2010

Defending Privacy: the Development and Deployment of a Darknet

Conor McManamon
Technological University Dublin

Fredrick Mtenzi
Technological University Dublin, Fredrick.Mtenzi@tudublin.ie

Recommended Citation

McManamon, C. & Mtenzi, F. (2010). Defending Privacy: The Development and Deployment of a Darknet, *Proceedings IEEE 5th Internet Technology and Secured Transactions (ICITST)*, pg.1-6. doi:10.21427/sv5v-6029

Defending Privacy: The Development and Deployment of a Darknet

Conor Mc Manamon and Fredrick Mtenzi
Dublin Institute of Technology, Ireland
Conor.McManamon@gmail.com; Fredrick.Mtenzi@dit.ie

Abstract

New measures imposed by governments, Internet service providers and other third parties which threaten the state of privacy are also opening new avenues to protecting it. The unwarranted scrutiny of legitimate services such as file hosters and the BitTorrent protocol, once relatively unknown to the casual Internet user, is becoming more obvious. The darknet is a rising contender against these new measures and will preserve the default right to privacy of Internet users.

A darknet is defined in the context of file sharing as a network which operates on top of another network such as the Internet for the purpose of secure and private distribution of digital material. While there are other darknet applications in existence, such as Freenet, WASTE again, and Relakks, they harbour some caveats. Whether they be proprietary solutions, depend on other services, are prone to feature creep or have security shortcomings, there is room for improvement.

The aim of this paper is to address and improve on some of the problems of these alternative darknet clients with the development of a lightweight darknet application suite - Umbra. It is then demonstrated how its deployment can circumvent or defeat the draconian measures currently threatening privacy in the public domain.

1. Introduction

In recent years, the Internet has become heavily policed with the aim of quashing the distribution of illegal materials whether they infringe copyright or are of another dubious nature. Applications such as peer to peer software and companies providing file hosting services are being heavily scrutinised, held responsible for the distribution of these materials and in some cases even made illegal. The measures imposed by Internet service providers and the courts on behalf of the firms responsible for the protection of copyrighted material have had major repercussions on the state of privacy in the public domain.

Two predominant measures threatening the state of privacy are filtering and traffic shaping. File hosting companies have been the subject of litigation by parties representing several large record labels. Mandates were passed down to remove all copyrighted material and implement a scheme of filtering measures to ensure the prevention of repeat uploads of any such materials to their servers. While this may seem a fair avenue of action from the perspective of the record labels, it is also imperative to take into account the impact of such actions on privacy.

Traffic shaping on Internet service provider networks is a relatively new public concern. Some Internet service providers openly advertise they throttle traffic associated with the BitTorrent protocol. While legislation surrounding traffic shaping and net neutrality is under current global review, introducing tiered services dependant on traffic types does present a violation of privacy. The content of Internet users' traffic is not directly analysed yet, but the type of traffic is certainly categorised and routed accordingly. This has ramifications in terms of network neutrality of course, but from a privacy perspective it also constitutes an overstepping of bounds as users traffic is subject to analysis and may be hindered due to choice of protocol.

There are several legitimate usage scenarios for these sorts of services. Home users sharing photographs or videos, small businesses distributing company documents or their products to customers, distributed software development projects or educational schemes hosting material for remote participants are all prime examples. The BitTorrent protocol has provided a medium for citizens of oppressed countries where censorship is a major obstacle to provide free media coverage of the situation within their countries and to communicate with the outside world.

One possible remedy to the current trend of privacy intrusion is the deployment of a darknet. A darknet, in the context of file sharing and this research, is defined as a network operating on top of another network, such as the Internet, in a closed or private manner. The content of data exchange is hidden from users outside the network or lacking sufficient authorisation. This can be achieved in a number of ways, the most common being encrypting all traf-

fic between nodes. In other networks, anonymity provided through proxy servers or onion routers is used alongside encryption. While there are some darknet applications in existence, it seems most of them either lack features, such as a complete security scheme or stand alone capability, which may be deemed important or, at the other extreme, are prone to feature creep. Other problems include the lack of portability, dependence on other services or are propriety solutions.

The aim of this paper was to develop Umbra, a darknet suite, which incorporates the good features of some of the current solutions in existence, improves upon their some of their weaknesses and introduces new features. Umbra will address the myriad of problems facing privacy on public domain networks such as the Internet and will target scenarios where interception and analysis of data is probable.

2. Related Work

The word privacy is generally an ill defined but widely understood term. A common correlation amongst the varied definitions is the availability, or lack there of, of information on a given subject or individual whether they be primary or derived via analysis [1]. In some countries, privacy is guaranteed by law or constitution, such as human rights acts or bills concerning electronic data protection. Many cases however see these documents superseded by more recent mandates and edicts passes down by higher echelons like the EU parliament or global agreements like the Anti-Counterfeiting Trade Agreement. In this research, privacy is defined as the right of Internet users to not have their electronic data in any form analysed or processed by any third party. This will include packet inspection, traffic analysis, and any means to discern the methods of privacy assurance such as cryptanalysis or the defeat of tools used to provide privacy. As the current situation stands, there are a number of core problems facing privacy. These range from technical problems such as filtering or traffic shaping measures imposed by governments and Internet service providers to social problems such as a lack of awareness surrounding the subject of privacy.

Filtering is a relatively new problem facing privacy in the public domain. Its introduction largely coincides with the spate of litigations by firms representing the major recording labels against hosting companies, Internet service providers and to some extent BitTorrent tracker sites for the hosting or trafficking of copyrighted music. These services themselves are, of course, not responsible for the copyright infringement however, but the people who are copying these materials and uploading or distributing them. The copyright firms, however, see the service providers as a less diffuse target and therefore focussed their legal attack on them. This has had a unilateral effect on the services, penalising users across the board; both legitimate and illegitimate. For example, Rapidshare's .de domain were the subject of legal action instigated parties representing Sony, Universal, Warner and many other major record labels to remove copyrighted material and implement content filters to prevent copyrighted material from being uploaded. The company also agreed to submit the user information of the uploader to avoid paying hefty fines [2]. The issue has not affected file hosting services only. The onslaught of the record labels has extended to BitTorrent tracker sites and Internet service providers also. Websites such as Mininova have installed a content recognition system which detects and removes torrent files linking to copyright infringing material in a move to mitigate the looming legal ramifications from BREIN, an anti-piracy foundation. Largely the problem with these content filtering systems is that they are prone to producing false positives. Without human interaction, copyright free data such as free documentaries, royalty and copyright free music or even home movies could be flagged and removed merely because they have a title similar to a box office movie or top of the pops hit. Noteworthy is the fact that the installation of this software did not satiate BREIN's legal pursuit against the website [3].

Network neutrality, an increasingly discussed topic as of late, is closely related to privacy in the public domain. The incorporation of traffic shaping measures into Internet service providers networks has been highly controversial and constitutes a violation of network neutrality and privacy. The concept of traffic shaping essentially entails analysing protocol headers of data packets to determine their priority given a particular set of heuristics, then routing them according to their ranking by the outcome of this analysis. What results is a tiered service, where the traffic of protocols they deem of high priority are routed before those of other protocols. This has become common practice for many Internet service providers to implement what they stipulate is a scheme to ensure quality of service to all customers. At times of peak network usage, however, this hierarchical scheme begins to become a hindrance to lower prioritised services.

American Internet service providers, such as Comcast, openly defend their use of traffic shaping, with particular reference to the BitTorrent protocol [4]. They insist that this is a measure to promote unilateral quality of service to all customers at all times of the day, and to a lesser extent, help combat the distribution of copyrighted material. BitTorrent, the now notorious peer to peer protocol geared for the distribution of large amounts of data in peer to peer networks, became a natural target for discrimination in traffic shaping schemes. Not only has it become increasingly associated with the distribution of pirated software, films and music, but also accounted for about fifty to ninety percent of Internet traffic in 2008/2009 [5, 6]. To Internet service

providers, this means congestion. While the content of traffic is not yet directly analysed, the filtering of the protocols being used means a violation of privacy for users.

In countries such as China and Iran, citizens are subject to heavy electronic censorship measures at the hands of their governments. There are a number of factors which contribute to the success of censorship policies. An overbearing government presence seems to be centric to begin with. Countries where people can be imprisoned or executed for merely expressing their opinion in public usually coincides with a national censorship scheme to silence any would-be opposition to the government. These schemes extend from vetted television broadcasts to the installation of government surveillance software on computers, as is the case in China, to the blocking of sites where anti-government propaganda and a platform for free speech reside, as is the case in Iran. State sponsored privacy violation is not limited to the Eastern hemisphere, however. The Culture of Fear [7] spouts phrases such as "national security" and "terrorism" allow governments to pass dubious anti-privacy legislation in countries such as Ireland, the UK and the USA. While these issues are non-trivial, they are abused to spread fear, uncertainty and doubt amongst citizens to allow the installation of increasingly pervasive censorship measures.

3. Umbra DarkNet System Design

The darknet developed in this project uses a hybrid centralised network topography. This consists of a server, whose sole purpose will be to authenticate clients and grant them access to the darknet, and peer to peer clients which will make up the darknet itself (See Figure 1). A centralised network suits real world deployment scenarios; A single entry point into the darknet is advantageous as it negates the requirement for clients to know the IP address of an online node as is the case in a purely decentralised topography. Using a central authentication server results in only trusted users being able to gain access to the darknet and also facilitates for simple addition and removal of trusted users should the web of trust break down. Registered darknet accounts also prevent spoofing as is possible on WASTE Again meshes.

The design of this darknet provides some immunity to server failure. The client application is a hybrid client/server. This means that independent from the server, the client can receive darknet connections so in the event of server failure the darknet stays up. Server failure may be detected when an offline client attempting to enter the darknet is presented with a network error indicating the server is unreachable.

The security of a successful darknet is multi-faceted. Trust and privacy are two of the core areas of the design

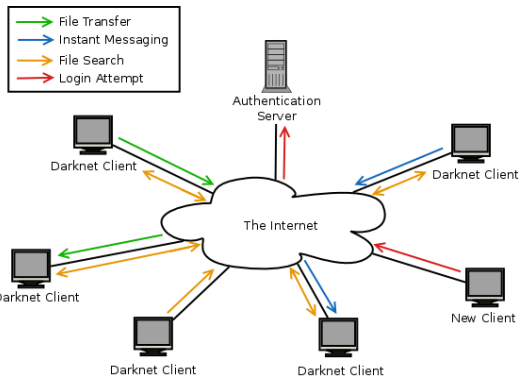

Figure 1. An overview of the Umbra darknet.

in this regard. As mentioned in section 1, cryptography is used to ensure the private and secure transmission of data between darknet nodes, but it also lends itself to the integrity of trust on the darknet as a whole. Darknets can be described as being built on a web of trust. While users at nodes may have personal relationships and trust one another, it is important from a security point of view to defend against subversion of the darknet and mitigate the impact of node compromise. Umbra makes provisions for security on these two fronts by implementing a hybrid cryptosystem consisting of a symmetric cipher for bulk data encryption and a public key algorithm for symmetric key distribution and peer verification. A more detailed discussion of this system is presented in section 4.

To ensure secure transmission of inter-nodal data, the design of Umbra includes the use of end to end encrypted streams. This makes only the sender and receiver privvy to the content of the data being transmitted. This is an improvement over the WASTE Again client which uses link level encryption, whereby any node on the darknet can intercept and interpret data, which creates a massive vulnerability to a "break-once-break-everywhere" (BOBE) attack [8, 9]; the compromise of a single darknet node compromises the entire darknet. Further improvements are made on the WASTE Again client with respect to the cryptographic layer of the darknet. The design of Umbra darknet improves security further by implementing the symmetric cipher used in the cryptographic layer in CBC mode, as opposed to PCBC mode which WASTE Again uses and is known to have vulnerabilities [10].

In designing Umbra, the breakdown of the trust relationship in the darknet is accommodated for. The popular WASTE Again client does not provide a means of removing compromised or untrusted nodes from a mesh. Instead, the entire mesh must be torn down and rebuilt. Umbra, however, allows the simple addition and removal of registered users' accounts from the authentication server, thus preventing their future entry into the darknet.

Security, although paramount, is not the only concern of a successful darknet application; usability plays an important role also [11]. As already discussed, the addition and removal of trusted nodes is made simple by the addition and removal of users' accounts on the server (a set of bash shell scripts aid the performance of these tasks in the current prototype). This is not the only area in which Umbra delivers usability, however. An intuitive user interface is vital to make the functionality of the client available to a wide demographic of users. Umbra boasts a standalone built in GUI which sought to improve on the shortcoming of the Freenet darknet client requiring a web browser to serve this purpose. Much time has been dedicated to the design of the frontend. The use of tooltips, meaningful informational messages and an overall clean, lightweight interface delivers a user friendly experience to users of varying technical proficiency.

In their paper, Biddle *et al* [8] suggest the functionality of a darknet should facilitate the introduction of objects into the darknet, the ability to search remote users' repositories for files of interest and the copying of these files across the darknet to their local machine. The design of Umbra is based on these prerequisites. Users have a shared directory on their local machine through which they make files available to the remainder of the darknet. These directories are remotely searchable by other users on the darknet and the files therein may be copied to a user's own machine in a secure and private manner via the darknet. In addition to these functions, an instant messaging feature has also been included in Umbra and which is secured in the same fashion as searching and file transfers.

4. Implementation and Evaluation

Umbra is written in C and C++ with the current suite of user space tools written in bash shell script. The OpenSSL cryptographic library and Nokia's Qt framework provide the cryptographic functionality and GUI respectively.

The cryptographic and networking layer is central to the secure functionality of Umbra. Not only does it ensure the secure transmission of data across the darknet, but it also preserves the integrity of the web of trust. By encrypting data, it's content known only to the sender and receiver and the absence of plaintext application layer packet headers in data sent over the Umbra darknet provides some immunity to application layer based traffic inspection schemes. Planned future improvements in the area of immunity to traffic analysis are discussed in section 6.

As previously mentioned in section 3, Umbra uses a hybrid cryptosystem consisting of a symmetric cipher for bulk data encryption and a public key algorithm for key distribution and peer verification. The latter functionality is handled in a secure handshake, which is discussed later in this section. Symmetric ciphers are far less computationally expensive than public key algorithms making them suitable for encryption of large amounts of data. Umbra uses symmetric encryption for all traffic once the secure handshake has taken place. Early prototypes used the Advanced Encryption Standard (AES) as a symmetric cipher. In the duration of development, the news of the cryptanalysis of AES192/256 [12] led to the long term security decision to switch to the Blowfish cipher of which, currently, there is no publicly known cryptanalysis.

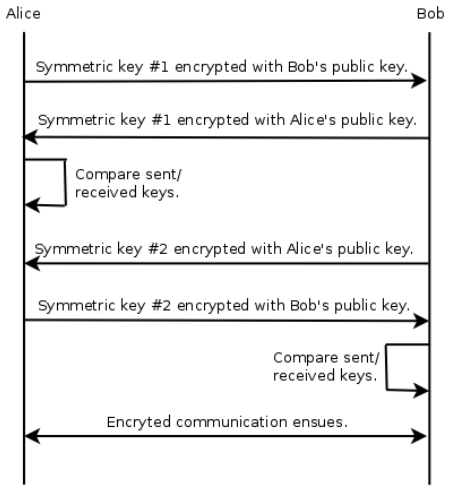

Figure 2. The Umbra connection handshake.

A public key algorithm, RSA, is used for key distribution and peer verification. Both tasks are carried out as part of the secure handshake, which is negotiated when a new connection is made on the darknet, either between clients on the darknet, or between client and authentication server. The premise of the handshake, illustrated in Figure 2, is simple. The symmetric keys and initialisation vectors used for symmetric encryption are derived from a cryptographically secure, high entropy source, namely the OpenSSL pseudo-random number generator (PRNG). Thus, these keys are extremely difficult to precompute which makes them suitable as tokens for peer verification. The parties exchange symmetric keys and initialisation vectors, one pair for each direction of communication, using one another's public RSA keys. If the returned copy of either matches the sent copy, both parties can determine one another's identity and continue the transaction securely using symmetric encryption with the exchanged keys.

The Umbra server application acts as the entry point into the darknet. To gain access to the darknet, clients must first authenticate successfully with the server. The login mechanism is two phase. Clients must first successfully handle the handshake challenge from the server to verify their identity and then password authenticate to the server with the user-

name and password associated with their registered account on the server. This two phase mechanism mitigates the risk of node compromise as an attacker would need to obtain a user's password in addition to a copy of their private RSA key in order to gain access to the darknet and masquerade as that user. As an additional security measure, passwords are stored as SHA256 hashes to prevent attackers obtaining this information in the event of a successful breach of the machine running the server. When clients successfully authenticate to the server, the server then synchronises the client with the details of the other online darknet clients and notifies the other online clients of the new client's entry into the darknet.

The Umbra client is the application through which the user interfaces with the rest of the darknet. After successful authentication with the server, the client becomes a peer to peer application which accepts and handles connections from other nodes on the darknet. The darknet functionality outlined in section 3 is implemented in the client application: Clients may share files with other users on the darknet by copying the file into their shared directory on their local machine; Clients may search one another's shared directories remotely for files of interest; Clients may copy files of interest to their own machine via the darknet; Clients may send instant messages to one another. All of this functionality is performed on top of the cryptographic and networking layer which ensures all data is transmitted securely and privately over the darknet.

Testing and evaluation of Umbra was carried out in both LAN and WAN environments on a relatively small darknet consisting of 20 nodes. A range of tests were carried out from the domain of security, usability, performance and functionality in the most objective manner possible. Evaluation was based on the results of these tests combined with user feedback.

Overall, Umbra performs quite well. The security of the cryptographic and networking layer and the server authentication mechanism has proven resilient to attack. The extent of the security tests included masquerade attacks, simulated node compromise and in the case of the server, replay attacks. The strength of the security of the cryptographic algorithms used is of course assumed secure until a breakthrough in quantum computing arrives or cryptanalysis of these ciphers is successful.

From a performance perspective, the overhead of the cryptographic and networking layer on the client side is not noticeable under average load, but becomes more significant under heavy load, for example, when several file transfers are being executed concurrently in the client side application. The threading architecture used is currently under review in the aim to remedy the impact of heavy load on the host machine.

5. Discussions

This paper has outlined the development and deployment of the Umbra darknet and demonstrated it's potential effectiveness against the encroaching threats to privacy on public domain networks. Umbra aims to continue to improve on the shortcomings of current darknet solutions and introduce new features to broaden it's defence against the privacy threats highlighted in sections 1 and 2.

There are, of course, social ramifications to the use of darknets on public domain networks. While the Culture of Fear does spread fear, uncertainty and doubt and uses these as tools to perpetuate the violation of privacy, such issues as child protection and terrorism are legitimate concerns. There is the possibility that child pornography rings may use darknets as a tool for the distribution of illegal materials or terrorists using them to disseminate instructions to carry out attacks. This possibility, however, should not herald the ban or stigmatisation of darknets; These criminals may also use the same technologies that the public use to carry out financial transactions online, or access online medical records yet these technologies are not subject to the same scrutiny as darknets.

The so-called rise of piracy in the advent of high speed Internet connections has been used as a scapegoat for the current economic climate to enforce unfair and legally dubious sanctions on legitimate services such as third party file hosters and BitTorrent trackers. Darknets provide a suitable replacement for these services and allow users to share data with only the people they trust. This implies the suitability of darknets in the corporate, education and home domains.

6. Future Research

Currently, Umbra is still in alpha stage of development. With just a few fixes and optimisations to be implemented, the project is nearing a beta release. There are many planned additional features to come in future versions of Umbra. As mentioned in section 4, immunity to various forms of traffic analysis will be implemented in a later version. These include link saturation and protocol obfuscation, two features implemented in the WASTE Again darknet. The planned link saturation mechanism will consist of false connections and junk data exchanges as well as junk data sent over legitimate darknet connections. This will prevent the ability of using traffic analysis to determine when legitimate data is being sent or deriving the size of a file or message as it is transmitted across the darknet. The ability for Umbra to obfuscate it's protocol will prevent it's detection by packet inspection mechanisms and avoid throttling by traffic shaping schemes.

Restrictive firewalls which allow access to only specific services present a problem to darknet users such as those

in countries ruled by oppressive regimes who implement pervasive censorship schemes. Reverse connections would allow users to instead make outgoing connections to other darknet clients wishing to connect to them. This feature will be implemented in a future version of Umbra.

Another major feature which would add an additional layer of privacy is support for anonymous proxy services such as the Tor onion router. Alongside the highly secure cryptographic and networking layer in Umbra, support for anonymity would further mitigate the impact of node or server compromise as the real IP address of a node would be masked by the proxy service. Combined with the other proposed features in this section, Umbra will deliver a robust, highly secure and private darknet solution to the open source community and make the preservation of privacy accessible to the widest demographic possible.

7. Conclusion

Recently, privacy has become scarce on public domain networks such as the Internet. Threats such as filtering, traffic shaping and large scale pervasive censorship schemes have left many users without refuge for securing their data as it is transmitted over these networks. While it is foolish to forecast the future as the privacy and security landscape is ever changing, the current trends in online privacy indicate that darknets will play a role as a "last bastion" for users who feel their right to privacy is threatened.

Umbra aims to be a valuable asset in the toolkit of anyone concerned with proactive protection of their privacy amongst the emergence of global anti-privacy action and legislation. These threats are ever evolving and inevitably open new avenues of privacy protection. The development of Umbra will be ongoing and will strive to defeat or circumvent these new threats as they arise. As a wise man once said;

> *"The Net interprets censorship as damage and routes around it."*
>
> – John Gilmore, Co-founder of the EFF.

8. References

[1] J. Waldo, H. Lin, and L. Millet. *Engaging Privacy and Information Technology in a Digital Age*. The National Academies Press, 2007. ISBN 978-0-309-10392-3 (hardcover) ISBN 978-0-309-66732-6 (pdf).

[2] EDRI. RapidShare needs to check every file for copyright infringement. EDRI-gram - Number 6.19, Last accessed: Monday November 9 2009, October 8 2008. http://www.edri.org/edrigram/number6.19/rapidshare-hamburg-decision.

[3] "Ernesto". Mininova Filters Copyright Infringing Torrents. Last accessed: Monday November 9 2009, May 6 2009. http://torrentfreak.com/mininova-filters-copyright-infringing-content-090506/.

[4] D. McCullagh. Comcast to FCC: We block only 'excessive' traffic. Last accessed: Monday November 9 2009, February 13 2008. http://news.cnet.com/8301-13578_3-9871287-38.html?tag=nefd.top.

[5] Ipoque. Internet Study 2008/2009 - The Impact of P2P File Sharing, Voice over IP, Instant Messaging, One-Click Hosting and Media Streaming on the Internet. Last accessed: Monday November 9 2009, 2009. http://www.ipoque.com/resources/internet-studies/internet-study-2008_2009.

[6] P. Svensson. Comcast blocks some Internet traffic- Tests confirm data discrimination by number 2 U.S. service provider. Last accessed: Monday November 9 2009, October 19 2009. http://www.msnbc.msn.com/id/21376597/.

[7] J. Giraldo. *Colombia: The Genocidal Democracy*. Common Courage Press, 1996. See Introduction by Noam Chomsky. ISBN-13: 978-1567510874.

[8] P. Biddle, P. England, M. Peinado, , and B. Willman. The Darknet and the Future of Content Distribution. Last Accessed: November 23 2009, 2002. http://msl1.mit.edu/ESD10/docs/darknet5.pdf.

[9] A. Acquisti. Darknets, DRM, and Trusted Computing: Economic Incentives for Platform Providers. Presented at TPRC 2004, 2004.

[10] C. J. Mitchell. Cryptanalysis of Two Variants of PCBC Mode when used for Message Integrity. *Spinger Lecture Notes in Computer Science: Information Security and Privacy*, 3574/2005:560–571, July 2005. ISBN-978-3-540-26547-4.

[11] J. Bethencourt, W. Y. Low, I. Simmons, and M. Williamson. Establishing Darknet Connections: An Evaluation of Usability and Security. *Proceedings of the 3rd ACM International Symposium on Usable Privacy and Security*, 229:145–146, 2007. ISBN:978-1-59593-801-5.

[12] A. Biryukov and D. Khovratovich. Related-key Cryptanalysis of the Full AES-192 and AES-256. Last Accessed: Sunday 14 March 2010, 29 May 2009. https://cryptolux.uni.lu/mediawiki/uploads/1/1a/Aes-192-256.pdf.