UNIVERSITY OF MASSACHUSETTS AMHERST

Department of Computer Science
140 Governors Drive
Amherst, MA 01003

February 12, 2012

Brian Neil Levine, PhD
Professor of Computer Science

Statement to United States Sentencing Commission
Public Hearing on Federal Child Pornography Offenses

Summary

- **Offenders engaged more seriously in child pornography file trafficking can be distinguished from lesser traffickers in part by their online actions and the technology used to access and share images of child exploitation on the Internet.**
- CP offenders are members of online communities that are supported by various mechanisms including p2p file-sharing network, web sites, web services, and chat rooms. The value of these communities to the offenders is determined by several factors: the number of peers involved, the amount of content peers share, the amount of time peers devote to the community, and the resources (i.e., bandwidth) peers contribute to meet demands for content. Users that have contributed a great deal of value to a community in these terms are more serious offenders. Counting the number of files shared by an offender is necessary but not sufficient.
- In some venues, offenders receive non-pecuniary benefits and incentives for their actions. Offenders that take advantage of these benefits are more serious offenders. In some cases, the benefits are related to improved network performance, which additionally improves the value of the network for all offenders. But in other more serious cases, the benefits can include training and encouragement that may lead from trafficking to contact offenses.
- Network details, such as the IP addresses of the user's computer, form the basis of online criminal investigations. Offenders that intentionally use mechanisms to mask network addresses and other information as part of these crimes should be viewed as significantly more serious offenders. Masked offenders can participate fully in open communities, making content available internationally, but stonewall justice and thwart investigators' abilities to put a stop to these communities and rescue exploited children.
- The above are important modern aspects of this crime and its offenders that are not taken into account by current guidelines. Assessments of offenders by Congress, sentencing judges, and the federal sentencing guidelines can use these differences to distinguish serious offenders.

1. Introduction

Current sentencing guidelines for child pornography offenders include enhancements that have been introduced in part because of the prominent role of the Internet and technology in these crimes. For example, offenses that involve the use of a computer have an increase of 2

levels. Similarly, the Internet has eased the ability of offenders to acquire massive collections, and offenses involving a greater number of images see increased levels.

Current sentencing guidelines do not take into account many modern aspects of online child pornography (CP) trafficking, some of which are employed by more serious offenders. The seriousness of child pornography offenders can be better distinguished through examination of their *online actions* and the *technology used* to access and share images of child sexual exploitation on the Internet. In this testimony, I provide factual information about trafficking technology and suggestions for distinguishing more serious online offenders, with the goal of informing changes to the statutory and guideline structure. Though the focus of this document is on the technology that supports online child exploitation, I expect the content to be considered within the context of other testimony that speaks to the legal and sociological aspects of these offenders and this crime.

In sum, more serious offenders can be distinguished in three major ways. First, they take actions that improve the value of online trafficking communities for other offenders. They share more files, for longer periods of time, and lend more bandwidth to support demand, ensuring others find success in trafficking as well and increasing the community. Second, they take advantage of non-pecuniary benefits of trafficking. These benefits may be related to improved network performance. More dangerously, the benefits may involve training and encouragement that can lead to contact offenses. Third, they intentionally mask the network information that forms the basis of most investigations. Masked offenders can participate fully in open communities, making content available internationally, but stonewall justice and thwart investigators' abilities to put a stop to these communities and rescue exploited children.

2. Background

There are many Internet venues supporting the search for and sharing of child pornography (CP) and content related to child sexual exploitation. These venues include: web sites; chat rooms; and peer-to-peer file sharing networks (p2p, for short). In addition, some Web-based services, such as backup and file hosting services, can be leveraged to create an ad hoc space for sharing web links to files and for creating ad hoc private communities around a set of files.

There are a great number of p2p file sharing networks. They include Gnutella,[1] eMule,[2] Ares,[3] BitTorrent,[4] OneSwarm[5,6] and GigaTribe,[7] among many others. Each is supported by one or more pieces of software that enable easy sharing of files by users to peers running the same software across the Internet. P2P networks are widely used and regularly provide new features beyond what is possible with a web site.

Each p2p network has a community of users, with some overlap among the networks in terms of the users involved and in terms of the content shared. Generally, users are able to join any network or introduce content to any network, though some contain private subgroups that require an invitation. The more important differences relate to the functionality offered by a particular p2p network to the user. For example, some (less popular) p2p networks directly attempt to mask the offender's network address from observation by third parties, and some thwart third

[1] T. Klingberg and R. Manfredi. *The Gnutella RFC, version 0.6*. `http://rfc-gnutella.sourceforge.net/src/rfc-0_6-draft.html`. 2002.

[2] Y. Kulbak and D. Bickson. *The eMule Protocol Specification*. `http://www.cs.huji.ac.il/labs/danss/p2p/resources/emule.pdf`. 2005.

[3] See `http://aresgalaxy.sourceforge.net/`.

[4] Bram Cohen. *The BitTorrent Protocol Specification*. `http://bittorrent.org/beps/bep_0003.html`. Version 11031. 2009.

[5] T. Isdal et al. "Privacy-preserving P2P data sharing with OneSwarm". In: *Proc. of ACM SIGCOMM Conference*. 2010, pp. 111–122.

[6] S. Prusty, B. Levine, and M. Liberatore. "Forensic Investigation of the OneSwarm Anonymous Filesharing System". In: *Proc. ACM Conference on Computer & Communications Security (CCS)*. 2011, pp. 201–214.

[7] See `http://www.gigatribe.com/`.

parties from enumerating of all content that a specific user is sharing. Some p2p networks are starting to integrate ideas from social networking sites, including the notion of deeming another peer as a "friend", which may lead to greater access to content or better network performance.

There are many technical details that distinguish popular networks, and these venues and programs are always changing. Detailing every difference here is not practical. Therefore, below, I focus on features of online trafficking that are common across programs, rather than focusing on programs themselves. In particular, I focus on features that I expect to remain available beyond the life cycle of any particular program or community.

3. Distinguishing Serious Offenders

Each online venue for CP trading is a *community*. At one end of the spectrum, the communities are an impersonal group-effort of people online that share content but never communicate. At the other end, the communities are comprised of offenders that chat and have stronger relationships involving training or encouragement, while also exchanging content. Each community has a value, and when the value is high, it attracts and helps more offenders.

The seriousness of an offender's actions online can be considered in the context of how it helps increases the value of their CP community. It is simpler to ask only whether the offenders downloads or uploads content, but in fact, there is more to being online than uploading and downloading. For example, one offender might share 900 files for 1 day; another offender might share 9 files for 100 days. As I describe below, these two actions both have benefits for the community, though current guidelines increase sentencing levels only for the former offender. One offender might be one of many sharing a file; another offender might be the only person sharing a file, ensuring the file's availability. One offender might be overt on the network; another offender might intentionally mask their network information, thwarting eventual traceback to their geographic location and the efforts of law enforcement.

With these examples in mind, consider the following distinguishing features and actions.

3.1. **Content diversity, availability, and redundancy.**

The value of CP-based communities of content to offenders is determined by several factors: the number of peers involved, the amount of content they share, the time they devote to the community, and the resources they contribute to meet demands for content. Offenders can be distinguished by what they have brought to a community along three dimensions.

- **Content Diversity.** The value of a CP-based community is strongly tied to the amount of content available. Communities survive and grow when they are a good resource for content, new and old. The number of files that an offender shares that are available only from their library (and not from other peers on the network) can vary quite a bit from user to user. Two offenders may each have 1,000 files each. But if one has unique content, it greatly increases the value of the network to others because the others are more likely to find what they seek.

  When an offender increases the diversity of files available to a particular network they increase not only the value of the network. It also means, generally, that they are increasing the number of victims whose images remain online well past the captured event and well past the introduction of the associated images to Internet venues. If an offender is the only person who is sharing a file, they are not unlike (though not the same as) the first person to introduce it: it is only for their actions that the file is available for others to download.

  Offenders with larger libraries (e.g., an offender with 2,000 files rather than 1,000) tend to have larger amounts of unique content made available to others online. The offenders with even more massive collections (well over several thousand) are, in general, able to gather such a library by acquiring content that isn't widely available otherwise. Hence, there is a significant difference between offenders with 2,000 CP files and those with 10,000 CP files, although current guidelines offer no enhancements above 600 CP files. But, it's not just the

size of the library: what matters as well is the amount of content offenders keep alive in the community through only their own actions.

- **Content Availability.** The value of CP-based communities is also tied to the amount of time that an offender's content is available. When demand continues to be met in terms of the percentage of time that files are available, the community survives and grows. Not all offenders are online at the same time in a peer-to-peer network. Some join for minutes or hours and return only days later. More active offenders will keep their sharing software online and available for longer durations, perhaps most of the hours of a day, most days of the week. Even if such an offender offers only a small shared library, being online longer increases the window of opportunity for more occasional offenders to find and download content. Thus, if one offender offers content constantly, then that helps ensure that any new users that come online at an arbitrary time can find content they seek. In that sense, an offender that shares 900 files but only a single day, is perhaps not adding as much value as an offender that shares 9 files for 100 days straight. Under current guidelines only the former offender is subject to an enhancement.
- **Content Redundancy.** Finally, the value of a CP-based community is also based on the resources available to meet demands for content. When an offender offers a file, even if for a short time and at a time that others offer it, they are lending resources to the network. Most p2p networks perform *parallel downloading*. In short, the downloader gets pieces of a specific file from multiple sources at once. Because most Internet Service Providers (ISPs) offer asymmetric service (the download rate is much higher than the upload rate), the most effective download strategy is to download from multiple remote peers at once. When an offender is one of many to offer CP, they are directly improving the network performance of the network in a way that a single offerer cannot provide. In other words, when offenders share even popular files, other offenders can typically download files faster, with larger bandwidths. When users can download files faster, the value of the network is higher, and the network thrives and grows.

3.2. **Community building and benefits from trafficking.** In some venues, offenders receive non-pecuniary benefits and incentives for certain actions. Offenders that take advantage of these benefits are more serious offenders. There are at least two types of benefits.

First, these benefits can be related to only network performance. For example, in OneSwarm p2p networks, peers can mark each other as trusted friends. If peers trust another, they reveal to each other an explicit list of shared files (that remain hidden from unknown peers), and they also boost the rate of file transfers between each other. In BitTorrent networks, peers receive significantly higher download rates when they upload data they have just received to others seeking the same content. In these systems, the software manages the incentives.

Second, these benefits can be social in nature. For example, an offender may have agreed to upload new content in order to gain access to a closed group. Many CP-based communities have grown from simple file transfer mechanisms to distribution networks, but increasingly have aspects of social networks. For example, some communities encourage chat and communication among members and not just file transfer. The most dangerous aspect of these types of communities of CP traders is that they offer training and encouragement among offenders.

Offenders that take advantage of benefits typically have taken actions that benefit the availability of content on the network, at the expense of the victims. These actions improve the value of the networks to others. When these benefits include training and encouragement, the offenders actions have ensured that the offending persons thrive regardless of and beyond the online network and programs used. Most seriously of all, the training encouragement might involve moving from image trafficking to contact offenses.

3.3. **Masking network addresses.** The ability of law enforcement to investigate these crimes is partly based on network information that can be observed remotely, including the remote peer's *IP address*. Some offenders use technology that masks this information. Offenders that

use network masking mechanisms are significantly more serious because the peer can participate in fully open networks (e.g., Gnutella or BitTorrent), making content available globally, while thwarting investigators abilities to stop their actions and shutdown these CP-based communities.

Applications on the Internet are able to route data to remote computers based on network-level IP addresses[8]. An ISP assigns these addresses to a subscriber. (Although it is unique while it in use, IP addresses can be reassigned to another subscriber at a later date.) Typically, the ISP keeps a log of these assignments, which can be subpoenaed by investigators, along with billing addresses that link network activity to a residential address. Some network applications have additional identifiers that can link the actions of the user across many IP addresses over time. For example: IRC users have a username; Gnutella users have a GUID; email users have an email address.

The use of IP addresses is fundamental to the Internet, but there is a basic method of obscuring the information: users can make use of a *proxy* that relays packets sent by the source to a destination. The main role of a proxy is to take incoming packets that have the source's IP address and send outgoing packets with only the proxy's IP address. There are many types of proxies available. As we note below, proxying is a mechanism that has been developed for many legitimate reasons outside of CP trafficking.

- There are hosts on the Internet that will proxy packets for free and others that do so for a fee. With sufficient technical knowledge, a person can also (illegally) hack into a remote host and use it as a proxy.
- *Virtual Private Network* (VPN) services also perform this function. The difference as compared to a simple proxy is that the traffic between the sender and the VPN proxy is encrypted. In that case, observers on the sender's network, or on any networks passed to get to the VPN proxy, do not know neither the final destination of the packets nor the content of the packets. Once the packets leave the VPN proxy, the sender's IP address is replaced with the VPN's IP address, and the content is overt. Variations of this model exist on the Internet (for example, ssh proxies) but the basic idea is the same.
- In more advanced networks, called anonymous communication systems, a sophisticated chain of three (or more) proxies is used, and all traffic is encrypted until between the third proxy and the final destination. This use of a proxy actually obfuscates the sender's IP address in a much more serious fashion. The first proxy in the chain knows the source but neither the destination nor the third proxy; the final proxy in the chain knows the destination but neither the source nor the first proxy; the second proxy knows only the two other proxies but neither the source nor destination. This setup is much stronger than relying on a single proxy, as above, that knows the source and destination. Tor[9] (`http://www.torproject.org/`) implements versions of this model, with variations of the idea in such applications as OneSwarm (`http://oneswarm.cs.washington.edu/`)[10], and Freenet (`http://freenetproject.org/`).
- *Network Address Translation (NAT)* is also a type of proxying, but typically it is one designed not for obfuscation but for sharing a single IP address among many users. In fact, all cellular users are behind NATs and are unintentionally masking their true network address.

Ordinary proxies are sometimes used as a bridge for linking two services, though they are also used for obfuscation. Because it encrypts data packets, one might use a VPN to gain privacy from local eavesdroppers (for example, while working in a cafe). However, many VPN services

[8]Kurose and Ross (J. Kurose and K. Ross. *Computer Networking: A Top-Down Approach Featuring the Internet*. Addison-Wesley, 2000) offer an excellent textbook on IP networks that details these advanced concepts.

[9]R. Dingledine, N. Mathewson, and P. Syverson. "Tor: The Second-Generation Onion Router". In: *Proc. USENIX Security Symposium*. 2004.

[10]See also `http://forum.oneswarm-fr.net` and `http://oneswarm.ru/`.

advertise themselves as an obfuscation service, preventing a destination from knowing the true sender. There is largely no reason to use Tor other than to obscure one's IP address, though there are many reasons to use Tor other than CP offenses. In general, proxying is a mechanism that has many legitimate uses, just as p2p file sharing networks contain content other than CP.

When a CP offender intentionally masks their IP address via a single or multi-chain proxy, it is analogous to robbing a bank with a mask. It does not increase the harm to victims, but it ties the hands of investigators. Investigating CP trafficking online is critical because it is the only effective *proactive* method of finding persons who directly and physically abuse children. Wolak et al.[11] have observed that 16% of investigations of CP possession ended with discovery contact offenders.

In some venues, the existence of the community is simply not advertised. It's not proxies that maintain the secrecy of the group, it's something as simple as a password protected website. In these cases, the network and the users are not anonymous, they are just hidden.

4. Intent

It is important to take these comments as part of a larger context. Described above are mechanisms only; the intent of a offender must be determined as well to apply these factors to sentencing. For example, some programs make it more obvious than others that sharing is involved. Some filenames are more obvious than others as being involved in child sexual exploitation and many programs share files before they can be viewed. The difference between a peer (in a p2p network) and a client (browsing a web site) is that a peer fulfills the roles of both client and server; some offenders may not understand that they have a role that is equivalent to a web server.

[11] J. Wolak, D. Finkelhor, and K. J. Mitchell. *Child-Pornography Possessors Arrested in Internet-Related Crimes: Findings From the National Juvenile Online Victimization Study.* Tech. rep. National Center for Missing & Exploited Children, 2005.