Content Identification Technologies

Business Benefits for Content Owners

By Bill Rosenblatt

April 15, 2008

Table of Contents

Introduction and Executive Summary

Content owners, including film studios, cable and broadcast programmers, television networks, and music companies, have been building out various business models for their content in the emerging world of digital distribution. In this world, content owners find themselves less and less in control of downstream entities – including content packagers, distributors, and consumption devices – than they have been in traditional value chains such as terrestrial broadcast, CD/DVD distribution, and theatrical motion picture release.

Certain enabling technologies are needed to make connections among online content retailers, content delivery networks, advertising and media buying agencies, consumer device manufacturers, anti-piracy service providers, audience measurement agencies, user-generated content sites, social networks, and so on. Among the most critical of these enabling technologies are those known collectively as *content identification technologies*: specifically *watermarking*, *fingerprinting*, and *content identifier standards*.

This white paper provides contextual background on these technologies and discusses several specific types of business benefits they can bring to content owners. These include: deterring piracy, increasing online ad revenue, tracking content usage, managing media assets, integrating content owners' internal systems, and monetizing so-called transformational uses of content. We show several examples of how watermarking or fingerprinting can combine synergistically with content identification standards to enhance business benefits.

Although the focus of this white paper is on content owners, it is clear that content owners cannot bring about the business benefits described here without cooperation from various types of businesses. New types of service providers are needed; these include *content identifier registration agencies* and *network monitoring service providers*.

In the last section of the paper, we review third-party incentives that are required to create a viable market for content identification technology providers. These include revenue opportunities for new types of service provider businesses, ways of reducing costs to participate in the market (or spreading those costs more equitably), and adoption of standards where they make sense.

This paper is sponsored by Sun Microsystems' Communications and Media Industry Practice. The impetus to research this topic arose from meetings of Sun's executive Media Advisory Board (MAB) during 2007. Led by Sun and leading media industry executives, MAB is focused on identifying practical solutions to shared business challenges of cable and TV programmers, broadcast networks, film studios, and other major content owners. The paper was first presented to participants of a Sun MAB roundtable on content identification held at the 2008 NAB conference in Las Vegas.

In making this work freely available, Sun MAB's intent is to foster industry dialog on certain technologies and practices that have potential to accelerate realization of digital distribution business benefits, particularly as they pertain to video content. The paper is available on www.sun.com and at www.giantstepsmts.com.

Content Identification Technology Backgrounder

The need to implement a wide range of business models for online content has led to the development of three key enabling technologies: watermarking, fingerprinting, and content identifier standards, collectively called content identification technologies. In this section, we describe each of these, provide brief histories, and summarize the current state of the market.

Watermarking

Digital watermarking, or *watermarking* for short, is one of two techniques for examining a file in order to determine its identity. It involves modifying the "noise" portion of the content in a file so that it contains some data, called the *payload*, in such a way that the user's perception of the content is not impaired[1]. Digital watermarking works with still images, audio, and video content[2].

There are two principal operations in watermarking: *insertion* (also called *embedding*) of the payload into the file, and *detection* of the watermark by special hardware or software that examines the file, searches for the watermark, and extracts the payload. Although the most typical watermark payload is information about the identity of the content or its owner, watermarks can contain any data at all – subject to limitations on payload sizes, which depend on various factors but are generally limited to a few dozen bytes.

Watermark designs incorporate tradeoffs among several criteria, including:

- **Capacity**: the maximum payload size, usually a few dozen bytes.

- **Robustness**: the ability of the watermark to survive various operations on the content, such as excerpting, cropping, compression, format conversion, or analog conversion.

- **Imperceptibility**: the lack of effect on the user's perception of the content when it is played or shown.

- **Security**: imperviousness to removal without affecting the content's appearance or sound.

- **Efficiency**: the speed or ease with which the watermark can be inserted and detected.

As we will see in the next section, it is becoming more popular to embed data other than content ownership information as watermarks, such as the identity of a user or device that downloads the content. It is also possible (within limits) to insert more than one watermark in a single file.

---

[1] There are also applications for perceptible (visible or audible) watermarks, such as on sample content from stock agencies, but we don't discuss those here.
[2] Watermarking is a special case of steganography, or "covered writing." An example of text steganography is in the Arthur Conan Doyle story *The "Gloria Scott,"* in which Sherlock Holmes decodes a secret message by looking at every third word in an innocuous-looking text.

Digital watermarking techniques have been in existence since the mid to late 1990s[3]. The first commercial solutions applied to digital still images, and watermark insertion tools come standard with many popular image processing software applications.

Watermarking was sold to the media industry as an anti-piracy tool for digital images during the first Internet bubble, but it wasn't very effective, and some content owners had concerns over losses in perceptual quality. In the ensuing years, the technology has spread to audio and then video, and a number of more sophisticated applications and business models have been developed that have revitalized the watermarking field. One sign of its new vitality is the establishment of an advocacy organization, the Digital Watermarking Alliance (DWA), in 2006.

Fingerprinting

The term "fingerprinting" has, unfortunately, had a number of overlapping and confusing meanings. It has been used to describe a watermarking technique in which when a user downloads a file, her identity (or that of her device) is inserted into the file as a "fingerprint." This is now generally called transactional watermarking (see p. 11).

Meanwhile, "fingerprinting" has come to mean examining the bits of a file and intelligently determining the identity of the content. That is, fingerprinting now means "taking the file's fingerprints" rather than "putting a fingerprint in the file." This is the meaning we use here.

The basic idea of fingerprinting is to examine a file, compute its fingerprint as a set of numbers, and look them up in a database of fingerprint values to determine the identity of the content. Fingerprinting can be used with audio, video, and (with somewhat different techniques) text.

A fingerprint is a special type of a mathematical construct called a *hash.* A hash is a kind of mathematical abbreviation or shorthand for a large amount of data, such that two files containing different data are virtually guaranteed to yield different hashes. Hashes are often used to check the integrity of data: a large amount of data is sent along with its hash value, which is recomputed on the receiving end; if the recomputed hash does not match the hash supplied with the data, then the data must have been altered in transit.

Yet standard hashing algorithms are not very useful for identifying content files: differences between two files can be perceptually trivial, but they will yield different hash values. For example, if you take a file containing a music track, make a copy, and change just one bit, then the two files will probably sound identical, but they will compute completely different hash values.

Fingerprinting solves this problem by providing a hashing scheme that is impervious to certain types of changes to content. In particular, all files that "sound" or "look" like the same content should yield the same fingerprint, while files containing different content should yield different fingerprints. Fingerprinting is thus also known as *robust hashing* or *perceptual hashing*. Analogously to watermarking, good fingerprinting technologies can account for operations on content such as cropping, rotation, color-space shifting, audio equalization, format conversion, and so on.

Fingerprinting is newer technology than watermarking. It is more complex, relying on heuristic and psychoacoustic or psychovisual techniques on top of the basic signal-

---

[3] For example, the landmark U.S. patent no. 5,745,604 to Geoffrey Rhoads of Digimarc Corp. was filed in 1996 and issued in 1998.

processing techniques that are also used in watermarking. Fingerprint computation therefore generally requires more computing resources than watermarking.

The first serious explorations of fingerprinting techniques for music – so-called acoustic fingerprinting – took place after a federal court shut down the original Napster file-sharing network in 2001 and the company sought ways of going "legit" and respecting copyright. Yet the music industry did not embrace fingerprinting until a few years later.

At this writing, several fingerprinting techniques are emerging for video and are being tested on film and television content.

Watermarking vs. Fingerprinting

Watermarking and fingerprinting have overlaps in the kinds of applications they support – as we will see in the remainder of this paper – but they are generally viewed as complementary technologies. Table 1 summarizes the most important differences and tradeoffs between them.

	Watermarking	Fingerprinting
Changes to Content	Watermark must be inserted	None
Process	Insert watermark in every file on server and/or consumer device; detect later	Compute fingerprint once for each content item and deposit in vendor's master database; re-compute later for lookup
Identification Accuracy	100% accurate by definition	Less than completely accurate
Data Flexibility	Can store any data, up to capacity limitations; files with identical content can have different watermarks	Cannot store any information; identical content files compute identical fingerprints
Costs of Implementation	Spread fairly evenly throughout the value chain	Primarily fall on network service providers
Device Support	Insertion and detection possible on some consumer devices as well as on servers	Computation can only currently be done on servers

• Table 1: Key differences between fingerprints and watermarks.

These differences have important implications for how the cost of implementing the two technologies is apportioned among different types of entities in the content value chain. We will explore this issue later on; see p. 25.

In the next section, we will see how these two technologies apply to the various parts of the content value chain under different scenarios.

Content Identification vs. DRM

Many of the business benefits of content identification technologies illustrated in this paper have also been ascribed to digital rights management (DRM) technologies – particularly in the areas of anti-piracy and usage tracking. Although the term DRM has had a range of meanings over the years, its narrower definition appears to have become generally accepted: technology that relies on keeping digital content in an encrypted state until it is safely ensconced in a secure end-user environment, such as a consumer device or software application.

In terms of its age, DRM technology is roughly contemporaneous with watermarking[4]. Both were offered as piracy remedies during the first Internet bubble of the late 1990s; large media companies viewed DRM as more promising.

There is a key conceptual difference between DRM and watermarking or fingerprinting: the former is normally proactive while the latter are reactive. DRM technologies (like Microsoft's Windows Media DRM and Apple's FairPlay) assume that users aren't allowed to do anything to content without explicit permission. Conversely, content identification technologies assume that users can do anything with content unless they are specifically forbidden (as in filtering; see p.12) or caught forensically later.

With regard to anti-piracy applications, legal experts sometimes say that content identification technologies mirror copyright law more closely than DRM technologies. For example, copyright law in most European Union countries[5] gives people rights to make copies of legitimately obtained content for private purposes, though some types of copies are not legal (such as those made for resale without authorization). Content identification technologies facilitate catching such exceptions rather than only allowing uses that are known in advance to be legal. On the other hand, DRM technologies can prevent illegal uses before damage has been done, whereas content identification technologies (by themselves) cannot.

There have been a few examples of integration of DRM and watermarking technologies. The first well-known one was the Secure Digital Music Initiative (SDMI) standards initiative, which began in 1999 as a reaction to the Napster peer-to-peer phenomenon but lost momentum in 2001. Another example is Blu-ray discs, which use watermarking as a backup to the AACS encryption-based DRM scheme.

In the future, video content owners will likely view DRM and content identification as complementary technologies whose synergies have largely yet to be explored – in contrast to the music industry, which is now turning away from DRM while increasing its interest in content identification technologies (see the watermarking example on p. 12).

Content Identifier Standards

The final technology related to content identification that we explore in this paper is *content identifier standards*. How should content be identified? On the surface, this seems like a simple enough question, but the subtle complexities of content identification have led to a proliferation of standards.

---

[4] For example, Stefik and Casey's 5,629,980 patent at Xerox PARC was filed in 1994 and issued in 1997; Ginter et al's 5,892,900 patent, the first "DRM" patent from the company now known as Intertrust Technologies, was filed in 1996 and issued in 1999.

[5] With the notable exception of the UK.

Identifier standards for online/digital content have their roots in legacy identifiers for offline/analog content, including standards such as ISBN (International Standard Book Numbers) and proprietary identifiers such as the old Schwann numbers for music albums. Digital/online content identifiers can (and need to) be more flexible than their legacy progenitors, but with that flexibility comes ambiguities, such as:

- **What to identify**: content items can be identified as abstract pieces of intellectual property, as products (analogous to SKU numbers in retail), as individual units (analogous to serial numbers), or as services (a content item bundled with some representation of rights granted to the user).

- **Level of granularity**: this is straightforward for some types of content, such as music tracks or scientific journal articles, but more ambiguous for other types, such as feature films, television news programs, or books.

- **Level of uniqueness**: unique within an enterprise or workflow, or globally unique.

Identifiers for digital content can have a number of interesting properties[6], including global uniqueness, persistence (identifiers last forever), location independence (unlike a web URL), and interoperability with legacy identifiers or with media companies' internal content identifier schemes. Some properties trade off against one another: for example, a requirement that an identifier be used as a watermark payload trades off against its length, and thus against the size of the universe of possible identifiers; ascribing semantics to certain fields or components of an identifier trades off against its flexibility.

Identifier Standard	Registration Authority	Content Types	Identifies What?
ISAN: International Standard Audiovisual Number	ISAN International Agency	Audiovisual	Intellectual property
UMID: Unique Material ID	SMPTE	Audiovisual	Products
ISWC: International Standard Musical Work Code	CISAC	Music	Intellectual property
GRid: Global Release ID	IFPI	Music	Products
DOI: Digital Object Identifier	International DOI Foundation	Any[7]	Anything
PURL: Persistent URL	OCLC	Any[8]	Anything

- Table 2: A sampling of content identifier standards.

---

[6] For more on this, see Rosenblatt, B. The Digital Object Identifier: Solving the Dilemma of Copyright Protection Online. In *Journal of Electronic Publishing*, University of Michigan Press, , Vol. 3, No. 2, December 1997, available at http://quod.lib.umich.edu/cgi/t/text/text-idx?c=jep;view=text;rgn=main;idno=3336451.0003.204.
[7] DOIs can apply to any type of content, but their most widespread use has been in scientific journal publishing.
[8] PURLs are most typically applied to library references.

As a result, many different content identifier schemes have been introduced, typically along the lines of media industry segments or anticipated applications. Table 2 lists a representative sampling of these.

Virtually all content identifier standards have *registration authorities* or *registration agencies* associated with them. These are entities that are responsible for issuing unique IDs that conform to the given standard and for tracking their ownership (which may change over time as media properties are bought and sold). They also typically maintain databases, accessible online, that contain information about the item to which an identifier has been assigned, such as bibliographic metadata. Such entities have long existed for offline identifier standards, such as ISBN agencies in different countries: R.R. Bowker in the United States, Nielsen BookData in the UK, AFNIL in France, etc.

Some identifier standards are set up so that multiple registration authorities can exist and offer competitive features and services; ISAN, ISWC, and DOI are examples of this. In such cases, the entity listed in Table 2 under Registration Authority is actually a governing body that certifies third-party registration authorities and manages them over a distributed architecture. Later in this paper we will discuss the roles that registration authorities can play in realizing the business benefits of content identification.

Business Benefits of Content Identification

In this section, we examine five categories of business benefits to content owners that arise from adoption of content identification technologies. For each one of these, we describe the benefits and how content identification technologies are being (or can be) used to achieve them.

The common theme among all of these business benefits is the integration of content identifier standards with watermarking, fingerprinting, or both. Each of these technologies is useful in isolation, but only up to a point; as we will show, the combination of watermarking or fingerprinting with content identifier standards is powerful and synergistic in enabling business benefits.

Figure 1 shows the basic media value chain, to which we will allude throughout this section.
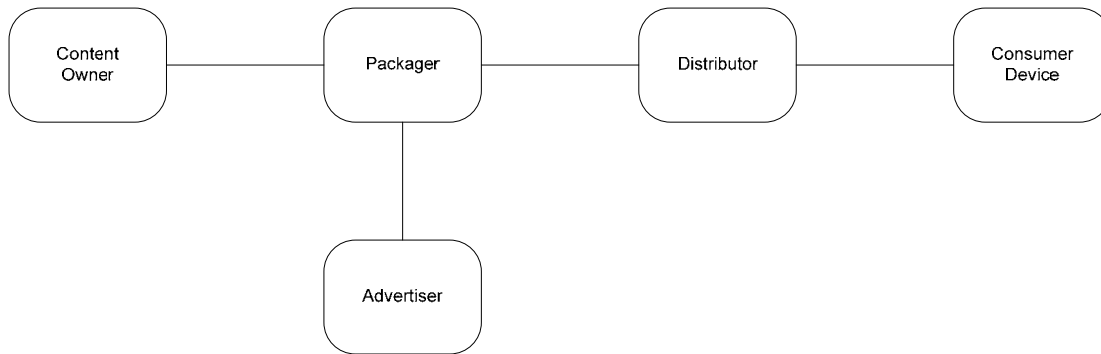
- Figure 1: The basic media value chain.

Figure 1 shows the following value chain participants:

- **Content Owners**: suppliers of content, such as film studios, record labels, and television programming companies.

- **Packagers**: businesses that package content for consumption by users and their devices. These can include online content retailers, websites, cable TV networks, etc.

- **Advertisers**: for these purposes, we really mean agencies that create ads and buy media for advertisers. Agencies provide ads to packagers, which insert them into content (as with TV programs) or juxtapose them alongside content (as with contextual ads on web pages).

- **Distributors**: network operators that take content and distribute it to consumers, including broadcasting networks, network service providers, and wireless carriers.

- **Consumer Devices**: these can include PCs, set-top boxes, mobile handsets, and other devices.

Of course, there are many examples of businesses that combine more than one of these types, such as the major television broadcast networks, which often function as content owners, packagers, and distributors all in one. But in the world of digital distribution, it is more likely for each of these to be separate entities.

Deterring Piracy

The first business benefit we discuss is the most commonly associated with content identification technologies and generally the first application envisioned for them: deterrence of content misuse.

Watermarking

As mentioned above, digital watermarking appeared as a purported cure for Internet content piracy during the first Internet bubble of the late 1990s. For example, Digimarc's MarcSpider crawled the Internet to find digital images with watermarks in them and determined which ones weren't where they were authorized to be. The technology was not very effective for that purpose (although it survives and is widely used for other purposes today), and a backlash resulted against watermarking as an antipiracy tool.

The demise of SDMI in 2001 (see p.7) represented the end of the first wave of watermarking applications for piracy deterrence. More recently, piracy deterrence applications of watermarking have become more sophisticated, leading to a resurgence of interest in the technology.

One of the most interesting newer watermarking applications is known as *transactional watermarking*, in which the user's device inserts its own identity (or that of the user) into the file as a watermark payload. A variation on this is *media serialization*, in which the payload is simply a serial number that is used to access the same information stored in a database on a server. Either of these techniques marks a downloaded file with the identity of the downloading user or device, so that if the file is found in an unauthorized place, such as on a file-sharing network, it can be traced to the downloader. Working examples of this technique for Internet content include Bitmunk's modified peer-to-peer network and Fraunhofer's LWDRM (Light Weight DRM), both of which embed user IDs as watermarks in content files.

Transactional watermarking requires that the client device have the functionality to insert watermarks and that this functionality cannot easily be disabled. This is not hard on a PC, where it can be done in software and the main concern is to guard against substitution hacks (i.e., swapping in client software that does not insert the watermark or inserts a false one). But for a consumer device, it requires that watermark insertion functionality be present on the device, which can add to its cost. Cinea, a division of Dolby Labs, recently developed a technology called Running Marks, which supports transactional watermarking on consumer devices, such as set-top boxes, with no added hardware costs.

Transactional watermarking has potential for piracy deterrence but comes with privacy as well as cost concerns. As a practical matter, it is currently limited to applications where content is distributed to a relatively small number of clients. For example, the vendor Activated Content provides a transactional watermarking-based solution for distribution of pre-release music to radio stations, music critics, and so on, so that if any files are leaked to P-to-P networks, the record company will know who leaked them. Cinea has created a watermarking-based technology for distributing Oscar screeners, which are specially-encoded DVDs of movies being considered for Academy Awards.

Now that the major music companies are in the process of abandoning encryption-based DRM for music downloads in favor of unencrypted MP3s, they are beginning to use watermarks to insert the identity of the retailer through which a track is purchased, such as Amazon or Wal-Mart. They can detect those watermarks in files that appear on P-to-P file-sharing sites or other unauthorized locations to get information about which retailers' users (if any) tend more towards piracy.
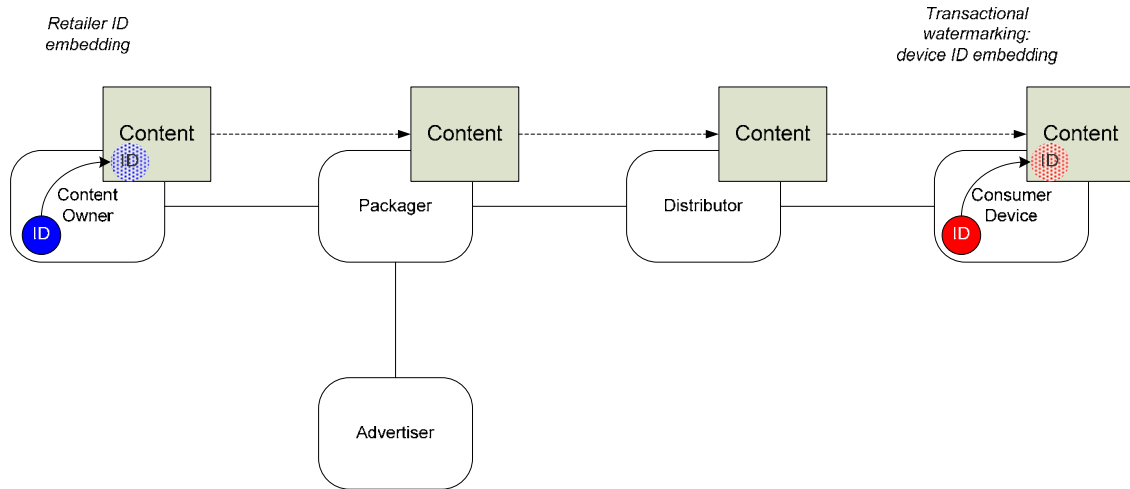
- Figure 2: Watermarking-based piracy deterrence techniques.

Figure 2 shows both of these watermarking-based antipiracy approaches. The hashed "ID" circle on the left shows the content owner embedding an identifier denoting the packager (retailer) before sending it the content; the spotted circle on the right denotes transactional watermarking.

Fingerprinting

Acoustic fingerprinting, meanwhile, emerged as an antipiracy tool after the shutdown of the original Napster file-sharing network in 2001, as mentioned above. In 2005, the major music companies licensed the iMesh file-sharing network to use acoustic fingerprinting technology from Audible Magic to block uploads of major-label content to the network. Instead, users who want major-label music on iMesh can purchase DRM-encrypted versions of it that are stored on its own servers. This use of fingerprinting is widely known as *filtering*, although watermarking can be used for filtering as well.

Since then, the use of acoustic fingerprinting to filter uploads of copyrighted music (and music videos) from P2P networks and websites has become fairly common; most major social networking sites use it, including YouTube, MySpace, and DailyMotion. Audible Magic and Gracenote are the leading vendors of the technology.

More recently, video fingerprinting technologies have emerged to address the increase in uploads of clips from movies, TV shows, and other copyrighted video material. Google is developing video fingerprinting technology for YouTube in house, to be used in conjunction with Audible Magic's acoustic fingerprinting. Audible Magic, meanwhile, is marketing a video fingerprinting technology called Motional Media ID, which is separate from its music fingerprinting technology; Philips, Thomson, and several startup vendors are also entering this emerging market.

The next step in the use of fingerprinting as a piracy deterrent is to use it at the ISP level instead of at the website level. AT&T, at the behest of Viacom and other media companies, has agreed to start testing video fingerprinting on its Internet service network; it has chosen the startup Vobile as its fingerprinting vendor. If the technology works, then AT&T will be able to filter out unauthorized copyrighted material whenever an ISP subscriber sends it over the network.

France will soon be a venue for fingerprinting – and possibly watermarking – at the ISP level, as mandated by law. The so-called Olivennes Agreement, which was announced in November 2007, will require French ISPs to implement content identification technologies to issue warnings to users who download unauthorized content, followed by ISP account terminations for repeat offenders.

ISP-level content identification is a bold step forward; some experimentation will be necessary to determine how successful it can be. Some of the unknowns include impact on network performance and the ability to distinguish between legitimate and illegitimate uses of content. As an example of the latter, a user may send a file containing a copyrighted work in an email message to a friend or relative in such a way that might be considered legal under such doctrines as Fair Use in the United States or Private Copying in many EU countries.

The use of fingerprinting or watermarking together with content identifier standards could provide significant synergies in the industry's ability to use the technologies to deter piracy. An important example of this is in so-called "takedown notices".

Currently, in the United States a content owner must ask a network service, such as an ISP or social networking site, to remove content that it owns and has not authorized. It must locate the content on the service and then send the service a message called a takedown notice in order to get the material removed. The service is legally obligated to comply[9].

Although U.S. law specifies the kinds of information required in takedown notices, each network service operator has its own ways of accepting them, so content owners have to submit different takedown notices to different operators. With new network services starting up all the time, this all adds up to quite a burden on content owners – as Viacom is claiming in its lawsuit against Google over copyrighted material on YouTube.

One way to make this process more scalable is to require that standards be defined for the formats of takedown messages. Then content owners could send out a single message to multiple service providers. At one level, this could mean simply creating a standard format, such as an XML DTD, that contains a tag for each type of information required in a takedown notice. Site operators would need to accept messages in that format. But that would still require content owners to create separate messages for each operator with the names or locations of files that they want taken down.

With standards for content identifiers in place, content owners could send the same message everywhere for each given content item. This is shown in Figure 3. Site operators would need to establish the identifiers for each file uploaded to them – by fingerprint computation or watermark extraction – and store them in a database that maps content ID to location on the site. Then when they receive a takedown notice, they can use the database to find the file and remove it.

---

[9] This is provided for under Section 512 of the Digital Millennium Copyright Act (DMCA). Accordingly, lawyers sometimes call such notices to site operators "512 notices."
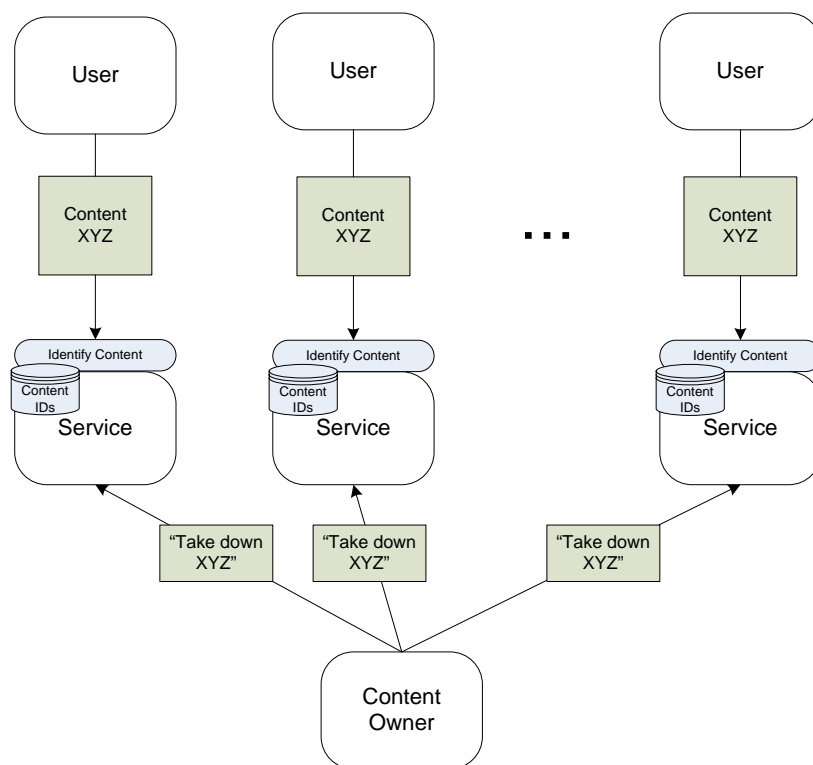
- Figure 3: Using content identification and standard content identifiers to make issuance of DMCA 512 takedown notices more efficient.

Increasing Internet Ad Revenue

The ubiquity of free content on the Internet has called into question many business models based on consumers paying directly for content. At the same time, the Internet is dramatically enhancing the ability of businesses to get marketing exposure by associating themselves with content or activities that appeal to a well-defined target audience. There are many Internet technologies that help advertisers determine which users might be interested in their offerings. As wireless networks become more open and pervasive, these technologies should apply there as well.

An increasingly important subset of this trend is *contextual advertising*: the ability to target an audience by its affinity to content. The best example of this is Google's wildly successful AdSense contextual advertising technology, which was enabled through Google's 2003 acquisition of Applied Semantics. Applied Semantics' computer-aided indexing technology analyzes text and assigns keywords based on its meaning. Google offers ad placements based on those keywords even if they aren't explicitly present on the web page where the ad appears.

The next logical extension of contextual advertising technology is beyond text to music and video. Content identification technologies can play a vital role in building out such capabilities, and real-world examples are starting to appear.
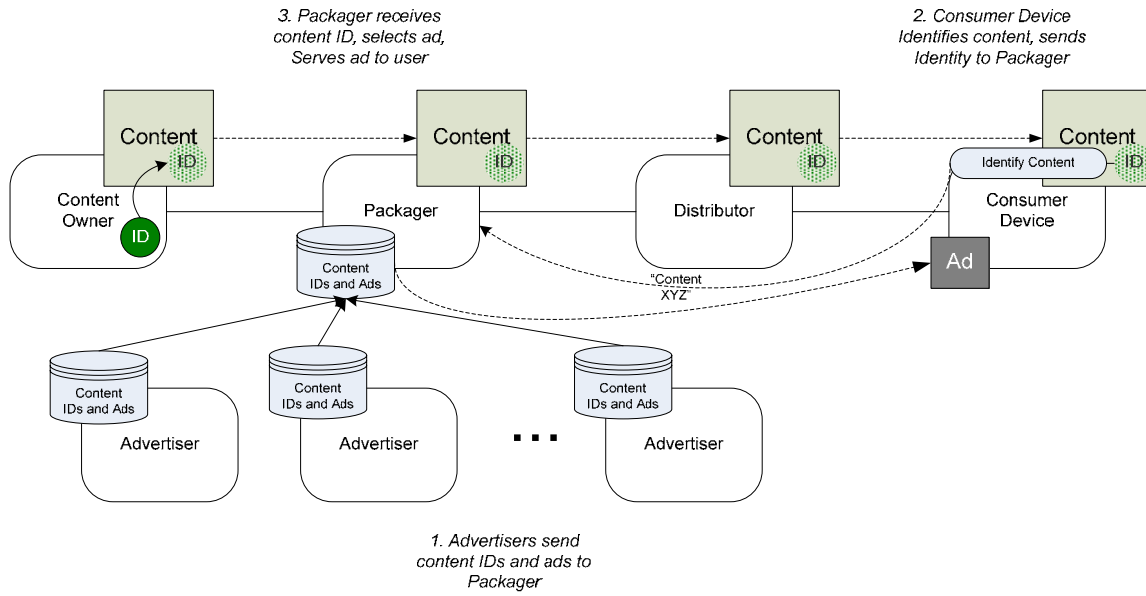
One of the most interesting examples of contextual advertising applied to non-text content is imeem's music uploads. Imeem, a social networking site focused on content sharing, lets users upload music, videos, or photos to their personal profiles or to group pages. When a user uploads a music track to imeem, acoustic fingerprinting technology (from Gracenote) determines the music's identity and looks it up. If the rights holder has granted rights to the track, then other users can stream the track from the user's profile and be shown a contextual ad. The record label shares in the ad revenue. Warner Music Group was the first major music company to license its content to imeem in this way; the other three majors followed.

Yet the imeem example is somewhat limiting, because it involves a service that does not know in advance the identity of content being offered. Other types of services, such as streaming web radio, know what content they are playing so they don't need content identification technologies.

The use of content identification technologies for contextual advertising makes even more sense when used with PCs or consumer devices, where the user chooses music to listen to or video to watch from her own collection. A retailer could provide an application that detects the content that a user selects and inserts targeted ads in exchange for discounts, free content, bonus content, etc. For the foreseeable future, this could only work with watermarks, because the processing power required to compute fingerprints is too great for most consumer devices.

Content identifier standards can enhance these schemes by providing a common, interoperable way of associating content with ads. With standards in place, ad-delivery services can maintain databases of standard identifiers associated with ads to serve alongside those content items. When the application installed on the consumer's device detects a watermark, it can retrieve the content ID[10], send it to an ad-delivery service along with the device's IP address, and the appropriate ad will be served. Figure 4 shows this.

---

[10] This could be done directly from the watermark if the content ID is small enough to fit the watermark's payload capacity. Otherwise, the payload could be a serial number or hash value, which would be looked up on a server that would return the content ID.

3. Packager receives
content ID, selects ad,
Serves ad to user

2. Consumer Device
Identifies content, sends
Identity to Packager

1. Advertisers send
content IDs and ads to
Packager

• Figure 4: Using content identification on consumer devices to serve targeted ads to the user.

In this scheme, a media buying agency can create a list of content items to be associated with an ad campaign. It can then deliver the same list to multiple ad delivery services or other service providers, along with the ads themselves, so that the service providers know to serve a certain ad when they detect a certain content item. Conversely, as shown in Figure 4 (step 1), multiple advertisers could send ad/content ID lists to a service provider so that it will know, when a certain content item is identified on a consumer's device (step 2), which ad to serve to the device (step 3)[11].

Tracking Content Usage

Measurement of television and radio audiences has been a huge business for decades; audience measurement is the engine that drives the broadcasting industry. But when content is made available in a "nonlinear" fashion – that is, without reference to specific times or places where consumers can find content – then measurement of access to content becomes that much more difficult.

Content identification technologies not only enable a broad range of usage tracking techniques but also hold the great promise of providing accurate, measurement-based tracking instead of the sample-based techniques that are typical nowadays.

The most well-established usage tracking application based on content identification technology is watermarking for television programming. It is currently being used in two general ways: for *broadcast monitoring* and for *audience measurement*.

---

[11] Service providers would need ways of resolving disputes over multiple requests to associate ads with a given content item, such as a very popular song or video clip. For example, multiple ads could be rotated, or an advertiser could pay extra to establish exclusivity.

Broadcast Monitoring

Broadcasts are monitored to confirm that programs and ads on television actually ran as scheduled, as part of content licensing and ad affidaviting processes. Experiments have also been done with broadcast monitoring as a means of reporting radio airplay of music for royalty calculation purposes.

Traditional broadcast monitoring has required rooms full of people, one in each geographic market, to listen to radio or watch TV and log when ads and programs come on. With watermarking, receivers connected to watermark detectors could replace this costly and error-prone human process.

Watermarking-based broadcast monitoring businesses launched during the first Internet bubble but did not get much traction[12]. Teletrax, launched in 2006, is a newer example of a service that uses watermarking (from Philips) to confirm that content such as ads and direct-response programming (infomercials or ads with "call this number now") ran according to contractual obligations.

It is possible to apply similar technology to monitoring video content on the Internet. We mention one example of this below; see p. 18. But this could also include the ability to ensure that ads and promos scheduled to run with a given piece of content actually do so. Broadcast networks can facilitate this by inserting watermarks into ads (as they come in from agencies) as well as into program content. Networks can also keep records of program rundowns that contain identifiers for each of the components – including ads and promos – that are embedded into the content itself.

Audience Measurement

Nielsen and Arbitron have been using watermarking for audience measurement in television and radio. Arbitron developed its watermarking-based Portable People Meter (PPM) system for radio in the early 1990s and is currently rolling it out to the major broadcasting markets in the United States. Nielsen already uses watermarking for its TV audience measurement service and encodes the vast majority of television content in the United States.

In the "nonlinear" world, the concepts of broadcast monitoring and audience measurement fuse into the overarching concept of content usage tracking. We know that – unlike on conventional broadcast media – content is available from many places at any time; it remains to measure who is using it, on what devices, and so on. Audience measurement devices like Arbitron's PPM and Nielsen's in-home devices have been used as statistical proxies for real data and only with the cooperation of "panel members" who are paid for their trouble.

The Internet and other digital distribution mediums hold the potential for precise, detailed tracking of personal content usage, just as they do for advertising. User-level content usage tracking can be done through content identification technologies as well as through encryption-based DRM. This is a controversial area, not least because royalty agreements vary widely by geography, artist, content type, and other dimensions.

Currently, in many countries – especially continental Europe and Canada – copyright owners receive royalties for content usage on various types of consumer devices through

---

[12] One example was Verance's ConfirMedia service, which launched in 2001 and ceased operations in 2006.

levies paid by hardware manufacturers to royalty collecting societies[13], which distribute them to the copyright owners.

Levy schemes differ widely from one country to the next, not only in magnitude but also in how they are calculated. And as consumer electronics become cheaper, levies that the hardware makers must pay become proportionately higher. As a result, consumer electronics companies have been lobbying the European Commission and other regulatory bodies to accept automated usage tracking schemes such as DRM and content identification technologies in lieu of levies.

Technological usage tracking is also controversial regarding consumer privacy. Users can rebel against consumer devices that actively report usage data, even if it is aggregated or anonymized[14].

Yet there are successful examples today of Internet usage tracking schemes that employ content identification technologies. One is Digimarc's ImageBridge, which is based on the original MarcSpider technology that Digimarc offered as an antipiracy technology during the first Internet bubble. ImageBridge tracks watermark-embedded digital images throughout the web and reports on usage, whether authorized or not. Brand owners use it to track the impact of their logos and trade dress, while others use it to gather evidence for litigation for copyright or trademark infringement.

Somewhat related to usage tracking is technology for using image recognition through watermarking as a direct marketing tool. For example, ad agencies and consumer product companies are experimenting with watermark detection on camera-equipped mobile phones: if a consumer takes a picture of an ad in a magazine or in a public space, some software on her handset can detect a watermark in the image and send her a text or email message with a coupon or other targeted offer.

Another recent example is Attributor's text fingerprinting technology, which crawls the web looking for text passages that are identical to those in its database, as well as for proper attribution of that text. Newsgathering organizations like AP, Reuters, and Canadian Press use Attributor to track the "play" of their news content on affiliates' websites as well as unauthorized sites.

Nielsen and Digimarc are developing a solution called Nielsen Digital Media Manager for tracking television content on file-sharing networks, user-generated content sites, and social networks. The solution employs Nielsen's existing infrastructure of watermarking and fingerprinting, the latter as backup to watermarks where they do not exist or cannot be detected. Nielsen Digital Media Manager will need cooperation from network service providers (see pp. 23-25) and others to become viable.

Standard content identifiers combine synergistically with content identification technologies for usage tracking in many ways. One is to make it easy to integrate data feeds from services that track usage of content in different media or distribution channels, e.g., terrestrial broadcasting, digital cable, and various Internet services (P-to-P networks, social networking sites, etc.). With common content identifiers, it would be easy for a

---

[13] Such as GEMA, SACEM, and Buma-Stemra for music royalties in Germany, France, and the Netherlands respectively.
[14] For example, a technology for protecting music CDs when they are copied onto PCs famously failed in the market because (among other things) of the revelation that the technology was using consumers' Internet connections to report usage data.

service provider to create integrated usage reports to content owners for a given content item. This is shown in Figure 5.
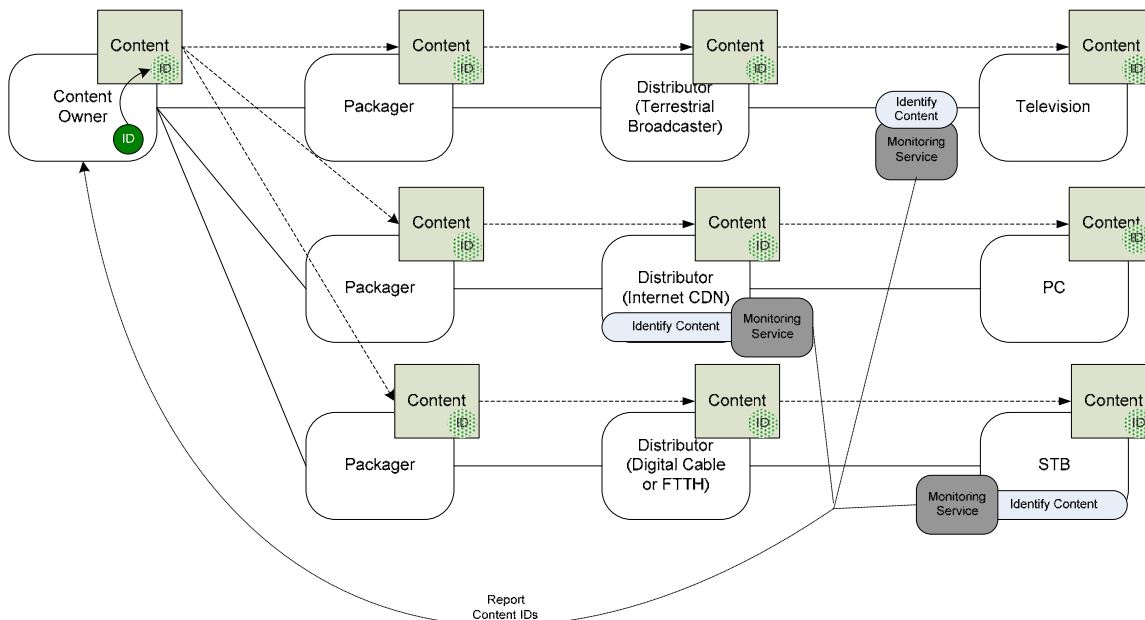
- Figure 5: Standard content identifiers enable integrated reporting of content usage across multiple delivery modalities.

Figure 5 shows a content owner distributing content (through packagers) to three different channels: a terrestrial broadcaster, an Internet content delivery network (CDN), and a provider of IPTV through digital cable or fiber to the home (FTTH). We show the content owner inserting a watermark before the content is distributed, but this example (or parts of it) could work with fingerprinting as well.

Each of the distribution channels may have a different way of monitoring usage. For example, terrestrial television may have a traditional broadcast monitoring service that is independent of both the broadcaster and the consumer's TV set. An Internet distributor may employ a CDN, which could monitor content on its servers; and a digital cable/FTTH IPTV service may employ set-top boxes with built-in monitoring capabilities. All such services can feed data to the content owner using the same content IDs – resulting in an integrated feed of usage data for the content owner.

Managing Assets and Integrating Systems

Content identification technologies can be instrumental in integrating the internal systems that touch on content within a media company. The most obvious such system is a media asset management (MAM) system or a digital asset management (DAM) system[15]. Other systems that handle content-related information can include DRM packaging, ad tracking, licensing, and royalty payment.

---

[15] For purposes of this paper, we'll use DAM to refer to metadata management, and MAM to refer to content management infrastructure including storage and movement of the actual files or "essence" with their associated metadata.

The term *enterprise content integration*[16] (ECI) has been used to describe an approach to integrating multiple content-related systems in a media company, as an alternative to creating a single centralized asset management repository. The term originally meant integrating multiple DAM/MAM systems so that users could search, browse, and retrieve content from across the enterprise through a single interface. But ECI is easily extended to encompass other systems such as marketing, product management, rights information, and the other types of systems mentioned above.

Standard content identifiers are the "glue" that can make it easier to integrate multiple internal systems. They can also be used to integrate with a content owner's own distribution systems or as metadata in feeds to external distributors or retailers.

The basic process in implementing ECI is as follows:

- Determine the smallest unit of content that you currently manage and that has value to your business or to consumers. This is the unit of content that will be identified.

- Set up a directory of standard identifiers along with basic metadata about the content that the identifiers denote. This is the *content catalog*.

- For each system to be integrated, find records that reference each content item. Store that system's identifiers for those records in the content catalog. For example, a royalty payment system will store information about content item XYZ. The record(s) in the system pertaining to XYZ will have their own identifiers, which could be as simple as internal tuple identifiers in a relational database. Store these identifiers in the content catalog under XYZ.

- Create functionality to update the above information periodically, so that the content catalog is always up to date with references to every system that stores information about content in the enterprise.

- Users can use the content catalog as a gateway to all information that the enterprise manages about content items of interest. Programming interfaces (APIs) can also be defined so that software can automatically access all of this information.

In other words, the ECI concept incorporates a process for normalizing content identifiers used – whether explicitly or implicitly – throughout the organization. Using common identifiers for content across multiple systems can lead to various business benefits, including cost efficiencies as well as scalability of product development. At the same time, we note that the above is not intended to be a complete "recipe" for enterprise DAM or MAM. Many other issues must be considered, such as the ability to group together assets that relate to the same piece of intellectual property (e.g., foreign soundtracks, letterbox versions, versions with subtitles, etc.).

---

[16] Rosenblatt, B. and Mark Walter, Enterprise Content Integration: Next Step Beyond DAM? The Seybold Report, Vol. 1, No. 20, January 21, 2002, pp. 10-14. The term was coined independently around the same time by a vendor of the technology, Venetica, which was subsequently acquired by IBM.

One organization that has used identifiers for internal assets in a very ECI-like way is Scripps Networks. Scripps uses an LDAP server to store its content catalog (the Scripps Asset Registry) and the widely-used Dublin Core standard for its basic metadata[17].

Monetizing Transformational Content Uses

The final business benefit of content identification technologies we examine here is one for the future. Many pundits predict that in the future online media world, so-called transformative uses of copyrighted material will become as important as the original content. Digital technology and applications such as Apple's GarageBand and iDVD make it easier than ever to take raw content materials and create something new through sampling, "mashups," and other examples of "cut and paste culture."

Some artists are even embracing transformation in their creations, such as Nine Inch Nails' recent album *Ghosts I-IV*, which was intended for remixing and mashups[18]. Another example is Lucasfilm, which made clips of "Star Wars" along with editing tools available last year on the Starwars.com website. This makes it easy for users to mash-up and use authorized Star Wars imagery in their blog posts or social networking pages[19].

Content owners have mixed feelings about such transformative uses of their content, and definitions of "transformative" are certainly subject to disagreement amid rapidly changing tools and technologies. But United States case law generally holds that transformative uses of content count as fair use[20] and thus do not qualify as copyright infringement.

Yet by now, at least some content owners – including major media companies – are becoming comfortable with transformative uses of their content under certain conditions. A simple example of this is Hulu.com, the joint venture of NBC Universal and News Corp., which allows users to embed links to video clips in their blogs, emails, web pages, etc.; Hulu and content owners share revenue from the ads embedded in the clips.

One of the most common criticisms of encryption-based DRM technologies is that rights-holders cannot reasonably allow such uses of content in conjunction with DRM, even though they may be legal. It is virtually impossible to imagine a DRM system that allows extraction of certain portions of a copyrighted work under certain conditions, for certain types of uses, while disallowing others. Among other things, editing tools would need to become "content aware" so that edit decision lists can track content genealogies.

Yet it is possible to use content identification technologies to track some types of transformative content usages through the value chain, using the same techniques as those described above (see pp. 16-19). The use of standard content identifiers would make this process more efficient, particularly if watermarks were used.

Here's how this would work: Let's say a user-generated content site accepts uploads of user-generated video clips. In addition to detecting full copyrighted works, as imeem does with uploaded music (see p. 15), it could detect snippets of copyrighted works that the user has included in the uploaded file. It could first search for watermarks in the snippets,

---

[17] Hurst, Chuck, *Assigning Identities to Enterprise Assets*. Presented at Digital ID World 2007, San Francisco, CA.

[18] Trent Reznor of Nine Inch Nails released the album, without a record label, under the Creative Commons "noncommercial-attribution" license, which means that anyone is free to create derivative works from the content as long as they attribute it to Nine Inch Nails.

[19] Make-It-Yourself 'Star Wars': Lucasfilm Will Post Clips From Film Saga on the Web, Inviting Fans to Edit at Will. *The Wall Street Journal*, May 24, 2007, available at http://online.wsj.com/article/SB117997273760812981.html.

[20] See for example the Second Circuit Appeals Court's 2006 ruling in *Bill Graham Archives v. Dorling Kindersley Limited*, or the 2003 Ninth Circuit ruling in *Kelly v. Arriba Soft*.

then compute fingerprints for those portions of the video clip where it could not detect watermarks. The server could invoke rules for the bits of content that it recognizes, such as to show targeted ads, block the upload if the snippet is too long, or simply note the user's tastes to drive recommendations.

The server would need to search for multiple types of watermarks, but if they all contained standard content identifiers, it would be easier to determine the identities of the snippets. This would also help enable the content owners embedding the watermarks to use a watermarking scheme of their choice, without having to be locked into the watermarking vendor's own database.

We have discussed a number of business benefits to content owners from adoption of content identification technologies, and shown various synergies between standard content identifiers and watermarking or fingerprinting. The question is not whether content owners can benefit from these technologies; the question is how to bring them to market.

As media packaging, distribution, and consumption go digital, functions become decentralized, and it becomes less and less possible for single entities to control the value chain. Furthermore, if common identifiers used internally within media companies conform to industry standards, they can also be used to feed external distribution schemes[21]. That is the essence of many of the services that Secure Path, for example (see below), offers to video content owners[22].

Therefore the key principle in bringing content identification technologies to market is ensuring that all of the players in the value chain have incentives to participate. Incentives come in three interrelated forms: revenue opportunities, ways of reducing participation costs, and the adoption of standards to increase confidence that the market will develop broadly. We conclude this white paper by looking at each of these three incentives in turn.

Revenue Opportunities

The success of content identification technologies is predicated on the buildout of certain types of services, which represent opportunities for entrepreneurial businesses. These services fall into two major categories: registration agencies for content identifiers and network monitoring service providers.

Content Identifier Registration Agencies

We discussed the key organizational roles that registration agencies play in content identifier standards (see p. 9). As mentioned, several identifier standards are set up so that multiple registration agencies can exist and offer unique sets of services, pricing structures, and so on.

Many registration agencies are (or are offshoots of) royalty collecting societies; for example, ISWC registration agencies for musical works include ASCAP in the United States and MCPS-PRS in the UK; and ISAN agencies include offshoots of film and television collecting societies such as GWFF in Germany and CFTPA in Canada. But they can also be for-profit businesses. One example is Secure Path, which is one of two current ISAN agencies in the US (the other being Microsoft).

Registration agencies' basic functions are to assign standard unique identifiers to content and to maintain directories that list identifiers along with information about the content they denote. Registration agencies can compete on the basis of the services they offer, which is invariably related to the metadata they store in their directories and the interfaces they offer to their data.

---

[21] For more on this, see Rosenblatt, B. *Enterprise Content Integration with the Digital Object Identifier: A Business Case for Information Publishers*. GiantSteps Media Technology Strategies white paper, June 20, 2002, http://www.giantstepsmts.com/ECIwhitepaper.pdf.

[22] As an example from another industry, many academic journal publishers use DOIs to identify journal articles so they can be accessed through the CrossRef reference linking service.

For example, Secure Path can register a content owner's ISAN numbers along with several different types of metadata for feeding information to distribution platforms, retailers, collecting societies, media metadata services (e.g., Tribune Media Services, Macrovision), fingerprint databases, anti-piracy service providers (see below), and so on; in other words, it enables a media company to "connect the dots" with various external service providers in the same way that ECI can connect the dots among different internal systems (see p. 19). Secure Path can translate among the myriad formats that these various entities use for medatata, mapping them all to ISAN numbers.

Network Monitoring Service Providers

The other type of function necessary to ensure many of the business benefits of content identification is *network monitoring*. Many of the business benefits enumerated in the previous section depend on some sort of network monitoring function:

- **Piracy Deterrence**: fingerprinting vendors like Audible Magic and Vobile are network monitoring service providers. They install their software, often delivered on a server as an "appliance," at a content service provider such as a user-generated content or social networking site. They analyze content that comes through the network, compute fingerprints, and look them up in databases. Beyond that, third-party antipiracy services detect the presence of copyrighted works on peer-to-peer networks, either to gather data for infringement actions (like BayTSP and MediaDefender) or to compile chart statistics (like BigChampagne). Companies like these will adopt the latest and greatest in content identification technologies as they seek to improve their services. MovieLabs, the R&D joint venture of the major film studios, has been conducting a series of tests of video fingerprinting technologies, which should help elucidate their strengths and weaknesses over time.

- **Contextual Advertising**: imeem's contextual advertising for music uploads depends on a fingerprinting service provider. Imeem uses SNOCAP's music e-commerce infrastructure, which incorporates Gracenote's audio fingerprinting engine as well as the ability to process business rules on fingerprint matches. Imeem acquired SNOCAP in February 2008; SNOCAP's other customers include MySpace. As the music industry continues to experiment with ad-driven free content, service providers for contextual advertising and other types of targeted marketing should proliferate.

- **Online Usage Tracking**: Nielsen's Digital Media Manager, which the company expects to launch in the second half of 2008, is a network monitoring service provider for television content that will require cooperation from websites and content networks to be successful. It uses Nielsen's own watermarking engine, and if it cannot detect a watermark, it uses video fingerprinting as a backup. Digimarc's ImageBridge (image watermarking) and Attributor (text fingerprinting) (see p. 18) are other examples of network monitoring service providers.

As content owners, packagers, advertisers, and other media value chain entities gain more experience with business models that depend on content identification, there will be more opportunities for network monitoring service providers to support those business models as well as for registration agencies to provide data feeds that knit all the pieces together.

Reducing Participation Costs

On p. 6 we discussed the differences between watermarking and fingerprinting technologies. Underlying these differences are questions about who pays for what, in terms of both dollars and effort. All else being equal, content owners prefer fingerprinting over watermarking because it only requires minimal effort on their part: they need only register their content with fingerprinting engines once per content item. They typically do not pay for this; on the contrary, they may charge fingerprinting vendors nominal service fees to feed them content for fingerprint computation and database entry.

Instead, fingerprinting vendors get their revenues from websites and other network service providers, who must install fingerprinting "appliances" in their infrastructures and incur the computational costs of analyzing content on their networks.

With watermarking, costs are spread more evenly throughout the value chain. Content owners must introduce watermark insertion into their production and release processes, which requires additional effort and cost -- of both watermark insertion per se and integration with DAM and other systems that supply payload data. Packagers, such as retailers and aggregators, may also – depending on the watermarking model – need to pay for watermark insertion. Website and network service operators need not pay as much in hardware or software costs, because watermark detection is computationally much cheaper than fingerprint computation.

Yet watermarking has functional benefits over fingerprinting, such as accuracy and flexibility, as explained on p. 6. Content owners should bear incentives and cost-functionality tradeoffs in mind when exploring arrangements with various types of technology vendors.

Standards Adoption

Adoption of standards is a way of reducing costs; it is also a way to increase confidence in the growth of an emerging market. For example, consider the piracy deterrence application. Copyright registration is not mandatory in the United States but is a prerequisite to litigation for infringement. If content owners had a uniform way of depositing standard content identifiers with the U.S. Copyright Office (or equivalents in other countries) as part of copyright registration, then anti-piracy service providers could use those identifiers to drive their services. This would reduce infrastructure and administrative costs for everyone[23].

Standardization of content identifiers ought to be obvious, but as discussed earlier (see pp. 7-9), subtleties about what to identify and differences in requirements across media industry segments have led to proliferation of content identifier standards[24]. This is the case despite the fact that ISO and other well-established standards bodies have ratified some of these standards. The industry needs to work together to reduce the number of viable content identifier standards.

---

[23] For more on this, see Rosenblatt, B., Thoughts on Notice, Takedown, Fingerprints, and Filtering. DRM Watch, March 15, 2007, http://www.drmwatch.com/special/article.php/3665921. Based on a talk given by the author at a Progress and Freedom Foundation Symposium, "What Goes Up Must Come Down: Copyright and Process in the Age of User-Posted Content," March 16, 2007, Washington, DC, slides available from the author.
[24] "The nice thing about standards is that there are so many of them to choose from." -- Andrew S. Tanenbaum, Vrije Universiteit, Amsterdam.

Should watermarks and fingerprints be standardized? There would be advantages to doing so, but we suspect that the answer is no – at least not yet.

Watermarks have tradeoffs (as mentioned on p. 4) that designers weigh based on the application and media type. For example, Cinea's Running Marks (see p. 11) feature zero computational cost to insert watermarks on a client device; this can be attractive to device makers who want to minimize costs. Yet Running Marks may not be appropriate to use in an environment where cost of watermark insertion is less of an issue and more capacity or flexibility is desired.

Watermarking vendors invest considerably in "secret sauce" that involves digital signal processing techniques. It would be nice to conceive, for example, of a standards-based watermark detection and payload extraction tool that could detect all types of watermarks. But requiring watermark insertion technology vendors to accommodate such a tool would inevitably lead to disadvantageous compromises in their designs.

Fingerprinting vendors are even more invested in "secret sauce," because their techniques for interpreting digitized content rely more heavily on heuristics and proprietary knowledge bases, not just on signal processing techniques. Fingerprinting is also less mature than watermarking, especially on the video side.

Fingerprinting technology vendors need to be able to compete on the basis of accuracy in content identification and efficiency. They will also need to adapt their technologies over time in order to respond to the "arms race" of content transformations that will appear in order to fool fingerprint-based schemes.   Hackers will surely develop simple tools to process or insert noise into content in a way that fools fingerprinting algorithms while keeping the content quality acceptable – and, somewhat ironically, tools to encrypt content to evade fingerprinting.

As fingerprinting applications continue to proliferate, content owners will get more information about these factors and will decide how much accuracy and robustness is acceptable to support different business models. Therefore it is important to preserve a competitive market for fingerprinting technology as well.

Sun Microsystems in Media & Entertainment

Sun's innovative, eco-friendly solutions are built to be open, standard, and interoperable – consistently providing the lowest total cost of ownership in the industry. Sun's solutions and partners – running on platforms and storage designed for video – improve the economics of production, storage and delivery of digital media in all its forms, to all manner of device, without sacrificing quality. Sun's business model is based on innovation and choice, providing a wide array of products, solutions, and architectural frameworks to enable media companies to be responsive to new revenue opportunities.

This paper is sponsored by Sun Microsystems' Communications and Media Practice as part of the Sun Media Advisory Board (MAB) program. Led by Sun and leading media industry executives, MAB is focused on identifying practical solutions to shared business challenges of cable and TV broadcasters, film studios, and other major content owners. In making this work freely available, Sun MAB's intent is to foster industry dialog on technologies and practices that have potential to accelerate realization of digital distribution business benefits. The paper is available on http://sun.com.

About Sun Microsystems, Inc.

Sun Microsystems develops the technologies that power the global marketplace. Guided by a singular vision -- "The Network is the Computer" (TM) -- Sun drives network participation through shared innovation, community development and open source leadership. Sun can be found in more than 100 countries and on the Web at http://sun.com.

About the Author

Bill Rosenblatt, president of GiantSteps Media Technology Strategies, is a recognized authority on digital media technologies, including digital rights management, content management, cross-media strategy, and content production systems, as well as on issues related to intellectual property in the online world. Before founding GiantSteps in 2000, Bill was 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. 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. GiantSteps' clients have included branded content providers, digital media technology vendors ranging from early-stage startups to Global 500 firms, and technology public policy entities in the United States and Europe.