Putting Network Anonymity in the Network

Barath Raghavan, Tadayoshi Kohno, Alex Snoeren, David Wetherall

University of California, San Diego and University of Washington

Goal

- Internet addresses directly identify long-lived hosts; anonymity remains a critical feature currently absent in the Internet
- We aim to provide a high-performance and deployable anonymity service by building it into the network, without re-routing

Background

What are the goals of network anonymity?

- Disassociating the user's identity from her observed network traffic
- Hiding the user's identity from a remote party
- Hiding the remote party's identity from the user

How can network anonymity be achieved?

- Via a set of re-encrypted overlay hops (a Mix network)
- Via a broadcast channel (a Dining Cryptographer network)
- Via masked forwarding through another host or gateway

What are the advantages and disadvantages of anonymity?

- Increased personal privacy; less potential for identity theft (+)
- Promotes a free exchange of ideas and information (+)
- Decreased personalization of content and services (–)
- Potential loss of accountability (–)

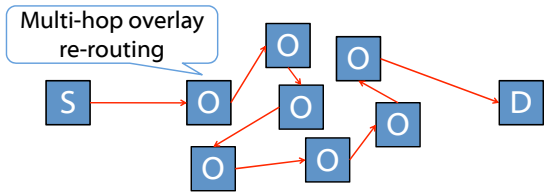

Prior systems

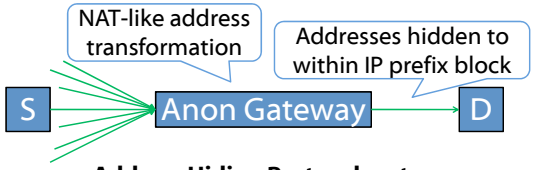

Address Hiding Protocol system

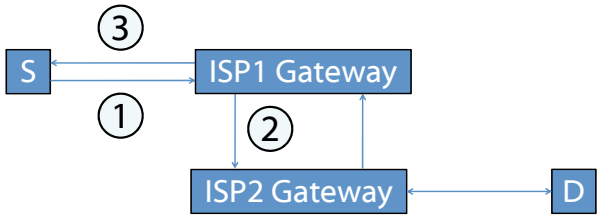

Address Hiding Protocol operation

1. Packet sent S → D, implicit nonce (src port and timestamp bits)
2. Source address encrypted: E_p(suffix of S) where p = E_k(nonce)
3. Destination address decrypted: D_p(suffix of S)

Related Work

Tor (The Onion Router)

- Socket-layer anonymity system; mix-net based
- Re-routing and resource constraints lead to poor performance
- Operates on ~ 1000 volunteer hosts worldwide

The Freedom Network

- Application-layer anonymity system; mix-net based
- Cover traffic, dedicated hosts, and data filters: strong anonymity
- Not in operation (out of business); bandwidth costs prohibitive

Crowds

- Web browser anonymity system; forwarding based
- Probabilistic forwarding, no encryption
- Weaker anonymity, higher performance

Challenges

Performance and Scalability

- Forwarding remains on default paths: low delay
- Simple cryptographic operations: high bandwidth
- Constant state at forwarding gateways

Anonymity sets

- Large ISPs have large IP address blocks for hiding
- Route advertisements can be aggregated to aid

Deployment

- Network service providers have unilateral incentive
- Users suffer little performance degradation
- Compatible with forensic requirements
- No client-side changes; DNS changes for server-side

System Design Considerations

Threat Model: Goals and Assumptions

- Prevent discovery of communication between unsuspect parties
- Hosts and their anonymity provider have a trust relationship
- Protect against a partially-global passive adversary

System Non-Goals

- To prevent insider attacks by those within an ISP
- To prevent long-term intersection attacks
- To provide data privacy or authenticity

External Factors

- Application-layer attacks require application-layer protection
- Timing attack prevention requires cover traffic
- DNS lookup protection requires an external, secure DNS resolver

Current Status

- In-line anonymity has the potential to deliver high performance with clear deployability incentives
- Our architecture can easily be composed with overlay-based anonymity approaches, enabling defense in depth
- No prior systems provide in-network anonymity; design space provides fertile ground for network architecture research

This work is supported by National Science Foundation grants CNS-0722000 and CNS-0722004