Herd: A Scalable, Traffic Analysis Resistant Anonymity Network for VoIP Systems

Stevens Le Blond, David Choffnes, William Caldwell,
Peter Druschel, and Nicholas Merritt.

Proc. of ACM SIGCOMM '15, 45(4):639-652, Oct. 2015.

Jack and Nitish

Voice Over IP (VoIP)

Transmit voice communication over Internet

Differences from telephone network: circuit vs. packet switch

Characteristics:

- Latency sensitive
- Somewhat loss tolerant
- Bursty

VoIP anonymization is desired for increasingly monitored, censored Internet

Freedom of the Net 2014:

"Between May 2013 and May 2014, 41 countries passed or proposed legislation to penalize legitimate forms of speech online, increase government powers to control content, or expand government surveillance capabilities."

"Since May 2013, arrests for online communications pertinent to politics and social issues were documented in 38 of the 65 countries"

https://freedomhouse.org/report/freedom-net/freedom-net-2014#.WERHrXUrI8o

Existing anonymity networks are insufficient to support VoIP

Tor:

Anonymity: vulnerable to adversaries at entry and exit
98.3% of calls can be traced

Performance: long RTT's

Existing VoIP services cannot be easily anonymized

Skype, etc. develop network protocols specifically for VoIP
- Designed for performance, not for anonymity

Can we add anonymity by using other services?
- VPN: obscure IP address, but can be compelled by auth.

Private VoIP services protect content, but can expose callers

RedPhone, etc. provide e2e encryption of VoIP

Rendezvous services can be compelled by auth.

VoIP services do not resist traffic analysis

Traffic analysis: infer comm. endpoints by correlating traffic
- Available data: time series of packets
  user activity

- Statistical analysis can lead to endpoint inference

Why is avoiding traffic analysis important?
- In Tor, 98.3% of calls can be traced

Avoiding traffic analysis: tradeoff between delay and bandwidth

Batch messages together
- Pro: limits diversity in packet time series
- Con: delay packets while waiting to send batch

Add chaff (nonsense) messages to stream
- Pro: garbage packets pollute time series
- Con: waste bandwidth

Herd System Model

- Clients - caller/callee
- Mix
  - Well-provisioned relays
  - Fully connected with other mixes
  - Assumed to be trusted if the client connects to it
- Superpeers
  - Highly available untrusted clients with low latency to the internet
  - Help reduce bandwidth and CPU usage on infrastructure

Herd System Model

- Adversary
  - Seeks to infer caller/callee, contents of call, time and duration of call
  - Able to observe time series of all packets on all global Herd links
  - Able to compromise mixes and network components in a local area
- Herd ensures zone anonymity
  - Knowing a call originated at a zone only tells you that the caller could be anyone in the zone
  - Inter-zone calls maintain zone anonymity independently for caller and callee

Layered Encryption

- All Herd traffic transferred over Datagram TLS links
- Herd uses both hop-by-hop and layered encryption
  - VoIP content is encrypted end-to-end between the caller and callee
  - Provide bitwise unlinkability
  - Hides content and routing information from both individual mixes and eavesdroppers
  - Mixes only know the previous and next hop on the circuit
- Borrows a lot of its cryptographic protocol from Tor

Rendezvous

- Directory server in each zone
- At least one mix for each caller and callee
- For each mix, a random and not necessarily distinct rendezvous mix
- To connect to a client:
  - Caller builds a circuit consisting of a mix and rendezvous mix in a trusted zone and registers the rendezvous zone with the zone directory
  - Callee does the same
  - Caller looks up callee rendezvous mix in directory and starts communicating
- Rendezvous helps ensure zone anonymity as you still don't know the mix that the client is connected to

Herd Architecture

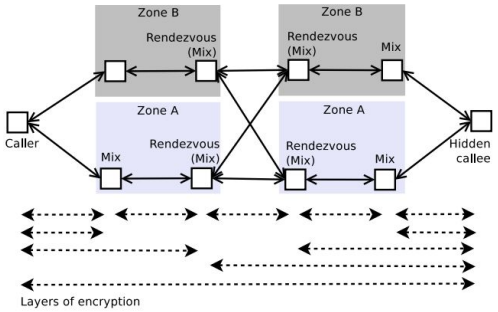

Traffic-analysis Resistance

- Adversary can observe time series of packets on different links and possibly statistically guess a caller or callee
- To defeat this, Herd adds chaff traffic as needed to pad flows so that an adversary can't figure out who is communicating
- On client links, Herd maintains constant chaffing at a rate sufficient for a single VoIP call
- Links connecting superpeers and mixes carry traffic at a rate that is a multiple of the VoIP base rate
  - This can be dynamically reconfigured by the zone directory

Superpeer Architecture

- Helps reduce client-side bandwidth load of mixes by a factor of N / A
  - N is the number of online clients in a zone
  - A is the maximum number of active clients
- Superpeers cannot decrypt the packets it forwards between the client and mix
- Also cannot determine when a client is active

Superpeer Downstream Network Coding

- Superpeer can carry r calls
- Partitions clients into r channels
- Receives packets in rounds of r packets, each packet in a round is forwarded to a different channel
- Active client decodes the packet, rest of the clients in channel consider it chaff

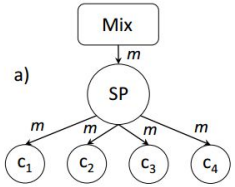

Superpeer Upstream Network Coding

- Superpeer receives packets from each client in a channel
- Forwards to the mix the XOR of all the clients
- Mix can recover the the payload packets from the XOR data as it knows which client is active
- Clients also send manifest data to help make it easier for the mix to decode data in the case of lost packets

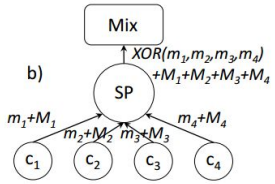

Superpeer Signaling and Channel Allocation

- Signaling
  - For incoming call, mix chooses an available channel which callee belongs to, and encrypts downstream packets in the channel
  - For outgoing calls, client sets the signaling bit in the manifest packet, and the mix responds on one of the channels
- Channel allocation
  - The mix allocates each new client to k distinct channels
  - Herd uses static allocations of clients to channels

Herd Against Various Attacks

- Passive traffic analysis attack
  - Traffic rates on links do not depend on payload flows
  - Chaff helps avoid this attack
- Active traffic analysis attack
  - Mixes can ban superpeers and clients that misbehave
- Compromised caller/callee
  - Rendezvous mechanism ensures that compromised clients can't learn anything other than the zone of their communication partner
- Long-term intersection attacks
  - Ineffective as clients tend to be online most of the time
  - All clients also send chaff data

Herd Against Various Attacks

- Sybil attacks
  - Adversary controls a large number of clients or superpeers
  - Herd is susceptible to this kind of attack if the adversary controls a large portion of the clients

Evaluation Metrics

Anonymity

Call Quality

Scalability/Bandwidth Consumption

Cost of Maintenance

Evaluation Setup

Use traces from mobile calls, Twitter, and Facebook

Compare Herd against Drac and Tor

Drac: anonymous comm. by using social network
- Rendezvous packets through exposed friends
- One chaffing conn. for each link in social network

Anonymity

Anonymity set: number of clients that could be the party

Drac: number of clients reached through *H* hops in SN
- Focus on *H* = 1, 2, 3 (*H*=0 → anon. set of 1)

Herd: number of clients assumed in single zone

Tor: number of clients determined through intersection

Herd provides comparable/better anon. than Drac

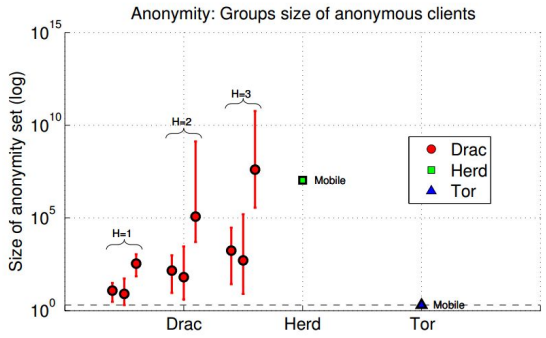

Herd provides comparable/better anon. than Drac

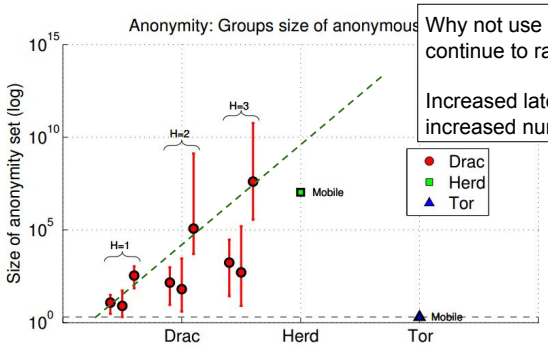

Why not use FB and continue to raise *H*?

Increased latency for increased num hops

Herd provides comparable/better anon. than Drac

Unrelated to this paper: Does Facebook have more-connected social networks?

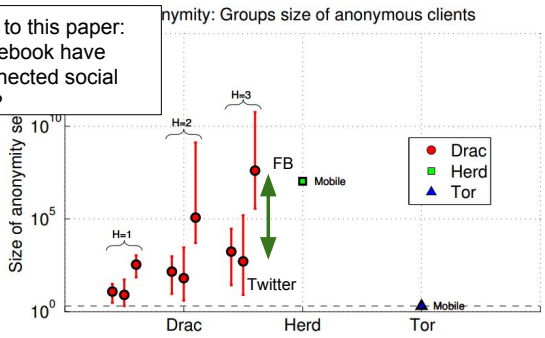

Herd has lower bandwidth requirements than Drac

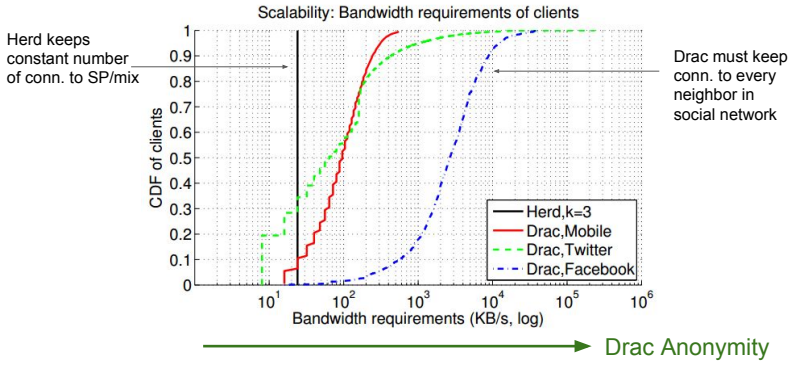

SPs reduce mix utilization, cost of operation

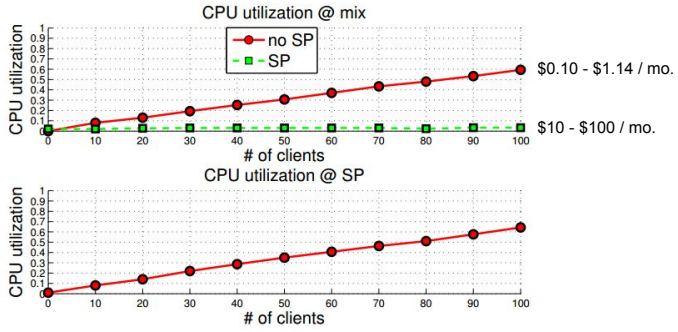

Herd incurs additional latency over Drac (H=0)

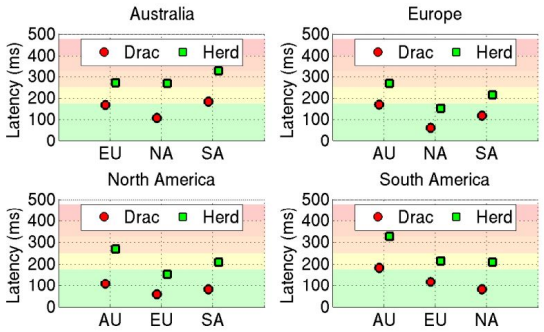

NOTE: With Drac H=0 means that there is no anonymity, no intermediate peer hops.

Summary

Herd provides traffic analysis-resilient anonymous VoIP system

Route calls through trusted zones selected by caller/callee

Offload computation to untrusted, well-connected superpeers

Lower bandwidth, higher anonymity, comparable call quality to existing systems

Discussion

- Authors give illusion that superpeers cause bandwidth/CPU util to “disappear”
  - CPU utilization is simply transferred from mix to superpeer
  - Cost reduction assumes that superpeers charge very little
  - How likely is it for superpeers to be willing to participate?
  - Seems that real contribution w.r.t. superpeers is using untrusted infra. to maintain anonymity
- Make a lot of general assumptions
  - Client can always access mix in another zone
  - Governments can collude to access significant parts of the internet