The Thinking Person's Guide to Document Rights Management

FileOpen Systems

THE THINKING PERSON'S GUIDE TO DOCUMENT RIGHTS MANAGEMENT

PART 1: WHY DOCUMENT RIGHTS MANAGEMENT?

PART 2: CHOOSING A DRM SOLUTION

PART 3: RESOURCES

INTRODUCTION

Welcome to **The Thinking Person's Guide to Document Rights Management.** If you've downloaded this guide, you've taken an important first step toward understanding and implementing a document security solution for your business—deciding to look and learn before you leap. Perhaps you've heard stories of extra security doing more harm than good, or of vendors over-promising and under-delivering (or just going out of business). However you got here, you know that choosing a rights management technology for your company's documents is too important not to do your homework first. You are, after all, a "Thinking Person!"

Who This Guide Is For

(Or rather, "For Whom Is this Guide?" in Thinking Person-ese)
This guide aims to give an introductory background to rights management to anyone looking to control what happens to digital documents after they are sent and delivered. When we talk about "documents," we mean the files that come out of the most popular business productivity and publishing applications—PDFs, Word and Excel files, etc. We are *not* talking about "ebooks" that are destined for consumption on Amazon Kindles or Apple iBooks (they have their own proprietary rights management over which you have no control).

So, if you are looking to prevent unauthorized sharing of confidential or private information contained in a document, or are publishing documents independently and want to prevent piracy, this guide is for you. You don't need to be technical to grasp the concepts we discuss in this guide, but we hope to give you a framework for understanding the critical technical considerations involved in choosing a rights management solution.

Who are We?

We founded FileOpen Systems, a software developer of document security and rights management tools, back in 1997. We've been witness to all the sea-changes, setbacks and lurching progress in our industry for almost twenty years, and are eager to share our learnings (and those of our customers) with you. We won't be touting our products here, and at the end we'll arm you with key questions to ask any vendor of document rights management technology.

PART 1: WHY DOCUMENT RIGHTS MANAGEMENT?

We've all seen the ads by security software vendors…your would-be attackers are out there, in dark rooms, with your website or backend systems in their crosshairs. Only the strongest firewall will keep them at bay and your data safe.

In this age of Anonymous, WikiLeaks and state-sponsored hacking, the very same ads that used to look like scare tactics now seem prophetic. But what's getting lost in the hysteria over hacking is that more often than not, document leaks can be traced back to a user who had legitimate access to them (e.g. Edward Snowden), and most incidents of content piracy start with someone who originally paid for it.

Today we have a global network that makes sharing, cloning and republishing documents faster and easier than ever before. The Internet sends first and asks questions later. The Web was never actually designed to know about the identity of users accessing its servers, or for browsers to limit access to documents stored in HTML or XML.

The ubiquitous image of the hoodied hacker belies the fact that most document leaks are perpetrated by insiders, whether knowingly or unwittingly

Partly for this reason, the digital document survives, in the form of PDFs that can be viewed in Adobe Reader or many other 3rd party viewers, and Microsoft Office files for viewing and editing in their native applications. Digital documents can be controlled in a way that Web files cannot, because their viewing applications support encryption and turning "off" certain functions to the end-user.

By storing valuable or sensitive information in a digital document, you can encrypt it not only in transit, but deny access to to all but those users to whom you explicitly grant access. If organizations encrypted their documents as standard practice, they would have a lot less to worry about from hackers and document leaks.

Why Passwords Don't Work

The most basic form of document security is setting a password, specific to the file, that is required to open it. Since all the major document applications enable passwords, this is a common (and free) method of controlling access. The limitations of password security are fairly obvious—the password must be communicated to the intended recipient, and the viewing application knows nothing about the identity of the user when it accepts the right password. The recipient can easily share the password with others. Furthermore, because the "key" to open the document resides in the file itself, password security is vulnerable to rudimentary and widely available cracks.

"WHEN YOU RESET MY PASSWORD, DID YOU HIT 'REPLY ALL?'"

Document-level passwords also break down with any kind of volume document sharing or publishing. If each document must have a unique password, keeping track of those quickly becomes an administrative headache.

Proprietary Viewers and Secure Containers

Another approach, which has largely fallen out of favor, has been for a security vendor to package documents into a secure "envelope," which requires the end-user to download an application to open it. Or, for a security application to convert a standard file format such as PDF into a proprietary format with a different file extension, which only their viewer can open (again, requiring the end-user to download and install a separate viewing application). As consumer tolerance for downloading special viewing apps has diminished, and corporate environments increasingly block installations of unknown applications, document security by means of proprietary containers has become an untenable method.

Secure the Perimeter!

In the enterprise security space, the focus has been on keeping the bad guys out, and confidential documents inside the firewall. The problems with this approach should be apparent to anyone who's been following news of document leaks. Documents stored in unencrypted form inside a firewall are completely unprotected once they leave the firewall, or in the event that firewall security is breached. To make matters worse, BYOD ("Bring Your Own Device") has exploded the number of devices on which confidential documents reside and can be shared.

Digital Loss Prevention (DLP) solutions, an offshoot of antivirus and malware products, try to mitigate the risk by scanning email attachments as they leave the firewall and flagging or blocking them from leaving. But DLP solutions don't adequately address the many use cases in which corporate insiders need to share sensitive documents with outside recipients such as as attorneys, accountants and other consultants. There is no mechanism in DLP systems for granting access to users outside the firewall.

20 Years of DRM

1995 **1995-2000: Format Wars**
Adobe pushes PDF as open standard for digital documents. Proprietary security containers die off as Web/HTML usage grows. First DRM products for PDF emerge.

2000 **2001-2005: The "Golden Years"**
Adobe opens Acrobat/Reader to DRM partners. Standard viewing environment and stability of OS/browser market make life easy for publishers.

2005 **2006-2010: Battleship Enterprise**
Adobe, Microsoft, Oracle and heavily funded startups duke it out for market share in Enterprise Rights Management. None of this really helps publishers.

2015 **2011-Now: The Race to Be Everywhere**
The explosion in mobile devices, competing platforms and viewing apps, plus a splintering of the PDF standard, present new challenges for publishers.

Enter Document Rights Management

Document Rights Management, or "DRM," seeks to control access to encrypted documents based on a user's identity and on specific usage policies put in place by the document owner. DRM solutions are designed to lock *all* copies of a file, opening them only for users who authenticate and meet the owner's permission settings. DRM should be thought of as separate from ERM, "Enterprise Rights Management," which enables every user in a corporate environment to protect their own files within the firewall. The key differentiator is that DRM is for **one-to-many** document distribution, as opposed to **many-to-many** document sharing. For example, an HR department may author and encrypt a confidential handbook, distributing it only to corporate officers.

There is a wide range of DRM solutions out there, both in terms of cost and technical capability. All of them provide a degree of security greater than using passwords or relying on the company firewall. The way these solutions achieve that security varies greatly, and can make or break your relationship with your intended audience.

In the next section we'll discuss the important considerations involved with choosing a DRM solution. By the end of this guide, you'll be armed with a thorough set of questions you can use to hone your requirements, and narrow the list of potential solutions you'll want to investigate.

The Race to Be Everywhere

The days when all a publisher had to worry about was delivering digital documents to Mac or Windows desktops are long over. Today users expect to access documents on iOS or Android tablets, on their smartphones, and also on their home laptops and work desktops. That's a tall order for a publisher wanting to control access to their documents, but support their legitimate users on a myriad of devices and platforms.

DRM solution providers have been racing to keep up with these demands, ditching OS-specific plug-ins and viewers and in some cases using the web browser as the primary document viewer. As we'll see below in the *Security - Usability Spectrum*, the move to device-agnostic viewers comes with security trade-offs you'll want to consider. Other DRM providers have developed methods of identifying and "fingerprinting" users on multiple devices, so their credentials and permissions follow them around.

As you think about your goals with DRM, be sure to consider how your current users use your documents. Are they always at a desktop? Will they always have access to high-speed internet to authenticate? Or might they want to view the documents on their phone or table, on a train or airplane? If so, you'll want to look at DRM solutions that support multiple devices and even multiple authentication methods for the same user (e.g. depending on whether they are inside a corporate firewall).

DRM in a Cloudy Landscape

DRM for documents is not a new idea or technology—far from it—but adoption by the enterprise has been slow. The high cost of enterprise document workflow platforms, of which DRM is just one feature, has deterred midsize companies and small businesses. Others are simply playing catch-up with a splintering IT environment which now includes mobile devices, cloud storage platforms, and increasing adoption of online productivity apps such as Google Docs. How can a company protect its information if that information is all over the place, and in many cases controlled by employee's private online identities (such as their Gmail account)?

The answer lies in DRM solutions that are able to function with, but also independently of these disparate environments. While many cloud platforms such as Box enable document security, access to those documents can only be granted within Box, and so on for all the other cloud platforms. This "silo'd security," has many of the same limitations of those secure containers we discussed, except now the cloud vendors function as the containers.

Cloud storage providers tout their strong security, but files can only be shared within the confines of each cloud service

Just as document management is moving to the cloud, so have some DRM solution providers. This migration brings its own set of concerns to publishers. Will you have to upload your unencrypted, vulnerable files to a remote server for encryption? Could they be compromised in transit? What happens to the unencrypted files—does the DRM vendor now own them? In fully cloud-based DRM systems, you may also have to upload your customer data so the server can authenticate your users.

If those are the kinds of questions that keep you up at night, you'll want to find solutions that allow you to perform some or all of the encryption and authentication functions locally, on your desktop or company server. In fact, the more dispersed the encryption, authentication, and document storage functions of your DRM implementation, the more secure it will be, because access to one part will not mean access to the others.

It's Not Just About You

If we could share just one lesson with you in this guide, this would be it:

No DRM solution, no matter how strong the encryption, will work if it is rejected by the legitimate end-users of your documents.

Placing too many authentication hurdles, forcing downloads and installations of heavy viewing applications, or failing to support the viewers and devices your users want, will compromise your security goals and may lead your document owners to abandon DRM altogether. The best DRM is the one that your intended users don't even realize is there.

Depending on the degree of sensitivity or value of the documents you are trying to protect, an important first step is deciding where you want to land on the *Security - Usability Spectrum.* We've developed this spectrum as a way of explaining to our customers how the way you implement document security is always a trade-off between security and usability. Just like a law of physics, this is one of those immutable laws of the technology world: as you ratchet up the security requirements on a document, the ease-of-use for the legitimate end user will go down. Likewise, if you optimize for usability, you'll have to trade off on some security features.

In this graphic, we've plotted some common document security features along the spectrum in terms of their impact on security and usability. In general, the more you want to control what happens to a document once it reaches an end-user's device, the less ease-of-use you're going to get. For example, if you want to control printing to a very granular degree, specifying certain page ranges and which kind of printer, you'll need a DRM solution that has client software sitting on the end-user's machine and can monitor OS functions.

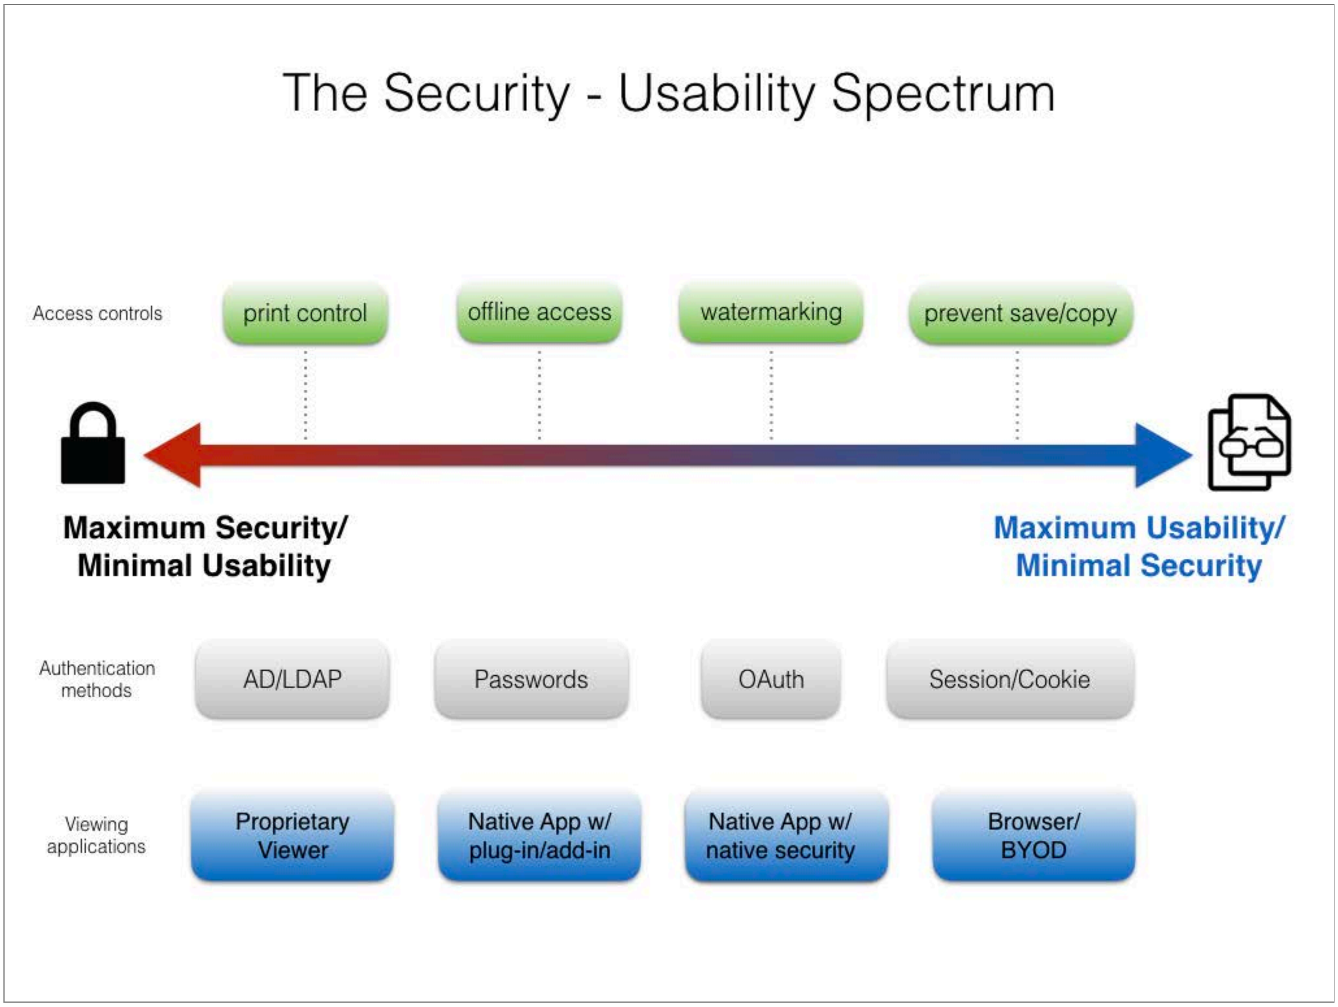

Certain access controls, authentication methods and viewing applications offer or require a higher degree of security, and place demands on the end-user to install special software.

Alternatively, if you want a very easy experience for your end-users, with no downloads or installations, you could use the web browser as your viewer and grant access via a session cookie. But, there are ways that a user could capture a document that way and share it with others. What you gain in ease-of-use, you generally give up in strong security. Where you want to end up on the *Security - Usability Spectrum* will determine what kind of DRM solution you choose, and which features you choose to implement.

Identity Management

Before you even consider implementing DRM, it's important to look at how you'll be managing your user data going forward. Most sophisticated DRM systems will allow you to import or interface with your user database in order to assign permissions to and authenticate your users. Identity management is what separates DRM from file-based password security. Rather than asking "Is this the right key to open this document?," DRM asks the question, "Does this user, at this moment, on this device, have the right to open this document?"

Good identity management makes it possible to change user permissions even after a document has been delivered to them, or to remotely revoke access completely (such as when an employee leaves the company).

In the movie "The Island," Ewan MacGregor's character adopts the identity of his human clone

There are degrees to how much a DRM system will integrate with your user data. Some enterprise solutions control both the user database, policy creation and authentication processes, and are therefore completely interoperable (within the company firewall). Other DRM solutions will "talk" to your user database via APIs, and can even use that data to grant access in an active directory setting. The more basic DRM offerings will allow you to import a spreadsheet of your users, so you will need to set up a process around keeping your user data up-to-date in two different places.

Open Standard or Walled Garden?

The Web is built on open standards. HTML, HTTP, SSL and many other Internet protocols are based on published specifications which international committees have agreed upon after years of research and debate. The advantages of standards are obvious to anyone working with technology—if we're all speaking the same language, we can innovate and interoperate faster, and deliver to our customers the functionality they want. Technology based on open standards levels the playing field for developers, and also gives consumers more choice and the freedom to migrate from one platform to another without too much disruption.

By contrast, a "walled garden" is a technology that is proprietary from top to bottom. Apple's iOS is a good example of a walled garden, because despite the App Store and Apple partner ecosystem, Apple exercises absolute control over what software runs on their operating systems and devices. Steve Jobs famously argued that this was necessary to ensure the highest quality customer experience, but it also results in a lock-in to a single vendor.

To date, there is no open standard for DRM in the way that SSL is an open standard. This stems in part from the inherent openness of the Web we referred to earlier—HTML pages were never designed to be encrypted. It may also be due to the security risks involved with publishing encryption and authentication methods for all the world to see.

That said, publishers can still choose DRM solutions that use standard components, and would be well advised to. DRM products that rely on proprietary wrappers or document formats require publishers to convert their existing documents to the new format, which becomes inextricably linked to the vendor's encryption method. A standards-based alternative would be a DRM system that supports PDF documents in their native state, so that they can be viewed in the PDF viewers that users' already have on their devices (although Adobe Reader is not an "open standard," it is an industry standard, and the PDF format is in fact a published open standard).

Using standards as much as possible gives publishers the most flexibility going forward, should they ever decide to change their security methods. It also benefits end-users, who won't have to continually upgrade a piece of proprietary viewing software.

In the enterprise space, IT professionals should carefully consider whether a one-stop shop document management platform with security features gives them enough flexibility going forward. If a corporation is using a single vendor to manage their productivity software, identity management, cloud storage, encryption and authentication services, that may offer consistency and convenience, but it also presents a single point of failure for essential operations. Furthermore, it locks a company into a very long and potentially expensive relationship with the vendor upon whom they depend.

Choosing modular solutions with standard components, which "speak" standard protocols to other systems via published APIs, gives publishers and corporations autonomy and flexibility when it comes to their most essential digital assets. Companies who adopt this strategy won't be brought to their knees of one component fails or a vendor goes out of business.

Controlling Access: Key Features

Once you've considered the broad strokes of your DRM requirements, you'll want to think about the specific access controls you'd like to be able to place on your users. In this section we'll look at the features many publishers use to limit usage, enforce subscriptions and corporate policies, and control the spread of valuable information.

As you evaluate DRM solutions, you'll want to think about *how* you would want to implement these features. Will they apply to a set of documents, or rather to a set of users? Will each user have their own unique permission set? Some DRM products will give you the option to apply permissions to users rather than to the documents themselves. This gives you the flexibility to change those permissions after you've delivered the document to the end-user.

Access Control Feature	Why It's Important
Copy/Save Prevention	Prevent copying of text, and saving/printing to another file format, so users cannot generate a version with no security. Some products can also block screenshots.
Expire/Embargo/Revoke	Expire access after a set period of time, or number of uses. Prevent access before a certain date/time (embargo), or remotely revoke access after document has been delivered.
Device Limits and Changes	Enable users to authenticate from specific devices. Limit the total number of additional devices, or grant access to a new device remotely.
Print Restrictions	Prevent printing altogether, or limit to a number of print operations, or for a period of time. Limit printing to certain page ranges, or lower resolution of printed version. Specify certain printers.
Watermarking	Overlay a watermark to the digital and/or printed version of a document. Set custom text for your watermark and/or include data about the user, time, date, IP address and device used.
Offline Viewing	Enable users to view documents when they are disconnected from the Internet (e.g. on a plane). Limit duration of offline viewing privileges. Specify devices for offline viewing.
Usage Tracking	Gather data about who is viewing and printing documents, their location via IP address, their specific device, etc. Track the flow of printed documents with watermarked user data.

These access control features impact what an end-user can do with your documents. For a reference to the features you'll want to look for on the admin side of DRM solutions, please see **Appendix B, Feature Checklist.**

The Outlook for DRM

The need for placing security on digital documents is not as controversial as it once was. Consumers have become accustomed to paying for valuable content, and understand the role of encryption in protecting both their and the publisher's investment. In the corporate world, executives are increasingly aware of their duty to protect confidential information, as well as to safeguard the privacy of their customer data.

What used to be a decision of whether or not to deploy document security has become a decision about *how* to implement it. We hope this guide has been helpful in illustrating some of the different ways you can make DRM work for you and your end-users. In Appendix A, we provide a list of questions to ask yourself and your team members about your goals and requirements. Appendix B is a "cheat sheet" of features you can use to evaluate and compare DRM solutions.

If you'd like to delve further into any of the topics in this guide, we invite you to subscribe to our blog at https://www.fileopen.com/blog. There you'll find an archive going back several years, including how-to guides, DRM news and industry perspectives.

About the Author

Diana Holm is co-founder of FileOpen Systems (www.fileopen.com), a software developer of document rights management systems since 1997. As the marketing lead and blogger-in-chief, she has been following the digital publishing industry from its earliest days, through the growth of the Web and now the mobile revolution. Prior to founding FileOpen Systems, Diana wrote the first book on publishing online with PDF, *Web Publishing with Adobe Acrobat and PDF* (Wiley, 1996). A graduate of Columbia University in New York, Diana now lives with her husband and two children in Santa Cruz, CA.

Appendix A: Questions to Ask Your Team

1. What is our business goal for implementing DRM? (Mitigating risk? Protecting revenues? Or simply gathering data/analytics?)

2. How do our users currently view our documents? In widely available applications such as Adobe Acrobat?

3. Would our users tolerate downloading a plug-in or separate viewer to access our documents?

4. Do our users primarily access our documents at work on a desktop, or also on mobile and tablet on the road or from home?

5. How do we manage user identities? Are they in a database that we control? Can our new DRM solution integrate with it?

6. How will we authenticate users? Have we already issued user logins that we could use to authenticate access to documents?

7. Are we using a cloud service to store our documents? What is our comfort level with uploading unencrypted copies of our documents to a cloud server?

8. Do we have the resources and IT expertise to set up a server on our premises, to ensure the highest level of security?

9. What other systems does our DRM solution need to work with? (E-commerce? Active Directory? Delivery workflows?)

10. Where do we want to land on the Security - Usability Spectrum? Do we need to implement features that require end-users to download client software?

11. Is it important to use to adopt a DRM system based on open standards? What would be the impact of choosing a top-to-bottom proprietary solution?

12. Which access controls will we use, and why? (Expiration? Printing restrictions? Watermarking?)

13. Do we want to be able to revoke access after a document has been delivered to the user?

14. What would we like to learn about our users from data about their behavior using our documents? How can we use that knowledge to improve our business?

15. Does the DRM vendor own or license the security technology they are using? If a license, from whom? Does the vendor develop their own core technology or act as a service provider?

16. What would be the potential cost/risk to our business of not implementing any DRM?

Appendix B: DRM Feature Checklist

DRM Feature	Yes/No
On-premise encryption option	
Web-based admin dashboard	
Active Directory integration	
Import/interface with identity management system	
APIs for custom integrations	
Support multiple devices per user	
Support smartphones and tablets (iOS, Android)	
Support common file formats in native viewers (PDF in Adobe Reader, .DOCx in MS Word)	
Support viewing in Web browser	
Authenticate via existing login	
Authenticate via device	
Authenticate via domain	
Authenticate via session cookie	
Expire access (by absolute date/time, time lapsed, # uses)	
Embargo access before a date/time	
Revoke access, even after delivery	
Restrict printing (time duration, # of prints, page range)	
Watermarking (static, dynamic, digital/print versions)	
Block screenshots (system, 3rd party tools)	
Block editing of text	
Enable offline permissions	
Enable search within encrypted documents	