

Freenet: A Distributed Anonymous Information Storage and Retrieval System

Ian Clarke¹, Oskar Sandberg², Brandon Wiley³, and Theodore W. Hong^{4*}

¹ Uprizer, Inc., 1007 Montana Avenue #323, Santa Monica, CA 90403, USA
`ian@octayne.com`

² Mörbydalen 12, 18252 Stockholm, Sweden
`md98-osa@nada.kth.se`

³ 2305 Rio Grande Street, Austin, TX 78705, USA
`blanu@uts.cc.utexas.edu`

⁴ Department of Computing, Imperial College of Science, Technology and Medicine,
180 Queen's Gate, London SW7 2BZ, United Kingdom
`t.hong@doc.ic.ac.uk`

Abstract. We describe Freenet, an adaptive peer-to-peer network application that permits the publication, replication, and retrieval of data while protecting the anonymity of both authors and readers. Freenet operates as a network of identical nodes that collectively pool their storage space to store data files and cooperate to route requests to the most likely physical location of data. No broadcast search or centralized location index is employed. Files are referred to in a location-independent manner, and are dynamically replicated in locations near requestors and deleted from locations where there is no interest. It is infeasible to discover the true origin or destination of a file passing through the network, and difficult for a node operator to determine or be held responsible for the actual physical contents of her own node.

1 Introduction

Networked computer systems are rapidly growing in importance as the medium of choice for the storage and exchange of information. However, current systems afford little privacy to their users, and typically store any given data item in only one or a few fixed places, creating a central point of failure. Because of a continued desire among individuals to protect the privacy of their authorship or readership of various types of sensitive information[28], and the undesirability of central points of failure which can be attacked by opponents wishing to remove data from the system[11,27] or simply overloaded by too much interest[1], systems offering greater security and reliability are needed.

We are developing Freenet, a distributed information storage and retrieval system designed to address these concerns of privacy and availability. The system operates as a location-independent distributed file system across many individual

* Work of Theodore W. Hong was supported by grants from the Marshall Aid Commemoration Commission and the National Science Foundation.

computers that allows files to be inserted, stored, and requested anonymously. There are five main design goals:

- Anonymity for both producers and consumers of information
- Deniability for storers of information
- Resistance to attempts by third parties to deny access to information
- Efficient dynamic storage and routing of information
- Decentralization of all network functions

The system is designed to respond adaptively to usage patterns, transparently moving, replicating, and deleting files as necessary to provide efficient service without resorting to broadcast searches or centralized location indexes. It is not intended to guarantee permanent file storage, although it is hoped that a sufficient number of nodes will join with enough storage capacity that most files will be able to remain indefinitely. In addition, the system operates at the application layer and assumes the existence of a secure transport layer, although it is transport-independent. It does not seek to provide anonymity for general network usage, only for Freenet file transactions.

Freenet is currently being developed as a free software project on Sourceforge, and a preliminary implementation can be downloaded from <http://www.freenetproject.org/>. It grew out of work originally done by the first author at the University of Edinburgh[12].

2 Related Work

Several strands of related work in this area can be distinguished. Anonymous point-to-point channels based on Chaum's mix-net scheme[8] have been implemented for email by the Mixmaster remailer[13] and for general TCP/IP traffic by onion routing[19] and Freedom[32]. Such channels are not in themselves easily suited to one-to-many publication, however, and are best viewed as a complement to Freenet since they do not provide file access and storage.

Anonymity for consumers of information in the web context is provided by browser proxy services such as the Anonymizer[6], although they provide no protection for producers of information and do not protect consumers against logs kept by the services themselves. Private information retrieval schemes[10] provide much stronger guarantees for information consumers, but only to the extent of hiding which piece of information was retrieved from a particular server. In many cases, the fact of contacting a particular server in itself can reveal much about the information retrieved, which can only be counteracted by having every server hold all information (naturally this scales poorly). The closest work to our own is Reiter and Rubin's Crowds system[25], which uses a similar method of proxying requests for consumers, although Crowds does not itself store information and does not protect information producers. Berthold *et al.* propose Web MIXes[7], a stronger system that uses message padding and reordering and dummy messages to increase security, but again does not protect information producers.