



Global Information Assurance Certification Paper

Copyright SANS Institute
Author Retains Full Rights

This paper is taken from the GIAC directory of certified professionals. Reposting is not permitted without express written permission.

Interested in learning more?

Check out the list of upcoming events offering
"Security Essentials: Network, Endpoint, and Cloud (Security 401)"
at <http://www.giac.org/registration/gsec>

Practical Use of Digital Rights Management in the Enterprise

GIAC Security Essentials Certification (GSEC)
Practical Assignment, Version 1.4b, Option 1
John C. Jarocki
19 August 2003

Abstract

Digital rights management (DRM) is a hotly debated topic in security circles. Vendors promise protection while security pundits warn of “snake oil” (Byfield). The truth, as with most things, is somewhere in the middle. DRM is challenging to get right, and even when done properly, it still leaves some risks to be accepted and/or dealt with. Recently, enterprises have become interested in DRM technology as a way to prevent information leakage of their own confidential and proprietary information. With well-publicized and spectacular failures of content protection schemes in the publishing industry, Information Technology organizations are wary (as they should be) of implementing such schemes. Sales, marketing, and engineering departments, meanwhile, are desperate for a tool to meet the requirements applied to them by the legal and audit organizations.

The goal of this paper is to discuss the components of an effective DRM system, to explain how to assess the need, and to provide some guidance on the implementation. The IT professional, armed with this information, should be able to make informed decisions and at least recognize the snake oil when it appears.

© SANS Institute 2003. All rights reserved. No part of this publication may be reproduced, stored in a retrieval system, or transmitted, in any form or by any means, electronic, mechanical, photocopying, recording, or by any information storage and retrieval system, without the prior written permission of SANS Institute. Author retains full rights.

1 What is Digital Rights Management?

Employees of high tech firms spend a good deal of their working hours entering information into computers. They transfer the creative output from their brains right into the computer systems they spend so much time staring at. From there, this raw information is molded, annealed, and polished into the Intellectual Property of the company they work for. That information is sometimes sold directly to customers, used as blueprints for a physical product, or entered into a database where it can be mined for later use in distilled form.

This process of gathering, storing, correlating, and outputting proprietary information is key to the success of many companies today. Since the protection of this information – from destruction, unauthorized access, and leakage – has become so important, digital rights management (DRM) vendors have begun to focus on this new market. Simply put: DRM systems aim to control who can access the content, when they can access it, and under what conditions.

2 How Does It Work?

Protecting electronic content in a non-trivial fashion requires attention to three major dimensions of access and control – confidentiality, integrity, and availability. A good system should also ensure utility, authenticity, and possession (Parker, p.230). Furthermore, multiple layers of protection should be employed – including authentication, fine-grained authorization, tamper protection, logging, and auditing.

It is possible to protect a document through encryption using a simple shared passphrase, but the challenge of distributing this shared secret may be difficult. Public key encryption may be used, but the management of an additional public and private key pair is a burden on both the administrators and the end users. Ideally, the system would employ an authentication method that the end user is already familiar with and at the same time spare the administrators the effort of provisioning new accounts manually. Additionally, a system that leveraged the existing password (and or security token) for the user community would smooth the bumps in implementing a new process.

Documents can be protected so that the secured version, the software to decrypt it, and the key are the only things needed to access the document. However, more sophisticated systems may require a network connection back to a set of servers so that access and authorization can be centrally controlled. This connection provides an opportunity to log access to the document. It also makes it possible to revoke or alter access permission after the document has already been delivered to the recipient. It is just this sort of centralized document protection system that will be the focus of the remainder of this paper.

2.1 Confidentiality and Integrity

In a proper document protection system, a document must be able to be encrypted and digitally signed so only the intended recipient can view it and so that the recipient is assured they are viewing the correct and original document. In order to decrypt the document, the recipient must have a way to provide decryption keys to the crypto engine (most likely by authenticating with a unique username and passphrase). Then, if the document is to allow some recipients access and not others (perhaps even at a finer granularity than the entire document), there must be a way for the system to decide who gets access to what content.

One important consideration when choosing an encryption system is that the encrypted documents may be compromised at some future time as the field of cryptanalysis evolves. In fact, that time might come sooner than expected (Clayton & Bond). An electronic document protection system that does not include a physical security component should be used only to protect information that has a limited shelf life of usefulness or that would not cause catastrophic damage if compromised.

Systems are now available that support the emerging Advanced Encryption Standard (AES) crypto system. 256-bit AES, for example, would be a good choice at the current time while 56-bit Data Encryption Standard (56-bit DES) is of dubious value (Clayton).

2.2 Availability and Utility

Information security now recognize the “utility” of information as an important component of any security framework (Parker). An encrypted document that cannot be decrypted to its original form and read by the intended recipient is not very useful, so one may expect any adversary to try to create this scenario. In a traditional crypto system, this sort of denial of service (DOS) attack can be carried out by garbling the encrypted version of the document, by stealing the keys, and even by surreptitiously swapping the keys with an invalid set.

With a centralized system such as the one we have been describing, there is yet another attack vector defend against: the communications with the server itself. The connection to the server can be blocked, the server can be shutdown or compromised, and (more subtly) a session can be hijacked or monitored in stealth using man in the middle attacks with tools such as dsniff (Song).

2.3 Authentication

There are several ways to provide authentication, but an emerging standard is the use of a centralized LDAP (Lightweight Directory Access Protocol) directory as a repository for account information such as usernames, passwords, and certificates. If the users are members of a single organization, the problem of creating accounts may already be solved. If password controls are in place, and if LDAP groups can be used as access control lists for the secured documents, then much of the hard work is already done.

Most environments, unfortunately, are not that simple. Corporations, for example, often have business partners and customers who need access to sensitive information. A good solution may be to maintain a separate directory (or tree) for these partners. In this scenario, the solution would have to provide a way to consult both directories – either by placing referrals inside the directory structure, or by specifying multiple LDAP servers as authentication sources.

If LDAP is used, there are several other important considerations:

- Should LDAPS (LDAP over SSL) be employed?
- Should directory lookups be available to anonymous (guest) users?
- If authentication to the LDAP server is used, what type should it be?
- If passwords are used, how will they be stored (and updated!) in the document protection system?
- Will the protection system store a copy of the LDAP directory’s information, or will it pass look-ups and authentication requests along?
- Where will the LDAP server be located, and how will it be protected from both external and insider attacks?

The details of interfacing properly with LDAP are outside the scope of this paper, but it is clear that there are many issues to consider. The document protection architecture will need to address these issues from the start to be successful. . *Understanding and Deploying LDAP Directory Services* is an excellent book for more information on LDAP (Howes, Tim).

2.4 Logging and Auditing

One of the most powerful tools for measuring the effectiveness of a DRM system is a good logging and auditing process. In the case of a document protection system, the logs can be used to determine who is accessing the content, and possibly whether it is still in possession of the intended user.

The logs should capture as much detail as possible. They should be centrally stored in an immutable fashion to ensure they can be protected from tampering, and a method of distillation should be available so the details contained in the logs can be extracted quickly and easily (Eaton). A good logging system will provide canned reports that can be reviewed periodically, but should also have the flexibility to produce custom reports to illuminate the unique details of a specific event. Real-time alerting is also useful, and a tool such as SWATCH (Atkins) may be used to email or page the appropriate parties when an exceptional event occurs (e.g., alerting the help desk staff after repeated unsuccessful login attempts).

Finally, the reports and alerts produced from this logging system should be audited periodically to ensure that the overall system is effectively providing the desired level of protection.

3 Assessing the Need

Now that the document protection system has been described, a more fundamental question needs to be addressed. Is it really needed? It is easy to describe such a system, and it may even be easy to design and build it, but if it is not a right fit for the target organization, it is doomed to failure. One needs to examine policies, culture, and infrastructure before dedicating resources to an implementation project.

3.1 Document Classification

In order to protect a document, it must be possible to determine how sensitive the information it contains really is. The effort of assigning a value to categories of documents is called document classification. Some organizations, such as the US Department of Defense, have very strict documentation classification policies (DoD 5220.22-M), while others have no formal policy at all. If the target organization is closer to the latter category, chances are that a document protection system will have to be preceded by a document classification project.

3.2 Paper versus Electronic

The paperless office has been just around the corner for many years. The problem is that many people still prefer documents printed on paper. Paper requires no power source, is light, thin, and viewable in many lighting conditions, and is still more reliable than the average electronic device (Sellen & Harper). However, some documents may actually be easier and more convenient to navigate in electronic form. Additionally, the savings from not printing and distributing a paper copy may outweigh the costs.

3.2.1 Process Improvements

One of the easiest ways to provide return on investment (ROI) numbers to favor a document protection system is to focus on the savings in paper, ink, and distribution. If the document was formerly printed thousands of times and distributed around the world, migrating to a completely electronic version will add up to tangible savings quickly.

Electronic documents also have other advantages over paper. Hyperlinks and bookmarks can take the reader quickly to the section he or she wants to view, and a good search engine can make it easier to find content related to particular keywords. If the document is large enough, and the computer used to view the electronic version small enough, people might prefer the electronic version – especially if they need access to several of these documents.

Information that changes rapidly is also better handled in electronic form. If the system can be built such that the source data (e.g., production and marketing costs) propagates automatically to an updated secure document (e.g., a product price list) that immediately replaces the former version, the enterprise will benefit immensely from the increased responsiveness to market conditions.

3.2.2 Culture Shock

Regardless of the advantages of electronic documents, a conversion will still meet resistance. Some people will always simply prefer paper, while others will be resistant to the very idea of change. Depending on the culture of the organization, it may be necessary to provide the option to print the document (either permanently or only during the transition period). If this option is chosen, watermarks (discussed in Section 4.2) can be used to lessen the risk.

On the other hand, at least a segment of the user population is going to like having an electronic version of the document. If it was previously distributed in paper format (and especially if it contained 500 pages full of cross-references), the electronic version will be easier to transport and navigate. However, these same users may also experience a bit of culture shock when told that the cut-and-paste feature is not allowed. A table of data might tempt them with the possibility of running calculations or generating graphs in a spreadsheet program, but without cut-and-paste they will have to manually input the data. Determining whether this activity should be forbidden is yet another difficult policy detail that will require serious consideration.

3.3 Supporting Infrastructure

Finally, the viability of a document protection system will depend on the existence certain key pieces of underlying infrastructure already being available. While it is possible to implement the entire infrastructure at one time, projects with such a large scope will be easily mired in delays and politics.

3.3.1 Centralized User Account Database

If there is no centralized database of user accounts, implementing a centralized document protection system will be challenging, to say the least. LDAP has been mentioned already because it is standards based and supported by many vendors and software packages.

LDAP is not the only way to provide an account database. Microsoft's Active Directory (essentially Microsoft's implementation of LDAP) and NT domains are other possible places to centralize accounts. For UNIX-based environments, NIS domains can also provide this function.

Typically, vendors of DRM software will provide integration with one or more of these. They may also have the ability to store accounts in their own internal database. A combination may be useful in cases where the majority of users (for example, all employees) will be located in a central LDAP directory, but a few exceptions (perhaps consultants) will be located in the internal account database. The differences in the way vendors handle these different account databases can be subtle but important. If the document protection system needs to group both internal and external users into a single policy group, for instance, it is important to make sure the product you choose supports groups containing members from multiple authentication domains.

3.3.2 Authorization (Role) Mapping Policy

The previous section mentioned policy groups. A policy group is a way of specifying the "who" in a document protection system. Most DRM products support some method of specifying who can view, save, print, or modify specific documents.

While it is possible to manually build such groups each time a new document is protected, this would destroy any hope of creating a scalable system in the long run. Instead, it would be nice to have a way to

specify groups once, and then re-use them for multiple documents. It would also be nice to attach roles to individuals. If these roles can be assigned automatically based on, for example, department membership, employment status, and location; the document protection system will suddenly become much more powerful and less of a burden.

3.3.3 Account Provisioning and Termination Procedures

If policy groups will be created in a mostly automated fashion as described above, then we want to ensure that the set of users is as accurate and up-to-date as possible. The process of creating an account for a new user is known as user account provisioning and is something of an interesting art. It turns out that most vendors provide provisioning tools for their software when accounts are needed. It also turns out that those provisioning tools are often not sufficient to do everything that a sophisticated organization requires. The author has personally asked many vendors about the common usage of their provisioning tools only to find out that most organizations eschew the provided ones and write their own to handle the details.

In order to make sure an organization is ready to pursue a document protection system that leverages an existing account database, a review of the provisioning process should be performed to make sure that any additional work to be done is identified as early as possible.

Even more importantly, the question: “What exactly happens when someone leaves the organization?” needs to be answered. If the document protection system relies on only current valid employees, it won’t pass an audit if the normal termination process removes accounts a week after they leave. This can be made even more complex if the accounts belong to employees of business partners who have no formal termination process in place.

4 Understanding the Risks

The decision to implement any security system should include analysis of the possible vulnerabilities of the system and whether the level of associated risk is acceptable.

This section provides some insight into ways that a document protection system might be compromised. It is not intended to be an exhaustive list, but should provide some food for thought.

4.1 Screen Shots

The first vulnerability most people think when presented with a system to protect documents is taking screen shots once the document is decrypted and viewable.

Vendors of DRM software have implemented various schemes to defeat “screen scraping,” but, barring esoteric approaches such as Tempest Fonts (Kuhn & Anderson pp.137-139), it is impossible to completely prevent this sort of attack. Instead, watermarks can be used to make the job of extracting the useful information more challenging (and also as a visual deterrent). Activity logs can also aid in a forensic investigation of the circumstances surrounding the information leak.

4.2 Printing

If the document can be printed, then it can be re-distributed. Even if the individual printing the document has no nefarious intent, physical security now becomes an issue. A document that has been printed can be also be scanned back into electronic form and Optical Character Recognition (OCR) tools used to extract any text for easier “laundering” and re-distribution. Again, logs and watermarks can provide some mitigation, but the risk remains.

4.3 Password Sharing

If the only way a user of the system is identified is via username and password, what guarantees it's not someone else using those credentials? This could happen if the username and password are stolen, but it might also be simply a matter of laziness: if a remote office is required to one person with an account, but someone else needs access later, will they go to the trouble of applying for another account? It probably depends on how hard the account is to obtain as well as what the punishment is for impersonating someone else in the system.

A safer way to go would be to adopt a multi-factor system using biometrics or security tokens. The latter could still be shared, but it at least becomes more inconvenient.

Certificate-based authentication is also another possibility, but if the certificate is based on a particular client computer, consideration needs to be given to how often the user community roams to different computers.

4.4 Connection Tapping/Hijacking

This is a more esoteric attack (i.e. more fun to think about), but is definitely possible (if the connection is using a weak TCP sequence numbering scheme, for example). Of course, the application protocol itself provides an additional layer of protection, but to someone motivated enough to use this sort of attack, it is probably trivial to compromise.

An interesting note here is that tapping (such as with `dsniff`) could be carried out in complete stealth, which might make this an attractive option to an adversary despite the higher skill and amount of effort required.

4.5 Tempest Attacks

Another possible way to acquire the information in the secured document is to wait until some valid recipient is viewing it and then capture EMF radiation from the computer display and re-assemble the content on a remote system. This form of attack is known as a Tempest attack, and has been shown to have been successful under certain conditions (van Eck). If your organization is really the target of such techniques, electronic document protection without physical security is probably not a good idea anyway.

4.6 Social Engineering

Social engineering is often one of the most effective ways to subvert any security system. If accounts are being provisioned manually by an operations staff, an adversary could gain access to the system by convincing a staff member they should have access. In a small organization, or one where employment/membership status can be verified independently, this may not be much of a risk. However, if the system is deployed globally and with business partners, there is a good chance that the operations staff will have no idea who is calling them – and possibly no good way to verify the caller should have access. The mitigation for this sort of social engineering technique is to make sure a good process is in place to verify the identity of potential users and the level of authority they should have.

4.7 Manual Transcription

In the end, if someone can see the information, they can copy it down. If the original format of the document is not important, or if only a small amount of data is really needed from the document, then manual transcription will defeat any protection scheme. A malicious insider who has access to the document can perform the transcription, or an adversary can “shoulder-surf” or take images with a well-placed hidden camera while an unaware user views the document.

Because of this simple exploit, and unless the document can only be viewed under controlled conditions (i.e., at a secure location), only documents that are impractical to transcribe should be protected in this way.

5 It's About Mitigation – Not Absolute Control

A document protection system is not a panacea. When the benefit of distributing a sensitive document electronically outweighs the potential cost of losing that information to an adversary, such a system can still provide some assurance that the document reached the intended recipients and was not inadvertently available to the wrong people. Specifically, it is a good way of keeping honest people from making honest mistakes. For example, a busy salesperson with price lists from a dozen different suppliers might accidentally send a copy of Seller X's price list to Buyer Y, who just happens to be good friends with the CEO of Seller X's main competitor.

An otherwise honest person may rationalize, if sensitive information from the competition appeared on the desk, he or she would be doing nothing wrong by using it. After all, it was not actively sought out. A document protection system can make that scenario more unlikely.

Even in the case where an adversary is specifically attempting to acquire sensitive and proprietary information, providing some protection is useful as long as the risks are understood and accepted or mitigated.

5.1 Communicate the Policy

Especially for keeping the honest people honest, a well-defined and communicated policy is a valuable tool. In addition, the policy will at least provide a basis for litigation if sensitive information is acquired purposefully by someone who clearly should not have it.

5.1.1 Document Classification

The information classification policy should explain the various levels of sensitivity, how to handle each, and what the repercussions of not complying with the policy will be. Some, typical classification levels include Confidential, Secret or Classified, and Top Secret (DoD 5220.22-M).

5.1.2 Terms and Conditions

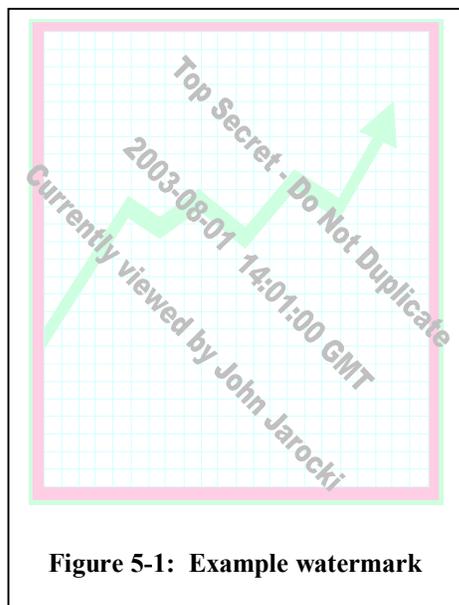
Before allowing access to a secured document (and perhaps even in the preamble of the document itself), the user should be presented with a set of Terms and Conditions and be required to accept them before being granted access. The Terms and Conditions should explain the sensitivity of the document and that the unauthorized access is forbidden. It might also explain that activities are being logged, and that perpetrators will be prosecuted.

5.1.3 Enforcement

If an employee can be terminated because they leaked sensitive information, this fact needs to be disclosed in the policy. Likewise, if a business partner can have their partnership terminated, or if an unrelated party can be sued for damages, they need to be notified of actions that may be taken against them.

5.2 Use Watermarks

Visible watermarks can provide a modest deterrent to improper use of a sensitive document by clearly stating the classification and possibly a phrase such as "DO NOT DUPLICATE." Adding the time the document was viewed and the name of the person viewing it sends the message that the system is keeping track of activities related to the document. It might also provide a way of tracking the source of leaked copies. Additionally, if the watermark is placed diagonally (as in Figure 5-1), OCR systems will have a harder time extracting useful text from the document.



Although users may complain about the watermark obscuring the information in the document, in practice they will find that the watermark is relatively easy to ignore.

5.3 Analyze the Logs

Logging details about access to protected documents is a good start, but the logs need to be analyzed in order to make them useful.

The defined policy should lead to a set of rules that can be developed to filter the data so that reports and alerts can be produced to highlight any problems.

5.3.1 Rules

Rules state what to do when a particular type of log entry is encountered. A rule, for instance, might be implemented to increment a counter when a specific document is accessed. A different rule might cause an alert to be generated when too many login failures for a specific user account are logged. If the policy states that user accounts should not be shared, the logs can be watched for access via a single user account from multiple different clients in close succession.

5.3.2 Reports

The set of reports that are necessary will again depend on policy, but some that might be interesting to produce include:

- Total Number of Unsuccessful Login Attempts
- Total Number of Document Accesses (with Details by Document)
- Number of Copies Printed (by User and/or Document)
- Connectivity Problems / Connection Resets
- Client Hosts With Connections Using More Than One Account

Providing the user community with feedback on details gleaned from these reports periodically lets them know that the system is being actively monitored. For example, during a project the author was recently involved with, an email was sent to the user community with the number of times a particular document had been printed in the previous week, along with a suggestion of how to use the document more effectively in electronic form.

5.3.3 Alerts

With a tool such as SWATCHe (Atkins), logs can be monitored in real time for extraordinary conditions that require immediate attention. Alerts can then be generated to notify technical support and/or security staff that action needs to be taken. For example, the system could alert on the following conditions:

- Multiple unsuccessful login attempts from the same account
- Access to a large number of accounts in quick succession from the client address
- Access from specific hosts or network ranges
- Attempts to access accounts of terminated employees

5.3.4 Aggregation

If multiple servers are being logged, or many documents are protected by the system, it is useful to watch for patterns at a higher level. Events from many client connections can be pieced together into patterns, and events separated in time can be correlated using even more advanced tools such as the Simple Event Correlator (SEC) (Vaarandi).

5.4 Enforce the Policy

In order to provide a deterrent against theft or misuse of the protected information, the usage policy must be enforced. Accidental violations will occur, and these obviously should be handled on an individual basis with fairness and consideration. However, if someone knowingly and maliciously ignores or subverts the policy, he or she should expect to be terminated, sued, or pursued under whatever conditions have been set forth in the policy. Ignoring this step will invite others to follow the steps of the first perpetrator and the policy (and the entire protection system) will eventually mean nothing.

6 Conclusion

Digital Rights Management is currently a hot (and contentious) topic in the security world. While it does provide an opportunity to be lulled into a false sense of security, it can also help control information leakage when implemented with attention to the important details. The key to a successful implementation is careful assessment of the need and understanding of the risks involved. The proper policies and infrastructure need to be in place, and access to the system needs to be monitored constantly to detect unauthorized access. Finally, fair and consistent enforcement of the policies is crucial for a successful enterprise DRM implementation.

7 References

Atkins, E. Todd. *SWATCH: The Simple WATCHer of Logfiles*. URL: <http://swatch.sourceforge.net/>

Byfield, Ted. *Security through Absurdity*. Disappearing Inc. Repackages Key Escrow. 19 Oct. 1999. URL: <http://www.heise.de/tp/english/inhalt/te/5395/1.html>

Clayton, Richard. *Brute force attacks on cryptographic keys*. 29 Oct. 2003. URL: <http://www.cl.cam.ac.uk/~rnc1/brute.html>

Clayton, Richard and Mike Bond. *Experience Using a Low-Cost FPGA Design to Crack DES Keys*. URL: <http://citeseer.nj.nec.com/539403.html>

DoD 5220.22-M. *National Industrial Security Program Operating Manual (NISPOM)*. Chapter 4: Classification and Marking. URL: <http://www.dss.mil/isec/chapter4.htm>

Eaton, Ian. *The Ins and Outs of System Logging Using Syslog*. SANS Reading Room. 14 Aug. 2003. URL: <http://www.sans.org/rr/paper.php?id=1168>

Howes, Tim, Timothy A. Howes, Mark C. Smith, Gordon S. Good. *Understanding and Deploying LDAP Directory Services* (2nd Edition). Addison Wesley, 2nd Edition, 2 May 2003.

Kuhn, Markus G. and Ross J. Anderson. *Soft tempest: Hidden data transmission using electromagnetic emanations*. In David Aucsmith, editor, *Information Hiding: Second International Workshop*, volume 1525 of *Lecture Notes in Computer Science*, pp.124-142. Springer-Verlag, 1998.

Parker, Donn B. *Fighting Computer Crime*. John Wiley and Sons, Inc., 1998.

Sellen, Abigail J. and Richard H.R. Harper. *The Myth of the Paperless Office*. Cambridge: MIT Press, 1st Edition, 2001.

Song, Dug. *dsniff*. 27 May 2001. URL: <http://www.monkey.org/~dugsong/dsniff/>

Vaarandi, Risto. *SEC - Simple Event Correlator*. 23 Jul. 2003. URL: <http://www.estpak.ee/~risto/sec/>

van Eck, W. *Electromagnetic Radiation from Video Display Units: An Eavesdropping*, in *Computers & Security* v 4 (1985) pp 269—286. 1985. URL: <http://citeseer.nj.nec.com/vaneck85electromagnetic.html>

© SANS Institute 2003, Author retains full rights.