The Free Haven Project
Distributed Anonymous Storage Service

Roger Dingledine, MIT
Michael Freedman, MIT
David Molnar, Harvard

Differences from Freenet

- Content-neutral
- Focus on storage and persistence
- Stronger anonymity?

Kinds of Anonymity

- Publisher should not be known
- Readers should not be identified
- Location of documents unknown
- Server deniability
- Private information retrieval

Other Aspects of Anonymity

- Linkability: Anonymity vs Pseudonymity
- Partial anonymity
- Computational vs Information-Theoretic Anonymity
- Perfect Forward Anonymity
- Formal models

Some Requirements

- Robust (k of n reconstruction)

- Content-neutral

- Configuration preferences set by each node operator

- Free/Opensource

Some related works

- Gnutella
- Freenet
- Publius

Supported operations

- Inserting (Information Dispersal algorithms)
- Expiring (expiration chosen at publication)
- Trading (currency)
- Retrieving (via public key for document)

Reasons for trading

- Publisher anonymity
- Dynamic (easy to join/exit)
- Longer expiration dates
- Accomodates ethical concerns
- Difficult to track shares

Trading Protocol

- Simplicity vs Accountability (fair exchange protocol, third party)

- Four-message exchange:
  - Two to exchange shares
  - Two to exchange receipts

Accountability

- Shepherd system (one fixed share, one roving)

- Buddy system (both shares roving)

- Receipts

Trust Network

- Broadcast opinions (or changes in opinions)
- Believe broadcasts based on trust
- Accept trades based on trust
- Only risk as much as you've already gained

Mixnet

- Mixmaster III (SMTP)
- Low-latency onion-routing
- ‘Meeting’ servers, keyservers, other extension ideas

Open Problems

- Ideal mixnet
- Accountability and trust algebra
- Modelling and metrics
- Formal definitions of anonymity
- Convenient/useful interfaces

Conclusion

- Implementation
- Legal issues/Education
- No systems are good enough yet