

**DIGITAL-RIGHTS-MANAGEMENT TECHNOLOGY IS  
THE NEXT STEP IN PROVIDING REAL PROTECTION,  
BUT AT WHAT COST AND TO WHOM?**

# The war on copying

ILLUSTRATION BY MIKE O'LEARY

**T**HE COMMUNICATIONS INDUSTRY is ready for an infusion of data, such as digital video, to drive it to recovery, but music, video, and other digital-content owners continue to keep a tight rein on their growing mass of IP (intellectual property) while waiting for a se-

ecure DRM (digital-rights-management) scheme to materialize. The complexity of DRM, however, makes it a nontrivial addition to a system, especially a consumer device with a low cost threshold.

Several standards are under development to define how to encrypt material and distribute it over public networks. Rather than define the policies themselves, they define a foundation framework that can support a variety of DRM policies. For example, the ISMA (Internet Streaming Media Alliance) 1.0 Encryption and Authentication spec, scheduled for approval this month, describes how to apply the Advanced Encryption Standard (AES) to content and packetize content for distribution in such a way as to prevent a late packet from disrupting the stream.

Intentionally, these specs do not define what keys you should use or how to man-

age particular rights. One reason for this omission is that the industry is unclear on how you should manage keys or express rights. Key management and expression of rights, however, is where DRM gets the most complex.

The TCP (Trusted Computing Platform) from the TCG (Trusted Computing Group) addresses key management by specifying a trusted module that applications can use to protect content. The module is actually a processing subsystem; all encryption and decryption happens on the module, so keys are never in the clear. However, to decrypt a 2-Mbyte/sec video stream, the module needs significant processing ability. The \$4.25 (1 million) AT97SC3201 TPM from Atmel, for example, performs a 2048-bit RSA sign in 500 msec.

Support for TCP should appear in Microsoft's next version of Windows (for-

*At a glance*.....42  
*Forensics*.....42  
*Protecting content* .....44  
*For more information* .....46



merly known by the code name “Palladium”). However, Microsoft has made public that it intends to introduce changes that will make the operating system incompatible with chips that follow the current version of the TCP spec.

Several proprietary DRM schemes are under development. Sony and Philips, among a plethora of hopefuls, have their own architectures. However, the lack of a consistent model for how to pay for content and use it will easily confuse users. Different business models that may restrict use to a single person or device or a specific time period will constrain the use of similar content, such as music files. Such flexibility is great for content owners but requires users to interact differently with every piece of content. Confusion will arise when users think they are buying one form of license and getting another. Buying and using content needs to be as easy as saying “I want it” and clicking on a button. It shouldn’t mean deciding how you want it and having to read pages of small print to understand what you are buying.

### PROTECTION AT ALL COSTS

Most companies mistakenly believe that content protection is about protecting content. Consider that renting a new

#### AT A GLANCE

- ▶ Content protection is really about protecting the release window.
- ▶ A serious impediment to content protection is the lack of a consistent model for how to pay for and use content.
- ▶ Any place you can patch around software authentication, you can patch around hardware authentication as well.
- ▶ You don’t have to copy every bit exactly to get a usable copy.

video release for a single night costs \$4 to \$5, but renting an old video for five days costs \$5. Most content makes the majority of its revenue in the first few weeks of release. Thus, content protection is really about protecting the release window.


Many companies mistakenly focus on the technology when trying to understand DRM and fail to consider the real social issues that managing content involves. For example, DRM schemes that tie content to a single PC fail to address the needs of, say, a child of divorced parents who lives in two homes. Even more common is the

person who wants to play music at home, at work, in the car, on a portable player, and at a friend’s house. The killer app for digital content is the connected home, yet most DRM schemes undermine consumers’ ability to easily move content between devices. Protection isn’t just about security; you need to consider convenience, as well.

It is in these real-world individual issues that DRM will meet the most resistance. Content owners want strict definitions of stealing and honesty. However, schemes that are too difficult to work with or appear insensitive to individual circumstances may drive “honest” people to reconsider what honesty really means.

### TYING DOWN CONTENT

A popular protection mechanism is to tie content to a hardware “dongle,” trusted module, or system component that cannot be copied. Traditional dongles hold all keys to content (or the master key that protects all other keys). You need the dongle to access content, and you can tie only one dongle to any particular content, a fairly limiting prospect. Consider that if you lose or misplace the dongle, you lose access to all the content it protects.



## FORENSICS

The most effective DRM (digital-rights-management) schemes will adjust to compromise. Therefore, content-protection systems will need a programmable component that allows new or existing content to address exposed vulnerabilities.

Key to adjusting to compromise is being able to determine how compromises are made. If all the intelligence is in the content player, when the player is compromised, the device will mask the compromise, and no one will be able to learn how the compromise was executed. Thus, the broken content must itself carry the details of how it was compromised.

Several companies have considered watermarking to identify the source of compromise. For example, before downloading

content, the distribution server will insert a watermark into the content, identifying your account.

Although watermarks can mark a downloaded file, they can’t mark a DVD bought off the shelf. Pirates can also reverse-engineer public watermarks or break secret ones to remove them. Another attack uses successive approximation to make user-unperceivable changes to content until the watermark is no longer detectable.

Cryptographic Research proposes a new strategy that embeds executable code within content. When a player tries to decrypt the content, it must present to the executable code-system information—such as the player, mixer, and speakers in

use. The content determines the risk of releasing unencrypted content to the system. For example, it considers a player that provides a serial number less risky than a player that doesn’t. For a risky configuration, the code might decode the content to produce a lesser experience in the form of lower resolution, for example.

The code can also dynamically embed the system information into the content. If the content is broken, this information travels with the content, enabling content owners to track the kind of system that broke the content and use this information in their next content release. You can also use the system information to build a case against an attacker, showing that the same equipment

broke several titles.

To make it effective, you must secretly embed forensics information in a nondeterministic manner. For example, a scene’s background color or the actual transition between scenes could hold information. Each system configuration would produce a set of subtle markings. Thus, 10 players might show Frame A, but an 11th player would show Frame A’. With a variety of mixed marks, pirates may be able to tell that the content was marked, but they will be unable to tell how the content was marked. Overcoming the secret of the marks is not critical to copying the content; it is for forensics, not protection. However, because it’s only a mark, you won’t know when you’ve successfully defeated it.



Also, you can access content from only one location at a time; in other words, you can't listen to a song on your MP3 player while your spouse uses the stereo to listen to a different song tied to the same dongle.

A two-tiered scheme allows more flexibility. A master security dongle, perhaps on the motherboard of a PC, holds the keys to content. To access content, you use an individual dongle. The master can either pass keys to the individual dongle, which introduces vulnerabilities, or decrypt and then re-encrypt the content to tie it to the individual dongle. The advantage of this system is that several users can access the same content. However, you need to have mechanisms in place, such as a limit on the number of dongles supported, to prevent users from re-en-

crypting content to a friend's dongle.

Flash modules are currently a popular method for transporting content. Several schemes tie content to individual flash modules. However, because an individual can have multiple modules, it's not unreasonable to consider that friends might share content by exchanging flash modules. Thus, tying content to the flash alone is insufficient.

Requiring users to occasionally "check in" over the Internet to authenticate access to content addresses the issue of content running on an approved list of registered devices. If a user wants to transfer rights to another device, the old device can be unregistered, and access to content will fail after the next scheduled check-in. Scheduling check-ins requires diplomacy, because users may perceive

that you are tracking their usage. Alternatively, a device may not have access to the Internet each time a user wants to access content, such as when a user watches a movie on a plane. You need to be able to postpone check-in by a reasonable window.

You also need a plan to account for the failure of a tied component. Users may tie personal data to the component that they need to be able to recover. You can't have a backdoor in the scheme, because others can compromise it. Thus, you need some way to register the material a user owns or to back up keys. Backing up keys is not a preferred mechanism, however, because it requires you to take the keys off the secure dongle. A rogue entity's authenticating itself to the secure dongle compromises the keys. Additionally, you have

## PROTECTING CONTENT

Given that nonvolatile memory, such as flash, is necessary to internally secure keys, no one expects Intel or AMD any time soon to integrate key management onto PCs' main CPUs; flash requires a different manufacturing process from that of the CPU and would disproportionately drive up the cost of a CPU with flash.

Protection is not a one-size-fits-all endeavor. Not all content is equal. More important, not all users are equal. If you want to support multiple licensing models, including one-time use, timed use, evaluation use, and so on, all of these mechanisms must be in place and working together.

The value of content defines how worthwhile it is to protect. Delivering movies to theaters in a digital format represents a value on the order of millions of dollars from a single hack. The number of theaters to protect is in the thousands. Compare those figures with a 99-cent song that millions of users download. You cannot use the same mechanisms to protect both these types of data.

Some DRM (digital-rights-management) schemes use

expiration dates. For example, you can use the content for 30 days, or you need to check in your device on a scheduled basis. To support such a function, however, you need access to a reliable clock, not a user-managed time and date. Maintaining or accessing such a clock raises device cost. You also need to consider mechanisms for protecting the clock.

One potential issue with open but proprietary DRM schemes is that large companies could lock smaller vendors out of the market by making the price of entry too high.

If successful, DRM will make today's multimedia equipment obsolete. How will consumers react to their HDTVs' becoming useless after only a year or two?

Content owners are foolish to try to force OEMs to develop a secure system when they have no real stake in its success. The OEMs don't lose revenue when their products are broken, and they may actually sell more when they are!

If a device can run user-created content, you can use it to run broken content. However, if a device can't run user-created content, users won't buy it.

(Videocamera OEMs aren't going to allow content owners to protect content at the expense of their lucrative market.)

Note that personal information is digital content, as well. The difference between a movie and medical records, however, is that the user has a vested interest in keeping the medical records secure, and so will cooperate with schemes to protect them.

A trusted platform will not actually protect user privacy. Users need to pay for content, requiring some kind of credential, account, or identity, which you can then use to track usage. Alternatively, content owners could use driver and application code to track usage.

One potential "abuse" involves users tying limited content to a dongle. When the user is finished with the content, he or she can sell the dongle and the rights associated with it. There's not much difference between selling dongles and selling used CDs.

Note that a trusted platform might actually make tracking "dishonest" users even more difficult. Protecting content in case a device is stolen means that all content is encrypted. Thus, even

stolen content will be encrypted using the trusted module. If you don't have the password, you can't identify what has been stolen. Forcing a user to give up a password requires evidence of wrongdoing, but you can't get the evidence until you have the password, thus stalling any legal approaches. And, you could use any mechanisms that you could use to break the stolen content to break protected content.

The problem with approaches such as limiting the number of times you can use a ripped file before you have to rip it again is that you can easily circumvent them. Consider a utility that would keep your files ripped for you automatically. Does such a utility circumvent protection technology by the DMCA (Digital Millennium Copyright Act)?

Hackers will test the DMCA to its limits. Imagine a virus that breaks content on a user's computer and broadcasts it over the Internet. Is the user liable for that virus? How would you know if a user intentionally accepted the virus? Blurring the line between intentional and unintentional hacking will make separating the innocent from the "guilty" a costly legal challenge.



to protect against the accidental destruction of keys. If it is possible to erase the module via software, then a virus could trigger an erasure and wreak havoc.

Many users don't back up their computers. Therefore, some mechanism must be in place to recover licenses after a system crash. One strategy dictates that the user has to contact the site from which they purchased content, but having to contact a multiplicity of vendors about a variety of content could be a herculean task. Consider also that when you tie content to a specific component, this component will not be present when the user upgrades equipment, forcing the user to download and rekey all content.

A robust DRM scheme needs to automate the transfer of rights between devices. For example, Turbo Tax 2002 offered a flexible license that allowed you to install the software on multiple machines, so that you could work on your taxes from any machine in the house, but only the designated primary machine supported the key functions of printing and electronic submission. If you wanted to designate another machine, you had to make a phone call. This process was inconvenient for users and ate into profits.

Some vendors have suggested a clearinghouse that manages all of the rights and licenses an individual has purchased. If the computer crashes or the user upgrades, the clearinghouse manages the rights-transfer process. Additionally, with a clearinghouse, you can use a portable component, such as a smart

card, to exercise rights at a temporary site, such as a friend's computer.

#### THROUGH THE OPEN WINDOW

A simple means of defeating rule-based protection is patching. A user who patches the application so that it ignores a flag or dongle defeats a do-not-copy flag or dongle. Tying the rules directly into the content is more effective but more complex (see sidebar "Forensics").

An improved defense is to execute critical-function code from a protected domain. You can execute code on a back-end server, which requires a network connection, or a trusted coprocessor, such as a smart card. For example, Sospita and Schlumberger have teamed up to allow you to execute critical code on a smart card. Note, you can't run high-performance code; response time for a first execution can take around 270 msec or as little as 8 msec if the function is cached on the card. Card readers might be pricey for some applications: They cost \$15 to \$25 each, depending upon volume.

Trying to prevent every instance of copying will probably require schemes so complex that users will be unwilling to use them (see sidebar "Protecting content"). Even if you can create a secure scheme, managing all the special cases, including lost tokens, legitimate transfer of keys, and others, may eat away more profit than those cases protect. Consider supporting a reasonable level of protection. For example, allowing users to install software twice lets a user to create a backup, transfer to a new computer, quickly handle the loss of a dongle, and

so on. You can avoid most service calls with this one strategy. Certainly, two friends might copy the software, but compare that lost revenue with the cost of managing legitimate transfers.

#### LAW VERSUS TECHNOLOGY

Protected content yields a questionable benefit. You can't share content, access it anywhere you like or resell it after you are finished with it. Additionally, understanding DRM licenses and dealing with restrictive use are annoying issues. If no perceivable benefit exists, then users have no incentive for cooperating with DRM.

OEMs face the issue of having to allocate more development resources and equipment budget to DRM functions, resulting in a more expensive device that is more complex to use but provides no extra value to the consumer to offset the cost. Compound this expense with devices needing to support several DRM schemes.

In other words, the users and OEMs, not the content owners, bear the cost burden of DRM. Neither users nor OEMs gain any real benefit from cooperating to create secure DRM; in fact, it's in their best interests if such schemes never succeed.

Movie and record labels also seem to be counting on legal instead of technical protection. The DMCA (Digital Millennium Copyright Act) of 1998 criminalizes any action that breaks copy protection, regardless of the reason for it. For example, if you copy a movie to your laptop hard drive so that you don't have to run the DVD drive on your battery, you

### FOR MORE INFORMATION...

For more information on products such as those discussed in this article, contact any of the following manufacturers directly, and please let them know you read about their products in *EDN*.

**Apple Fairplay**  
www.apple.com/itunes/

**Atmel**  
www.atmel.com

**Cryptographic Research**  
www.cryptography.com

**DHWG (Digital Home Working Group)**  
www.dhwg.org

**Digital Media Project**  
www.chiariglione.org/project/

**EFF (Electronic Frontier Foundation)**  
www.eff.org

**Infineon**  
www.infineon.com

**ISMA (Internet Streaming Media Alliance)**  
www.isma.tv///

**ISSA (Information Systems Security Association)**  
www.issa.org

**Liberty Alliance**  
www.projectliberty.org

**Microsoft**  
www.microsoft.com

**MPIAA (Motion Picture Association)**  
www.mpiaa.org

**Philips**  
www.philips.com

**Real Networks**  
www.realnetworks.com

**RIAA (Recording Industry Association of America)**  
www.riaa.com

**Schlumberger**  
www.smartcard-drm.com

**SDMI (Secure Digital Music Initiative)**  
www.sdmi.org

**Smarte Solutions**  
www.smartersolutions.com

**SmartRight**  
www.smartright.org

**Sony**  
www.sony.com

**Sospita**  
www.sospita.com

**TCG (Trusted Computer Group)**  
(successor to Trusted Computing Platform Alliance)  
www.trustedcomputing-group.org

**Wave Systems**  
www.wave.com

**Windows Media**  
www.windowsmedia.com

have committed a crime, even though there have been no real damages or lost revenue.

Recently, the US Supreme Court ruled on DVD protection, stating that publishing trade secrets circumvents protection schemes as covered under the DMCA. If they take this idea to the extreme, vendors could protect content using any “trade secret,” even one as simple as shifting bits by one place. Traditionally, trade secrets were valuable only in the sense that they were secret. Now, the law protects them. And content owners are pushing for laws that allow prosecution of individuals. Effectively, the government and users will bear the burden of protecting content through the court system.

Content owners fail to see that it simply doesn’t make financial sense to take someone to court for “stealing” \$20 of content. Just imagine the RIAA (Recording Industry Association of America) asking district attorneys to fill courtrooms with 13-year-old

“criminals.” In the end, legal means alone will not protect content. If enough people ignore the legal barriers, the system will collapse under its own weight. The only effective protection scheme is one that actually prevents action.

The truth is that DRM is not for the benefit or protection of users, no matter what content owners or standards groups say. For DRM to succeed, content owners must take responsibility and bear the burden of its implementation. They must recognize that users already expect certain capabilities, such as the ability to use purchased content as and where they like.


Owners must also recognize that people will accept a definition of “honest” only if it doesn’t unreasonably inconvenience them.

Both sides—owners and users—must benefit in real ways for DRM to take hold. DRM must be seamless, non-intrusive, and painless if content owners expect users to cooperate. The main challenge

for DRM is that an efficient system, at least from a user and an OEM perspective, is in place. DRM’s big promise is to make listening to music or watching a movie a more complex endeavor.

It’s easy to lose sight of the true goal of copy protection. It is not about preventing copying. In most cases, a perfect digital copy is unnecessary; many DVD-copying applications make good-enough copies—copies that users can’t tell from the originals— from analog outputs. Effective copy protection, rather, is about making distribution of pirated material difficult enough that you can turn most nonpaying pirates into paying users.

Of course, you can always try charging a reasonable price and trusting people to be honest. Just think of all the money you’ll save not having to implement DRM. □



You can reach  
Contributing Technical  
Editor Nicholas Cravotta  
at editor@  
nicholascravotta.com

---

#### ACKNOWLEDGMENTS

*Special thanks to Jay Srivatsa, a principal analyst at iSuppli Corp, and Susan Kevorkian, a senior analyst at IDC, for their contributions to this article.*

---