Persistent Information Security – Beyond the e-Commerce Threat Model

Merv Matson
President
RightsX Inc.
Calgary, AB Canada
(403) 617-0703
MatsonM@RightsMarket.com

Mihaela Ulieru
Canada Research Chair
Faculty of Computer Science, Univ. of New Brunswick
Fredericton, NB Canada
(506) 458-7277
Ulieru@UNB.ca

ABSTRACT

This paper introduces a new class of information security solution. The core technology, to authorize and track the use of digital files, was originally developed in eCommerce applications, there known as Digital Rights Management (DRM). In applications to non-commercial confidential records, such as health and safety documents, we call the solution "Persistent Information Security". We distinguish it from DRM because the threat models of the fields of application differ significantly. An implementation, RightsEnforcer, is described to clarify some concepts of operation. A simple model for a cost-benefit study of deploying a security technology is suggested and illustrated.

Categories and Subject Descriptors

H.4.3 [Information Systems Applications]: Communications Applications---Electronic mail; E.3 [Data Encryption]

General Terms

Security

Keywords

Persistent Information Security, document security, threat model, Digital Rights Management, DRM, risk remediation, insecurity expense

1. Introduction: Digital Mobility – Blessing and Curse

Our new digital communications networks are an enormously valuable facility for copying and distributing information as digital files. Most conspicuous of all is the Internet, publicly accessible worldwide, delivering huge benefits referred to as the 'digital advantage'. However, together with these benefits comes a series of bad side effects which offset the advantages in this blessing/curse duality.

Blessing – communication efficiency and reach, the 'digital advantage' of computer processed and distributed information, including speed, accuracy, cost, search capability, and valuable new services – consider eMail, the Web and search engines

Curse – invasion of privacy, broken confidentiality and loss of trust, cyber warfare and terrorism, loss of intellectual work, theft of information, identity and money; the same network that lets us reach out lets the enemy reach in

Persistent Information Security, a class of file security, promises to help lift the curse for confidential records.

2. Persistent Information Security and DRM

Digital Rights Management (DRM) is an information protection and control technology developed primarily for eCommerce applications, particularly eBooks and eMusic [http://www.publishers.org/press/pdf/DRMExecutiveSummary.pdf]. DRM is built on a technology to authorize and track the use of digital files, wherever they are used. It is important to distinguish DRM from Persistent Information Security, sometimes referred to as Enterprise DRM, which is built on the same core technology (Figure 1).

The distinction is important because DRM and Persistent Information Security are deployed in different applications and different threat models apply. DRM is used to protect commercial properties (eg eMusic and eBooks) that are made generally available, typically through an eCommerce Web store. Persistent Information Security is applied to protect confidential records (eg health records) which are to be made available to the right person at the right time, only.

The labels can be confusing. Often both the core technology and an eCommerce system built upon it are called 'DRM'. What we call 'Persistent Information Security' is often called Enterprise DRM. But DRM is associated with the eCommerce applications by long common use. We must be careful to avoid indiscriminately populating the confidential records threat model with all elements of the eCommerce threat model.

ICEC'06, August 14–16, 2006, Fredericton, Canada.

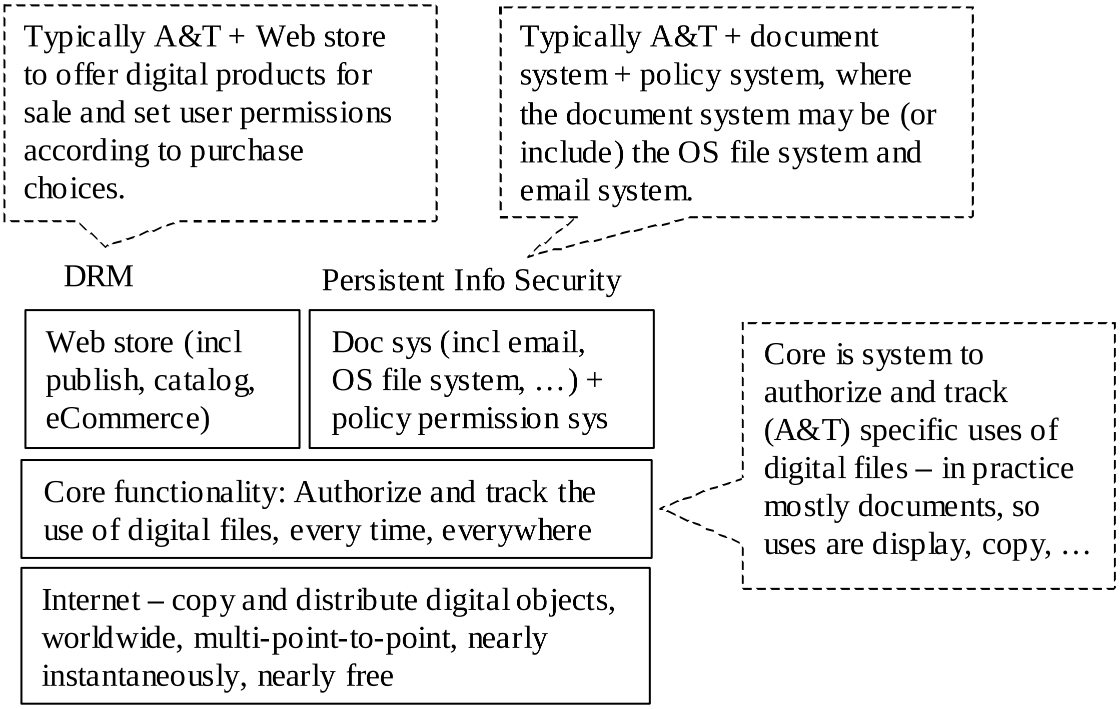

Typically A&T + Web store to offer digital products for sale and set user permissions according to purchase choices.

Typically A&T + document system + policy system, where the document system may be (or include) the OS file system and email system.

DRM

Persistent Info Security

Web store (incl publish, catalog, eCommerce)

Doc sys (incl email, OS file system, …) + policy permission sys

Core is system to authorize and track (A&T) specific uses of digital files – in practice mostly documents, so uses are display, copy, …

Core functionality: Authorize and track the use of digital files, every time, everywhere

Internet – copy and distribute digital objects, worldwide, multi-point-to-point, nearly instantaneously, nearly free

Figure 1. Persistent Information Security vs. Digital Rights Management

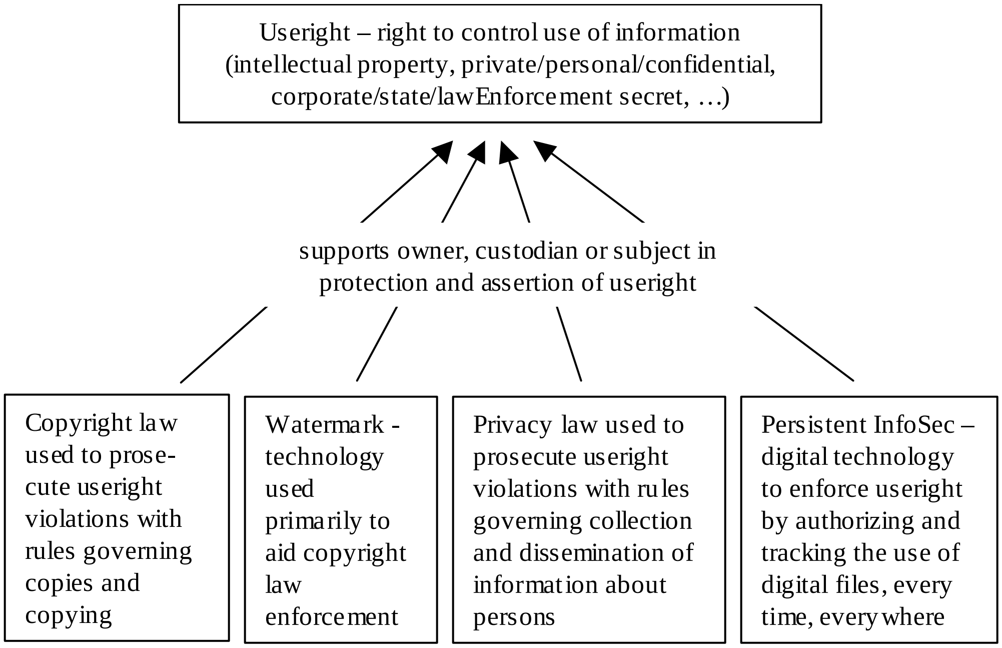

Useright – right to control use of information (intellectual property, private/personal/confidential, corporate/state/lawEnforcement secret, …)

supports owner, custodian or subject in protection and assertion of useright

Copyright law used to prosecute useright violations with rules governing copies and copying

Watermark - technology used primarily to aid copyright law enforcement

Privacy law used to prosecute useright violations with rules governing collection and dissemination of information about persons

Persistent InfoSec – digital technology to enforce useright by authorizing and tracking the use of digital files, every time, everywhere

Figure 2. Useright vs. Copyright

Table 1. Persistent Information Security vs. Copyright

Copyright	Law	Copies	Print	Reactive
Persistent InfoSec	Technology	Uses	Digital	Proactive

The main non-commercial applications for the core technology are:

- Personal information, eg electronic health records, employee records, customer data – the focus of all new privacy legislation [1, 5].
- Non-personal confidential information, eg corporate, R&D, police, military, government – the focus of all hackers [2].

3. Persistent Information Security as Useright Defense

The Internet is an extraordinary mechanism for copying and distributing digital files, nearly instantaneously anywhere in the world, nearly free. (Thus we experience endless email spam.) It's hopeless to depend on copyright (or any law) to protect rights to information in digital form [3]. Instead of trying to control the making and possession of copies we must depend on controlling the *use* of copies. Figure 2 illustrates four 'useright' defenses while Table 1 contrasts two of them: Persistent Information Security and copyright.

The reactive-proactive dissimilarity is significant. To enforce copyright one must catch someone doing something wrong and then employ ponderous legal machinery for redress, if it's worth it and one can afford it. In practice this is so seldom done it's often ineffective for enforcing 'useright', even in the print world.

Persistent Information Security is used to prevent rights violations proactively. This does not mean we should abandon copyright and watermark. The law and other enforcement technologies are still needed for cases where the technology for legitimately enforcing useright is threatened or defeated.

4. Persistent vs Channel and Lock-Unlock Information Security

Figure 3 illustrates the difference between three concepts for securing information files. The file source is shown at top, then being delivered through the network to a user's machine (the oval). Shaded elements are protected. Two of the three are file level technologies.

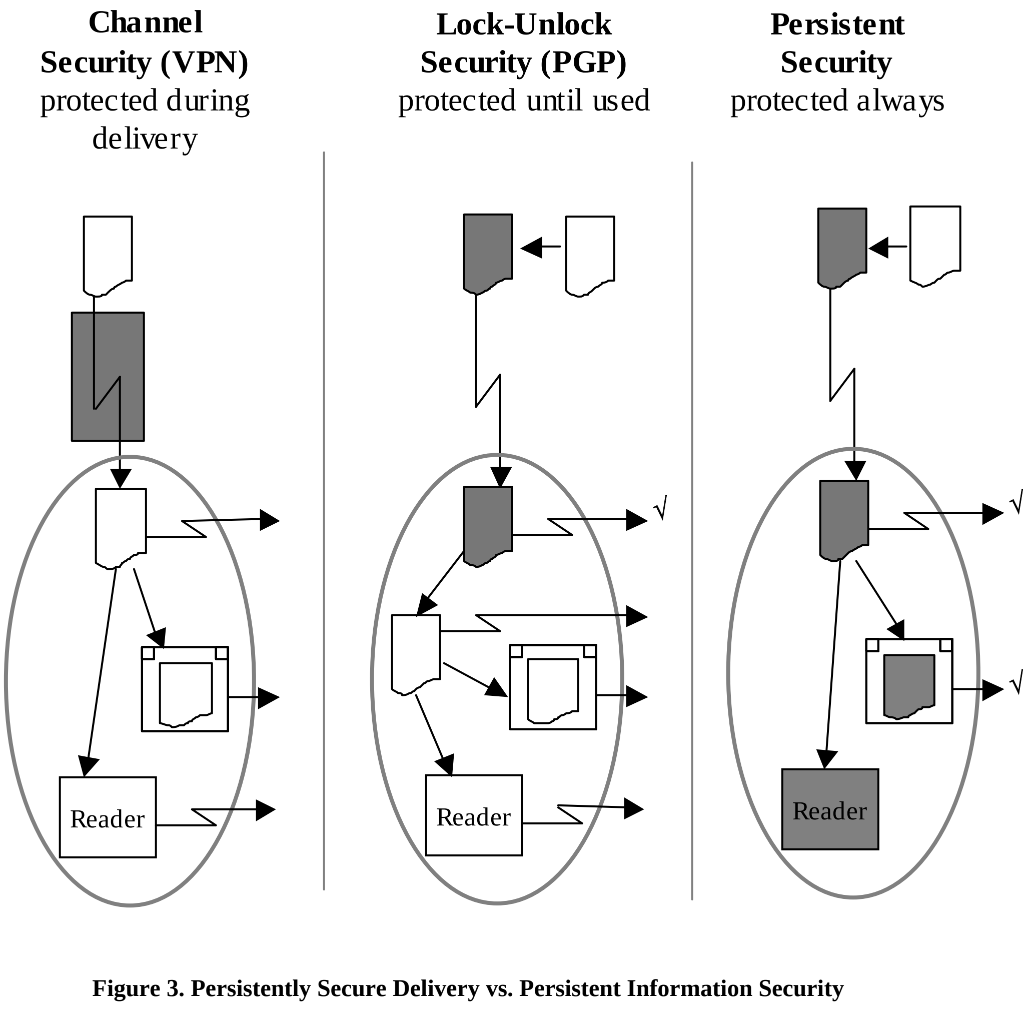
Channel Security (VPN)
protected during delivery
Lock-Unlock Security (PGP)
protected until used
Persistent Security
protected always
Reader
Reader
Reader
√
√
√

Figure 3. Persistently Secure Delivery vs. Persistent Information Security

Channel security, typically implemented in a VPN (Virtual Private Network), protects only during delivery over the network (but it protects more than files, eg commands from a remote user). The file is automatically unprotected when it arrives and can leak in many ways invisible to the information custodian. Lock-unlock security is explicitly applied to the file at source and explicitly removed at destination. Then it can leak as for channel security.

Persistent Information Security differs from *repository/perimeter* (eg firewall) and *delivery* (eg VPN, PGP) security (Figure 3) in that the information stays protected in all domains. Persistent security is always 'ON' once applied at source (unless the receiver is also an information custodian, a co-author for example, and is given the right to make unprotected copy). The protected file can be copied at its destination and redistributed, but any pass-along recipient is subject to the information custodian's control over use. It's no use to have a copy. You must have useright to use it.

From this perspective, Persistent Information Security is a *file level access control method* that enforces the information custodian's terms of use every time, everywhere, even after distribution and use by remote legitimate users.

5. A Threat Model for Confidential Records

5.1 Confidential Records

The 'confidential records' information domain is relevant to organizations and individuals who are custodians of confidential information. Information is confidential for one of three usual reasons.

- It's private. The information is about individuals and they trust (have confidence) the custodian will protect their privacy.
- It's valuable if it's undisclosed. It's (more) valuable if it's only known to the owners, for example hard won evidence about the location of an oil deposit, or customer buying intentions. The owners or benefactors of the intelligence trust the custodian will protect their valuable information.
- It's dangerous in the wrong hands.

Table 2 summarizes several classes of threats that can affect confidential records.

Table 2: Threat Classifications: Increasing Degree of Consequence

Threat class	Description	Example	Freq, Consequence
Legitimate user accident or carelessness	Legitimate user unintentionally exposes confidential information to illegitimate user.	eMail accident (attach wrong file, answer 'reply all' instead of 'reply', wrong addressee); misplace in commonly accessible location; lost laptop or CD	Most frequent of all, threatens to become even more frequent as use of info systems grows. Consequence: sometimes, but not too often, very bad.
Legitimate user mischief	Legitimate user of information system accesses information without need to know.	Because he can and he is curious and naughty	Quite frequent. Consequence: usually not too bad
Legitimate user attack	Insider accesses info with malicious intent.	Disgruntled associate embarrasses organization with info leak. Legitimate user recruited by criminals.	Infrequent, but criminal activity and espionage growing. Consequence: very bad
Illegitimate user attack	Hacker defeats perimeter defenses of server or desktop (if any on desktop). Hacker steals identity of legit user. Invader gains physical access to system.	Hacker gains interactive control of machine over Net. Via social engineering, or hacking. Social engineering or theft of laptop or CD.	Infrequent, but ruthless targeting. Consequences: very bad, perhaps worst of all, since intent is to harm and attacker has attack resources.

5.2 eCommerce

The threat model for the digital properties of eCommerce is very different. It's relevant to those with information to sell, including entertainment products in the form of information (eg eMusic, eBooks). It's not confidential. The owner wants everyone in the world to have the information, but pay for it. Even though many of the threat classes (eg Legitimate user attack) and defenses (eg firewall) are the same as for the confidential information threat model there are essential differences, as illustrated in Table 3 – contrast with confidential records and the comparable entry in Table 2.

In the case of confidential records, given that the community of legitimate users is closed and controlled by employment or other relationship one can consider that 'the enemy is at the gates'. There is rarely any unchecked malice inside that community. A malicious outsider can usually get in only by stealing an identity or co-opting an insider. In eCommerce anyone in the community of malicious hackers can become a legitimate user ('the enemy is

Table 3. Legitimate User Attack for eCommerce

Threat class	Description	Example	Freq, Consequence
Legitimate user attack	A talented hacker (or several of them) picks an eCommerce target and immediately becomes a legit user.	Hacker buys an eBook, or picks up the loss leader for free, thus becoming a legitimate user of the DRM system.	Very frequent. Consequence: very bad for business

among us' - on the consumer end of operations). It suffices that one legitimately acquires one eBook or eSong and with it all of the end user software for analysis and experimentation. That's a lot easier to hack than just working with a copy of the eBook alone.

6. Risk and Remediation – The Burden of Security

In today's digital world absolute security is impossible to achieve (unless one chooses to not use digital equipment, but then they do not live in the 'digital world' anymore…). Therefore a risk management approach to security appears the only viable solution. To strengthen the 'security belt' it is essential to minimize risk induced by long exposure of confidential information in a vulnerable state or place. To minimize the cost one has to balance the need for security with the difficulty to access the confidential data by legitimate users. As illustrated in Figure 4, various security measures lead to different trade offs, as follows:

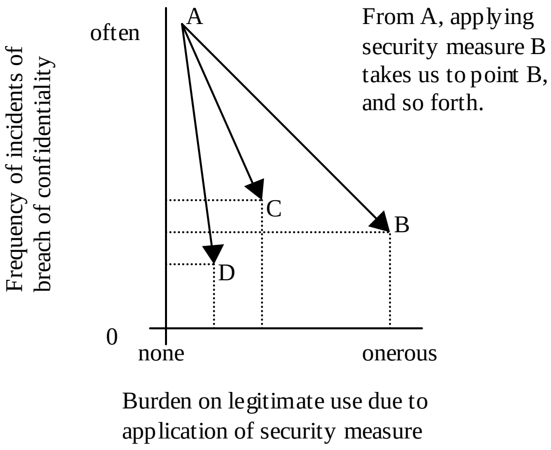

Figure 4. Mitigating the 'Burden of Security'

A – No security measure is applied, therefore there is no burden but too many unfavorable incidents occur.

B – After applying security measure B we have a huge reduction of incidents, but at a huge cost.

C – Security measure C is not as good as B in reducing incidents, but it's a lot less burdensome, so may be the more acceptable of the two.

D – Of course this is the security measure we really want: effective and not too onerous.

It is important to recognize that sometimes simple-minded criteria around security may lead to the wrong choice. In Figure 4 solution C is preferable to B for all but the most secret or valuable information, given that the increase in frequency of incidents will be outweighed by the huge reduction in usability cost for the legitimate users. Therefore a thorough analysis is worth doing when deciding to settle on a solution.

To help decide between candidate solutions we propose a value function, $V(S_i)$ to be computed for each solution – defined as the benefit $B(S_i)$ minus the cost $C(S_i)$ of *using* solution S_i without acquisition or deployment costs, as per relation (1) below. Based on this selection criterion the best solution is the one with the highest value, but none would be chosen if all were negative.

$V(S_i) = B(S_i) – C(S_i)$, where (1)

- $V(S_i)$ is the value of using solution S_i,

- S_i is an element of the set S of n candidate security solutions,

- $B(S_i)$ is the "insecurity expense", the value of using S_i, quantified as probabilistic avoidance of expensive incidents due to insecurity,

- $C(S_i)$ is the "use cost", the cost of using S_i, quantified as the burden (increase in work) on legitimate use due to S_i

$B(S_i)$ is directly proportional to the effectiveness of security measure Si, ie more effective implies greater benefit and therefore more value.

$C(S_i)$ is directly proportional to the burden imposed on the users of Si, ie more burden implies higher cost and therefore less value.

Based on definition (1) above we claim that its low burden and high benefit will make Persistent Information Security the most valuable solution. In the sequel we will demonstrate this using as example RightsEnforcer, a specific Persistent Information Security system developed by RightsMarket [http://www.RightsMarket.com/].

7. RightsEnforcer – A Solution for Persistent Information Security

RightsEnforcer is a suite of software modules designed to integrate and/or interoperate with existing systems to protect documents every time use is attempted. Existing security policies, user directories and content management systems are all utilized by or with RightsEnforcer to provide "persistent protection" while minimizing the need for training and administrative overhead.

To illustrate RightsEnforcer's client-server architecture and communication between the components let's consider an e-Health application (Figure 5) and follow the messages illustrated by solid arrows (dashed arrows depict document flow).

1. **Permissions** A document file is extracted or generated from the Health Information System. The RightsClient (RightsEnforcer client) is used to 'wrap' (encrypt, attach identifying metadata) the document and set permissions for use. Permissions (or terms of use) can

be set explicitly for individuals (peer permissions) or the document can be given a type so policy can express terms of use [4]. The file is then distributed, by email, download, CD, … RightsEnforcer does not control distribution nor take possession of the documents, protected or unprotected. It interoperates with email and file systems, including the Windows file system.

2. **Permissions Query** A user attempts to use the file. The file is unusable without the assistance of the RightsClient which alone can decrypt it. RightsClient will only work with applications that have been made 'trusted' to enforce the information custodian's terms of use (eg display and print document from now for one week). RightsClient authenticates the user and seeks an answer to the query "What rights does this user have to use this document?" If the user is online to the RightsServer (RightsEnforcer server) then the query will be directed there, even if offline use is permitted. If the user is offline, and offline use has been permitted, RightsClient may be able to answer the query from secure local store.

3. **Permission Answer** The user's permitted operations (typically for documents display, print, clear copy-paste, clear file copy) govern what the user can do with the document.

4. **Audit track** An audit trail is built of permission responses, as well as many other operations by end users and system administrators. In case of offline use the track is stored locally and opportunistically uploaded to the server.

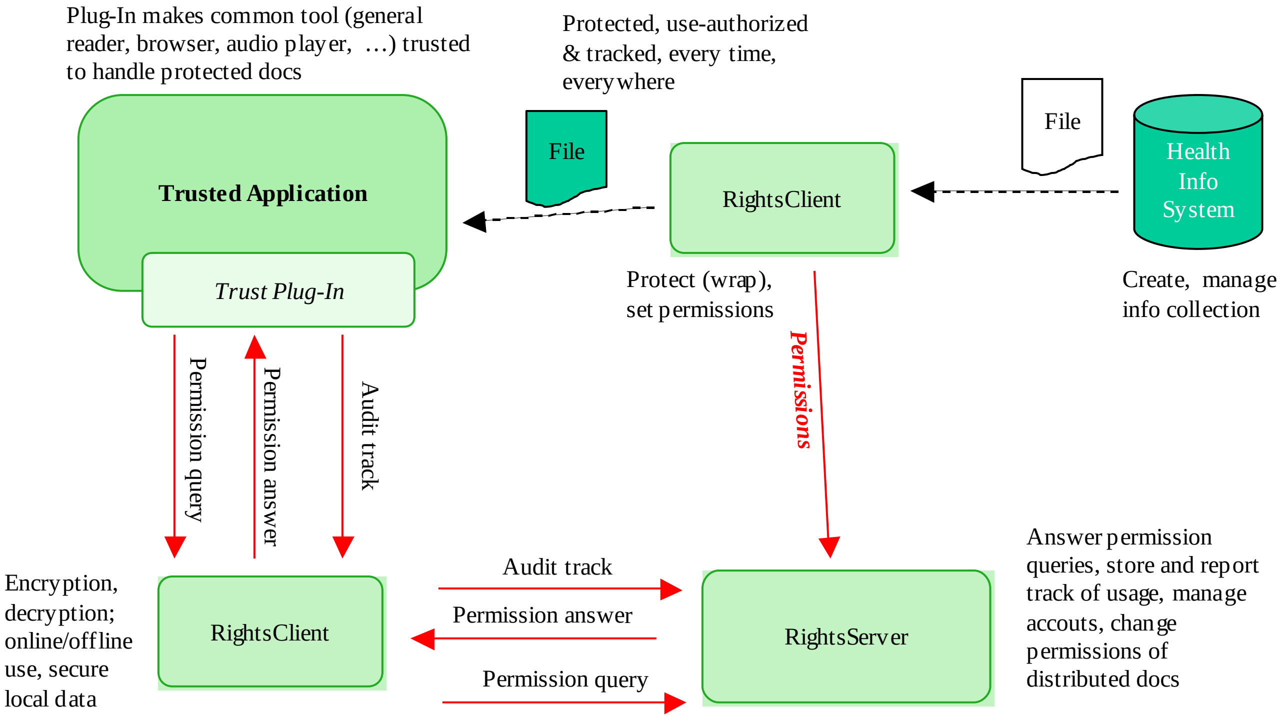
Plug-In makes common tool (general reader, browser, audio player, …) trusted to handle protected docs
Protected, use-authorized & tracked, every time, everywhere
Trusted Application
Trust Plug-In
File
RightsClient
File
Health Info System
Protect (wrap), set permissions
Create, manage info collection
Permission query
Permission answer
Audit track
Permissions
Encryption, decryption; online/offline use, secure local data
RightsClient
Audit track
Permission answer
Permission query
RightsServer
Answer permission queries, store and report track of usage, manage accouts, change permissions of distributed docs

Figure 5. RightsEnforcer Process Flow

8. The Value of Deploying Persistent Information Security

Consider the benefit-cost of RightsEnforcer (RE), according to the definitions from Section 6:

$$V(RE) = B(RE) - C(RE), \text{ where} \tag{2}$$

- V(RE) is the value of using solution RE,

- B(RE) is the "insecurity expense", the value of using RE, quantified as probabilistic avoidance of expensive incidents due to insecurity,

- C(RE) is the "use cost", the cost of using RE, quantified as the burden (increase in work) on legitimate use due to RE

Deploying RightsEnforcer avoids many exposures and therefore many incidents. The benefit-cost analysis depends on access policy, so for example let's consider two cases, internal and external control, and plausible policy.

8.1 Internal Control

- **Policy**: All sensitive documents and emails are protected. For all but the most sensitive documents, all internal users are permitted to display and print. Only creators and editors can copy-paste and unwrap.

- **Benefit**: It’s vastly more difficult for an insider to make collections of unprotected digital documents. Questionable use can be questioned so it is inhibited. An outsider breaking in will usually only get protected documents. If the hacker has stolen Joe’s identity he will only be able to clear copy the documents that Joe can clear copy, and that activity is noticeable, by Joe for one. If Joe’s identity is compromised there is recourse: disable that identity (account) and give Joe another, thus cutting off the hacker from the protected documents.
- **Benefit Quantification**: Picking some numbers for illustration,
    - Exposure incidents: 1,000,000
    - Exploited exposure incident rate: 1/1000
    - Exploited exposures: 1000
    - Frequency distribution of expense: 1/1000 – 100k$, 10/1000 – 10k$
    - Total expense of insecurity: 1 x 100k$ + 10 x 10k$ = 200k$
- **Cost**: To avoid exposure incidents the information custodian needs to authenticate and wrap and the information consumer needs to authenticate to RightsEnforcer. Also, occasionally, because clear copy is inaccessible legitimate work is hindered, ie the policy is not perfect.
- **Cost Quantification**: Making some assumptions about average cost of labor and picking some numbers for illustration (weakly related to numbers above; the analysis is just missing),
    - Wrap instances: 10,000
    - Cost to wrap: 1$
    - Access instances: 100,000
    - Additional cost to access: 0.25$
    - Total cost of usage: 10,000 x 1$ + 100,000 x 0.25 = 50k$

8.2 External Control

- **Policy**: All sensitive documents and emails are protected. By default external users get only display permission, but if it’s reasonable, print and even copy permission can be granted.
- **Benefit**: Documents are protected and rarely exist in an unprotected state. Exposures are greatly reduced, even if the outside recipient is careless or gets hacked. It’s vastly more difficult for an invading hacker to find anything usable. Incorrect documents can be killed and replaced. Users can be cut off completely or from access to selected documents.
- **Benefit Quantification**: Picking some numbers for illustration,
    - Exposure incidents: 10,000
    - Exploited exposure incident rate: 1/100
    - Exploited exposures: 100
    - Frequency distribution of expense: 1/100 – 1000k$, 10/100 – 100k$
    - Total expense of insecurity: 1 x 1000k$ + 10 x 100k$ = 1100k$
- **Cost**: To avoid exposure incidents the information custodian needs to authenticate and wrap and the information consumer needs to authenticate to RightsEnforcer. The external cost to use is not an expense to the organization deploying the security.
- **Cost Quantification**: Making some assumptions about average cost of labor and picking some numbers for illustration (weakly related to numbers above; the analysis is just missing),
    - Wrap instances: 1,000
    - Cost to wrap: 1$
    - Total cost of usage: 1000 x 1$ = 1k$

9. Conclusions

Persistent Information Security protects information every time, everywhere, not just behind the firewall or in the VPN tunnel. The information owner or custodian controls its use even if a legitimate end user is careless or malicious, even if the public Internet (eMail, Web download) is used to distribute information. An essential technology in the important field of information security, Persistent Information Security protects uniquely at the end points of legitimate use where most leaks occur. It costs less than other technologies and non-technological protections it displaces. Because it uniquely plugs the end-point gaps, it permits the deployment of systems that otherwise would be blocked by privacy or risk restrictions.

Along with perimeter defense (eg firewall) and malicious code defense (eg anti-virus), Persistent Information Security is deployed from data centers to websites to home computers. It may supersede existing network transmission defenses (eg VPN, PGP). It is or can be integrated with most information management systems (eg clinical health and human resource systems) and information distribution tools (eg eMail).

In deploying Persistent Information Security one must distinguish that the ‘enemy’ is not (as often) inside the gates as in the case of ordinary eCommerce applications.

10. REFERENCES

[1] Government of Canada. *PIPEDA - Personal Information Protection and Electronic Documents Act* http://www.privcom.gc.ca /legislation/02_06_01_01_e.asp,

[2] CSI/FBI. *2005 CSI/FBI Computer Crime and Security Survey*. http://www.gocsi.com/forms/fbi/csi_fbi_survey.jhtml. (free download)

[3] Matson, M. *Digital Property Protection – Useright not Copyright*, IEEE Canadian Review, 1999 Summer

[4] Matson M. *A policy engine for granting access to persistently secure EHRs*. eHealth 2002 Conference, Vancouver, Canada, 2002-Apr-21

[5] United States Government. HIPAA - Health Insurance Portability and Accountability Act. http://www.hhs.gov/ocr/hipaa/privacy.html