



GiantSteps
Media Technology Strategies

1841 Broadway, Suite 200
New York NY 10023
212 956 1045
fax: 212 258 3286

<http://www.giantstepsmts.com>

Integrating DRM

with

Peer-to-Peer Networks

Enabling the Future of Online Content
Business Models

By **Bill Rosenblatt**

November 22, 2003

Contents

Contents.....	1
Executive Summary	2
Background.....	2
The Rise and Importance of Peer-to-Peer.....	2
The Rise and Importance of DRM.....	3
The Gulf between P2P and DRM.....	5
Integrating DRM with P2P: Needs and Opportunities.....	5
DRM Technology Features for P2P Networks	7
Reasonable Usage Support	8
Lightweight Superdistribution.....	8
Standards Support	10
User Experience.....	12
Gaps in Existing DRM Technology	13
Cost-Related Functionality Limitations.....	13
Device Tethering	13
Lack of Superdistribution Support	14
Complexity of Integration	14
Conclusions: Developing the Market	15
About the Author	17
About GiantSteps Media Technology Strategies	17
About DigitalContainers LLC	17

Executive Summary

The rise of peer-to-peer (P2P) networks has been an inevitable outgrowth of the rise of the Internet. Unfortunately, P2P networks have grown from useful tools in information sharing to havens for trafficking in unauthorized copies of intellectual property (IP). Owners of IP, meanwhile, have been pushing for digital rights management (DRM) technologies to control distribution of IP so that it does not fall into the wrong hands.

Supporters of P2P networks appear to be at odds with DRM-supporting IP owners, but P2P networks offer a lot to users as well as other participants in content business models, and they are here to stay. Integration of DRM into P2P architectures is inevitable, as IP owners try to walk the fine line between embracing functionality that users want and maintaining control over their IP.

This white paper explains the motivation for and inevitability of integrating DRM with P2P. After briefly reviewing how both DRM and P2P came into being, we explain the need and opportunity to integrate DRM functionality into P2P networks. We discuss features of DRM technology that make it especially appropriate for integration with P2P, and we summarize shortcomings of many existing DRM solutions with respect to those features. We conclude with some suggestions for how to develop the market for DRM solutions that are optimal for integration with P2P networks.

Background

Both DRM and P2P are creatures of the Internet era, but they came into being at different times and for different reasons. Here we will examine the origins of and motivations for each.

The Rise and Importance of Peer-to-Peer

Technologies for peer-to-peer data exchange over networks have been in existence virtually since the beginning of computer networking in the 1980s. Nowadays, in its most generic form, the term peer-to-peer is used to distinguish a network architecture from client-server, which has been a dominant architecture in both pre-Internet network applications and on the Internet itself.

The idea of client-server is that resources (such as files) are on a server computer, and clients can only obtain resources through servers. If Client C_1 wants to get Resource R from Client C_2 , then it needs to go through a server to do so, thereby requiring the server to have a list of resources that includes Resource R and C_2 as its location. In contrast, peer-to-peer networks allow clients to exchange resources directly among each other.

Peer-to-peer architectures came into being in the pre-Internet age about ten years ago with technologies such as Microsoft's Windows for Workgroups (WFW), which enabled PC users to access files on each others' PCs. Sun Microsystems's Network File System (NFS), which emerged even earlier and enabled all computers on a network to make their file directories available in a network-wide hierarchy, can also be considered as a form of peer-to-peer. When the commercialization of the Internet began in the early-to-mid 1990s, File Transfer Protocol (FTP) – particularly the variation called "anonymous FTP" that does not require a file user to identify itself to the file owner – became the most important antecedent to P2P as we know it today.

Internet P2P networks provide services similar to the likes of NFS, WFW, and FTP, though with more sophisticated searching and browsing functionality, over the public Internet instead of institutional networks. Most of the early commercial development of the Internet centered on the World Wide Web, which is very much a client-server model. P2P networks needed to build on Internet-based protocols other than the HTTP protocol that powers the Web. The important thing to understand is that P2P networking is not a new model; at its core, it is simply an application of well-known networking models to the Internet.

P2P networking is not a new model; at its core, it is simply an application of well-known networking models to the Internet.

The first well-known P2P service on the Internet was, of course, Napster, which came online in June 1999. Napster was actually not a pure P2P network, because it relied on a central server to act as a catalog of files on the network and their locations. (Napster's server-based architecture ultimately led to its shutdown by a judge a year after it started.)

The Napster phenomenon gave rise to post-Napster P2P networks, such as the proprietary FastTrack network used by KaZaA and Grokster, and the open-source Gnutella network used by LimeWire and Morpheus. Both of these networks were designed without central servers so as to avoid Napster's legal fate, but even FastTrack is not a pure peer-to-peer service: it relies on so-called supernodes, which constitute the first level of connectivity in the network and help make request routing decisions. Gnutella, in contrast, is purely peer-to-peer, with no clients having special distinctions of any kind.

Owners of copyrighted intellectual property (IP) have seized upon P2P networks because they embody a set of attributes that make them ideal for unfettered distribution of files:

- Unlike local-network file sharing technologies such as NFS, WFW, and their successors, they are accessible throughout the Internet, not just on an institutional network.
- Unlike sending file attachments in email messages, they do not require that the source of a file actually send it or even know the identity of the recipient.
- Unlike duplication of physical media such as CDs or DVDs, P2P networks allow files to be copied instantaneously, with maximum automation, and without the cost of physical media.

Of course, the same attributes that frighten IP owners make P2P networks attractive to those who genuinely want to publish information as easily and widely as possible.

The Rise and Importance of DRM

Although P2P on the Internet did not come into being until 1999, IP owners were concerned with digital networks as conduits for unauthorized file copying long beforehand. Most industry observers identify 1994 as the year when digital rights management began to emerge as a field on its own¹ – the same year as the beginning of the

¹ See, for example, Proceedings: Technological Strategies for Protecting Intellectual Property in the Networked Multimedia Environment, 1994, <http://www.cni.org/docs/ima-ip-workshop/>.

commercialization of the Internet, although early contributors to the DRM field did not necessarily see the Internet as being as dominant as it has become.

IP owners in the mid-1990s looked at online rights management primarily as a question of emulating business models from the offline world. As a crude example, the “rights management” properties of a printed book result directly from its physical characteristics, e.g., it is difficult to copy books in their entirety and virtually impossible to change their contents in place. Publishers sought technologies that would bring similar behavior to the online digital world, and early DRM solutions, such as IBM’s Cryptolope and EPR’s (later InterTrust’s) DigiBox, attempted to provide this.

Just as P2P is an Internet application of preexisting network architectures, DRM technology is really an extension of techniques long used in operating systems to control users’ access to system resources. There are many different types of DRM implementations, but they tend to conform to a common architecture². In this architecture, the user receives an encrypted file, containing the content, along with a *license* that stipulates what rights the user has to the content. A piece of software or hardware on the user’s client device interprets the license and, if authorization is successful, decrypts the content and does what the user intends (play, view, print, copy, etc.).

Variations on the canonical DRM architecture involve such issues as:

- Whether the authorization is done on the basis of a user’s identity, a device’s identity, or both.
- Whether the software doing the authorization is built in to the playback device or software, built in to the platform on which it runs, or independent of those.
- Whether the license is bundled in with or separate from the content.
- How much fine-grained control the IP owner has over specification of rights.
- Whether or not the user is required to be connected to the network at all times.
- How financial transactions are integrated with the authorization process.

IP owners have been using DRM to implement new business models, which are not just analogs of existing offline models. Such models represent the brightest future for online content distribution. However, they have only been modestly successful, because it takes a lot of time and effort to get users comfortable with new ways of consuming content.

As a result, DRM is starting to take off as a component of online content models in niche markets, such as the online music distribution of Apple’s iTunes, RealNetworks’s Rhapsody, and Napster 2.0; eBooks and ePeriodicals from various publishers; and online film download services like MovieLink and CinemaNow.

DRM is starting to take off as a component of online content models in niche markets.

² For more details, see Bill Rosenblatt et al, *Digital Rights Management: Business and Technology*, John Wiley & Sons, 2001, Chapter 5.

DRM has had a hard time developing as a market, for several reasons. Online emulations of offline content models have been very rough from the perspectives of user convenience and support for some usage modes that are legally protected or that users have come to expect, which we will examine later. There is also an ingrained notion in consumer behavior (and, some feel, in legal precedent as well) that people should be allowed to do what they wish with digital content products, without fear of being controlled or monitored – as DRM technology can do.

Yet at the same time, the networked digital paradigm has opened up the possibilities of “do what they wish” to include rampant, unrestricted, perfect copying, and IP owners need to be able to control that. Therefore, DRM continues to develop toward giving users convenient, seamless experiences along with guarantees of privacy.

The Gulf between P2P and DRM

The way various advocacy groups portray it, DRM and P2P are polar opposites. To IP owners, P2P offers open invitations to copyright infringement and rampant theft of intellectual property, while DRM is the only way to keep the Internet from killing the media industry. To consumer advocates and some others, P2P is natural outgrowth of the “open” functionality of the Internet, while DRM represents the media industry’s attempts at playing “Big Brother” and controlling user behavior in ways that are inconsistent with the balance of interests embodied in intellectual property law.

As a result, there is a lot of posturing on both sides of the issue, as people from both sides work to get sympathetic ears from technology implementers, legislators, and the news media.

We can hope that everyone will see both DRM and P2P for what they are *and* are not, and get on with the business of using both to their advantage.

The reality, of course, is that both DRM and P2P are sets of capabilities, and they are far from mutually exclusive. As we will see, P2P functionality is key to implementing important new business models for content – models that IP owners ignore at their long-term peril. At the same time, DRM is necessary to close at least the larger holes that P2P creates in IP owners’ ability to profit from their IP. We can hope that once both sides finally get past the rhetoric, everyone will see both DRM and P2P for what they are *and* are not, and get on with the business of using both to their advantage.

Integrating DRM with P2P: Needs and Opportunities

IP owners need to consider how they can integrate DRM functionality with P2P networks so that they can offer their customers P2P functionality while also protecting themselves from copyright abuse. If we look at the history of IP owners’ business models on the Internet, we draw the inevitable conclusion that P2P is really an evolutionary extension of the user-oriented features that IP owners have been obliged to offer since the beginning of Internet commercialization.

The early digital products (including CD-ROMs and websites) produced by publishers and other IP owners were “shovelware” that merely contained repurposed content from the companies’ analog output. User contributions to websites were confined to the kind of

letter-to-the-editor sections typically found in print products. As publishers needed to find ways to overcome the liabilities of screen reading and slow, expensive dialup connections, as well as to compete with one another, they added various interactive features that included user control or user-originated content. Examples of the former include personalized content filtering, choices of look and feel, and rich search functions; examples of the latter include “community” features like discussion groups and chat rooms. Despite initial resistance from both editorial and legal departments, these features flourished and are standard on virtually all name-brand media websites today, while product formats with limited interactivity, such as CD-ROMs, are in permanent decline.

Another important step beyond shovelware is IP owners looking beyond their own websites and sending their content to places where they may find more audiences for it. Two important manifestations of this development can be viewed as precursors to P2P: syndication and affiliate programs.

Syndication is one IP owner sending content to other publishers on a formal, regular basis – for example, a restaurant reviewer syndicating its content to travel websites. Several vendors created tools for automating syndication relationships, and the open standard protocol ICE³ (Information and Content Exchange) was developed to promote interoperability. ICE has constructs that let publishers describe the rights they are conveying to subscribers, but no mechanism to enforce those rights; ICE-compliant (and other) syndication software merely automates publish-subscribe relationships that are controlled by legal contract terms. With syndication, in other words, the publisher has to know and trust the subscribers.

Affiliate programs provide another key step forward for IP owners, although they are more closely associated with general online retailing than with content. Perhaps the most famous user of affiliate programs is Amazon.com; in addition, many other retailers – including IP owners like Sony Music and Scholastic – have affiliate programs through the LinkShare network. In an affiliate program, operators of special-interest websites “stock” products from retailers by placing special links on their sites; for example, a website devoted to stamp collecting may feature various books and have links to their pages on Amazon.com. If a user clicks on one of those links and purchases the book, then the stamp collecting website gets a commission.

When used with content products, affiliate programs approximate *Superdistribution* – which is similar to peer-to-peer but is more controlled and implies that there is an IP owner that originates the distribution⁴ – except that the Superdistribution is only done to two levels. Although most products purchased through affiliate networks are physical, there is no reason why they can’t be digital – that is, delivered in digital form to the buyer as part of the purchase process.

If one were to look at extending both syndication and affiliate networks for content, one may well want capabilities that allow distribution of content to arbitrary (and arbitrarily many) parties, with technological controls over usage supplanting contractual ones because the trustworthiness of the other parties is unknown. One would also want the ability to facilitate e-commerce among all levels in the network. Put these requirements together, and you get P2P with integrated DRM.

³ See www.icestandard.org.

⁴ For more information, see Cox, Brad J. *Superdistribution: Objects as Property on the Electronic Frontier*. New York: Addison Wesley, 1996.

Participants in a P2P network can bring lots of legitimate value to both IP owners and users. IP owners have been bemoaning the need to “compete with free,” but they are coming to realize that there is much more to a positive user experience than merely claiming to have a certain item of content available for those who specifically look for it. There is a huge difference, for example, in looking for a song on KaZaA and having to put up with decoys, poor-quality encodings, spyware, and so on, versus finding the same song on a legitimate online music service, playing it in its entirety with good quality, finding artist information and recommendations for similar music, getting technical support, and having one’s privacy respected.

There are many types of services around content that might appeal to users. It’s unrealistic to think that IP owners’ websites will provide them all; it’s also unrealistic to think that all desirable content-related services even fit the business-to-consumer model in general. For example, DRM-integrated P2P networks make a lot of sense for certain types of corporate applications, such as knowledge management and collaboration, where maximum dissemination of data is paramount but so is security, and in the distribution of video and other large-sized content, where it’s desirable to offload corporate servers.

In general, peer-to-peer data exchange models provide IP owners with more ways to add value to content, including ways that the IP owner may not think of on its own. At the same time, P2P networks do provide large, scalable opportunities for abuse. The architecture that solves this problem, while scalably facilitating value-added services for content, is P2P with integrated DRM.

Peer-to-peer models provide IP owners with more ways to add value to content, including ways that the IP owner may not think of on its own. At the same time, P2P networks do provide large opportunities for abuse. The architecture that solves this problem, while facilitating value-added services for content, is P2P with integrated DRM.

At the same time, DRM technology must meet certain criteria to be acceptable to P2P participants. The two most important requirements are easy to state, if not to implement:

- The DRM technology must support users’ reasonable usage expectations. At a minimum, this includes a user’s right to use content in any format on any device she owns. Ideally, it also includes legal fair-use rights such as copying for research or criticism purposes.
- The DRM technology must also be seamless and unobtrusive. This includes installation with no extra effort on the user’s part, no adverse effects on the user’s device or platform, and operation in the background as much as possible.

In next section, we will look at technical requirements of DRM technologies that enable them to meet the above criteria and others.

DRM Technology Features for P2P Networks

DRM technology has been around for almost a decade. There are many types of DRM solutions on the market today, some of which have found success in niche markets, as mentioned above. What are the specific features of DRM solutions that make them attractive for integrating with P2P networks? Here we suggest several.

Reasonable Usage Support

The term “fair use” is a loaded one; it has a specific meaning under U.S. copyright law (its analog in the U.K., Canada, and Australia is “fair dealing”), but consumer advocates and others have extended it to stand for content consumers’ reasonable expectations of usage rights. The legal term refers to uses of content that are valid defenses to charges of copyright infringement. Uses must conform to broad legal guidelines, but ultimately a judge and jury make decisions about whether uses are fair. Therefore it is impossible to create any kind of automated system that proactively decides whether to allow a use based on legal fair-use criteria.

However, reasonable usage expectations are another matter. If a user buys a piece of content, she may well expect to be able to render (display, play, or print) that content on any device she owns⁵. The paradigmatic example of reasonable usage expectations in the analog world is to play a music CD in one’s car in addition to one’s home stereo, perhaps by taping it onto a cassette.

DRM systems should be able to support a user’s reasonable content usage expectations; this should include acting independently of individual formats and playback software or devices, and facilitating any necessarily format conversions or transcoding. DigitalContainers is an example of a DRM system that facilitates reasonable usage support: it is cross-platform, works with a multitude of media formats, and does not require a client application that the user must download and install.

DRM systems should be able to support a user’s reasonable content usage expectations.

The most important precondition to supporting reasonable usage expectations is interoperability of identification schemes for both users and devices. Currently, and with few exceptions, each DRM scheme has its own notion of identity and its own way of authenticating identities. A user’s identity in one scheme (e.g., for an Adobe eBook) is only coincidentally related to her identity in another scheme (e.g., for an online music subscription service). Attempts to create universal online identification schemes have been thwarted by a combination of technical complexity and concerns over privacy.

An ideal DRM scheme for integration with P2P networks should at least offer some degree of identity interoperability among popular formats, devices, and services; existing technology for aggregating personal information online (such as Yodlee in financial services) might apply. Yet consumer rights advocates tend to concur that identity schemes – such as DigitalContainers’ – that are based on users, not devices, offer a first approximation to reasonable usage support.

Lightweight Superdistribution

Superdistribution has been mentioned in the same breath as DRM since the early days of DRM, when a few DRM technology vendors attempted to support it. The complexity of a DRM and e-commerce scheme that allows every participant in a content Superdistribution scheme to make its own economic offers is prohibitive. For example, one peer may want

⁵ This is sometimes known as “space shifting” content, a term that is related to “time shifting,” i.e., playing broadcast content after it was originally broadcast. The U.S. Supreme Court upheld the right to time shifting in *Sony v. Universal*, 1984 (the landmark “Betamax” decision); the *legal* right to space shifting is still being contested.

to sell content items individually at a profit, while another may want to sell them at cost, another may want to loan them, and yet another may want to make a repository of items available on a monthly subscription basis.

The nearest that most DRM schemes have gotten to “Superdistribution” is a URL included in encrypted files that takes users who are not authorized to access the content to a website where they can purchase rights. This is inadequate to the needs of a P2P network, in which peers should be able to define their own business models, as suggested above. At the same time, peers should not be expected to deploy cumbersome, expensive e-commerce systems in order to implement their chosen business models.

For Superdistribution to work well with P2P networks, DRM systems should provide simple ways to define and implement content business models, including rights specifications and commerce terms. Emphasis on defining individual users or classes of users for authentication purposes should be minimized, because one of the most important aspects of P2P, as mentioned above, is that the identities of participants in P2P networks are not known in advance.

DRM systems should provide simple ways to define and implement content business models.

DRM schemes can also facilitate Superdistribution by providing as much business model support as possible integrated with content objects, to minimize implementation complexity. This implies the ability to precisely specify details of content rights being offered, such as number and type of renderings, time limits, and so on; see Rights Expression Languages on p. 10. It also means the ability to handle certain functions directly onboard content items, and to interface with web services that handle external functions that make it easy for participants to implement their business models – such as billing, usage tracking, and subscription management.

For example, DigitalContainers is a DRM technology that supports lightweight Superdistribution through its Hybrid P2P architecture, which supports the ability to describe content rights in a fine-grained manner, the ability to facilitate on-the-fly server-based user authentication, and rich functionality for supporting a wide variety of business models, including payment processing, onboard the encrypted content objects rather than on a server.

For Superdistribution support that is too complex to handle onboard encrypted content objects, DRM schemes should support integration with web services through standard interfaces so that they can be developed by a multitude of vendors. Ease of integration with web services will encourage the development of such services and their adoption by P2P participants.

For example: assume peer P_1 makes content item C available as part of a repository through a paid-subscription service. When user P_2 obtains the object, it should have self-contained functionality to retrieve P_2 's identity, send it to a service for verification that P_2 is a subscriber to P_1 's service, and then receive a license L from that subscription service that enumerates the rights to which P_2 is entitled. This is shown in Figure 1.

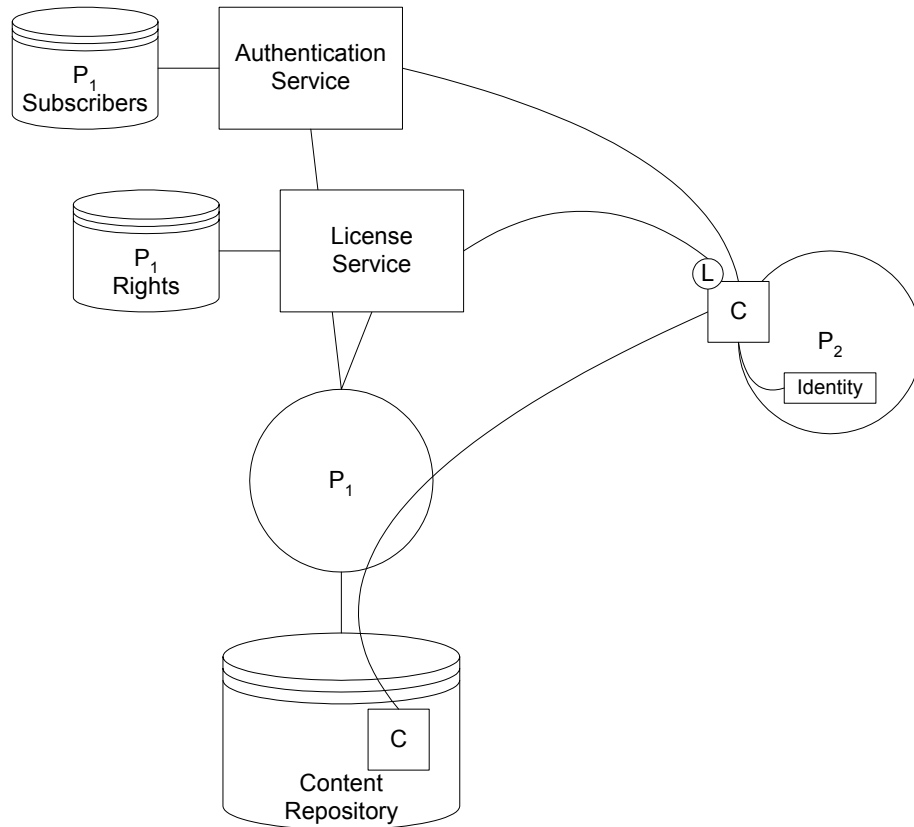


Figure 1: Two peers in a peer-to-peer architecture with DRM-packaged content. The content C has functionality for accessing web services. The Authentication Service authenticates P₂'s identity, and the License Service issues a License L for P₁'s content C.

Standards Support

The development of functions and services surrounding DRM in P2P networks would be best facilitated by DRM technology that supports various types of open standards, including:

Rights Expression Languages

To implement flexible, interoperable content distribution schemes on P2P networks, DRM schemes need to embrace standards for creating content rights specifications; these are usually known as Rights Expression Languages (RELs). RELs provide standard semantics for elements of rights specifications, such as those that would be stored in a rights database such as the one labeled *P₁ Rights* in Figure 1, including:

- The right being granted, such as Play or another render right.
- The entity to which the right is being granted, such a user or device.
- The terms under which the right is granted, such as payment or presentation of credentials (e.g., a valid subscription to a service).

The most prevalent standards in the REL area are MPEG REL⁶, from the Moving Picture Experts Group, which derives from XrML⁷ (eXtensible Rights Markup Language) from ContentGuard, Inc.; and OMA DRM⁸ from the Open Mobile Alliance, which derives from ODRL⁹ (Open Digital Rights Language) from IPR Systems Ltd. Other standards bodies, including OASIS (the XML and SGML standards body) and the Open eBook Forum, are also defining RELs.

RELs are especially important in Superdistribution networks. If P_1 passes some content to P_2 , then P_2 's rights to that content need to be a subset of P_1 's rights, and if P_2 passes the same content to P_3 , then P_3 's rights need to be a subset of P_2 's – or, if P_2 or P_3 want additional rights, they need to be able to define them with precision and acquire them from the original IP owners. A properly designed REL enables this.

Network Identification

As mentioned above, universal – or at least interoperable – identification of users and devices is a critical factor in supporting DRM ease of use and consumers' reasonable content usage expectations. The concept of a universal ID implies that a single entity controls all such IDs, which concerns privacy advocates and others. Microsoft's .NET Passport¹⁰ identification scheme, which allows users to use a single ID to access many different online services (including Microsoft's own services as well as many others), is the closest thing there is today to a universal ID scheme.

Universal – or at least interoperable – identification of users and devices is a critical factor in supporting DRM ease of use and consumers' reasonable content usage expectations.

Short of a universal ID scheme, the next best possibility is a standard for interoperability of ID schemes, sometimes known as federated network identity. In a federated ID scheme, there is no single repository of IDs, but organizations can use each others' IDs on a per-transaction or per-service basis as long as users give permission to do so. The Liberty Alliance¹¹, a consortium originated by Sun Microsystems, has defined a specification for a federated ID scheme based on the SAML¹² (Security Assertion Markup Language) standard from OASIS.

Meanwhile, Microsoft has announced that it will create a new version of .NET Passport that provides federated ID capability and uses the older Kerberos¹³ distributed authentication standard from MIT.

Web Services

Web services are the ideal way to foster the development of services that P2P network participants can use in conjunction with DRM schemes to create new types of content-related value added services with minimized cost and complexity. Two examples of web

⁶ See http://www.chiariglione.org/mpeg/standards/mpeg-21/mpeg-21.htm#_Toc23297977.

⁷ See <http://www.xrml.org>.

⁸ See <http://www.openmobilealliance.org/tech/docs/index.htm>.

⁹ See <http://odrl.net>.

¹⁰ See <http://www.microsoft.com/net/services/passport/developer.asp>.

¹¹ See <http://www.projectliberty.org>.

¹² See http://www.oasis-open.org/committees/tc_home.php?wg_abbrev=security.

¹³ See <http://web.mit.edu/kerberos/www/>.

services related to DRM shown in Figure 1 are the Authentication Server and the License Server; if P_1 gets these through service providers instead of through licensed software, then P_1 's implementation can be much cheaper and simpler.

Web services are the ideal way to foster the development of services that P2P network participants can use in conjunction with DRM schemes to create new types of content-related services with minimized cost and complexity.

There are several emerging standards in the web services area, the most important of which is WSDL¹⁴ (Web Service Description Language), from IBM, Microsoft, and Ariba, currently a draft W3C (World Wide Web consortium) specification. WSDL enables the definition of service descriptions through messages that service requesters pass to service providers and vice versa.

Other important web services-related standards include the W3C standard SOAP¹⁵ (Simple Object Access Protocol), for describing data objects, and the OASIS standard UDDI¹⁶ (Universal Description, Discovery and Integration), a directory service that enables listing and finding web services. There are many other web services related standards in various stages of development; these are beyond the scope of this white paper.

User Experience

Above all, a DRM scheme that is suitable for integration with P2P networks has to preserve a seamless user experience. In addition to providing for reasonably expected usage rights, such as time and space shifting, as mentioned above, the following are aspects of DRM that contribute to user experience:

- Installation of the DRM has to be seamless, including the initial installation of the software as well as maintenance. Ideally, the user should not do or even notice anything about the installation. This should be true for all platforms. Java, XML, and other cross-platform technologies, such as are used in DigitalContainers' Hybrid P2P architecture, should help achieve this.
- Payment processing should be integrated with ISPs and other service providers, so that users don't have their experiences disrupted by requests for payment information. Universal or interoperable ID schemes will go a long way towards facilitating this.
- The DRM should track content usage but do so in a way that respects privacy. This is a well-known problem – tracking software is often referred to as “spyware” – and solving it is largely the responsibility of service providers that process usage information. Service providers need to take steps to give users confidence that tracking information is not being abused.

¹⁴ See <http://www.w3.org/2002/ws/desc/>.

¹⁵ See <http://www.w3.org/TR/SOAP/>.

¹⁶ See <http://www.oasis-open.org/committees/uddi-spec/tcspecs.shtml>.

Gaps in Existing DRM Technology

DRM is complex technology, and in this early phase of its development, designers have chosen to focus on features of most immediate concern to their customers, mainly media companies, who have mainly been focusing on piracy prevention and simple distribution schemes that emulate physical media distribution models. Another concern at odds with the complexity of DRM is the cost of deployment, particularly any unit costs of bundling DRM functionality in with platforms and consumer devices.

Meanwhile, P2P networks have only recently come into being. Therefore, many DRM technologies in existence today have various gaps in their ability to be integrated into P2P networks. Here are some of the most important of those gaps, which should represent opportunities for DRM technology designers in the future.

Cost-Related Functionality Limitations

DRM schemes designed by consumer device makers typically have just enough functionality to satisfy IP owners while keeping unit costs at a minimum. The best-known example of this is the CSS (Content Scrambling System) for DVDs, which was designed by two consumer electronics makers (Toshiba and Matsushita) and accepted by movie studios with the promise of a stronger future solution that has yet to materialize.

Cost of DRM has also been an issue beyond the world of consumer devices. In general, DRM is virtually unique in the technology world, in that it is a complex technology that encumbers the user without any direct benefit (unlike, say, burglar alarms, which protect people from theft of their own physical assets); the biggest challenge in the market has been to find those participants in the content value chain who would be willing to pay for it. (The media industry in particular has been reluctant to make investments in DRM technology compared to, say, the software industry.) This problem should recede over time as DRM becomes more and more bundled into value-added services that have tangible benefits for users.

The problem of the cost of DRM should recede over time as DRM becomes more and more bundled into value-added services that have tangible benefits for users.

Device Tethering

Many DRM schemes permit access to content only on a specific device, instead of supporting space shifting and other reasonable usage expectations. For device and platform vendors, the reasons for this are obvious: why support usage of content on competitors' platforms? Some DRM schemes allow usage on up to a fixed number of devices or software of the same type but not of different types.

For IP owners, the reasons for supporting device tethering derive from the media industry's traditional product orientation: the principle that two different formats of the same content – for example, the DVD and VHS versions of the same movie, or the print and eBook versions of a book – are separate products and should be paid for separately. Most corporate IP owners, which use content for purposes of knowledge management, marketing, collaboration, etc., would not agree with this.

IP owners also feel that device tethering is sometimes necessary to curb infringement; for example, if a college textbook is published in eBook format using a DRM technology that allows reading on up to 10 eBook readers, then a class of 20 students is likely to collectively purchase as few as 2 copies. Admittedly, device tethering is a legitimate response to the imperfection of reasonably-priced user authentication, such as passwords that can be shared as opposed to more expensive but more effective biometric authentication devices.

Lack of Superdistribution Support

Most DRM schemes only support single levels of distribution, or they support the limited form of Superdistribution discussed above. As early DRM vendors found out, support for true Superdistribution requires far more complex technology than that required for single-tier distribution. Yet web services, cross-platform functionality, and other technologies that have appeared since the mid-1990s can ameliorate this problem. Standards for rights and web service descriptions will especially help remove the complexity of Superdistribution with DRM.

Complexity of Integration

Even with single-tier distribution schemes, a serious barrier to growth in the DRM market has been how expensive, time-consuming, and complex it is to integrate DRM technology with all of the necessary surrounding functions, including: content production and packaging, user identification, transaction processing, and CRM (customer relationship management). Launching a new content e-commerce initiative has been so complex that integration costs dwarf that of off-the-shelf software, including DRM packaging software. Furthermore, the capital outlay required even with off-the-shelf software is prohibitive for smaller IP producers, including many who might be interested in making content available over P2P networks.

A serious barrier to growth in the DRM market has been how expensive, time-consuming, and complex it is to integrate DRM technology with all of the necessary surrounding functions.

Once again, replacing licensed software with services, and standardizing the interfaces to those services so as to minimize integration effort, will help solve this problem. There have been many attempts to build DRM service provider businesses; most have failed because the kinds of prices that IP owners have been willing to pay for services did not measure up to the service providers' high cost structures. But the success of a handful of current DRM-related service providers in niche markets points the way to a brighter future for service-based architectures.

Conclusions: Developing the Market

We conclude this white paper with some thoughts on how DRM can grow to support integration with P2P networks. Some of the problems that must be solved are technological in nature; others are problems of perception rather than reality; but most are more matters of economics than anything else.

Of the technological problems, the largest one is Superdistribution. Early DRM technologies such as IBM's Cryptolope attempted to support Superdistribution but failed because of all of the functionality that needed to be built from scratch, on both the server and client platforms, to support it. Nowadays at least some of the required functionality (e.g., network authentication and e-commerce) is standard and widely available, and such functionality is becoming available through standard web-service interfaces. But it is still a daunting technical challenge to implement Superdistribution without undue complexity and disruption of user experiences – to say nothing of prohibitive cost. DigitalContainers is a DRM technology vendor that is addressing these challenges today.

Network identity is a problem of both perception and technology. Universal network identification schemes like Microsoft's .NET Passport have a "Big Brother" perception problem that may be exaggerated. The same is true for user tracking technology vis-à-vis privacy. On the other hand, federated (interoperable) network identity, à la the Liberty Alliance, is seriously difficult to implement in today's heterogeneous trust environments.

Technology problems related to meeting reasonable usage expectations, such as device and format portability and rights specification interoperability, derive largely from economic considerations. As mentioned above, one of the biggest challenges in the development of DRM has been getting participants in the content value chain to pay for it.

The two types of participants most closely involved in designing DRM schemes have been platform/device vendors and IP owners. IP owners, as mentioned above, have long thought in terms of "products" instead of "content," leading them to feel that purchasing content in one format should not allow the purchaser to access that content in other formats. And device vendors are not at all motivated to create DRM technology that allows users to access content on other types of devices.

Aside from advocacy groups like the Electronic Frontier Foundation and DigitalConsumer.org, who attempt to assert content usage rights through lobbying of legislators, case support, and other such activities, third-party DRM vendors are the ones that are actually motivated to build technology that supports users' reasonable usage expectations as well as other features that promote the integration of DRM with P2P networks.

Third-party DRM vendors are the ones that are actually motivated to build technology that supports users' reasonable usage expectations as well as other features that promote the integration of DRM with P2P networks.

There have been countless "standalone" third-party DRM technology vendors over the past several years, but very few of them have succeeded, due to several factors, including unrealistic revenue expectations, lack of understanding of content business models, and of course, inadequate technology. To understand how these vendors might find

customers, we should answer the question: who stands to gain from the proliferation of DRM-enabled P2P networks?

The answer lies in the fact that P2P network usage promotes use of network bandwidth, equipment, and services; in fact, numerous recent statistics have shown that the majority of bandwidth on the Internet is used by a small percentage of users who mostly engage in file sharing. Therefore, we suggest that network hardware/software makers and internet service providers (ISPs) are the best potential sources of interest in and funding of DRM development for P2P networks.

Network hardware and software makers' interest in embracing DRM is hampered somewhat by the open nature of the Internet and the W3C's lack of interest in DRM, but network equipment makers have been looking at DRM, though they have yet to become active in the market. For example, Cisco designed a DRM protocol called OCCAM (Open Conditional Content Access Management) in 2001, but the company appears to have no interest in developing products around it¹⁷.

The major Internet service providers have largely avoided DRM; one reason for this is that noninvolvement in DRM has enabled them to stay aloof from various legal liability issues. However, recent activity related to the Digital Millennium Copyright Act, such as the music industry's subpoena of Verizon over the name of a Verizon Online user suspected of music piracy, may be sending a signal to ISPs that noninvolvement breeds liability too, therefore they should participate in the market and start looking at the service provider opportunities it can afford. Major ISPs are naturals to support DRM-enabled P2P networks and provide value-added services to their participants, as a way of garnering revenue from their heaviest users instead of (or in addition to) monitoring bandwidth usage and charging tiered pricing, as a few have begun to do over strong user objections.

Major ISPs are naturals to support DRM-enabled P2P networks and provide value-added services to their participants, as a way of garnering revenue from their heaviest users.

Of course, the development of some of the technologies mentioned in this white paper should also encourage startups, as well as more established vendors, to build various types of new content-related services that can grow the market.

Finally, we should emphasize that standards can help hasten and lower the cost of solutions to many of the problems mentioned above. There are several existing standards efforts that related to DRM integrated with P2P, as previously mentioned. The problem with many of them is that they are design with much broader areas in mind than that of DRM and P2P networks, meaning that approval processes take longer and applicability is not as straightforward. P2P-related trade associations (e.g., P2P United) are beginning to appear; unfortunately, they engage in anti-DRM posturing for political purposes. Such groups need to get beyond polemics, understand the opportunities for everyone available in DRM integration with P2P networks, and start representing the P2P community in relevant standards initiatives alongside the independent DRM vendors that can most directly impact the market.

¹⁷ The OCCAM white paper, no longer available on Cisco's web site, is available by request from the author.

About the Author

Bill Rosenblatt, president of GiantSteps Media Technology Strategies, has 20 years of experience in technology architecture, business development, and marketing; publishing; new media; and online education. He has been a business development executive at a leading technology vendor, an IT executive at major publishing companies, and chief technology officer of an e-learning startup. He has expertise in digital media technologies such as content management, digital rights management, streaming media, and publishing systems. Bill is the author of several books, including *Digital Rights Management: Business and Technology* (John Wiley & Sons, 2001), and he is Managing Editor of the Jupitermedia newsletter DRM Watch (www.drmwatch.com).

About GiantSteps Media Technology Strategies



GiantSteps Media Technology Strategies is a management consultancy focused on the content industries that help its clients achieve growth through market intelligence and expertise in business strategy and technology architecture. Clients have included publishing companies, news, entertainment, and professional information providers, and digital media technology vendors ranging from early-stage startups to Global 500 firms.

Contact:

phone: +1 212 956 1045

email: info@giantstepsmts.com

Web: www.giantstepsmts.com

White paper commissioned by DigitalContainers LLC

About DigitalContainers LLC



DigitalContainers provides patented digital rights management for use in peer-to-peer networks and the Internet. Digital Containers include self-contained file protection, authentication and e-commerce system,

allowing files and media to travel around the Internet, yet perpetually be tracked, controlled and audited by the content owners. This enables content owners to securely monetize their digital goods in peer-to-peer networks.

Contact:

phone: +1 703 208 1040

email: info@digitalcontainers.com

Web: www.digitalcontainers.com