors, based on machine readable licences that can be automated to a significant and therefore cost reducing degree. As a consequence, today's business models have limited appeal to general consumers.

One of the main aims of the MPEG REL has been to unlock the possibilities of innovative models, which will help to encourage consumers to engage in on-line commerce. For instance, the MPEG REL will facilitate consumer access to rich multimedia content on a limited time basis (similar to a loan) or on the basis of membership of a group. It would also be possible to create a business model where access to content is conditional on another action, such as the purchase of a physical item (e.g. purchased on a web site and delivered by carrier). The possibilities are seemingly endless and limited only by the imagination of the rights owner and the functionality of the DRM service.

Given that the MPEG REL has been specified in line with a very extensive list of requirements, it will be possible to design exceptionally complex, multilayered business models, with licences that intermediaries can further customise for end users. This will enable content owners to reward their distributors by the issue of a single licence upstream that can be progressively amended, while ensuring that the original terms set by the content owner are enforced. It will also enable securely, profitable peer to peer business models, where end users can reap benefit from distributing content to their friends and colleagues.

8 MPEG REL in a wider context

Like any standard before it, MPEG will be judged by the extent to which it is seen as useful and taken up by industry. Many factors will be considered, such as ease of implementation, integration with other standards and ultimate usability. To ensure that MPEG standards are what industry wants, MPEG has not only actively encouraged participation by all stakeholders, but it has also developed liaison relationships with other standards organisations that share its technical background. By this method, it exchanges information about on-going work and promotes its own specifications. For instance, it has recently started a joint activity with the ITU on video coding, which will bring together two organisations that have hitherto produced separate standards

The importance of broad stakeholder participation and liaison activity for the MPEG REL is clear and has become more urgent in the context of MPEG-21, given that its specifications are content agnostic. Already, other communities are building extensions. Towards the end of 2002, it was announced that the Open eBook Forum plans to develop its rights grammar based on the MPEG REL. And at the same time, the TV-Anytime Forum, which deals with technologies related to content distribution in the digital television industry, including metadata, content referencing and rights management and protection, is actively considering the development of an extension. In addition there is now a proposed liaison between the Learning Technologies Sub-committee (LTCS) of

the IEEE and MPEG, specifically in order to support the LTSC in selecting a rights language for the delivery of teaching and learning.

While it is expected that many communities will create extensions to the baseline REL, it is also possible to profile the language. This involves the creation of a subset of the language for a specific purpose within a content vertical. This will become extremely important as content moves through the value chain down to many different devices. For instance, for a mobile phone the full power of the language as required by a PC will not be needed and a slimmed down version – a profile – can be used to express usage permissions. In another instance, a profile might be created for delivery to public music kiosks with limited functionality. By this means, the MPEG REL can be exploited both for demanding applications on powerful equipment and for simple applications on lightweight consumer devices.

Beyond the content industries, other sectors will almost certainly require the functionality provided by a rights language. Almost all business communities, for example, insurance and banking services, and healthcare systems have a requirement for document security that may not involve financial exchange. In the case of business, a company may wish to restrict certain documents to certain employees or to communicate with third parties on the basis of certain conditions. Both these requirements can be met by the MPEG REL. In healthcare, the security of patient records on networks is a pressing need. The MPEG-21 Rights Expression Language would help enable a hospital or healthcare network to set very explicit rules for access to documents managed by a conditional access system. What is interesting is that neither of these examples involves payment of any kind, yet both benefit from the functionality of the MPEG REL. By developing its own, sector specific extension, based on the REL Core, each vertical can benefit from the functionality of the rights expression language.

The adoption of the MPEG REL beyond the content industry will have advantages in addition to those that accrue to any individual sector. The obvious advantage is that it will be possible to communicate content governed by rights expressions beyond the boundaries of a particular sector. This could be particularly useful in certain types of corporate publishing, such as the exploitation of data in the financial services industry, where a standard shared with commercial publishing may produce new business opportunities. However, an equally important advantage for consumers is likely to flow from the existence of a common REL standard for implementers of rights management and conditional access technologies. Such a standard will produce greater competition among vendors and implementers. Users will be able to choose “best of breed” based on a single standards, which will lead to lower prices and more customers. In turn this will enable implementers to spread their own costs and reduce prices still further. The start of the process of the MPEG REL moving outwards from the content industries is exemplified by the activity of the Oasis Technical Committee investigating the use of the MPEG REL Core Schema and Standard Extension as the basis for a new Enterprise Extension for business documents.

While true interoperability may still be a distant prospect, a common rights expression language, with extensions based on the MPEG REL, can incrementally bring many of the benefits true interoperability will eventually yield. As extensions are created in multiple content verticals, it will be possible to transfer content generated in one securely to another. This will lead to cross channel fertilisation and the growth of multimedia content. At the same time, a common rights language will also lead to the possibility of broader content distribution (by

enabling cross-DRM portability), thus providing more channel choice for consumers. It is this vision of the MPEG REL spreading out that is such an exciting prospect.

9 The future

Just as the telephone was validated by its own take-up (the more subscribers the more socially and commercially valuable it became), so the MPEG REL will become increasingly useful the more widely it is implemented. The history of MPEG standards would seem to suggest that implementers will start building to the specification in mid-2003, coincidental with the completion of the standard. This will be followed by extensive take-up within two or three years, so that by mid 2006, the MPEG REL will be a pervasive technology, implemented across many different digital rights management and conditional access systems, in both the content industries and in other, non-rights based industries. In association with other MPEG technologies, the REL will provide the “glue” for the networked economy, enabling parties to have trusted dealings with each other. It will be the basis of commercial licensing transactions, just as much part of the commercial infrastructure as many other, previously innovative technologies have been. In this way, the REL will ultimately become a “transparent” technology, as invisible to the user as the phone infrastructure is today.