Black Market Botnets

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

(Appeared in MIT Spam Conference, 2008)

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. A botmaster who controls a botnet can use technology that exists now to create an infrastructure for selling information to third parties in a new way, exploiting the so-called "Long Tail." This results in not one, not two, but three new markets and untapped revenue streams for botmasters. 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. Can users' private documents be *selectively* stolen by an adversary?[1]

Botnets play a critical part in answering this question. The zombie computers that comprise a botnet have access to the private documents of the people that use the zombie computers, to documents that would otherwise be inaccessible to adversaries. However, not all adversaries have the desire or capability to establish a botnet.

Enter the botmaster. The botmaster, the person who creates and operates a botnet, is not the adversary (at least not directly) in this paper. The botmaster is a businessperson, who has access to private documents through their botnet. They do not necessarily want these documents themselves, but there are many adversaries who would; the motivation for the botmaster is to sell the documents for a profit. More importantly, the key insight underlying our "black market botnets" is this:

> *The botmaster does not know what documents are valuable to the adversary.*

Notice the important distinction between this scenario and what happens in current botnets, where botmasters harvest information they know has value, like credit card and PIN numbers. In contrast, the botmaster in our black market botnets will *not* know what has value, and will compensate by allowing their customers – the adversaries – to search for what they want.

In business terms, the botmaster of a black market botnet is taking advantage of the "Long Tail" [2]. This is the name given to the observation that the total sales volume of specialized items, targeting niche markets, can compete with sales of mass-market items. The trick is that sellers must be able to provide specialized items efficiently, and customers must be able to find them easily. In botnet terms, the sales curve might look like Figure 1.

Besides being able to target new niche markets, there is another advantage to the botmaster. Because of the scarcity, or maybe even uniqueness, of the specialized documents being stolen and sold, adversaries would pay much higher prices than they would for run-of-the-mill information. The botmaster of a black market botnet may

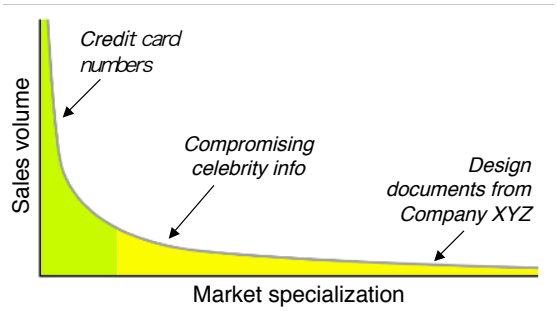

Figure 1: The Long Tail of a black market botnet

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

find the Long Tail good for both total sales volume and profit.

There are many possible scenarios in which a botmaster could sell specific 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.
- 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 [27–29].

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]. While currently restricted to a fixed type of data, Gozi demonstrates that the general threat we describe is beginning to evolve.

While we would like to put a dollar amount on the markets we present, it is not yet possible to do so in any meaningful way. Unlike fledgling studies of current underground economies [12], there is no extant data to be gathered and analyzed, nor is it clear what value an adversary would place on an invaluable document.

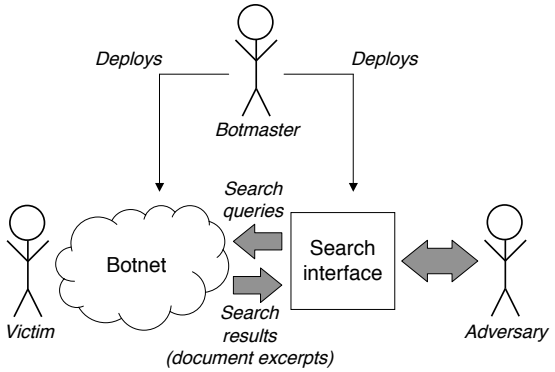

Figure 2: Basic architecture of a black market botnet

2 Basic Black Market Botnet Architecture

Conceptually, a black market botnet consists of two parts. First, there is the botnet that has access to salable documents. Second, there is a mechanism through which adversaries can search for documents of interest. In practice, the search mechanism may be hosted by machines in the botnet.

Because the total volume of documents on all of a botnets' computers would be large, we assume that it is not feasible to transfer all documents to a central location in their entirety. We therefore assume that adversaries' search queries would be injected into the botnet, and the botnet would return the search results. (We revisit this latter assumption in the next section.) This leads to the architecture shown in Figure 2. We note that the botnet does not even need to be excessively large for this to work, just well-targeted: for example, a botnet affecting only Fortune 500 companies would contain documents of interest to a number of adversaries.

To search documents, one option is for the botmaster to use existing peer-to-peer software and searching algorithms in their botnet (e.g., those used by the Gnutella or FastTrack networks [13]). Indeed, botnets have been known to employ P2P structures already [9]. Adversaries would use peer-to-peer clients to submit their search requests to the black market botnet. The search results would only contain a small, fixed excerpt 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 documents on zombie machines will almost certainly be required.

Another search option would be for the botmaster to

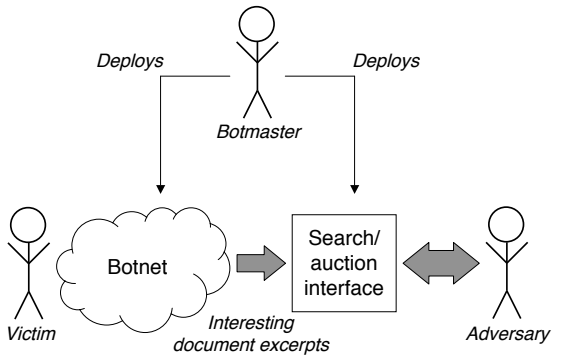

Figure 3: Advanced architecture with search term identification

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, safeguarding their websites using methods like bullet-proof hosting [23], fast flux [22], and levels of redirection [30].

Regardless of the mechanism used to perform searches, the question of *how* an adversary would search for documents of interest still remains. One possibility is that adversaries could use a traditional keyword-based search. However, this approach would be limited. In particular, the onus is on the adversaries to enter the magical incantation of keywords to find what they are looking for. 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 private documents of interest to only a few adversaries cannot be ranked the same way in terms of popularity. A ranking could be constructed of the popularity of various search terms rather than of the documents themselves, aiding the adversary in constructing useful searches, but revealing information to defenders who infiltrate the search interface. A search term ranking also assumes some degree of commonality in search terms, an assumption that may not be valid depending on where an adversary sits on the Long Tail.

How can a botnet be effectively searched for interesting documents? The simplest approach would be to search recently edited documents, assuming the adversary will want current information. However, only looking for recently edited documents would miss other potential targets, such as the previous year's tax returns. Therefore, more sophisticated indexing methods must be used. If a botmaster chooses to use existing software, Google Desktop 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. However, while convenient for the botmaster, this approach does not solve the problem of excessive query results. This issue is the subject of the next section.

For completeness, although it is probably obvious, an adversary who finds an interesting document excerpt would be able to purchase the entire document from the botmaster. Payment could be managed through established means, like WebMoney, e-gold, or PayPal.

3 e(vil)Bay: Advanced Black Market Botnet Architecture

In the last section, we assumed that the botmaster would not know what documents were interesting to an adversary; this assumption helped shape the basic black market botnet's architecture. This setup was limited in effectiveness, however. In this section, we explore how the botmaster can have zombies automatically identify potentially interesting search terms.

Dörre et al. discuss a system for text mining [10], the first part of which is to algorithmically extract what they call features, or "significant vocabulary items," such as "credit line." Amazon has followed along similar lines with their "Statistically Improbable Phrases" (SIPs) [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. Closely related is the concept of "inverse document frequency" (IDF) [24], a measure which can be used to retrieve individual documents containing a term which occurs in few of all the available documents.

The ability to automatically identify potential search terms suggests that a different architecture for the black market botnet is possible, shown in Figure 3. Upon identifying terms, a zombie could take excerpts of those terms with a small amount of their context in a document, and transmit the excerpts to a search interface. The search interface would perform adversaries' searches locally on the excerpts, rather than send queries into the botnet. For a botmaster, this architecture has the advantage of less telltale traffic to zombie computers, at the cost of a centralized target for defenders.

There are other advantages, too. A straightforward search interface would meet the needs of an adversary regardless of where their interests were on the Long Tail. Automatically collecting interesting excerpts from the black market botnet would allow the excerpts to not just be searched, but be advertised – they may be of interest

to more than one adversary, be sitting at the thicker part of the Long Tail. The "search interface" could then be an auction site, where adversaries can bid on the document excerpts. The adversary who wins the auction pays the botmaster, who then instructs the zombie computer to send the document to the buyer.

When document excerpts are posted for auction, a botmaster has two options as to how and where the listings are posted. First, the botmaster could create their own auction website, though as with the search portal described in the previous section, there is a risk that the site will be shut down, depriving the botmaster of income. This is the approach taken by Gozi [15].

Second, a botmaster could use an existing auction website. 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 – information hiding – has obvious application here (for more, see [19]). Many auction sites such as eBay allow sellers to post images of an item being sold, and this presents a good opportunity to apply steganography to hide information in the images. From a human's perspective, the auctioned items will look perfectly normal: a battered old teapot being sold, complete with a picture. However, adversaries 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 [32].

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 with which the items are posted. This list could be distributed through different kinds of channels, such as direct communication over IRC, which is already known to be used by fraudsters selling credit cards and other personal information [25]. 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, and provide the new information to their targets (in this case, the buyers). This is certainly not an ideal solution, but as long as it is good enough, a botmaster will be able to maintain operations.

Decoding the steganographically-hidden information could be accomplished if the botmaster provided adversaries with a program able to decode an image. Such programs are freely-available, mitigating trust issues between botmasters and adversaries. A more elaborate program could even search eBay using their API [11], 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, as before. Upon receipt of the payment, the botmaster would then provide the full document to the buyer. 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 Additional Revenue Streams

The auction or outright sale of documents to adversaries is one revenue stream for botmasters, but black market botnets present two more.

4.1 Bidding Wars

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. However, this is likely to erode any trust the botmaster enjoys from the adversaries.

Also, the nature of this activity is overt. It is conceivable that an adversary could sell many documents from one victim under normal circumstances, but once a bidding war started, the victim would be alerted to the zombie's presence and take counter-measures.

4.2 Query Persistence

As presented, adversaries' search queries in a black market botnet are ephemeral; if a document of interest appears only seconds after an adversary searches, the adversary will not see it. An outgrowth of the black market botnet scenario is for the botmaster to allow adversaries' queries to persist, to watch for new documents matching the search criteria.

There are definitely tradeoffs. For the adversary, persistence raises the risk of information exposure should the black market botnet be compromised by defenders. Such information could reveal the identity of a targeted person

or company, for example, which the adversary may wish to keep secret.

For the botmaster, query persistence provides another revenue stream, as a value-added service that adversaries can be charged for. Extending this idea, the botmaster could allow adversaries to inject "agents" into the botnet or the search interface. Adversaries could then not only have persistence, but could perform much richer searches, not limited by the botmaster's provided search interface.

Allowing agents might seem like the botmaster is placing a high degree of trust in adversaries, but this is not the case. A savvy botmaster would not allow arbitrary agents to run, but only safe, code-signed agents that the botmaster provides to adversaries, for a fee. It is thus not only safe, it is a business opportunity for the botmaster. The botmaster protects their investment in the black market botnet by maintaining control of agents and also the number of concurrent agents, to avoid the botnet being overloaded. The botmaster also opens up a new market with the black market botnet, selling agents to adversaries.

5 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, is the most effective defense. The usual suspects are helpful here: installing the latest patches, using firewalls, running up-to-date antivirus software.

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, nor how long such documents will remain available. This further motivates the use of agents that perform persistent queries in the black market botnet, as discussed in the last section. That said, this defense takes extra planning and effort to implement, and is not likely to be practical for everyday use.

A related issue is the current trend in retaining documents and other personal information for longer periods of time [4]. Government legislation mandating data retention for auditing purposes, such as the Sarbanes-Oxley Act, only provides more opportunities for a black market botnet to gain access to private documents [8]. 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). The idea is that, if a document is only readable on a particular zombie computer, then it's not usable even if the document is stolen.

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.

Various DRM schemes exist, and in 2003 the World Intellectual Property Organization published a comprehensive study on various DRM technologies [31]. In particular, implementations of DRM systems for documents already exist [17].

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. And, once a document is extracted from its steganographic cocoon, it is again vulnerable to persistent black market botnet queries.

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 [21].

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, how-

ever, 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 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 that 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.

6 Related Work

Schechter and Smith have discussed how a botmaster could sell access to infected computers [20]. However, for an adversary to correctly guess *which* computer holds interesting information is a hard problem, one that a black market botnet naturally solves.

Although a botmaster engaging in extortion is not good for business in the long run, the idea of compromised computers being used to extort payment from their owners has been examined. One way of accomplishing this is to encrypt files on the compromised computer, after which the adversary offers to decrypt them for a price [34]. This is not just academic; a number of cases have occurred in the wild [3, 16, 18, 26]. Young [33] presents yet another malware-based extortion attack, this one against brokerage firms, but is silent on how exactly malware is to automatically find sensitive documents to extort with. This is a common problem with extortion scenarios: how does malware decide what is valuable/sensitive/incriminating? Black market botnets, in contrast, harness the intelligence of human adversaries to find interesting documents, a mechanism which is both simpler and more reliable.

Bond and Danezis [5] 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 in their scale, and are much more plausible.

Chau et al. have examined the possibility of finding fraudsters on eBay [7]. 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" [35]. 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 (steps are taken to try and prevent screen capturing through software). In the case of a black market botnet, a display-only file server (DOFS) would prevent harvesting documents from an infected computer, because the documents simply do not exist there. A DOFS has clear advantages in a corporate environment, which could supply the necessary DOFS infrastructure, but would be of little help to the average home user. Furthermore, as Yu and Chiueh admit, it is not possible to prevent screen captures completely.

7 Conclusion

In this paper we have presented a scenario where a botmaster can use existing technology to create novel markets that exploit the Long Tail, where adversaries purchase private documents stolen from victims' computers. There is a clear motivation for creating these markets, both for adversaries who would like access to the documents, and for botmasters who would be able to profit from providing the access that adversaries desire. In addition to the primary revenue stream of document sales, there are additional revenue streams in document bidding wars and supplying agents for persistent queries.

But this is only the beginning. Researchers studying the Long Tail phenomenon have conjectured that there are "second-order effects" of the Long Tail on producers and consumers [6]. The unfortunate implication is that, as the market for private documents becomes popular, new niche black markets will emerge that are currently unfathomable.

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. The authors' research is funded in part by the Natural Sciences and Engineering Research Council of Canada; the first and third authors are also supported by the Informatics Circle of Research Excellence. The Long Tail curve used in Figure 1 is by Hay Kranen and is in the public domain.

References

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

[2] C. Anderson. *The Long Tail: Why the Future of Business Is Selling Less of More*. Hyperion, 2006.

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

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

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

[6] E. Brynjolfsson, Y. Hu, and M. D. Smith. From niches to riches: Anatomy of the Long Tail. *MIT Sloan Management Review*, 47(4), 2006.

[7] 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.

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

[9] 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.

[10] 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.

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

[12] J. Franklin, V. Paxson, A. Perrig, and S. Savage. An inquiry into the nature and causes of the wealth of Internet miscreants. In *14th ACM Conference on Computer and Communications Security*, pages 375–388, 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] LURHQ. Cryzip ransomware Trojan analysis, 2006.

[17] Microsoft. Microsoft Windows Rights Management Services for Windows Server 2003, 2005.

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

[19] 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.

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

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

[22] Spamhaus. What is "fast flux" hosting? Frequently Asked Questions (FAQ).

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

[24] K. Sparck Jones. Index term weighting. *Information Storage and Retrieval*, 9(11):619–633, 1973.

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

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

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

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

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

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

[31] 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.

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

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

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

[35] 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 Managment*, pages 31–39, 2004.