Security Issues and Solutions in Peer-to-peer Systems for Real-time Communications

draft-schulzrinne-p2prg-rtc-security-00

Henning Schulzrinne
Enrico Marocco
Emil Ivov

Overview

- Attacker motivations
- Attacker resources
- P2P for real-time (vs. file sharing)
  - more than just a DHT

Attacker motivations

- Disrupt communications
  - extortion, dislike, political, …
  - incumbent operator?
- Financial gain
  - impersonation
  - theft of service
  - spamming (SPIT)
- Fun & fame

Attacker resources

- Identities:
  - IP addresses
    - if used for DHT position
    - user subscription limitations
  - mobile phone #'s
  - email addresses, …
- Computational resources
  - botnets make proof-of-work largely useless

Attack timing

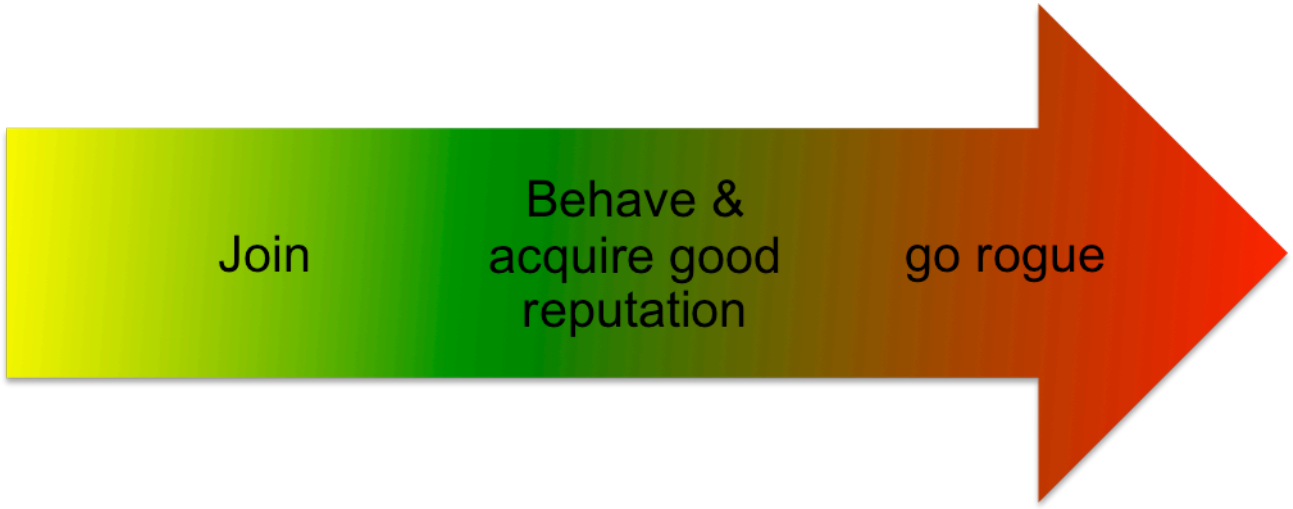

Review: P2P for real-time

- Map names to other identifiers
  - sip:alice@example.com → alice@128.59.16.1
- Provide (computational) services
  - proxying (registration, services)
  - relaying (NAT traversal)
- Store data
  - configuration data
  - voicemail

File sharing vs. real time

	File sharing	Real-time
Distributed database	file location hundreds or thousands per user	User locations: one per user
Availability	same file, hundreds of copies	each user is unique
Integrity	poison file store with bogus material → but no direct threat to user	impersonate user → compromise user communication integrity
Confidentiality	Files are public (may want to hide origin)	Communications is private (src/dest & content)

Admission control

- Goal: keep rogue percentage low
  - allows detection, voting, bypassing
- Group charter + group authority
  - authority certifies candidates compliance with charter
  - central authority or voting
    - how practical in semi-anonymous systems?
    - what information can votes be based on?
    - ballot stuffing by compromised nodes
- Use CAPTCHA to reduce impact of bots
- RELOAD (and Skype) uses central authority

Position in overlay

- Sybil attacks do not depend on identifier
  - but preventing nodes from choosing location randomizes attacks
- IP address or identifier provided by central authority
  - IP address doesn't work well for NATed devices
  - Allows attacker more choice
- Use temporary identifiers?
  - randomizes attack targets
- Use diametrically opposed IDs to avoid local collusion
  - rogue nodes can add neighbors

Identifying malicious peers

- Proactive
  - use test cases to detect misbehavior
  - “mystery shopper”
- Reactive
  - detect and report misbehavior
- Reputation management
  - mostly investigated for file sharing
  - difficult to prevent another denial-of-service attacks of rogue nodes
  - transitive trust

Real-time services are different

- Don’t need everyone to be a peer
  - just enough resources to get job done
  - just increases routing latency (log(N))
  - increases chances of corruption
- Typically, promote nodes from clients to peers
  - use invitation, rather than self-promotion
  - based on uptime, resources, public IP address, geographic need
- Why would a client want to become peer?
  - Skype: closed → (almost) no choice
  - Open systems: incentives ←→ randomized promotion for sybil prevention

Attack

- Denial of service
  - black hole signaling or media
  - fictitious error responses (“no such number”)
  - use iterative routing – getting closer?
- Integrity of location bindings
  - Identity-based crypto → non-intuitive identifiers
- Integrity of content (voice mail, …)
  - generally, only inserter needs access

Summary (& my take)

- P2P systems for real-time applications ≠ file sharing
  - more than just key → value mapping
- Identity scarcity is crucial
  - leverage existing hard-to-clone identities
- Reputation systems are unlikely to work
  - either central entity knows “good guys”
  - or they all look the same
- Avoiding centralization at all cost may not matter for real-time services
  - typically, don’t have Napster/PirateBay problem