

The Security Newsletter

In this issue

Be our Guest 2

The news

- Researchers uncover detectable chips 2
- New Jersey voting machine accuracy 3
- Cracking CSM encryption 3
- Physical access attacks with Firewire 3
- AES is promised, XOR is implemented: trust no one! 4

Are you sure it is gone? 4

Captcha 5

Freenet 5

Published Quarterly By:
Corporate Research - part of the
Licensing, Research & Innovation Division

Technical Editor:
Eric Diehl

Editors:
Sharon Ayalde
Nicholas de Wolff
Natalie Hamrick

Contributors:
Davide Alessio
Michael Arnold
Patrice Auffret
Jeffrey Bloom
Olivier Courtay
Eric Diehl
Marc Éluard
Ulrich Gries
Yves Maetz
Michel Morvan
Dekun Zou

VP and Head of Corporate Research:
Gary Donnan

LR&I Head:
Beatrix de Russé

Email and to subscribe:
security.newsletter@thomson.net

Copyright Thomson 2008

INTRODUCTION



This quarter, the battlefield of copyright infringements was extremely active. The RIAA changed its strategy by increasing the settlement fees. The later in the litigation process, the higher the requested fees will be. So, the RIAA expects to quickly suppress litigations. In the Viacom v. YouTube litigation, the judge ordered YouTube to make available user identities to Viacom as well as the list of videos said users watched. The French "Création et Internet" law, also called the HADOPI or Olivennes law, is soon to be approved by legislators, and French warez/release teams cinefox and CaRNAGE were arrested.

Nevertheless, two other pieces of news interested me even more: the death of TorrentSpy, and the rise of mininova. TorrentSpy was one of the biggest tracker sites. At the end of March, it stopped working. Under strong pressure from the studios, it decided to cease its activities, but that wasn't to be the end of the story. In May, a California federal judge ordered TorrentSpy to pay \$111 million (72 million euros) to the MPAA. This high penalty is mostly due to the assertion that TorrentSpy destroyed evidence. TorrentSpy refused to provide information about its "customers" and destroyed the corresponding data. This is an interesting parallel with the Viacom v. YouTube case. Of course, TorrentSpy will not be able to pay. The aim of the MPAA was to send a strong warning to other tracker sites.

The TorrentSPy story did not scare every tracker site. A few days later, mininova, another famous tracker site, announced that it had passed the symbolic threshold of 5,000,000,000 downloaded torrents. The distribution of the type of downloaded content is interesting - 39% are TV series, 22% are movies, with only 20% for music.

Who does the closing of TorrentSpy benefit the most? Is it content owners or illegal file sharers? If the closing of TorrentSpy would effectively deter users, then content owners would clearly be the winners. If the users of TorrentSpy will move to other trackers, then the results would be more seeders and leeches for the same content - in other words, better service for P2P.

In this issue, we will present what we see as the next threat in P2P technology: anonymity of peers. The battle is not over. The war has certainly not yet been won.



Eric Diehl
Domain Director, Security

Be our Guest: Ton Kalker



Thomson: How did you become involved with security?

TK: In 1992, I started to work on MPEG compression at Philips Labs. Then, in 1996, I heard about a cool new technology: watermarking, and I became interested in multimedia security.

Thomson: What are your current topics of investigation?

TK: I have several areas of interest. Of course, I am still involved in signal processing. The second area is DRM interoperability. Currently, DRMs are operating in vertical silos. With the wide spread of electronic distribution, consumers will suffer more and more from these isolations. For four years, I have participated in the design of CORAL's interoperability architecture. Also, I am currently involved in DLNA CPS and OpenMarket, and am focusing on interoperability and increasing the quality of the consumer experience.

Thomson: Tell us a little bit more about OpenMarket.

TK: OpenMarket is an initiative driven by studios, technology solution providers, content distributors and others. The big idea is to introduce the download equivalent of the DVD, aiming to enable the "buy once, use anywhere" experience. OpenMarket will enable the consumer to use a consistent domain of playback/recording devices, independent of the retailer. Any OpenMarket branded content will be allowed to play on any of your devices in your domain. It will also provide the infrastructure that enables the receipt of the right content (in terms of resolution, format, and protection) for your devices. OpenMarket will deploy a domain model that provides sufficient flexibility for the consumer and at the same time prevents infinite dilution.

Thomson: Many people claim that DRM is dead. What is your opinion?

TK: In some cases, it is true. Music, for example, has two potential business models: sell-through and subscription based. Most likely, music for sell-through will be DRM-free. However, if the subscription model ever proves to be viable, then DRM will be needed.

Video has a different usage model than music. Music is mostly consumed as a background experience, with a high repetition rate for any given item. Video consumption is completely the opposite: it requires the undivided attention of the viewer, and typically,

most items will be watched only a few times. As a result, the economic value of movies is higher than music, both per usage and per item. I therefore expect that DRM for movies will continue to be widely deployed, not only for rental but also for sell-through. Interoperability is currently the biggest issue for DRM acceptance. Consumers expect greater flexibility, as they want to consume content anytime and anywhere. This is where initiatives such as Coral, DLNA CPS, and OpenMarket play a role.

Thomson: What are the current trends in DRM?

TK: Identification technologies attempt to limit illegal distribution. It is not yet clear if the best solution is watermarking or robust feature extraction. Many critics claim that both solutions are too complex and expensive. Nevertheless, there is a strong push in that direction.

Thomson: What do you think are the future hot topics?

TK: Cost of travel is currently skyrocketing. There is an increasing need to communicate through electronic means. Today, solutions are not natural. The next generation of communication tools will remove physical barriers by becoming more natural. They will require new sensor networks, new coding schemes, network technology, and new rendering techniques. This is one of the main topics of my current research.

In the security field, identification technologies will become the hot topic: identification of content and people. Biometrics has to tackle many issues, among which privacy is not the least.

T. KALKER (IEEE fellow, HP Labs)
Interview by E. DIEHL

The News

Researchers uncover undetectable chips



The classical approach to hacking computer systems is to find bugs in computer software that give unauthorized access. A new way to hack the microprocessor.

This year, at the "Usenix Workshop on Large-Scale Exploits and Emergent Threats", researchers from the University of Illinois demonstrated how they altered the behavior of a computer chip to grant attackers back-door access to a computer [1]. It would take a large amount of work for this attack to succeed in the real world, but it would be virtually undetectable.

To launch its attack, the team modified the design of an open source processor and programmed it in a programmable chip (an FPGA). The modifications bypass some memory access protection and load malicious firmware in the processor. Researchers altered only a tiny fraction of the original processor circuits.

Their demonstration consists of a development board with their modified processor running the Linux operating system. They implement attacks such as logging into the machine as legitimate users (including the root) or stealing passwords and sending them over the network.

"This is like the ultimate back door," said Samuel King, one of the researchers. "There were no software bugs exploited." The researchers are now working on tools that could help detect such a malicious processor. This detection will be extremely difficult.

> M. ÉLUARD

New Jersey voting machine accuracy

On February 5, the state of New Jersey held its presidential primaries. In seven different counties throughout the state, there were more than 35 cases where the Sequoia AVC Advantage voting machines showed discrepancies on the summary tapes. While the discrepancies were small (one or two votes here or there), they are significant in that each tape represents the activity of a single machine. Discrepancies here imply a software or hardware bug in the design, thus raising questions as to the trustworthiness of the whole system.



The counties tried to start an investigation to determine the cause of the errors, but Sequoia threatened legal action, claiming that such an investigation would violate their copyrights and expose their proprietary designs. This raises a bigger concern: How can a democracy trust its most valued right - the right to vote - to a proprietary system that cannot be independently examined?

> J. BLOOM

Cracking GSM encryption

For at least a decade, GSM encryption protocol (A5/1 [2]) was known to be vulnerable, despite GSMA's (GSM Association) advertisements that it was very secure. GSMA underestimated the hazards because known attacks required large resources (computational power, time, money, etc.). Therefore, only a motivated organization would be interested in implementing and performing the attacks.



Recently, two researchers [3] (D. Hulton and S. Miller) proposed a new attack reducing the time consumption and computational power needed; they alleged that this was the first practical attack on GSM encryption protocol. Both researchers claimed that their

work was aimed only at pointing out the weakness of the algorithm (although it had already been demonstrated in the past). Interestingly, two companies that specialize in cell phone encryption employ these researchers.

The attacker needs to collect some information sent in clear text by the Base Station (BTS) as the user's identifier and the phone identification code. To get this information, attackers just need to call the targeted phone and intercept the answer given by the carrier, or just wait for the



user to originate a call while in the same reception area than the eavesdropper. With the data, the attacker can collect all the traffic (encrypted) generated by the targeted user. The second part can be done offline and it is similar to the most (in)famous WEP attack. The attacker looks up collected data in a pre-computed and fixed table and can finally get the clear version of the call.

Hulton and Miller claim [4] that a cheap attack can be achieved in about a half hour. If you are really interested in your neighbor's calls and ready to spend a little more money, you can decrypt the call in about thirty seconds. They expect to begin selling specialized hardware to achieve the attack very soon.

> D. ALESSIO

Physical access attacks with Firewire

A recently published physical attack - the so-called Firewire exploit - exposes Linux, MAC and Windows Operating Systems. The attacker PC, connected to the Firewire port (IEEE 1394) of a target machine, can gain reading and writing access to whole physical memory without any reboot operation. One of the first demonstrations bypassed the Windows Vista password verification. Through DMA access and its addressing scheme, Firewire provides access to the memory, which permanently stores NTLM's Vista password authentication procedure. The attack locates the corresponding portion of software using code signature techniques. Then, it disables the comparison instructions.



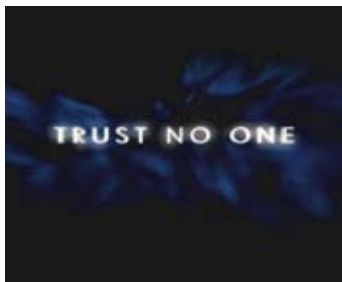
The exploited weakness is due to poor implementation of Firewire drivers that allow DMA access to be wide open. It is possible to forbid this access on Linux. As long as Windows has not proposed corrections, you should disable the automatic connection to Firewire port.

> M. MORVAN

AES is promised, XOR is implemented: trust no one!

The Easy Nova Data Box PRO-25UE-RFID [5] specifications sound promising. This external USB hard disk drive protects data using 128-bit AES hardware-based encryption. A RFID chip that contains the secret key controls its access.

The German publication, "Magazine für Computertechnik" discovered, however, that the promises were not true [6]. Indeed, the encryption uses a simple XOR. The key can be extracted with a simple command line.



The problem comes from the Innmax's controller chip IN7206 [7]. AES is not used for data encryption, but only to protect the ID of the RFID chip. This chip is used by many external hard disks. The hard disk's datasheet has been updated. It no longer mentions AES

encryption, but refers to "simple encryption." XOR is really too simple! If you want to know if your AES advertised protection is real, check your controller chip. Remember our fourth law: "trust no one" [8].

> Y. MAETZ

Are you sure it is gone?

Every year, the US government declassifies many documents for public access. These documents are redacted to keep sensitive information from public knowledge. For example, the name of an undercover agent is replaced with a black bar to protect this agent and his/her ongoing operations. During rump session of Eurocrypt04, David Naccache disclosed a method proving that improper redacting processes may not completely remove the sensitive information. He used the length of hidden words as an example. In 2005, his work was extended by using more information about the used fonts [9]. At the Information Hiding Workshop this year, a research group in Singapore presented a process to guess/recover a portion of the sensitive information [10] using a different approach.

In many cases, the documents are fed into a scanner and the outputs are compressed document images for electronic storage and archiving. Popular image compression methods such as JPEG and JPEG2000 use quantization in the transform domain. As a result, artifacts like Gibbs-effect are created. This introduces the correlation between pixels. Information in areas where sensitive data resides might spread



into pixels in the surrounding areas. The researchers in Singapore studied this phenomenon for JPEG and JPEG2000 compressed and then redacted the document's images. In cases where there is no recompression after the redacting process, or if the recompression has a higher bit-rate, the compression artifacts in pixels surrounding the redacted sensitive area can be exploited to deduce the removed information. A particular example discussed [10] is one in which the sensitive information contains a limited number of choices, for example, "YES" or "NO" in a questionnaire table (Figure 1). Two pattern images can be generated and filled into the redacted area. Then, they compress the images and compare which pattern generates compression artifacts that are closer to the original artifact. They demonstrated that success rate was higher than just a random guess.

The proposed attack method requires that the attackers have prior knowledge of the redacted sensitive information. In other words, the sensitive information has to be limited to a small set of choices in order for the attacker to build the trial template. Nevertheless, the threat is still valid because the attacker may be able to derive the candidate list from other sources. This article definitely raises alarms, and should compel government agencies to revisit their de-classification process. The authors also provide suggestions to prevent the attack. As simple counter measures, they suggest that a lower bit-rate compression or additive noise would help to completely remove this information so it would not run the risk of leaking.

Personal History Survey

In each of the boxes below, please answer either only YES or NO.

Do you like the interface of this website?	██
Do you like this company?	██
Are you concerned about your reputation?	██
Do you prefer smart wear over casual wear?	██
Do you like to have longer hair?	██
Do you believe in love in first sight?	██
Are you concerned about your weight?	██
Are you concerned about your height?	██
Do you like spicy food?	██
Do you like chocolates?	██

Figure 1: Example of a "black bar" process

> D. ZOU

CAPTCHA

(Completely Automated Public Turing test to tell Computers and Humans Apart)

Have you ever come across web pages with images like this?



This so-called CAPTCHA is currently used by many popular websites to prevent automated access to services by bots. In order to gain access, the user has to read the distorted text in the image and type it into a dialog box. This access control method relies on the assumption that the problem of extracting the text in the image is easily solvable by a human, but not a computer program. The term was coined in 2000 by Luis von Ahn and Manuel Blum (see [12]): researchers from Carnegie Mellon University and IBM. They developed this method to verify that the user is in fact a human by requiring the user to read distorted text from a bitmapped image.

Application areas:

- Preventing comment spam in blogs
- Protecting website registration
- Protecting email addresses from scrapers
- Online polls
- Preventing dictionary attacks

The requirements of an ideal CAPTCHA are the following:

- Today's computer programs should not be able to read the text
- Humans should be able to easily read the text
- New challenge-response tests should be generated automatically

The difficulty in developing ideal CAPTCHAs lies in fulfilling the first two contradictory requirements. In order to prevent computer programs from decoding the CAPTCHA, the information must be heavily distorted. This, in turn, leads to poorly decipherable text. The de-facto standard is to use warped text in images, but they are poorly recognizable. Also, visually impaired people cannot read the text at all without additional means and are thus excluded from the service. Nevertheless, sophisticated optical recognition systems already have the ability to extract text from images (example projects are PWNtcha [11], Anti-Gimpy [13] and aiCaptcha [14]).

This area of conflict between generating a secure and human readable CAPTCHA may render CAPTCHAs obsolete in the future once superior pattern matching algorithms are developed [15].

Alternatives

Several techniques are available that prevent the misuse of such services. While they are as effective as CAPTCHAs for verification purposes, they are more accessible for impaired people. Some examples are listed below, but these also have some pros and cons:

- Test for logic e.g. math calculation
- Biometrics technology, such as fingerprint or iris scanning
- Limited-use account, making high-value sites unattractive to bots

In summary, CAPTCHAs are attractive in providing short-term security for registration and access problems, but are likely to be superseded in the long run.

> U. GRIES, M. ARNOLD

FREENET

One major problem that content owners face is Peer-to-Peer (P2P) technologies being used to illegally download copyrighted content. Content owners lobby governments to set up laws to fight against piracy. Two approaches exist: filtering P2P service and stopping the source of illegal content.

The first one is under the responsibility of Internet Service Providers (ISPs). They comply with laws by using filtering technologies based on source/destination ports to throttle communication. Each datagram throughout the Internet contains some information: the IP address of the source and destination, and the service port. By analyzing an exchange between two computers, they can guess information such as the used protocol. Most secure protocols start with clear packets before exchanging encrypted ones. This method is called deep inspection packet. Using this method, an ISP can banish some P2P protocols from its network. The drawback is that the filtering does not distinguish legal from illegal content.

The second approach looks for clients sharing copyrighted content and threatens to shut down their Internet account. Some companies specialize in uncovering P2P usage. They use three steps: The first step browses available content on P2P (or the web) and detects illegal content. The second step, by using tools dedicated to each P2P protocol, collects the IP address of sources. In the third step, ISP delivers the user name of the incriminated IP address to the legal jurisdiction. However, the first step is the most difficult one. There is a huge quantity of information to process due to the multiplicity of information sources: websites (and private websites), IRC, and P2P search engines.

The P2P community already adapted their protocols to these technologies. Current P2P approaches bypass filtering using

randomized ports and encrypt protocol exchanges. However, they have not completely changed the architecture. Traditional P2P protocols were designed for fast and scalable file exchanges. Security was not a concern. They failed to truly make the user's identity (IP address) anonymous.



Some P2P architectures were not designed for speed, but for a political goal. The first and most known of these is Freenet. The initial goal (in 1999) of Freenet was to allow the exchange of information, even in dictatorial regimes. Two main use cases drove Freenet's architecture.

The first use was for sending a message to a group of people without the possibility of anyone stopping its diffusion (ISP or government). The second use was so that somebody can store a secret document and that only this person can retrieve it from any computer connected to the Internet.

The P2P community reuses some of Freenet's components to respond to the new "threat." For example, Omemo is P2P software based on Freenet for file sharing. Omemo claims that the sources of files are anonymous.

To participate in Freenet, you agree to share an amount of your storage space. Freenet content is distributed and replicated on this shared storage space. There is a complex scheme of encryption and routing based on cryptography.

Stored data is encrypted. Thus, nobody can read stored data on his hard disk and know the nature of the data. Every piece of content uses a different encryption key.

To retrieve content, you must have a key: known as a Content Hash Key (CHK), which has three elements: the cryptographic hash of the content (that was a unique identifier on Freenet), the encryption key, and metadata about algorithms.

Following is an example of a CHK:

```
CHK@SVbD9~HM5nzf3AX4yFCBcA4dhNUF5DPJZLL5NX5BrS, bA  
7qLNJR7IXRKn6uS5PAySjIM6azPFvK~18kSi6bQ, AAEA—8
```

This key is mandatory for retrieving content. Only the person who knows this key has access to the clear content.

Documents can also be grouped together, for instance, a website. Thus, the set of corresponding CHK is accessible through a Signed Subspace Key (SSK), which has four parts: hash of the set of documents, the decryption key, a name, and a version of the set.

Following is an example of an SSK:

```
SSK@GB3wuHmt[.]o-eHK35w, c63EzO7u[.]3YDduXD, ,  
AQABAAE / mysite - 4
```

Updateable Subspace Keys (USK) provide an updateable entry point. If you know the USK, you may find the newest web version.

A computer connected to Freenet (or any other P2P platform) is called a node. On Freenet, each node is connected to a limited number of nodes. When a node searches content, the query is transmitted to one (or more) connected nodes. If a neighbor node has the response, it transmits the response to the querying node and it forwards the query to its neighbor. There is a chain of forwarded queries. Each node communicates only with its direct neighbors. The small world theory states that each node can reach any other node.



At each query, the intermediate node keeps track of the data. So, if content is often requested, it will just as often be duplicated. This is an important component to understanding the robustness against censorship. To remove content, its location is found through multiple queries, each generating copies of the content! It is impossible to list all content available in Freenet. With Freenet, the originator of the content is never known.

Freenet has two options. The open option (called OpenNet) shares content between all Freenet users. Nodes can potentially be connected to all other nodes. The closed option (called DarkNet) shares content only within a closed group of people. Two nodes are connected only if both users accepted the peer node. A Darknet should be invisible to the rest of the Internet.

In 2005, Freenet was flawed, and it led an ISP in China to filter Freenet. A deep inspection packet allowed China to identify Freenet protocol. Freenet coped with this flaw. Now, in the Darknet option, they guarantee that there is no longer a flaw of this type. OpenNet is theoretically prone to this threat. It mitigates the threat through steganography to hide some packets of Freenet protocol. A well-known Freenet site is <http://freenet-china.org/> - it was accessible on Freenet through this address: SSK@fjfkHAbxdwMyTMFgtZjcP2ge-AYPAgM/sites/

File sharing is the aim of Freenet. Nevertheless, other services have been built over this architecture, such as mail, newsgroups, or blogs.

Omemo is the newest solution for anonymity in P2P. It is strongly based on Freenet. Omemo eases access to content by indexing every uploaded file. Participating users can browse content using multiple criteria, and they can easily share content. The drawback of Omemo is that it does not protect as well against network filtering. Nevertheless, it is currently robust enough.

A counterstrike to P2P filtering is rising. Illegal copyrighted content distribution platforms are slowly switching to P2P anonymizing technologies. This is just the start of a new battle. P2P users will try to hide their identity.

> O. COURTAY, P. AUFFRET

Where will we be?



* International Workshop on the Arithmetic of Finite Fields (WAIFI 2008), Siena, Italy, July 6-9, 2008

Paper presentation: Fast point multiplication on elliptic curves without pre-computation, by Marc Joye

* Broadband World Forum Asia, Honk Kong, 15-17 July 2008

Hot seat panel: P2P threat or opportunity, chaired by Eric Diehl

* 5th Workshop on Fault Diagnosis and Tolerance in Cryptography (FDTC 2008), Washington DC, USA, August 10, 2008

Paper presentation: On the security of a unified countermeasure, by Marc Joye

* Eighth Smart Card Research and Advanced Application Conference (CARDIS 2008), Egham, UK, September 8-11, 2008

Paper presentation: SmartPro: A smart card-based digital content protection for professional workflow, by Sylvain Lelievre

References

[1] S.T. King et al., "Designing and implementing malicious hardware," San Francisco, USA: 2008; http://www.usenix.org/event/leet08/tech/full_papers/king/king_html/.

[2] "Attacks on A5/1 as used in GSM," Wikipedia; http://en.wikipedia.org/wiki/A5/1#Attacks_on_A5.2F1_as_used_in_GSM.

[3] M. Kassner, "Cracking GSM encryption just got easier," TechRepublic.com; <http://blogs.techrepublic.com/wireless/?p=206>.

[4] D. Hulton, "Intercepting GSM traffic," 2008; <http://www.blackhat.com/presentations/bh-dc-08/Steve-DHulton/Presentation/bh-dc-08-steve-dhulton.pdf>.

[5] "easyNova :: Produkte"; <http://www.easy-nova.de/index.php?siteID=18&productID=28>.

[6] C. Rutten, "Enclosed, but not encrypted," Heise Security, Feb. 2008; <http://www.heise-online.co.uk/security/Enclosed-but-not-encrypted--/features/110136/0>.

[7] "INNMAX Corporation"; <http://www.innmax.com/en/im7206.html>.

[8] THOMSON security labs, "Ten laws of security"; <http://eric-diehl.com/index.php?lang=En&page=lois>.

[9] D. Loprestia and A.L. Spitz, "Information Leakage Through Document Redaction: Attacks and Countermeasures," Document Recognition and Retrieval XII, 2005.

[10] N. Ho and E. Chang, "Residual Information of Redacted Images Hidden in the Compression Artifacts," Santa Barbara, USA: 2008.

[11] "PWNtcha - libcaya"; <http://libcaya.zoy.org/wiki/PWNtcha>.

[12] L.V. Ahn, M. Blum, and J. Langford, "Telling humans and computers apart automatically," Commun. ACM, vol. 47, 2004, pp. 56-60.

[13] G. Mori and J. Malik, "Breaking a Visual CAPTCHA"; <http://www.cs.sfu.ca/~mori/research/gimpy/>.

[14] "brains-N-brawn.com"; <http://www.brains-n-brawn.com/default.aspx?vDir=aicaptcha>.

[15] J. Yan and A.S. El Ahmad, "A Low-cost Attack on a Microsoft CAPTCHA," Apr. 2008; <http://homepages.cs.ncl.ac.uk/jeff.yan/msn.htm>.