Black Market Botnets

Nathan Friess and John Aycock
Department of Computer Science
University of Calgary
2500 University Drive N.W.
Calgary, Alberta, Canada T2N 1N4
{nfriess,aycock}@cpsc.ucalgary.ca

TR 2007-873-25, July 2007

Abstract

Botnets have yet to be exploited to their full potential, because they have yet to take advantage of all the information available to them. The zombie computers that comprise a botnet have access to the private documents of the people that use the computers. A botmaster who controls the botnet can harvest the documents and sell them to third parties, creating a viable – if illegal – online business. We outline motivations for such a business model, as well as the mechanics of a possible implementation. We then present a variety of defenses against this scenario.

1 Introduction

Many details about people are best retrieved at the source: the computer which a person uses to store their information digitally. Users enter information about almost every aspect of their lives. This is particularly true in a business environment, where email is a key form of communication, and internal documentation and reports are created constantly. This leads to a question. If all of this private data is "out there" already, can it be harnessed in such a way that an adversary[1] could use that information? In other words, why would one want to index private data on a global scale, is this even possible, and if so, what can be done to keep digital data saved on a desktop computer truly private?

The number of computers that have been gathered in botnets has grown steadily over the past few years [29]. Although botnets are already being used for identity theft [2] (among other things), larger botnets and the ingenuity of adversaries lead to more sophisticated approaches to harnessing the private information stored on zombie computers. The botmaster, the person who creates and operates the botnet,[2] would not

[1] We use the term "adversary" to generically refer to someone with malicious intent with respect to a targeted person.

[2] We distinguish between the botmaster and adversary here because they are not necessarily the same person.

necessarily want this information, but there are many adversaries who would, and so the motivation for the botmaster is to sell the documents for a profit. Botnets play a key role in this scenario. They cross organizational boundaries, providing access to documents that are otherwise be inaccessible to interested parties.

There are several possible scenarios in which a botmaster could sell documents to adversaries. For the sake of discussion, we ignore the fact that few of these scenarios would be legal:

- A company looking for information about current research projects of their competitors could search for internal documents from competing companies.
- A company could perform "market research" on customers.
- A private investigator could use private documents as another source of information in their investigations.
- Paparazzi could search for information on celebrities.
- Terrorists could search for security weaknesses by looking for classified documentation on a target facility.
- Counter-terrorism agencies would be able to search for intelligence to thwart terrorists' plans.
- An adversary that is planning a targeted electronic attack against an organization (e.g., using social engineering) could start by searching for insider information about the organization, making their attack more convincing.
- The botmaster could start a bidding war between the original owner of the document and an adversary. The document's owner would be extorted into paying the botmaster into keeping the document private. This would work particularly well against large corporations or celebrities.
- Police agencies could use private documents in an attempt to catch people who commit serious crimes, like people who produce child pornography. While the legal ramifications of this would require some consideration, evidence acquired by illicitly-conducted computer searches has been admitted previously [33–35].

As the market for private documents becomes more popular, new uses would undoubtedly emerge, producing niche black markets which are currently unfathomable. Brynjolfsson et al. provide insight into how reducing search costs promotes niche markets [7]. Previously, distributing products and services to these "Long Tail" markets wasn't possible due to the cost of connecting buyers and sellers. As this cost drops, more niche markets can be accommodated. This applies to the situation we describe, where the creation of a black market for private documents yields new opportunities to exploit the technology.

Assume for the moment that a botmaster knows how to identify "interesting" documents (we discuss this in the next section). The basic architecture of a black market botnet system could then be described as illustrated in Figure 1: the botmaster instructs

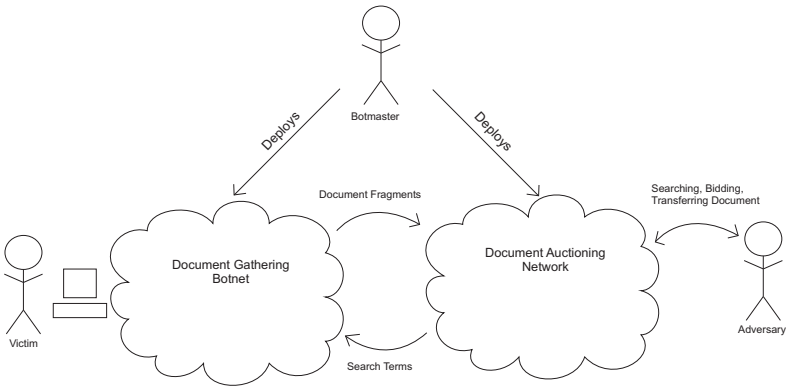

Figure 1: Basic architecture of a black market botnet

their zombie computers to search for interesting documents stored on the compromised computers. When a zombie computer finds a document that looks interesting, it posts a small excerpt from the document to an auctioning site, where adversaries can bid on the documents. The adversary who wins the bid pays the botmaster, who then instructs the zombie computer to send the document to the buyer. Therefore, the black market botnet system is composed of two main components: the zombie computers that gather documents from compromised computers for auction, and the auctioning system itself. The latter part may be hosted on a botnet to reduce the risk of it being shut down, or the botmaster may be able to use an existing auction site like eBay. Both of the two components will be outlined in this paper, as well as the possibility of using a well-known auction site to sell the documents.

It is important to stress that this threat is *not* just another "doom-and-gloom" scenario. Shortly after we completed the first draft of this paper in February 2007, the Gozi Trojan was discovered in the wild. This Trojan steals posted web form data and transmits it to the botmaster's web site; adversaries can search through the captured data and purchase the results [15]. This is a limited form of the more general threat we describe in this paper.

2 Gathering Documents

There are two major problems for a botmaster to solve in terms of gathering documents. First, how can interesting files be discovered automatically? Second, how can adversaries search for these interesting files?

There are various algorithms that could be used for building a searchable index of local files on a compromised computer, the simplest of which would be to find recently edited documents. However, only looking for recently edited documents would miss other potential targets such as the previous year's tax returns. Therefore, more indexing methods must be used. If a botmaster chooses to use existing software, Google Desk-

top already provides indexing and document searching capabilities. A zombie could download and install Google Desktop, let it index all of the documents on the computer, and then direct search queries to Google's software. A more talented botmaster could implement a text mining algorithm, aggregate the results, and use those to assist in searches. For instance, Dörre et al. present an algorithm that can be used for mining data [11]. They use basic pattern recognition and heuristics to find interesting pieces of information from a document. They extract what they call features, or "significant vocabulary items," such as "credit line." The next step is then to cluster or categorize the documents based on what features are found in them. This categorization would then give another parameter which adversaries could search by. Amazon has followed along similar lines with their "Statistically Improbable Phrases" [1]. SIPs are phrases gathered from books or other literary works that are mostly unique to each book. For example, if one particular book mentions "fuzzy bunnies on the beach at sunset" multiple times, but very few other books use this phrase, then this would be a SIP for the book.

To search documents, one option is for the botmaster to use existing peer-to-peer software and searching algorithms (e.g., those used by the Gnutella or FastTrack networks [13]) Indeed, botnets have been known to employ P2P structures already [10]. Adversaries – potential buyers – would use peer-to-peer clients to submit their search requests. The search results would only contain a small, fixed piece of the document, rather than the entire file. This would prevent adversaries from piecing together a document based on repeated search results. Recent peer-to-peer file sharing applications typically include search features to allow users to find files based on attributes like their name, type, and size. While these attributes might prove useful, more complex indexing of local documents will almost certainly be required.

Another search option would be for the botmaster to construct their own web search portal, allowing adversaries to use regular web browsers to search. Obviously the botmaster takes a risk here, because the search portal can be shut down. Spammers and phishers have dealt with this problem for years, using methods like bullet-proof hosting [27], fast flux [26], and levels of redirection [36].

A searchable botnet may provide some capabilities for connecting adversaries with documents, but it would be quite limited in capabilities. In particular, because adversaries would find documents primarily using a traditional keyword-based search engine, the onus is on the adversaries to enter the magical recipe of keywords to find what they are looking for. Unfortunately, it is doubtful that an adversary would find a document that says "our big financial scandal" in it, making keywords of limited use. Moreover, due to the vast numbers of documents that would be available to a large botnet, a search query that includes common words would return far too many documents to easily sift through. Web site search engines mitigate this problem by ranking results based on the popularity of a web site, but that is not an option for black market botnets. Private documents cannot be ranked in terms of popularity. In Section 6, we suggest an alternative approach to gathering documents, using a multi-agent system.

3 Selling Documents

When documents are posted for auction, a botmaster has a number of options as to how and where the listings are posted. The botmaster could create their own auction system, or use one of the existing auctioning systems on the Internet. When creating their own auction system, the botmaster could create a web site, or use a more covert medium such as IRC. IRC is known to be used by fraudsters selling credit cards and other personal information [28]. Another option would be to create a web site to facilitate the auctioning process, and, like creating a search portal, there is a risk that the site will be shut down, depriving the botmaster of income. This is also the approach taken by Gozi [15].

Given potential problems with setting up a custom auction network, a botmaster could instead choose to use an existing auction site. As the largest and most recognized auctioning site, we use eBay as an example. The black market botnet wouldn't be able to post fragments of documents directly on eBay, as it would be trivial for the site's operators to take down such listings. Instead, the listings would need to be obfuscated so that an average viewer would not know the illicit nature of the auction. Steganography has obvious application here. One of the most common places to hide information is in image files, with new techniques being researched in recent years [17, 21]. Many auction sites such as eBay allow for sellers to post images of the item being sold, and this presents a good opportunity to apply steganographic techniques. From a human's perspective, the auctioned items will look perfectly normal. However, anyone who knows what to look for will be able to uncover information about the real document being sold. This is similar to reported criminal activity, where drugs are being sold on eBay using cover items [38].

If a black market botnet were to use images containing hidden information in eBay listings, two main issues must be addressed: how does an adversary find the illicit listings, and how does an adversary extract the hidden information? To answer the first question, a botmaster can provide adversaries with a list of accounts which the items are posted with. This list could be distributed through different kinds of channels, such as direct communication over Internet Relay Chat, or by posting it on a darknet like Freenet. eBay has a search utility where one can limit search results to a specific seller. Adversaries can use this search utility to quickly locate listings which are about private documents. Of course, if the botmaster only has one or two accounts, eBay would eventually find the accounts and shut them down. However, phishers have also dealt with this problem for years. Their solution is to simply move on to other accounts, and provide the new information to their targets (or in this case, the buyers). This is certainly not an ideal situation, but as long as it is "good enough" a botmaster will be able to maintain operations.

The major advantage to this approach for the botmaster is that the listings themselves can be about anything, and in any section of eBay. One time the botmaster could add a listing describing an antique tea pot, and the next time a listing can be about cardboard boxes. The black market botnet could have a canned list of postings, modify them slightly each time, and rotate through the list.

Decoding the steganographically-hidden information could be accomplished if the botmaster provided adversaries with a program able to decode an image. Such pro-

grams are freely-available, mitigating trust issues between botmasters and adversaries. A more elaborate program could even search eBay using their API [12], and automatically decode images in the search results. The hidden bits could be checksummed, to allow the program to distinguish between actual hidden information and random garbage.

Once a listing is found, adversaries would use eBay like in any other auction. The successful bidder pays the botmaster using an payment system such as e-gold or PayPal. Upon receipt of the payment, the botmaster would then provide the full document to the buyer. The document itself could be mailed using an anonymous email account. If the document is small enough, another option would be to encrypt the document and include it in the auction posting, so that only the decryption key needs to be provided to the buyer.

4 Defenses

Several methods can be used to defend against black market botnets; some proactive, others reactive. Most of the defenses are not mutually exclusive, and therefore multiple defenses can (and should) be employed.

Preventing Infection. Preventing computers from becoming part of a botnet, i.e., avoiding infection by Internet worms and other malicious software will be the most effective defense. The scenario presented in this paper is based on the premise that a botmaster will have a large botnet under their control. Without the botnet, this scenario quickly breaks down. Installing the latest patches for operating systems, web browsers, and other frequently-used software will go a long way towards reducing exposure to malicious software. Installing firewalls and anti-virus software and keeping them up to date is also a good way to prevent malware from executing on the target computer. Unfortunately, these defenses are not in themselves complete, because they do not necessarily prevent unknown malware and exploits from being used. Other forms of defense are desirable as well.

Limiting Document Exposure. A cautious user might attempt to hinder a black market botnet by limiting access to private documents. This could be accomplished by moving infrequently-needed documents to offline storage; documents that are saved on a DVD and stored on a bookshelf simply won't be within reach of an adversary. However, if the user does require a document and inserts the DVD into their computer, then it instantly becomes accessible and vulnerable again.

While certainly not a complete defense, limiting document exposure through offline storage would act as part of a defense in depth. The implication is that a black market botnet would not know when interesting documents may become available, and how long such documents will remain available. A black market botnet would therefore need to copy interesting documents to counter this user practice, copying activity that could be used to detect the botnet's presence. That said, this defense takes extra planning and effort to implement, and is not likely to be practical for every-day use.

A related issue is the current trend in retaining documents and other personal information for longer periods of time [5]. Government legislation, such as the Sarbanes-Oxley Act mandating data retention for auditing purposes, only provides more oppor-

tunities for a black market botnet to gain access to private documents [9]. Archived documents must therefore be handled with great care.

Digital Rights Management. Another way to limit access to private documents is to use digital rights management (DRM). Various DRM schemes exist, and in 2003 the World Intellectual Property Organization published a comprehensive study on various DRM technologies [37]. A simple DRM technique, for example, would be to protect documents with a password. When combined with strong encryption, where the password is the encryption key, this scheme limits a black market botnet's access to private documents, much like saving the documents on removable media. However, a similar problem also exists in that documents are immediately accessible to a botnet once decrypted.

Trusted Computing. Software-based digital rights management techniques have been circumvented repeatedly, such as with DVDs [22], iTunes [22], eBooks [25], and more recently, HD-DVDs [19]. In order to combat these threats, DRM is now being integrated directly in the low-level components of computers, resulting in so-called "trusted computing." The approach taken here is that by adding a hardware trusted-computing module to every computer manufactured, it becomes possible to tie documents to specific computers [32].

An ideal trusted computing platform will protect documents right from the point that they are retrieved from disk to the point at which they are displayed on the screen. This means that the disk, operating system, video card, and monitor will all negotiate a encryption scheme such that nowhere along the line can an intruder view the contents of the document [31]. Currently, trusted computing platforms aren't widely available, so this defense against black market botnets is still a work-in-progress.

Use Steganography. Another possible defense against black market botnets would be to use steganography to hide very sensitive documents. Users could hide sensitive financial information on their computer inside a seemingly harmless image of their puppy. In this sense, all of the steganographic techniques that can be used by black market botnets to hide postings on auction sites can be applied to defend against them. The drawbacks to steganography are the limited storage capacity it offers, and extra steps required of the user to hide and retrieve their documents.

Document Fingerprinting. If one assumes that it will be impossible to fully protect every private document on a computer system, then the next best defense is to find out how documents are leaking, and plug those holes after the fact. Reactive defenses may not be an ideal approach, but it is unlikely that every single black market botnet scenario can be predicted and proactively defended against. Fingerprinting is a technique where each copy of a document contains some unique modification (fingerprint) so that the document can be examined later to determine who this copy belonged to [24].

Fingerprinting research is typically aimed towards multimedia content, where content distributors attempt to prevent piracy by linking copies of the multimedia content to specific "owners." A corporation could take a similar approach when releasing documents under a non-disclosure agreement. In the context of this paper, however, a corporation would need to not only fingerprint documents which are distributed to external organizations, but also documents which are distributed within the company. This way, if a document is harvested from an infected computer on the inside and later

appeared in some public forum, it would be possible to trace the document back to the computer that it leaked from.

Follow the Money Trail. The key motivation presented in this paper for a botmaster to gather private documents is the monetary incentive. Thus, money will be trading hands frequently in exchange for documents. If law enforcement agencies stumble across even a handful of buyers, they may be able to trace payments to their destination and catch the botmaster. More proactive law enforcement agencies may even conduct sting operations by purchasing documents themselves for the purposes of following deposits to the botmaster's bank account.

Unfortunately, laundering money is a well-developed process, and it is hard to weigh the possible success of tracing money to the botmaster against the deployment of proactive defenses.

Active Countermeasures. Similar to existing honeypots that are used to track spammers and Internet worms [14], systems can be set up to provide fake documents for a black market botnet to mine. If combined with fingerprinting, the owner of a document honeypot could gain extra insight into how black market botnets work, such as the characteristics that make a document interesting to adversaries. Given this extra knowledge, a large document honeypot with many fake documents could be established in order to decrease the signal-to-noise ratio in the auction. The "good guys" could even bid on the fake documents to throw off black market botnets who learn from previous sales. That said, this "defense" would be yet another arms race which does not address any of the underlying issues involved, and is perhaps best left to people with too much time on their hands.

5 Related Work

Schechter and Smith have discussed how a botmaster could sell access to infected computers [23]. While that may be somewhat profitable, selling the information stored on an infected computer while retaining control over it would be far more profitable. The idea of compromised computers being used to extort payment from their owners has been examined as well. One way of accomplishing this is to encrypt files on the compromised computer, after which the adversary offers to decrypt them for a price [40]. This is not just academic; a number of cases have occurred in the wild [3, 18, 20, 30].

Bond and Danezis [6] argue that people may willingly install and maintain malicious software on their machines, and one of the incentives for this is access to a compromised computer's files. They even suggest a search facility. However, their "Satan" virus extends the invitation only to people in a victim's social network; our scenarios extend far beyond this and are much more plausible.

At the same time, Young is interested in how smarter malware can collaborate to provide a directed attack against the owner of an infected computer [39]. In his example, three pieces of malware work together to extort a stock broker into performing transactions. One copy of the malware infects a computer inside the stock broker's network and finds some sensitive information. The other two copies are outside the stock broker's network, each of which will hold part of the sensitive information from

the stock broker. The malware inside then instructs the stock broker to perform a transaction. If the stock broker refuses, then the malware on the outside combines the pieces of the sensitive information, and posts them to a public forum.

Young's idea is related to this paper in a few different ways. First, Young brings up the idea of broadcasting sensitive information, which is similar to the harvesting and selling of private documents. Second, Young raises the bar on the level of intelligence of malware. In his scenario, the malware doesn't just encrypt a document and display a message indicating who to contact to have the document decrypted, the malware communicates with other systems automatically, reducing the amount of interaction required by the adversary. This paper takes Young's idea to the next level, based on the observation that if three pieces of malware can communicate to accomplish a goal, so can any number of peers in a botnet. More importantly, once a black market botnet is set up, auctioning off private documents is automatic, lucrative, and scales well. Young did not address scalability of extortion attacks in his paper.

Chau et al. have examined the possibility of finding fraudsters on eBay [8]. They were interested in mining various pieces of information from eBay's online records, examining them to spot fraudsters, as well as locate accomplice accounts. Part of their implementation was building a graph of which accounts traded with other accounts; eBay conveniently provides this information through their user feedback system. They were able to identify several accounts which were involved in fraudulent sales, and several other accounts which were presumably used by the fraudsters in order to boost the trustworthiness of the accounts used in the fraudulent sales.

Chau's research raises two interesting points. Like many papers in the area of computer security, results that are published with good intentions can be turned around and used for malicious purposes too. Therefore, if researchers can follow the graph of traders in order to locate accomplice accounts, adversaries in our black market botnet scenario could use the same technique to find other documents being sold. This could be an extension to the implementation outlined previously. A botmaster could open several accounts on an auction site, and post documents using all of the accounts. The botmaster creates some number of transactions between the accounts by selling and then buying their own (fake) items. Then, rather than giving out the complete list of accounts to adversaries, adversaries find the other accounts by looking through the feedback for the known accounts. This provides yet another way for the botmaster to obscure their activities for investigators.

Yu and Chiueh have proposed an interesting approach to digital rights management, where they never give users access to an underlying document, but still allow users to view and update the document. They call their system a "Display-Only File Server" [41]. Their belief is that by limiting the ways in which a user can interact with a document, they can prevent a user from stealing the document with digital means. That is, they acknowledge that a user can still take a photo of their monitor, but the user cannot simply copy the document using the computer.

Yu and Chiueh's implementation utilizes a Windows Terminal Server (WTS) and client application. When a user double clicks on a file in the Windows Shell, they intercept this action and start up their client application. This application then connects to the WTS and opens the document on the server, so that the document is never available directly on the client computer. They even attempt to prevent screen captures by

hooking into Windows events and "clear" a capture as soon as it is created.

Yu and Chiueh's approach to DRM provides protection against the scenario mentioned in this paper. Even if a desktop computer is infected by a black market botnet, the display-only file server would prevent harvesting documents from the infected computer, because the documents simply do not exist there. Using a DOFS will have clear advantages in a corporate environment, but an average home user is unlikely to have a WTS sitting about. Furthermore, as Yu and Chiueh admit in their paper, it is not possible to prevent screen captures completely. As they mention, a determined adversary who cannot use the high-level API calls will simply move to low-level attacks to gain access to the contents of the screen. This does not mean that a DOFS is not useful, because anything that increases the level of difficulty for an adversary is helpful in combating their activities.

6 Future Work

One avenue which we would like to explore is the possibility of a botmaster implementing a multi-agent system to gather documents for sale. As we mentioned previously, simply providing searching capabilities to adversaries is quite limiting, in that there may be a lot of search results to sort through. It will also be difficult for adversaries to gauge their interest in a document based only on a small snippet of text. Therefore, if a botmaster would like to create a viable black market for private documents, it will be necessary to allow adversaries to use better methods to search for the documents. Instead of providing direct searching capabilities to adversaries, a black market botnet could host document-gathering agents. These agents would search for documents on compromised hosts and flag them for the botnet, which would post the candidate documents to the auction system. For example, an agent could search for spreadsheets that contain dollar figures in the order of millions of dollars and have a calculated cell that shows a negative amount. As before, adversaries would bid on the documents based on the fragments posted, and the auction's winner could then claim a complete copy of the document. Some agents could be created by the botmaster, while others could be created by some trusted adversaries, provided that an adversary could demonstrate that their agent wasn't malicious.

In terms of actual implementation of the agents, there are many examples of previous work which could be applied to this situation. Kusumura et al. wade through auction postings using "domain knowledge," and for example, extract features from computer sale postings such as CPU, memory, and disk space from non-uniform descriptions [16]. The notion of domain knowledge is an important point here, because agents looking for private documents could exploit domain knowledge as well. For example, an agent can look for documents from a stock broker by including words that are in the brokerage domain. When looking at a table or spreadsheet of financial information, an agent could extract dollar amounts from the table, to determine whether this document is about a corporation's finances, or grandma's monthly bookkeeping. Each of these features are domain-specific, and agents can be tweaked in one or more domains in order to increase their usefulness in finding saleable documents.

Birukov et al. have examined an agent-based system for assisting users in finding

web sites related to a topic of interest [4]. In their implementation, they created an agent to run on each user's computer, which learns about what sorts of web sites that a user visited while researching a particular topic. The agents then cooperate in order to suggest other web sites to users looking for similar information. A similar level of co-operation would be useful for agents searching for "interesting" (saleable) documents. Once the mining and auctioning systems are in place, agents could learn from previous sales in order to automatically determine which keywords or document properties are most likely to sell. This is akin to how Birukov's agents learn from other user's browsing habits.

The approach of using a multi-agent system on the botnet will require additional thought into both the high-level design and implementation details. The main questions here are whether a black market botnet could feasibly employ a multi-agent system, and how such a system could be detected and disrupted. Even without a multi-agent system, it would be helpful to run various indexing systems on computers in a laboratory situation to determine the speed and effectiveness of gathering documents.

7 Conclusion

In this paper we have presented a scenario where a botmaster can use existing technology to create a novel market, where adversaries can purchase private documents stolen from victim's desktop computers. There is a clear motivation for creating a market, both in adversaries who would like access to the documents, and the botmaster who would be able to profit from providing the access that adversaries desire.

The ability to harvest the documents in a simplistic manner is well within the reach of a botmaster; they would only need to build a botnet and provide searching functionality. The documents can be auctioned using the botnet as well, or more likely, through an existing auction site such as eBay. While using an existing auction site might not provide the most stable venue to sell documents, a sufficient level of obfuscation would likely permit such a scheme to be "good enough" and therefore viable. There is also the possibility of the botmaster taking the next step and automating the harvesting and selling processes using an adaptable multi-agent system.

Finally, we have presented several forms of defense that can be deployed immediately and in large scale, as well as some that are part of on-going long term efforts. These roughly fall into the following categories:

- preventing computers from becoming infected;
- preventing or slowing down the harvesting of documents;
- tracking documents to determine from where they are leaking;
- deceiving document harvesters to decrease the signal-to-noise ratio of harvested documents.

Black market botnets highlight the importance of vigilance in computer security research. As defensive mechanisms against known attack vectors are improved, adversaries will continue to re-arrange existing techniques and exploit them in new and creative ways.

Acknowledgments

Thanks to Jörg Denzinger and Randal Acton for initial discussions about this paper. Rei Safavi-Naini gave some helpful comments on digital rights management and steganography. Both authors' research is funded in part by the Natural Sciences and Engineering Research Council of Canada; the first author is also supported by the Informatics Circle of Research Excellence.

References

[1] Amazon.com. What are statistically improbable phrases? http://www.amazon.com/gp/search-inside/sipshelp.html.

[2] P. Bächer, T. Holz, M. Kötter, and G. Wicherski. Know your enemy: Tracking botnets. http://www.honeynet.org/papers/bots/, 2005.

[3] J. Bates. Trojan horse: AIDS information introductory diskette version 2.0. *Virus Bulletin*, pages 3–6, Jan. 1990.

[4] A. Birukov, E. Blanzieri, and P. Giorgini. *Implicit*: An agent-based recommendation system for web search. In *Fourth International Joint Conference on Autonomous Agents and Multi-Agent Systems*, pages 618–624, 2005.

[5] J.-F. Blanchette and D. G. Johnson. Data retention and the panoptic society: The social benefits of forgetfulness. *Information Society*, 18(1), 2002.

[6] M. Bond and G. Danezis. A pact with the Devil. Technical Report UCAM-CL-TR-666, University of Cambridge Computer Laboratory, 2006.

[7] E. Brynjolfsson, Y. J. Hu, and D. Simester. Goodbye Pareto principle, hello long tail: The effect of search costs on the concentration of product sales. Available at SSRN: http://papers.ssrn.com/sol3/papers.cfm?abstract_id=953587, 2006.

[8] D. H. Chau, S. Pandit, and C. Faloutsos. Detecting fraudulent personalities in networks of online auctioneers. In *Principles and Practice of Knowledge Discovery in Databases*, pages 103–114, 2006.

[9] C. Crump. Data retention: privacy, anonymity, and accountability online. *Stanford Law Review*, 56(1):191–229, Oct 2003.

[10] D. Dagon, G. Gu, C. Zou, J. Grizzard, S. Dwivedi, W. Lee, and R. Lipton. A taxonomy of botnets. Unpublished, available at http://www.math.tulane.edu/ tcsem/botnets/ndss_botax.pdf, 2005.

[11] J. Dörre, P. Gerstl, and R. Seiffert. Text mining: Finding nuggets in mountains of textual data. In *Fifth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 398–401, 1999.

[12] eBay. What is the eBay API? http://developer.ebay.com/common/api, 2007.

[13] O. D. Gnawali. *A Keyword-Set Search System for Peer-to-Peer Networks*. MIT, 2002. M.Sc. thesis.

[14] The Honeynet Project. http://www.honeynet.org/.

[15] D. Jackson. Gozi Trojan. SecureWorks, 2007.

[16] Y. Kusumura, Y. Hijikata, and S. Nishida. Text mining agent for net auction. In *2004 ACM Symposium on Applied Computing*, pages 1095–1102, 2004.

[17] Y. K. Lee and L. H. Chen. High capacity image steganographic model. In *IEE Proceedings on Vision, Image and Signal Processing*, pages 288–294, 2000.

[18] LURHQ. Cryzip ransomware Trojan analysis, 2006.

[19] muslix64. BackupHDDVD, a tool to decrypt AACS protected movies. http://forum.doom9.org/showthread.php?t=119871, 2006.

[20] Panda Software. PGPCoder.A. Virus Encyclopedia, 2005.

[21] F. A. P. Petitcolas, R. J. Anderson, and M. G. Kuhn. Information hiding – a survey. *Proceedings of the IEEE*, 87(7):1062–1078, July 1999.

[22] L. Rowel. The ballad of DVD Jon. *netWorker*, 10(4):28–34, Dec 2006.

[23] S. Schecter and M. Smith. Access for sale: a new class of worm. In *2003 ACM Workshop on Rapid malcode*, pages 19–23, 2003.

[24] D. Schonberg and D. Kirovski. Fingerprinting and forensic analysis of multimedia. In *12th Annual ACM International Conference on Multimedia*, pages 788–795, 2004.

[25] D. Sklyarov and A. Malyshev. eBooks security - theory and practice. In *DefCon 9*, 2001.

[26] Spamhaus. What is “fast flux” hosting? Frequently Asked Questions (FAQ).

[27] Spammer-X. *Inside the SPAM Cartel*. Syngress, 2004.

[28] The Honeynet Project. Know your enemy: Profile - automated credit card fraud. http://honeynet.org/papers/profiles/cc-fraud.pdf, 2003.

[29] Trend Micro. 2006 annual threat roundup and 2007 forecast. http://fr.trendmicro-europe.com/global/products/collaterals/white_papers/-2006AnnualThreatRoundup.pdf, 2006.

[30] Trend Micro. TROJ_ARHIVEUS.A. Virus Encyclopedia, 2006.

[31] Trusted Computing Group. TCG specification: Architecture overview. https://-www.trustedcomputinggroup.org/groups/TCG_1_0_Architecture_Overview.pdf, 2004.

[32] Trusted Computing Group. Trusted platform modules strengthen user and platform authenticity. https://www.trustedcomputinggroup.org/specs/TPM/-Whitepaper_TPMs_Strengthen_User_and_Platform_Authenticity_Final_1_0.pdf, 2005.

[33] United States v Bradley Joseph Steiger. 318 F.3d 1039. Eleventh Circuit, United States Court of Appeals, 2003.

[34] United States v Ronald C. Kline. 112 Fed.Appx.562. Ninth Circuit, United States Court of Appeals, 2004.

[35] United States v William Adderson Jarrett. 338 F.3d 339. Fourth Circuit, United States Court of Appeals, 2003.

[36] D. Watson, T. Holz, and S. Mueller. Know your enemy: Phishing. http://www.honeynet.org/papers/phishing/, 2005.

[37] World Intellectual Property Organization. Current developments in the field of digital rights management. http://www.wipo.int/documents/en/meetings/2003/sccr/pdf/sccr_10_2.pdf, 2003.

[38] WSOC. China Grove girl finds pot inside of Christmas present box. http://www.wsoctv.com/news/10624516/detail.html, 2006.

[39] A. Young. Non-zero sum games and survivable malware. In *Information Assurance Workshop, IEEE Systems, Man and Cybernetics Society*, pages 24–29, 2003.

[40] A. Young and M. Yung. Cryptovirology: Extortion-based security threats and countermeasures. In *IEEE Symposium on Security and Privacy*, pages 129–140, 1996.

[41] Y. Yu and T.-C. Chiueh. Display-only file server: A solution against information theft due to insider attack. In *Fourth ACM workshop on Digital Rights Management*, pages 31–39, 2004.