

The Darknet and the Future of Content Protection

Peter Biddle, Paul England, Marcus Peinado, and Bryan Willman

Microsoft Corporation*
Redmond, WA 98052, USA
(peterbi, pengland, marcuspe, bryanwi)@microsoft.com

Abstract. We investigate the darknet – a collection of networks and technologies used to share digital content. The darknet is not a separate physical network but an application and protocol layer riding on existing networks. Examples of darknets are peer to peer file sharing, CD and DVD copying, and key or password sharing on email and newsgroups. The last few years have seen vast increases in the darknet’s aggregate bandwidth, reliability, usability, size of shared library, and availability of search engines. In this paper we categorize and analyze existing and future darknets, from both the technical and legal perspectives. We speculate that there will continue to be setbacks to the effectiveness of the darknet as a distribution mechanism, but ultimately the darknet genie will not be put back into the bottle. In view of this hypothesis, we examine the relevance of content protection and content distribution architectures.

1 Introduction

People have always copied things. In the past, most items of value were physical objects. Patent law and economies of scale meant that small scale copying of physical objects was usually uneconomic, and large scale copying (if it infringed) was stoppable using policemen and courts. Today, things of value are increasingly less tangible: often they are just bits and bytes or can be accurately represented as bits and bytes. The widespread deployment of packet switched networks, and the huge advances in computers and codec technologies, have made it feasible (and indeed attractive) to deliver such digital works over the Internet. This presents great opportunities and great challenges. The opportunity is low cost delivery of personalized, high quality content. The challenge is that such content can be distributed illegally. Copyright law governs the legality of copying and distribution of such valuable data, but copyright protection is increasingly strained in a world of programmable computers and high speed networks.

For example, consider the staggering burst of creativity by authors of computer programs that are designed to share audio files. This was popularized

* Statements in this paper represent the opinions of the authors and not necessarily the position of Microsoft Corporation.

by Scour and Napster, but today several popular applications and services offer similar capabilities. In addition, CD-writers have become mainstream, and DVD-writers may well follow suit. Hence, even in the absence of network connectivity, the opportunity for low cost, large scale file sharing exists.

1.1 The Darknet

Throughout this paper, we will call the relevant items (e.g. software programs, songs, movies, books, etc.) objects. We will use the term to copy to refer to the duplication of objects in circumvention of copyright. The persons who copy objects will be called users of the darknet, and the computers used to copy objects will be called hosts. The idea of the darknet is based upon three assumptions:

1. Any widely distributed object will be available to a fraction of users in a form that permits copying.
2. Users will copy objects if it is possible and interesting to do so.
3. Users are connected by high bandwidth channels.

The darknet is the distribution network that emerges from the injection of objects according to assumption 1 and the distribution of those objects according to assumptions 2 and 3.

One implication of the first assumption is that any content protection system will leak popular or valuable content into the darknet, because some fraction of users – possibly experts – will overcome any copy prevention mechanism or because the object will enter the darknet before copy protection is applied.

The term “widely distributed” is intended to capture the notion of mass market distribution of objects to thousands or millions of practically anonymous users. This is in contrast to the protection of military, industrial, or personal secrets, which are typically not widely distributed and are not the focus of this paper.

Like other networks, the darknet can be modeled as a directed graph with labeled edges. The graph has one vertex for each user/host. For any pair of vertices (u, v) , there is a directed edge from u to v if objects can be copied from u to v . The edge labels can be used to model relevant information about the physical network and may include information such as bandwidth, delay, availability, etc. The vertices are characterized by their object library, object requests made to other vertices, and object requests satisfied.

To operate effectively, the darknet has a small number of technological and infrastructure requirements, which are similar to those of legal content distribution networks: The static hardware requirements to support a darknet are:

1. The *injection* requirement comprises technologies, devices and mechanisms that convert objects into a form, in which they can be transmitted and consumed in a darknet. Examples include audio and video compression algorithms and tools, CD and DVD readers, and programs that circumvent content protection systems (cracks). Injection provides darknets with new objects.

2. Mechanisms for *storage* and *replication* are required to allow users to make and keep copies of objects and to support the store and forward model of peer to peer networks. Examples include tapes, CDs, DVDs, and computer hard disks.
3. *Ubiquitous rendering devices* required to allow consumption of objects. Examples include portable music players, computers and consumer electronics DVD players and television sets.

The following core network related requirements correspond roughly to the components of the graph model outlined above:

1. Any darknet requires nodes that operate as object *sources*. These correspond to users who let at least some other users copy objects available to them.
2. Similarly, any darknet will contain *destination nodes* – users who want copies of objects. Often, nodes operate as both sources and destinations.
3. *Transmission links* are necessary to move copies of objects from source nodes to destination nodes. The Internet is the link that supports today’s peer to peer networks. The postal service and hand carried CDRs (sneakernet) support other darknets.
4. *Search engines* or other introduction mechanisms allow new and existing users to find objects on the darknet.

The dramatic rise in the efficiency of the darknet can be traced back to the general technological improvements in these infrastructure areas. At the same time, most attempts to fight the darknet focus on limiting or auditing one or more of the infrastructure items. Legal action has traditionally targeted search engines and source nodes. As we will describe later in the paper, this has been partially successful. The drive for legislation on mandatory watermarking aims to deprive the darknet of rendering devices. We will argue that watermarking approaches are technically flawed and unlikely to have any material impact on the darknet. Similarly, most content protection systems are meant to prevent or delay the injection of new objects into the darknet. However, no such system constitutes an impenetrable barrier; later, we will discuss the merits of some popular systems.

We see no technical impediments to the darknet becoming increasingly efficient (measured by aggregate library size and available bandwidth). However, the darknet infrastructure is under legal attack. In this paper, we trace the historical and current attacks on darknets and speculate on the technical and legal future of sharing technologies, concentrating particularly, but not exclusively, on peer to peer networks.

The rest of this paper is structured as follows: Section 2 analyzes different manifestations of the darknet with respect to their robustness to attacks on the infrastructure requirements described above and speculates on the future development of the darknet. Section 3 describes content protection mechanisms, their probable effect on the darknet, and the impact of the darknet upon them. In Sect. 4 and 5, we speculate on the situations in which the darknet will be effective, and how businesses may need to behave to compete effectively with it.

2 The Evolution of the Darknet

We classify the different manifestations of the darknet that have come into existence in recent years with respect to the five infrastructure requirements described and analyze weaknesses and points of attack.

As a system, the darknet is subject to a variety of attacks. While legal action, aimed at deterring widespread infringement, continues to be the most powerful challenge to the darknet, the darknet is also subject to a variety of other common threats (e.g. viruses, spamming) that, in the past, have led to minor disruptions of the darknet. They threaten to become considerably more damaging.

In this section we consider the potential impact of legal developments on the darknet. Most of our analysis focuses on system robustness, rather than on detailed legal questions. We regard legal questions only with respect to their possible effect: the failure of certain nodes or links (vertices and edges of the graph defined above). In this sense, we are investigating a well known problem in distributed systems.

2.1 Early Small Worlds Networks

Prior to the 1990s, copying was organized around groups of friends and acquaintances¹. The copied objects consisted mainly of music on cassette tapes and computer programs. The rendering devices were widely available tape players and the computers of the time (see Fig. 1). Content injection was trivial, since most objects were either not copy protected or, if they were equipped with copy protection mechanisms, the mechanisms were easily defeated. The distribution network was a “sneaker net” of floppy disks and tapes (storage), which were exchanged in person by members of a group or were sent by postal mail. The bandwidth of this network – albeit small by today’s standards – was sufficient for the objects of the time. The main limitation of the sneaker net, with its mechanical transport layer, was latency: It could take days or weeks to obtain a copy of an object. Another serious limitation of these networks was the lack of a sophisticated search engine.

There were some attempts to prosecute individuals who were trying to sell copyrighted objects they had obtained from the darknet (commercial piracy). However, the darknet as a whole was never under significant legal threat. Reasons may have included its limited commercial impact and the protection from legal surveillance afforded by sharing amongst friends.

The sizes of object libraries available on such networks are strongly influenced by the interconnections between the networks. For example, schoolchildren may copy content from their “family network” to their “school network” and thereby increase the size of the darknet object library available to each. Such networks

¹ Prior to this, some early computer users had access to ftp servers, usenet, and bulletin boards. These provided high bandwidth access to computer programs, and later to objects, such as images scanned in violation of copyright. However, the size of the communities served by these darknets was negligible.

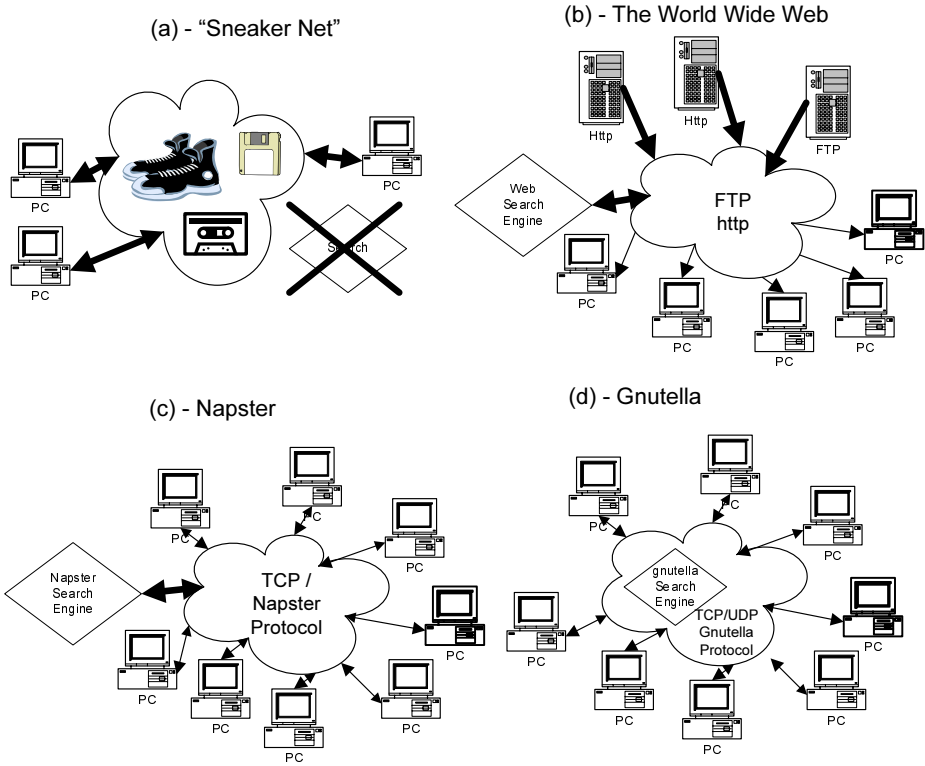


Fig. 1. Historical evolution of the Darknet. We highlight the location of the search engine (if present) and the effective bandwidth (thicker lines represent higher bandwidth). Network latencies are not illustrated, but are much larger for the sneaker net than for the IP-based networks

have been studied extensively and are classified as “interconnected small worlds networks” [1]. There are several popular examples of the characteristics of such systems. For example, most people have a social group of a few score of people. Each of these people has a group of friends that partly overlap with their friends’ friends, and also introduces more people. It is estimated that, on average, each person is connected to every other person in the world by a short chain of people from which arises the term “six degrees of separation.” These findings are remarkably broadly applicable (e.g. [2,3]). We suspect that these findings have implications for copying on darknets, and we will return to this point when we discuss the darknets of the future later in this paper.

The small worlds darknet continues to exist and indeed remains dominant for certain types of content. However, a number of technological advances have given rise to new forms of the darknet that have superseded the small worlds manifestation for some object types (e.g. audio).

2.2 Central Internet Servers

By 1998, a new form of the darknet began to emerge from technological advances in several areas. The internet had become mainstream, and could be used by anyone seeking to connect users with a centralized service or with each other. The continuing fall in the price of mass storage together with advances in compression technology had also crossed the threshold at which storing large numbers of audio files was no longer an obstacle to mainstream users. Additionally, the power of computers had crossed the point at which they could be used as rendering devices for multimedia content. Finally, “CD ripping” (from unprotected CDs) became a convenient, broadly available method for content injection.

The first embodiments of this new darknet were central internet servers with large collections of MP3 audio files. A fundamental change that came with these servers was the use of a new distribution network: The internet displaced the sneaker net – at least for audio content. This solved several problems of the old darknet.

Firstly, latency was reduced drastically. Secondly, and more importantly, discovery of objects became much easier because of simple and powerful search mechanisms – most importantly general purpose world wide web search engines. The local view of the small world was replaced by a global view of the entire collection accessible to all users. The main characteristic of this form of the darknet was centralized storage and search – a simple architecture that mirrored mainstream internet servers.

Centralized or quasi-centralized distribution and service networks make sense for legal online commerce. Bandwidth and infrastructure costs tend to be low, and having customers visit a commerce site means the merchant can display adverts, collect profiles, and bill efficiently. Additionally, management, auditing, and accountability are much easier in a centralized model. However, centralized schemes work poorly for illegal object distribution because large, central servers are large single points of failure: If the distributor is breaking the law, it is relatively easy to force him to stop. Early MP3 Web and FTP sites were commonly “hosted” by universities, corporations, and ISPs. Copyright holders or their representatives sent “cease and desist” letters to these website operators and web owners citing copyright infringement and in a few cases followed up with legal action [4]. The threats of legal action were successful attacks on those centralized networks, and MP3 web and FTP sites disappeared from the mainstream shortly after they appeared.

In the language of the model of Sect. 1, the centralized server darknet succumbed to a legal attack on its source nodes, whose small number made the attack tractable.

2.3 Peer to Peer Networks

The realization that centralized networks are not robust against attack has provided part of the impetus for the evolution of peer to peer networking and file sharing technologies. In this section, we examine architectures that have

evolved. Early systems were flawed because critical components remained centralized (Napster) or because of inefficiencies and lack of scalability of the protocol (gnutella) [5]. It should be noted that the problem of object location in a massively distributed, rapidly changing, heterogeneous system was new at the time peer to peer systems emerged. Efficient, highly scalable protocols have been proposed since then [6,7].

Early Internet Protocols. Simple peer to peer-like systems have existed on the internet for a long time. The main example is Usenet, which predates the central server darknets described above. While certain parts of Usenet have been and are still being used to distribute certain types of objects illegally, Usenet never became a mainstream darknet and never faced many of the attacks the more recent darknets are exposed to. We note, however, that the problem of endpoint anonymity arose in connection with Usenet. This resulted in work on anonymizing remailers and legal attacks on them.

Napster. Napster was the service that ignited peer to peer file sharing in 1999 [8]. There is little doubt that a major portion of the massive (for the time) traffic on Napster was of objects being transferred in a peer to peer model in violation of copyright law. Napster succeeded where central servers had failed by relying on the distributed storage of objects not under the control of Napster. This moved the injection, storage and replication, source nodes, network distribution, and consumption of objects to users.

However, Napster retained a quasi-centralized database with an index searchable on the file name. The centralized database itself became a legal target [4]. Napster was first enjoined to deny certain queries (e.g. “Metallica”) and then to police its network for copyrighted content. As the size of the darknet indexed by Napster shrank, so did the number of users. This illustrates a general characteristic of darknets: there is a correlation between the size and bandwidth of the object library and the appeal of the network for its users. This translates into positive feedback in the number of users: an efficient service quickly gains new users, and vice versa.

Gnutella. The next technology that sparked public interest in peer to peer file sharing was Gnutella. In addition to distributed object storage, Gnutella uses a fully distributed database described more fully in [9]. Gnutella does not rely upon any centralized server or service – a peer just needs the IP address of one or a few participating peers to (in principle) reach any host on the Gnutella darknet. Second, Gnutella is not really “run” by anyone: it is an open protocol and anyone can write a Gnutella client application. Finally, Gnutella and its descendants have substantial non-infringing uses. This changes its legal standing markedly and places it on a similar legal footing with email. Because email has substantial non-infringing use, it is not under direct legal threat in the jurisdiction of the authors of this paper, even though it may be used to transfer material unlawfully.

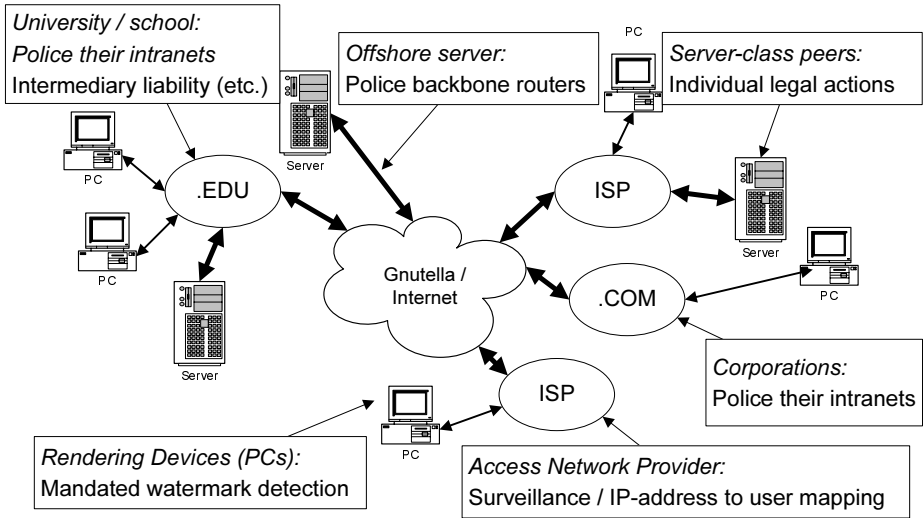


Fig. 2. Policing the darknet. Gnutella-style networks appear hard to police because they are highly distributed, and there are thousands or millions of peers. Looking more closely there are several potential vulnerabilities

2.4 Robustness of Fully Distributed Darknets

Fully distributed peer to peer systems do not present the single points of failure that led to the demise of central MP3 servers (injection) and Napster (search). It is natural to ask how robust these systems are and what form potential attacks could take. We observe the following weaknesses in Gnutella-like systems:

- Free riding
- Lack of anonymity

Free Riding. Peer to peer systems are often thought of as fully decentralized networks with copies of objects uniformly distributed among the hosts. While this is possible in principle, in practice it is not the case. Recent measurements of libraries shared by gnutella peers indicate that the majority of content is provided by a tiny fraction of the hosts which we term “super peers” [10]. Although gnutella appears to be a homogeneous peer to peer network of cooperating hosts, in actual fact it has evolved to effectively be another largely centralized system (Fig. 2). Free riding (i.e. downloading objects without sharing them) by many gnutella users appears to be main cause of this development. Widespread free riding removes much of the power of network dynamics and may reduce a peer to peer network into a simple unidirectional distribution system from a small number of sources to a large number of destinations. Of course, if this is the case, then the vulnerabilities that we observed in centralized systems (e.g. FTP-servers) are present again. Free riding and the emergence of super-peers have several causes:

Peer to peer file sharing assumes that a significant fraction of users adhere to a post-capitalist ideal of sacrificing their own resources for the “common good” of the network. Apparently, most free riders do not seem to adopt this ideology. For example, with 56 kbps modems still being the network connection for most users, allowing uploads constitutes a tangible bandwidth sacrifice. One approach is to make collaboration mandatory. For example, Freenet [11] clients are required to contribute some disk space. However, enforcing such requirements without a central infrastructure is difficult.

Existing infrastructure is another reason for the existence of super peers. There are vast differences in the resources available to different types of hosts. For example, a T3 connection provides the combined bandwidth of about one thousand 56 kbps telephone connections.

Lack of Anonymity. Users of gnutella who share objects they have stored are not anonymous. Current peer to peer networks permit the server endpoints to be determined, and if a peer-client can determine the IP address and affiliation of a peer, then so can a government agency. Users who share objects illegally face the threat of legal action. This appears to be another motivation for free riding.

2.5 Attacks

In this section, we analyze the robustness of distributed darknets with global databases. We consider how a variety of counter measures might apply to each of the technological and infrastructure requirements we identified in Sect. 1. These measures can be broadly classified as:

Legal: Filing lawsuits against users of the darknet or the operators of its infrastructure. Such attacks remove users from the darknet, but more importantly discourage participation of a much larger group of potential users.

Content protection: A collection of technical measures ranging from hindering injection (DRM) to attempts to make rendering devices reject darknet objects (watermark screening) and forensics (fingerprinting). These techniques are discussed in more detail in Sect. 3.

Network attacks: Like any other network, the darknet is subject to well known attacks, such as denial of service (DoS), spamming and viruses. We do not investigate the legal status of these attacks, but simply note that they are, in principle, possible and, to a very limited degree, appear to have taken place in the past.

Much of the static infrastructure (injection, storage, replication, rendering) has substantial non-infringing uses. Examples of such dual use technologies include audio and video compression tools, CD and DVD players, computers, monitors and television sets. These technologies appear largely immune to legal action. Furthermore, network attacks do not appear to apply in most cases. This leaves content protection as the main class of measures against the static darknet infrastructure. We analyze the effectiveness of these techniques in detail in Sect. 3. It appears unlikely that content protection measures alone will have a significant impact on the darknet.

The case of injection is different in the sense that injection tools that circumvent content protection mechanisms are subject to legal action – possibly under the Digital Millennium Copyright Act (DMCA). However, the most relevant recent example of such legal action appears to have been largely unsuccessful. DVD “ripping” tools that circumvent the CSS copy protection system are easily available on the internet.

Attacks against the network infrastructure of the darknet fall mostly into the categories of legal action and network attacks.

Sources. Source nodes of the darknet (i.e. hosts that make objects available to users in violation of copyright law) are subject to legal action. Lack of endpoint anonymity makes these hosts identifiable. Because of the prevalence of super peers the darknet depends on a relatively small set of powerful hosts, and these hosts are promising targets for attackers.

Darknet hosts owned by corporations are typically easily removed. Often, these hosts are set up by individual employees without the knowledge of corporate management. Generally corporations respect intellectual property laws. This together with their rational aversion to lawsuits, and their centralized network of hierarchical management, makes it relatively easy to remove darknet hosts in the corporate domain.

While the structures at universities are typically less hierarchical and strict than those of corporations, similar rules often apply.

If the .com and .edu OC-3 and OC-12 lines were pulled from under a darknet, the usefulness of the network would be impaired. Today, this would leave DSL, ISDN, and cable modem users as the high bandwidth servers of objects. We believe limiting source hosts to this class would present a far less effective piracy network today from the perspective of acquisition because of the relative rarity of high bandwidth consumer connections, and hence users would abandon this darknet. However, consumer broadband is becoming more popular, so in the long run it is probable that there will be adequate consumer bandwidth to support an effective consumer darknet.

The obvious next legal escalation is to bring direct or indirect (through the affiliation) challenges against users who illegally share large libraries of material. This is already happening and the legal actions appear to be successful [12]. This requires the cooperation of ISPs in identifying their customers, which appears to be forthcoming due to requirements that the carrier must take to avoid liability and, in some cases, because of corporate ties between ISPs and content providers. Once again, free riding makes this attack strategy far more tractable.

In addition to legal action, sources are subject to different kinds of denial of service attacks. These attacks become also more viable in the presence of widespread free riding.

Destination Nodes. Destination nodes suffer from the same endpoint anonymity problem as source nodes. In principle, similar legal attacks apply. In practice, destination nodes are better protected by their larger numbers.

Transmission. Attacks on transmission typically take the following forms. First, there have been attempts to identify and block darknet traffic on the internet. While such attacks may succeed with today's peer to peer systems, they are easily prevented by encrypting the darknet traffic. A second type of countermeasure is to limit the upload bandwidth of users who are suspected of providing large amounts of data into the darknet. While measures of this type may work against darknets with a relatively small set of super peers, they appear significantly less effective in darknet environments with more broadly distributed source nodes.

Search Engine. In Gnutella-style darknets, the search engine is integrated into the nodes. Thus, legal measures against the search engine are largely equivalent to legal measures against source and destination nodes, as described above. However, the global search engine has important implications for the feasibility of legal measures, as it removes endpoint anonymity and makes nodes globally identifiable. That is, the identity (IP address) of any source node is exposed through the global search engine to any client.

There are some technological workarounds to overcome the vulnerability presented by the lack of endpoint anonymity: anonymizing routers, overseas routers and object fragmentation complicate the effort required by law enforcement to determine the original source of unlawfully transferred bits. For example, Freenet tries to hide the identity of the hosts storing any given object by means of a variety of heuristics, including routing the object through intermediate hosts and providing mechanisms for easy migration of objects to other hosts. Similarly, Mnemosyne [13] organizes object storage such that individual hosts may not know what objects are stored on them. It is conjectured in [13] that this may amount to common carrier status for the host. A detailed analysis of the legal or technical robustness of these systems is beyond the scope of this paper. However, all such systems introduce the possibility of intermediary liability for the individuals who provide the "final hop."

Conclusions. The most relevant attacks we have identified exploit the lack of endpoint anonymity and are aided by the effects of free riding. We have seen effective legal measures on all peer to peer technologies that are used to provide global access to copyrighted material. Centralized web servers were effectively closed down. Napster was effectively closed down. Gnutella and Kazaa are under threat because of free rider weaknesses and lack of endpoint anonymity.

Should Gnutella-style systems become unviable as darknets, systems such as Freenet or Mnemosyne might replace them. It is hard to predict further escalation, but we note that the DMCA is a far reaching (although not fully tested) example of a law that is potentially quite powerful. We believe it probable that there will be ongoing technical efforts to sidestep existing laws, followed by new laws, or new interpretations of old laws, in the next few years. The rapid build out of consumer broadband, the decreasing price of storage, and the fact that personal computers are effectively establishing themselves as centers of home entertainment are technical developments that will continue to drive darknet demand.

Lack of endpoint anonymity is a direct result of the globally accessible global object database, and it is the existence of the global database that most distinguishes the newer darknets from the earlier small worlds. At this point, it is hard to predict whether the darknet will be able to retain this global database in the long term, but it seems clear that legal setbacks to global index peer to peer will continue.

2.6 Small Worlds Networks Revisited

In this section we try to predict the evolution of the darknet should global peer to peer networks be effectively stopped by legal or other means. The globally accessible global database is the only infrastructure component of the darknet that can be disabled in this way. The other enabling technologies of the darknet (injection, distribution networks, rendering devices, storage) will not only remain available, but will rapidly increase in power. We stress that the networks described in this section (in most cases) provide poorer services than the global network.

In the absence of a global database, small worlds networks could again become the prevalent form of the darknet. However, these small worlds will be more powerful than they were in the past. With the widespread availability of cheap CD and DVD readers and writers as well as large hard disks, the bandwidth of the sneaker net has increased dramatically, the cost of object storage has become negligible and object injection tools have become ubiquitous. Furthermore, the internet is available as a distribution mechanism that is adequate for audio for most users, and is becoming increasingly adequate for video and computer programs. In light of strong cryptography, it is hard to imagine how sharing could be observed and prosecuted as long as users do not share with strangers.

Students in dorms will establish darknets to share content in their social group. These darknets may be based on simple file sharing, DVD-copying, or may use special application programs or servers: for example, a chat or instant messenger client enhanced to share content with members of your buddy list. Each student will be a member of other darknets: for example, their family, various special interest groups, friends from high school, and colleagues in part time jobs (Fig. 3). If these small worlds are sufficiently well connected, we can anticipate that content will rapidly diffuse between darknets. Since the legal exposure of such sharing is quite limited, we believe that sharing amongst socially oriented groups will increase.

The limited exposure of sharing with strangers does not imply that such sharing will become universal. Non-technical admonitions will continue to discourage users from sharing. Such counsel may originate from parents, employers, or educators. The associated threats and possibility of discovery will factor into each individuals decision to share.

Small worlds networks suffer from the lack of a global database; each user can only see the objects stored by his small world neighbors. This raises a number of interesting questions about the network structure and object flow:




Fig. 3. Interconnected small worlds darknets. Threats of surveillance and prosecution may discourage participation in global darknets. In response, darknets form around social groups for which surveillance of illicit activity is unlikely. These darknets will use high bandwidth, low latency communications (intranets and the internet) and are supported by search engines. Custom applications, Instant Messenger style applications or simple shared file systems host the darknet. People’s social groups overlap so objects available in one darknet diffuse to others: in the terminology used in this paper, each peer that is a member of more than one darknet is an introduction host for objects obtained from other darknets

- What graph structure will the network have? For example, will it be connected? What will be the average distance between two nodes?
- Given a graph structure, how will objects propagate through the graph? In particular, what fraction of objects will be available at a given node? How long does it take for objects to propagate (diffuse) through the network?

Questions of this type have been studied in different contexts in a variety of fields (mathematics, computer science, economics, physics, and biology). A number of empirical studies seek to establish structural properties of different types of small world networks, such as social networks [2] and the world wide web [3]. These works conclude that the diameter of the examined networks is small, and observe further structural properties, such as a power law of the degree distribution [14]. A number of authors seek to model these networks by means of random graphs, in order to perform more detailed mathematical analysis on the models [15,16,17,18] and, in particular, study the possibility of efficient search under different random graph distributions [19,20]. We will present a quantitative study of the structure and dynamics of small worlds networks in an upcoming paper, but to summarize:

- For popular titles, small worlds darknets can be extremely efficient: very few peers are needed to satisfy requests for “top 20” books, songs, movies or computer programs. If darknets are interconnected, we expect the effective injection rate (injection from other networks) rate to be large. If darknet clients are enhanced to seek out new popular content, as opposed to the user demand based schemes of today, small worlds darknets could become very efficient.

- Less popular titles, will be harder or impossible to find, depending on the network parameters.
- Time sensitive objects will not be available.

For popular titles, small world darknets may provide a quality of service that matches that of peer to peer networks with global databases; for less popular titles, they may suffer from a reduced library size and latency.

3 Introducing Content into the Darknet

Our analysis and intuition have led us to believe that efficient darknet replication and propagation will remain a fact of life. In this section we examine rights management technologies that are being deployed to limit the introduction rate of content into the darknet.

3.1 Conditional Access Systems

A conditional access system is a simple form of rights management system in which subscribers are given access to objects based (typically) on a service contract. Digital rights management systems often perform the same function, but typically impose restrictions on the use of objects after unlocking.

Conditional access (CA) systems such as cable, satellite TV, and satellite radio offer little protection against objects being introduced into the darknet from subscribing hosts. A conditional access system customer has no access to channels or titles to which they are not entitled, and has essentially unencumbered use of channels that he has subscribed or paid for. This means that an investment of \$100 (at time of writing) on an analog video capture card is sufficient to obtain and share TV programs and movies. Some CA systems provide post unlock protections but they are generally cheap and easy to circumvent.

Thus, conditional access systems provide a widely deployed, high bandwidth source of video material for the darknet. In practice, the large size and low cost of CA-provided video content will limit the exploitation of the darknet for distributing video in the near term.

The same can not be said of the use of the darknet to distribute conditional access system broadcast keys. At some level, each head end (satellite or cable TV head end) uses an encryption key that must be made available to each customer (it is a broadcast), and in the case of a satellite system this could be millions of homes. CA system providers take measures to limit the usefulness of exploited session keys (for example, they are changed every few seconds), but if darknet latencies are low, or if encrypted broadcast data is cached, then the darknet could threaten CA system revenues.

We observe that the exposure of the conditional access provider to losses due to piracy is proportional to the number of customers that share a session key. So, cable operators are in a safer position than satellite operators because a cable operator can narrowcast more cheaply.

3.2 DRM Systems

A classical DRM system is one in which a client obtains content in protected (typically encrypted) form, with a license that specifies the uses to which the content may be put. Examples of licensing terms that are being explored by the industry are “play on these three hosts,” “play once,” “use computer program for one hour,” etc.

The license and the wrapped content are presented to the DRM system whose responsibility is to ensure that:

- The client cannot remove the encryption from the file and send it to a peer.
- The client cannot “clone” its DRM system to make it run on another host.
- The client obeys the rules set out in the DRM license.
- The client cannot separate the rules from the payload.

Advanced DRM systems may go further. Some such technologies have been commercially very successful – the content scrambling system used in DVDs, and (broadly interpreted), the protection schemes used by conditional access system providers fall into this category, as do newer DRM systems that use the internet as a distribution channel and computers as rendering devices. These technologies are appealing because they promote the establishment of new businesses and reduce distribution costs. If costs and licensing terms are appealing to producers and consumers, then the vendor thrives. If the licensing terms are unappealing or inconvenient or the costs are too high then the business will fail. The DivX “DVD” rental model failed on most or all of these metrics, but CSS-protected DVDs succeeded beyond the wildest expectations of the industry.

On personal computers, current DRM systems are software only systems using a variety of tricks to make them more or less hard to subvert. DRM enabled consumer electronics devices are also beginning to emerge.

In the absence of the darknet, the goal of such systems is to have comparable security to competing distribution systems – notably the CD and DVD – so that programmable computers can play an increasing role in home entertainment.

DRM systems strive to be BOBE (break once, break everywhere)-resistant. That is, suppliers anticipate that individual instances (clients) of all security systems, whether based on hardware or software, will be subverted. If a client of a system is subverted, then all content protected by that DRM client can be unprotected. If the break can be applied to any other DRM client of that class so that all of those users can break their systems, then the DRM-scheme is BOBE-weak. If, on the other hand, knowledge gained breaking one client cannot be applied elsewhere, then the DRM system is BOBE-strong.

Most commercial DRM systems have BOBE exploits, and we note that the darknet applies to DRM hacks as well. The CSS system is an exemplary BOBE weak system. The knowledge and code that comprised the De-CSS exploit spread uncontrolled around the world on websites, newsgroups, and even T shirts, in spite of the fact that, in principle, the Digital Millennium Copyright Act makes it a crime to develop or distribute these exploits.

A final characteristic of existing DRM systems is renewability. Vendors recognize the possibility of exploits, and build systems that can be field updated.

It is hard to quantify the effectiveness of DRM systems for restricting the introduction of content into the darknet from experience with existing systems. Existing DRM systems typically provide protection for months to years; however, the content available to such systems has to date been of limited interest, and the content that is protected is also available in unprotected form. The one system that was protecting valuable content (DVD video) was broken very soon after compression technology and increased storage capacities and bandwidth enabled the darknet to carry video content.

3.3 Software

The DRM systems described above can be used to provide protection for software, in addition to other objects (e.g. audio and video). Alternatively, copy protection systems for computer programs may embed the copy protection code in the software itself.

The most important copy protection primitive for computer programs is for the software to be bound to a host in such a way that the program will not work on an unlicensed machine. Binding requires a machine ID: this can be a unique number on a machine (e.g. a network card MAC – media access control – address), or can be provided by an external dongle.

For such schemes to be strong, two things must be true. First, the machine ID must not be “virtualizable.” For instance, if it is trivial to modify a network card driver to return a different MAC address, then the software-host binding is easily broken. Second, the code that performs the binding checks must not be easy to patch. A variety of technologies that revolve around software tamper resistance can help here [21].

We believe that binding software to a host is a more tractable problem than protecting passive content, as the former only requires tamper resistance, while the latter also requires the ability to hide and manage secrets. However, we observe that all software copy protection systems deployed thus far have been broken. The definitions of BOBE strong and BOBE weak apply similarly to software. Furthermore, once software is broken, the hacks or patched software are just as much subject to the dynamics of the darknet as passive content.

4 Policing Hosts

If there are subverted hosts, then content will leak into the darknet. If darknet propagation is efficient, then content will be available to all interested peers. In this section we evaluate technologies proposed for limiting output, or provide forensic information that allows users who inject objects in violation of copyright or contract to be identified.

4.1 Watermarking

Watermarking embeds an “indelible” invisible mark in content. A plethora of schemes exist for audio/video and still image content and computer programs.

There are a variety of schemes for exploiting watermarks for content protection. These schemes are implemented in output devices. Consider a rendering device that locates and interprets watermarks. If a watermark is found then special action is taken. For example, the output device may:

Restrict behavior: For example, a bus adapter may refuse to pass content that has the “copy once” and “already copied once” bits set.

Require a license to play: For example, if a watermark is found indicating that content is rights-restricted then the renderer may demand a license indicating that the user is authorized to play the content.

Such systems were proposed for audio content – for example the secure digital music initiative (SDMI) [22], and are under consideration for video by the copy protection technical working group (CPTWG) [23].

There are several reasons why it appears unlikely that such systems will ever become an effective anti-piracy technology. From a commercial point of view, building a watermark detector into a device renders it strictly less useful for consumers than a competing product that does not have one, and such detectors impose a “tax” in performance and cost on consumers who are using devices for perfectly lawful activities. Hence watermarking schemes are unlikely to be widely deployed, unless mandated by legislation. The recently proposed Hollings bill is a step along these lines [24]. Even with legislation, they are likely to meet severe resistance.

We contrast watermark based policing with classical DRM: If a general purpose device is equipped with a classical DRM system, it can play all content acquired from the darknet, and have access to new content acquired through the DRM channel. This is in stark distinction to reduction of functionality inherent in watermark based policing.

Even if watermarking systems were mandated, this approach is likely to fail due to a variety of technical inadequacies. The first inadequacy concerns the robustness of the embedding layer. We are not aware of systems for which simple data transformations cannot strip the mark or make it unreadable [25]. Marks can be made more robust, but in order to recover marks after adversarial manipulation, the reader must typically search a large phase space, and this quickly becomes untenable. In spite of the proliferation of proposed watermarking schemes, it remains doubtful whether robust embedding layers for the relevant content types can be found.

A second inadequacy lies in unrealistic assumptions about key management. Most watermarking schemes require widely deployed cryptographic keys. Standard watermarking schemes are based on the normal cryptographic principles of a public algorithm and secret keys. Most schemes use a shared key between marker and detector. In practice, this means that all detectors need a private key, and, typically, share a single private key. It would be naïve to assume that these keys will remain secret for long in an adversarial environment. Once the key or keys are compromised, the darknet will propagate them efficiently, and the scheme collapses. There have been proposals for public key watermarking systems. However, so far, this work does not seem practical and the correspond-

ing schemes do not even begin to approach the robustness of the cryptographic systems whose name they borrow.

A final consideration relates to the location of mandatory watermark detectors in client devices. On open computing devices (e.g. personal computers), these detectors could, in principle, be placed in software or in hardware. Placing detectors in software would be largely meaningless, as circumvention of the detector would be as simple as replacing it by a different piece of software. This includes detectors placed in the operating system, all of whose components can be easily replaced, modified and propagated over the darknet.

Alternatively, the detectors could be placed in hardware (e.g. audio and video cards). In the presence of the problems described this would lead to untenable renewability problems – the hardware would be ineffective within days of deployment. Consumers, on the other hand, expect the hardware to remain in use for many years. Finally, consumers themselves are likely to rebel against “footing the bill” for these ineffective content protection systems. It is virtually certain that the darknet would be filled with a continuous supply of watermark removal tools based on compromised keys and weaknesses in the embedding layer. Attempts to force the public to “update” their hardware would not only be intrusive, but impractical.

In summary, attempts to mandate content protection systems based on watermark detection at the consumer’s machine suffer from commercial drawbacks and severe technical deficiencies. These schemes, which aim to provide content protection beyond DRM by attacking the darknet, are rendered entirely ineffective by the presence of even a moderately functional darknet.

4.2 Fingerprinting

Fingerprint schemes are based on similar technologies and concepts to watermarking schemes. However, whereas watermarking is designed to perform a-priori policing, fingerprinting is designed to provide a-posteriori forensics.

In the simplest case, fingerprinting is used for individual sale content (as opposed to super-distribution or broadcast – although it can be applied there with some additional assumptions). When a client purchases an object, the supplier marks it with an individualized mark that identifies the purchaser. If the marked content appears on a darknet, a policeman can identify the source of the object and the offender can be prosecuted or other action can be taken.

Fingerprinting suffers from fewer technical problems than watermarking. The main advantage is that no widespread key distribution is needed – a publisher can use whatever secret or proprietary fingerprinting technology they choose, and is entirely responsible for the management of their own keys.

Fingerprinting has one problem that is not found in watermarking. Since each fingerprinted copy of a piece of media is different, if a user can obtain several different copies, he can launch collusion attacks (e.g. averaging). In general, such attacks are very damaging to the fingerprint payload.

It remains to be seen whether fingerprinting will act as a deterrent to theft. There is currently no legal precedent for media fingerprints being evidence of

crime, and this case will probably be hard to make since detection is a statistical process with false positives, and opportunity for deniability. However, we anticipate that there will be uneasiness in sharing a piece of content that may contain a person's identity and that ultimately leaves that person's control.

Note also that, with widely distributed watermarking detectors, it is easy to see whether a watermark has been successfully removed. There is no such assurance for determining whether a fingerprint has been successfully removed from an object because users are not necessarily knowledgeable about the fingerprint scheme or schemes in use. However, if it turns out that the deterrence of fingerprinting is small (i.e. everyone shares their media regardless of the presence of marks), there is probably no reasonable legal response. Finally, distribution schemes in which objects must be individualized will be expensive.

5 Conclusions

There are no inherent technical impediments to darknet based object sharing technologies growing in usability, library size, aggregate bandwidth and efficiency, but the legal future of darknet technologies is less certain. We have described successful or partially successful legal attacks on all network based object sharing technologies in widespread use today. We anticipate further escalation of attacks and of darknet technologies to remove the vulnerabilities that were exploited in previous attacks. We have analyzed the infrastructure components necessary to support arbitrary darknets, and have argued that, while some of the infrastructure components appear immune to legal or technological attack, some vulnerabilities will remain.

The largest vulnerability arises from the exposure of a user's identity, either directly or indirectly, to law enforcement masquerading as a peer. This vulnerability arises if users share with unknown or anonymous peers, and is a consequence of registering hosts and objects with a global database or other database without user access control. Should the threat of legal action make sharing among anonymous users too risky for average users, then we have argued that darknets will form around smaller, access controlled small worlds groups for which the risk of surveillance is smaller.

The reduced exposure afforded by small worlds darknets to their users may come at the price of diminished quality of service. The library size, availability, and latency of a small world darknet will always be inferior to that of a global darknet. This will almost certainly mean that small worlds darknets will be impractical for sharing less popular objects and time sensitive objects. On the other hand, even moderately efficient small worlds darknets are likely to provide high quality of service for the most popular objects.

It is our conjecture that darknets will survive, but the efficiency and size of these future darknets is uncertain. In the remainder of this section we speculate on the technical and business implications of the continued existence of darknets of varying levels of efficiency on the commerce of digital goods.

5.1 Technological Implications

Darknets replicate objects. An efficient darknet replicates objects rapidly, and makes the original and its replicas available to an expanding group of users. If the darknet is an efficient global darknet then all users can access an object immediately after it is introduced. If architectural deficiencies or attacks reduce the efficiency of a global darknet then significant time and effort may be required to obtain a copy of an object. If no global darknet exists, but a user is a member of one or more small worlds darknets then users must wait until an object reaches their small world – either by diffusing from an interconnected small world, or through direct injection.

Classical DRM systems inhibit the injection of objects into darknets. However, we must always assume that a fraction of DRM systems are subverted, or objects are introduced into the darknet through other channels. In light of the arguments in the previous paragraph we conclude that DRM systems will be effective in limiting the widespread availability of objects for isolated small worlds darknets, but will be ineffective security measures in the presence of efficient global darknets.

The interesting cases arise between these two extremes – in the presence of a darknet which is connected but in which factors such as latency, limited bandwidth or the absence of a global database limit the speed with which objects propagate. It appears that quantitative studies of the effective “diffusion constant” of different kinds of darknets and objects would be highly useful in elucidating the dynamics of DRM systems and the darknet.

Proposals for systems involving mandatory watermark detection in rendering devices try to impact the effectiveness of the darknet directly by trying to detect and eliminate objects that originated in the darknet appear flawed. In addition to severe commercial and social problems, these schemes suffer from serious technical deficiencies, which argue against their future value. We conclude that such schemes are doomed to failure.

5.2 Business in the Face of the Darknet

Darknets are a competitor to legal commerce, and the normal rules of competition apply. The level of competition of a darknet for an industry depends on its efficiency and effective price compared to the convenience and price of the competing legal channels (as well as other social factors like the price sensitivity and honesty of the users).

Historically, the efficiency of a darknet has been affected by the legal and technical attacks upon it. We have argued that global darknets have inherent vulnerabilities that can be exploited to reduce library size and aggregate bandwidth. Clearly, the level of competition provided by a darknet depends on the attacks it is exposed to, and we assume that businesses will continue to invest in such attacks. We have argued that these attacks may reduce the quality of service of darknets, even if they may not completely eliminate them.

A moderately efficient darknet will provide pressure on the price and convenience of legal channels for businesses. There are many technical and social

factors that determine the competitiveness of a darknet, and we will list those that seem particularly important. First, the size of the shared objects: Current peer to peer darknets appear adequate for audio, but are not adequate for video for most users. Second, the behavior of the customers: corporate customers are unlikely to engage in widespread sharing of digital objects in violation of contract or copyright. However, it appears that many people share audio files without compunction. Third, the distribution size: mass market media is widely distributed and widely interesting. This implies many potential injection hosts, and high demand driving darknet replication. In contrast, personalized documents or premium business reports are far less likely to be introduced and replicated. Fourth, the convenience of the legal channel: convenience can take many forms: a DRM-protected object may be less convenient than an unprotected object; a native digital representation of an object from a darknet may be more appealing to some users than an object embedded in a physical artifact (e.g. a CD). Fifth, time: if darknets are only moderately efficient then there will be a delay before a new object is widely available. Of course the price of the object is a huge factor, and there are many others.

We do not believe that darknets will drive the cost of all digital goods to zero, but it appears likely that the effects on some types of mass market digital commerce will be significant.

Acknowledgements

We are grateful to Cormac Herley, Rico Malvar, John Manfredelli and Yacov Yacobi, for many useful comments, ideas, and discussions.

References

1. Watts, D., Strogatz, S.: Collective dynamics of small world networks. *Nature* **393** (1998) 440–442
2. Milgram, S.: The small world problem. *Psychology Today* **2** (1967) 60–67
3. Albert, R., Jeong, H., Barabási, A.L.: Diameter of the world-wide web. *Nature* **401** (1999) 130–131
4. <http://www.riaa.com>
5. Javanović, M., Annexstein, F., Berman, K.: Scalability issues in large peer-to-peer networks – a case study of gnutella. Technical report, ECECS Department, University of Cincinnati (2001)
6. Stoica, I., Morris, R., Karger, D., Kaashoek, M.F., Balakrishnan, H.: CHORD: A scalable peer-to-peer lookup service for internet applications. In: *Proceedings of the ACM SIGCOMM 2001 Conference (SIGCOMM-01)*. (2001) 149–160
7. Dabek, F., Brunskill, E., Kaashoek, M.F., Karger, D., Morris, R., Stoica, I., Balakrishnan, H.: Building peer-to-peer systems with Chord, a distributed lookup service. In: *Proceedings of the Eighth IEEE Workshop on Hot Topics in Operating Systems (HotOS-VIII)*. (2001) 81–86
8. <http://www.napster.com>
9. http://www.gnutelladev.com/protocol/gnutella_protocol.html

10. Adar, E., Huberman, B.: Free riding on Gnutella. Technical report, Xerox-PARC (2000)
11. Clarke, I., Sandberg, O., Wiley, B., Hong, T.: Freenet: A distributed information storage and retrieval system. In: International Workshop on Design Issues in Anonymity and Unobservability. (2000)
12. Clarke, R.: A defendant class action law suit. <http://www.kentlaw.edu/perritt/honorsscholars/clarke.html>
13. Hand, S., Roscoe, T.: Mnemosyne: peer-to-peer steganographic storage. In: Proceedings of the first International Workshop on Peer-to-Peer Systems. (2000)
14. Barabási, A.L., Albert, R.: Emergence of scaling in random networks. *Science* **286** (1999) 509–512
15. Aiello, W., Chung, F., Lu, L.: Random evolution in massive graphs. In: Proceedings of the 42nd Annual IEEE Symposium on Foundations of Computer Science. (2001) 510–519
16. Cooper, C., Frieze, A.: A general model of web graphs. In: Proceedings of the 9th Annual European Symposium on Algorithms. (2001) 500–511
17. Newman, M.: Small worlds: the structure of social networks. Technical Report 99-12-080, Santa Fe Institute (1999)
18. Newman, M., Watts, D., Strogatz, S.: Random graph models of social networks. *Proc. Natl. Acad. Sci. USA* **99** (2002) 2566–2572
19. Kleinberg, J.: Navigation in a small world. *Nature* **406** (2000)
20. Kleinberg, J.: Small-world phenomena and the dynamics of information. *Advances in Neural Information Processing (NIPS)* **14** (2001)
21. Aucsmith, D.: Tamper-resistant software: An implementation. In Anderson, R., ed.: *Information hiding: first international workshop*, Cambridge, U.K. Volume 1174 of *Lecture Notes in Computer Science.*, Springer-Verlag (1996) 317–333
22. <http://www.sdmi.org>
23. <http://www.cptwg.org>
24. Hollings, F.: Consumer broadband and digital television promotion act
25. Kirovski, D., Petitcolas, F.: Replacement attack on arbitrary watermarking systems. In: *Proceedings of the 2002 ACM Workshop on Digital Rights Management*, Springer-Verlag (2003)