The Case of The Plan R* Network of Resistant Communication: Implications to Digital Forensic Investigations

George Chlapoutakis

Digital Forensics Group, School of Science and Technology, University of Teeside, Middlesbrough, TS1 3BA

Abstract

The Plan R* network of resistant communication is the project of the Autistici-Inventati collective, whose purpose is to create an anonymous and self-regulated environment where political and social activism can be discussed without fear of reprisals. Essentially, the free Plan R* network provides a number of services such as mailboxes and remailers, web-space, instant messaging and mailing lists operating in a distributed fashion through Virtual Private Network links, allowing for both user data and communication to be widely dispersed and replicated across the entirety of the network.

While the purpose behind the Plan R* network is benign in nature, its cost-free nature makes it a potential distribution medium for a number of illicit material and a potential staging ground for illicit activities. More importantly, the structure and mode of operation of this resistant communication network presents significant problems to digital forensic investigators, due to the dispersal of both data communication and data storage across a vast number of computers.

The purpose of this paper is to discuss the issues that can arise in a digital forensic investigation involving this network and the implications of those issues to the integrity and validity of the forensic evidence gathered.

Keywords: resistant communication, distributed computing, distributed storage, digital investigations, forensic evidence, anonymity

* Corresponding author.
Email address: g.chlapoutakis@tees.ac.uk (George Chlapoutakis).

1 Introduction

> ”The more you tighten your leash, Tarkin, the more star systems will slip through your fingers!” [1]

Anonymous and censor-resistant communication, publishing and storage services and networks have been a topic of much and very heated controversy involving a number of diverse parties.

Most notably used to provide political, social and ideological activists with a space in which to exchange documents, files and ideas, and organise activities without fear of government or corporate reprisals, such communication networks started attracting attention in the early 2000s with the creation of projects such as Free Haven [2] and FreeNet [3].

They, however, gained much more prominence with the addition of Peer-to-peer file-sharing services such as the nowadays infamous Napster and Gnutella networks. Later on and closer to our current time, the Tor anonymising proxy network [4] was written to allow political dissidents of oppressive regimes to bypass government-set restrictions on free speech on the Internet in the form of Internet firewalls in order to share and exchange information with the rest of the world.

However, and despite a number of different efforts from software designers and users alike to regulate such networks by themselves, those networks also quickly attracted the interest of a variety of offenders who saw the potential anonymous and censor-resistant communication networks have for both communicating with their peers and disseminating material that is considered by the legislature of numerous countries to be illegal to highly illegal and offensive in nature.

The Plan R* network of resistant communication is just such a network. The project of the Autistici-Inventati collective, its purpose is to create an anonymous and self-regulated environment where political and social activism can be discussed without fear of reprisals. Essentially, the free Plan R* network provides a number of services such as mailboxes and remailers, web-space, instant messaging and mailing lists operating in a distributed fashion through Virtual Private Network links, allowing for both user data and communication to be widely dispersed and replicated across the entirety of the network.

In Section 2 we take a look at the literature around censorship-resistant networks, paying particular attention to communication and storage networks, and briefly discuss a small number of examples of such networks. In Section 3 we outline and discuss to some detail the structure, topology, services and level of anonymity and privacy offered by the Plan R* Network. Then, in Section 4, we discuss how the use of the said network in the commitment of a digital crime will affect the investigation that will follow, and what potential issues the investigator will have to face in terms of the integrity and validity of the evidence gathered. Finally, Section 5 concludes the paper with a discussion of findings so far and possible future directions.

2 Censorship-resistant networks

Broadly speaking, censorship-resistant networks are considered to be an extreme example of the problem of network availability, or the ability of a network to withstand massive strains in workloads, such as those generated when a network attack takes place, and hardware and software failures.

In the case of censorship-resistant networks, massive strains in workloads may not necessarily be the result of a network attack but the effect of high accessibility of a piece of information stored on the network or the result of the rest of the network having to share the load of traffic due to a hardware or software failure that would occur when, for example, the UK Indymedia servers were removed by the UK Police in the much-publicised raid described in [5].

2.1 Censorship-resistant storage and publication networks

In terms of storage and publication, censorship-resistant networks primarily focus on storing information in as widely dispersed a manner as possible and in such a way as to ensure that information stored on the network will maintain high levels of persistance, robustness and availability as possible. This, in turn, requires the network to be robust both in terms of the information storage mechanism and in terms of data communication mechanism used for retrieval of the information.

Such networks can have a client-server architecture, as in the case of the Eternity service [6] and later the Free Haven [2,**?**], Publius [7], and Wikileaks [8] services, where the standard is for a list of well-known servers to store encrypted copies of files sent by client applications.

Cryptographic keys are used to both guard against faults or malicious tampering and as an identifier to allow for easy updating of the lists of files, which are published in the form of hierarchical collections.

Use of custom robust and censorship-resistant distributed filesystems such as the Least Authority Filesystem TAHOE [9] has also been observed. In the case of the TAHOE filesystem, Wilcox-O'Hearn and Warner designed the distributed storage grid to provide a cryptographically safe and secure file storage through the use of erasure coding and cryptography.

Censorship-resistant storage networks, however, can also have a peer-to-peer architecture, in which storage robustness and availability is handled by replicating the file across multiple peer-to-peer nodes, such as described in Fiat and Saia's research [10] and data retrieval requests are routed to the appropriate node depending on the geographical distance and network latency metrics each specific peer-to-peer storage network uses.

Examples of such storage networks include the FreeNet peer-to-peer storage network [11,3], the OceanStore network [12] and the Rosebud fault-tolerant storage network [13], as well as the Cooperative File System peer-to-peer network filesystem [14].

2.2 Censorship-resistant communication networks

Censorship-resistant communication networks have traditionally focused mainly on maintaining e-mail anonymity through a number of means, such as anonymous remailers, whose primary function is to conceal the sender's identify from the e-mail message itself, nowadays using techniques such as random or user-specified timing delays, or techniques based on the mix network design [15].

As censorship-resistant communication networks have been enlarged in scope to allow more services to be anonymised, such as Instant Messaging, Web-browsing and IRC (Internet Relay Chat) session services, the focus changed to assuring that not just the sender, receiver and the message were anonymised but also the route all the above took through the Internet.

This brought into prominence a network routing strategy that was originally developed in the 1990s but only became increasingly popular later on, as the computational power and complexity of modern computers increased. Onion routing [16] is a low-latency network that is based on the TCP protocol and delivers content between anonymous parties by tunneling them random-sequentially through multiple non-anonymous but encrypted links.

The nowadays very popular Tor anonymisation network [4] uses an updated and modified version of Onion Routing that increases entropy by inserting dummy packets partway along the tunnel or drops packets at random points of the tunnel in an attempt to make traffic analysis of the session harder.

Finally, I2P [17] is another anonymising network that is based on Onion Routing and Tor to ensure secure anonymity between parties. Its main difference to Tor, however, is in that it not only relies on circuit-based peer-to-peer networks (where communication is established between a close circle of ”friends” peer-to-peer nodes, each of which has other ”friends” as well, ad infinitum and communication passes through random ”friendly” nodes), which increases the randomness of the links, but also publishes all information including link node membership information, in a distributed cryptographic hash table under a public key, thus removing any linking information from the public eye.

3 The Plan R* Network

The Plan R* Network, as was outlined above, provides services such as mailboxes and remailers, web-space, instant messaging and mailing lists to interested parties. In order to do so it adopts a number of different types of structural and topological strategies and uses a combination of the censorship-resistant storage and communication methodologies and techniques we have discussed in the previous section.

3.1 Structure and topology

Based on Autistici/Inventati's plan explanation [18] and their how-to [19], it appears that the Plan R* Network's architecture adopts a combination of client-server and peer-to-peer architectures, depending on the type of data and the type of service offered.

The public data, such as web pages, is stored in a network that uses a peer-to-peer architecture, while private data such as mailboxes are stored in a client-server network.

Communication between the different servers, in a clear divergence from the previous methods discussed, is handled through VPN (Virtual Private Network) links between servers.

Thus far it can be seen that Autistici/Inventati's theoretical model for a censorship-resistant communication network attempts to combine storage, publication and communication architectures based on the strengths of each architecture in dealing with different types of data.

Configuration management, which is used for the administration of the Plan R* Network's infrastructure and is one of the most problematic aspects in terms of set-up of a distributed censorship-resistant communication network is handled centrally, through the use of the CFEngine [20] service, and changes are also stored in a CVS (Concurrent Versioning System) server. Those two

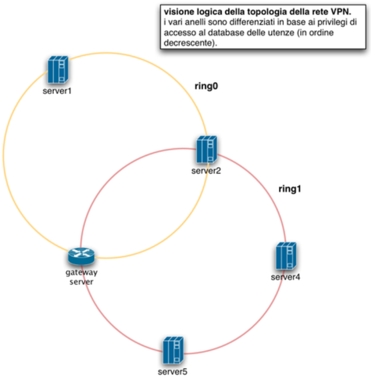
visione logica della topologia della rete VPN.
i vari anelli sono differenziati in base ai privilegi di accesso al database delle utenze (in ordine decrescente).
server1
ring0
server2
ring1
gateway server
server4
server5

Fig. 1. VPN connectivity in the R* Plan Network (http://dev.autistici.org/orangebook/html-en/images/rings500.png)

points of centralisation could be construed to represent relative weak points in the communication network.

3.2 Level of anonymity and privacy offered

Again according to the Autistici/Inventati's how-to [19], anonymisation of data occurs only at log-file level. Specifically, their anonymisation method consists of customising the services running on the network to either maintain log files lacking information that may identify users, such as IP addresses, or filtering the log files in memory to remove sensitive information.

With regards to the protection of the privacy of the data offered, now, there appears to be no user data encryption, as it was deemed impractical at this point in time due to the project's infancy, but encrypted partitions are used to store system/network-related information such as the SSL certificates and keys and the implementation of encrypted partitions is based on dm-crypt.

Otherwise, regarding the encryption and protection of the privacy of the users of the services offered, the Autistici/Inventati collective points the users to external sources discussing privacy and cryptography software they can use to secure their communication from their computers to the Plan R* Network servers.

4 Inferences, implications and issues for the digital forensic investigator

Based on the information provided in the previous two sections regarding censorship-resistant communication networks, and resilient and robust storage and communication networks in general, we can draw a number of inferences on the structure of this network that will then allow us to discuss the implications of its structure from a digital forensic investigator's point of view.

While the Plan R* Network utilises a combination of architectures in a clear effort to cater for both private and public data, it does so only at application level, through the use of an unencrypted LDAP database and a number of equally unencrypted rsync remote file synchronisation sessions, as well as a mysql database replication session. Encrypted communication between servers is, thus, handled solely at network level, through the use of VPN links between the servers.

Furthermore, the keys for the network encryption and the SSL certificates for the encrypted communication sessions are stored in an encrypted partition in the disk of each of the servers. However, no user-generated or stored content is encrypted in each of the servers. The Autistici/Inventati's howto suggests that all the services offered by the Plan R* Network all rely on SSL certificates, so the level of application-level protection is safeguarded by the encrypted storage space provided.

The process of data anonymisation occurs only on the log file level and involves reducing the level of log file detail, in a clear effort to minimise the amount of information that can be seized in the event of a government-sanctioned raid on any of the server locations.

It can also be seen that the configuration management of the entire network is centralised in the form of one server containing the CFEngine and another server containing the CVS versioning system. There is no evidence to suggest either that those two servers are geographically dispersed or that those two servers are utilising some form of system encryption or encrypted network communication. However, there is also no evidence to suggest that the data stored on these systems are user-generated/stored content.

Finally, there is no evidence of the current or intended use of any custom secure distributed filesystem, as in the cases examined in Section 2.

With regards to evidence integrity and validity, the implications to the digital forensic investigator come from the type of communication involved in the offense that is being investigated. If the offense involves personal data as defined by the Plan R* Network, for example e-mail messages, the difficulty lies in establishing a clear trail of communication between offenders or between offender and victim, given the fact that personal data are located in any number of different locations at the same time. Given that the nodes in the network all communicate in an encrypted fashion and the network services all implement the aforementioned forms of log file anonymisation, establishing a trail that can prove a connection or a link between offenders or between offender and victim via a specific network node becomes infeasible.

The same applies for the evidential integrity and validity of public data, which is transmitted in a peer-to-peer fashion, although that depends on the exact mechanism by which the peer-to-peer network works in the case of the Plan R* Network.

5 Conclusions and future work

Massively distributed communication and storage networks of the type of the Plan R* Network have only recently come into focus in terms of digital forensic investigation requests, mainly due to their distributed nature and the amount of time and effort required to investigate such networks, especially in the case of very advanced and state of the art distributed communication and storage networks.

The case of the Plan R* Network is particular in that, for one it is still in its infancy and not grown to its full capabilities. It demonstrates a number of interesting characteristics in the form of its network architecture and internal structure that show a wish to combine technologies from both client/server and peer-to-peer architectures in order to improve on the privacy, security and resiliency of the communication passing through it, and the use of Virtual Private Network point-to-point encryption between nodes shows it to be an experiment in the use of multiple architectures, structures and methodologies in order to better use their strengths.

In Section 2 of this paper we took a look at the literature around censorship-resistant networks, paying particular attention to communication and storage networks, and briefly discussed some of those networks, which allowed us to better understand the ways in which the Plan R* Network differs from other such networks by discussing its structure, topology and level of privacy and anonymity offered in Section 3. Then, in Section 4, we drew some inferences based on Section 3's analysis of the Plan R* Network in order to be able to discuss how the use of the said network in the commitment of a digital crime will affect the investigation that will follow, and what potential issues the investigator will have to face in terms of the integrity and validity of the evidence gathered.

A better and more concise analysis and discussion of findings certainly needs to be performed in the case of the Plan R* Network, but that will need to occur at a point in time where the network has matured to such degree so that we can see the interplay of those technologies and architectures in a clearer manner.

Censorship-resistant communication networks certainly merit a much more detailed research and analysis from the point of view of the digital forensic investigator. As much as privacy and anonymity should be protected and preserved, that self-same requirement can be and has been misused and misappropriated for the commitment of a number of diverse offenses in the past, and undoubtedly will continue to be misused and misappropriated well into the future.

References

[1] G. Lucas, Starwars episode 4: A new hope, 20th Century Entertainment, US (May 1977).

[2] R. Dingledine, M. Freedman, D. Molnar, The freehaven project: Distributed anonymous storage service, in: Proceedings of the International Workshop on Design Issues in Anonymity and Unobservability, Berkeley, CA, USA, 2001, pp. 67–95.

[3] I. Clarke, O. Sandberg, B. Wiley, T. Hong, Freenet: A distributed anonymous information storage and retrieval system, in: Proceedings of the International Workshop on Design Issues in Anonymity and Unobservability, Berkeley, CA, USA, 2001, pp. 46–66.

[4] R. Dingledine, N. Mathewson, P. Syverson, Tor: The second-generation onion router, 13th USENIX Security Symposium, San Diego, CA, USA.

[5] EFF, Indymedia server takedown, `http://www.eff.org/cases/indymedia-server-takedown/`, Last seen: 23/04/2009 (August 2005).

[6] A. Back, The eternity service, Phrack Magazine 7 (51), `http://www.cypherspace.org/adam/eternity/phrack.html`, Last seen: 23/4/09.

[7] M. Waldman, L. Cranor, A. Rubin, Publius, O'Reilly, 2001, Ch. Peer-to-Peer: Harnessing the Power of Disruptive Technologies, chapter 11.

[8] Wikileaks, `http://wikileaks.org/`, Last seen 24/04/2009.

[9] Z. Wilcox-O'Hearn, B. Warner, Tahoe - the least-authority filesystem, in: 4th International Workshop on Storage Security and Survivability (StorageSS 2008), Alexandria, VA, USA, 2008.

[10] A. Fiat, J. Saia, Censorship resistant peer-to-peer content addressable networks, in: 13th Annual ACM Symposium on Discrete Algorithms, San Francisco, CA, USA, 2002.

[11] I. Clarke, S. Miller, T. Hong, O. Sandberg, B. Wiley, Protecting free expression online with freenet, IEEE Internet Computing, 6 (1).

[12] J. Kubiatowicz, D. Bindel, Y. Chen, S. Czerwinski, P. Eaton, D. Geels, R. Gummadi, S. Rhea, H. Weatherspoon, W. Weimer, C. Wells, B. Zhao, Oceanstore: An architecture for global-scale persistent storage, in: 9th International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS 2000), Cambridge, MA, USA, 2000.

[13] R. Rodrigues, B. Liskov, Rosebud: A scalable byzantine-fault-tolerant storage architecture, Mit-lcs-tr-932, Computer Science and Artificial Intelligence Laboratory, Massachusetts Institute of Technology, `http://publications.csail.mit.edu/lcs/pubs/pdf/MIT-LCS-TR-932.pdf`, Last seen: 24/042009 (December 2003December 2003).

[14] F. Dabek, M. Kaashoek, D. Karger, R. Morris, I. Stoica, Wide-area cooperative storage with cfs, in: 18th ACM Symposium on Operating Systems Principles, Banff, Canada, 2001.

[15] M. Burnside, A. Keromytis, Low latency anonymity with mix rings, Proceedings of the 9th Information Security Conference (ISC 2006), Samos, Greece (2006) 32–45.

[16] D. Goldschlag, M. Reed, P. Syverson, Onion routing for anonymous and private internet connections, Communications of the ACM 42 (2) (1999) 39–41.

[17] I2p anonymous network, `http://www.i2p2.de/`, Last seen:24/4/2009.

[18] Autistici/Inventati, A/i orange book (1.0): An how-to for the realization of a resilient network of self-managed servers, `http://dev.autistici.org/orangebook/html-en/`, Last seen: 23/04/2009 (2005).

[19] Autistici/Inventati, The r* plan: What does it actually mean?, `http://www.inventati.org/en/who/rplan/what.html`, Last seen: 23/4/2009 (01 2007).

[20] Cfengine, `http://www.cfengine.org`, Last seen: 24/4/2009.