Eclipse and Re-Emergence of Anonymous P2P Storage Network Overlay Services

Marios Isaakidis
University College London, eQualitie
m.isaakidis@cs.ucl.ac.uk

George Danezis
University College London
g.danezis@ucl.ac.uk

1. INTRODUCTION

As soon as anonymous peer-to-peer (P2P) storage networks came to be in the early 2000s [8, 9, 11, 14], developers found themselves with a new playground for creating privacy enhancing tools. In many ways different to the now dominant client-server architecture, a whole generation of distributed overlay services emerged, offering a viable option for asynchronous messaging and bulletin boards.

Tor [10] onion services have rekindled interest in anonymous network overlay services and are spawning a new bunch of promising applications. Arguably, a significant subset of them has chosen Tor without properly evaluating the properties of alternative anonymity systems.

In an effort to guide the Privacy Enhancing Technologies developer community towards the best design decisions, we are revisiting the key features of anonymous P2P storage networks; we showcase overlay applications currently in use; and finally share our experience with CENO [1], a distributed censorship circumvention service on top of Freenet [8].

2. ANATOMY OF A P2P STORAGE NETWORK OVERLAY SERVICE

P2P storage networks are, as the name suggests, decentralized information storage and retrieval systems, where participating nodes act as equal peers in replicating the files and routing the requests in the network. Depending on the implementation, they can offer security guarantees such as anonymity for both the producers and the consumers of information, plausible deniability of the files being hosted at a node's datastore, and high availability and persistence of the information inserted.

The latter property supports the goal of censorship resistance; it is virtually impossible to remove every instance of a file from the network. Remarkably, once a user inserts a file, they may disconnect from the network assured that lookups for that resource will be carried out by the distributed cache. Furthermore, since no nodes play a special role in the network, there is no need for a central directory, making them easy to bootstrap from scratch, less prone to IP blacklisting and in overall reduce the Denial of Service attack surface.

The first P2P storage networks originated in Academia with FreeHaven [9], OceanStore [11] and Publius [14], to name a few. The current designs offer incentives in form of a cryptocurrency to the nodes providing resources to help run the system, usually for hosting either websites [12], or private data in decentralized cloud systems [6]. IPFS [4] and Zeronet [7] are recent implementations aiming to support a dynamic, decentralized and self-versioned web.

We focus on Freenet, due to the diversity of the existing overlay services [2], as well as for the security guarantees and application programming interfaces, that make it an ideal candidate for developing CENO.

Compared with the services overlaying low latency anonymous communication systems such as I2P [5] and Tor, which can forward traffic to a hidden server in the network, P2P storage network overlay services rely solely on file insertions and retrievals. This restriction makes the synchronous client-server architecture impractical. By applying public-key cryptography techniques though, developers can create signed subspaces [3], where only the owner of the private key may insert and update information. This method has been extensively used for looking up content inserted by a specific user or service provider.

3. A DIVERSE ECOSYSTEM OF SERVICES

Using the simple operations available, developers created the very first communication applications. Frost [2], a forum-like tool, made sharing Freenet lookup keys possible, thus advertising content which otherwise would have remained undiscoverable. Freemail [2] covered the need for anonymous asynchronous messaging. FLIP-IRC [2] was a synchronous multi-party chatrooms experiment, enabling users to connect to a virtual IRC server and fetch the discussion messages from the decentralized cache. Unfortunately FLIP-IRC experienced long delays, suggesting that onion/garlic routing [10, 5] suits better the real-time communication case.

Eventually, collaboratively editable Wiki systems [2] appeared, followed by the Infocalypse [2] source code management tool. Both empower truly anonymous collaboration and are prominent examples of exploiting the self-revisioning nature of Freenet-like networks.

The lack of centralized services further impedes the process of keyword searches in the distributed storage. Global indexes do exist, maintained by nodes who crawl Freenet sites. Yet the term matching is performed locally by the Library [2] Freenet plugin, after the indexes get retrieved from Freenet.

4. PSEUDO-IDENTITIES AND THE WEB OF TRUST

Distributed trust and spamming in decentralized systems are known issues: Freenet's safeguard, the Web of Trust (WoT) [2], is inspired by Levien's attack resistant trust metrics [13]. Each pseudo-identity, mapped to a public key, may publish a trust score for other identities: positive for those

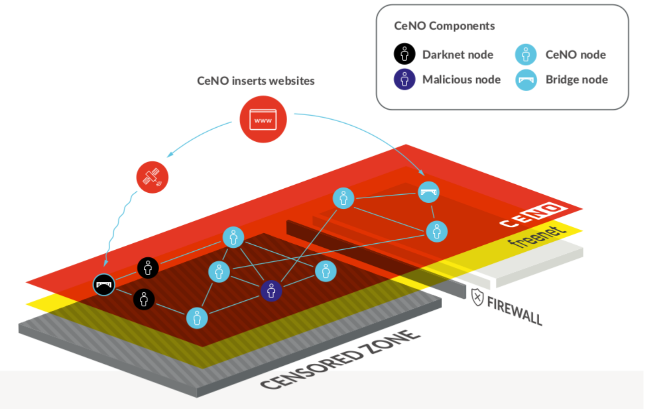

Figure 1: CENO service utilizing the Freenet network for propagating cached versions of websites in censored zones

they trust and negative for spammers. A weighted graph can then be constructed locally at each node.

A small set of globally trusted seed identities assist in introducing new identities to the rest of the graph. In detail, seed identities publish captchas for newcomers to solve, in order to get an initial score. Hereafter, nodes will avoid downloading content inserted by identities with a negative trust score, subsequently marginalizing spammers.

This pseudo-identity system is used as the cornerstone for building one-to-many publishing applications. Services that were heavily spammed, such as the Frost forum, got redesigned around this concept. Meanwhile, spam-resistant micro-blogging and social networking services have become the new standard.

5. CENO: A CLIENT-SERVER SERVICE

CENO proposes a new paradigm in censorship circumvention, relying on Freenet's resilience and security properties. It enables users to anonymously request from specific nodes that have access to the uncensored Web – from now on "Bridges" – to insert static bundles of websites into the Freenet cache, as it appears on Figure 1. Such queries need to be handled by the Bridge nodes only once; subsequent requests are served directly via the decentralized storage. In case of nationwide Internet throttling, users will still have access to otherwise censored content, given that a copy resides in the cache of some other nodes in the same country.

Unlike other censorship circumvention systems, CENO clients need no a priori knowledge on how to contact a Bridge node. The channel establishment mechanism does not require authentication or manual interaction by the users, is efficient compared to services like Freemail [2] and can scale horizontally so as to handle increasing demand. In order to speed up the service, CENO uses a well interconnected cluster of nodes for distributing the tasks as well as for getting advantage of the Small World routing phenomenon.

eQualitie[1] released CENO version 1.0 for Windows, Linux and OS X in April 2016. RSS news feeds from independent newsrooms, human rights NGOs and activists are continuously being cached and are available via the CENO portal.

6. CLOSING REMARKS

For over fifteen years, anonymous P2P storage network overlay services illustrate their capacity in providing sophisticated functionality, overcoming in genuine ways their inherent limitation to only a couple of fundamental operations: insertions and retrievals. Based on our experience, we believe that services being built today, such as whistleblowing and publishing platforms, should consider selecting systems like Freenet for the underlying network. All in all, they are resistant to global adversaries and traffic analysis attacks, while their versatility makes them easier to bootstrap and obviates the requirement of centralized directories. Remarkably, confronting the common sense, distributed services like CENO require fewer resources and become faster as they get widely adopted, considering that the probability of a website being already cached is higher and the respective resources become better replicated in the distributed cache.

On the other side, open challenges remain: supporting dynamic content and synchronous message exchange, achieving satisfactory performance, ensuring the availability of content that is not relevant to a large crowd and, specifically for distributed services, methodically blocking spammers while making provision for scaling up.

Acknowledgements. Marios Isaakidis and George Danezis are supported by NEXTLEAP (EU H2020 ref: 688722).

7. REFERENCES

[1] CENO. `https://censorship.no`.

[2] Freenet overlay services. `https://wiki.freenetproject.org/Projects`.

[3] Freenet Signed Subspace Keys. `wiki.freenetproject.org/Signed_Subspace_Key`.

[4] The InterPlanetary File System. `https://ipfs.io`.

[5] The Invisible Internet Project. `https://geti2p.net`.

[6] Storj. `https://storj.io`.

[7] ZeroNet. `https://zeronet.io`.

[8] I. Clarke, O. Sandberg, B. Wiley, and T. W. Hong. Freenet: A distributed anonymous information storage and retrieval system. In *Designing Privacy Enhancing Technologies*, pages 46–66. Springer, 2001.

[9] R. Dingledine, M. J. Freedman, and D. Molnar. The free haven project: Distributed anonymous storage service. In *Designing Privacy Enhancing Technologies*, pages 67–95. Springer, 2001.

[10] R. Dingledine, N. Mathewson, and P. F. Syverson. Tor: The second-generation onion router. In *13th USENIX Security Symposium*, 2004.

[11] J. Kubiatowicz, D. Bindel, Y. Chen, S. Czerwinski, P. Eaton, D. Geels, R. Gummadi, S. Rhea, H. Weatherspoon, W. Weimer, et al. OceanStore: An architecture for global-scale persistent storage. *ACM Sigplan Notices*, 35(11):190–201, 2000.

[12] N. Lambert and B. Bollen. The SAFE network: a new, decentralised internet. 2014.

[13] R. Levien and A. Aiken. Attack-resistant trust metrics for public key certification. In *Usenix Security*, 1998.

[14] M. Waldman, A. D. Rubin, and L. F. Cranor. Publius: A robust, tamper-evident censorship-resistant web publishing system. In *9th USENIX Security Symposium*, pages 59–72, 2000.

[1] `https://equalit.ie`