NAT Behavioral Requirements for TCP

Saikat Guha, Kaushik Biswas, Bryan Ford, Paul Francis, Senthil Sivakumar, Pyda Srisuresh

Presented by Philip Matthews

draft-ietf-behave-tcp-01

IETF 66

Document Scope

- Minimize NAT-hindrances to TCP
  - Public-server (most NATs are OK in this regard)
  - P2P (some NATs OK, some not)
- NOT make any changes to TCP
  - Must work with existing stacks
- NOT redesign NATs
  - Mostly small changes

. . . allow applications (especially P2P) to work consistently and fail gracefully

SYN Handling

Goal: Allow inbound SYNs whenever possible, allow diagnostics otherwise

- Allow inbound SYN:
  - For TCP S-O
  - For 3WHS to internal host w/ existing mapping (subject to NAT's security policy)
- Otherwise signal ICMP error:
  - Delay for 6s, give P2P apps a fighting chance to trigger S-O

Session Timeout

Goal: Guarantee at least a large idle timeout

- Guarantee idle TCP for 2h4m
  - Administrator configurable
- Handling of Time-Wait left unspecified
  - For connection-throughput reasons
  - Pointer to time-wait assassination hazards in [RFC1644]

TCPM Overlaps

ICMP errors during connection initiation

- NAT may send ICMP port-unreachable
- Non-P2P app can abort, report error to user
- P2P may persist in hopes of TCP S-O
- Stack may abort by default, but ultimately app should have the option to not abort in response to certain ICMPs

- Not strictly *necessary* (6s leeway just for this)
- Something to consider for Gont's ICMP draft

Things to come

Suggested way to setup P2P TCP

- Open multiple sockets at both ends (think ICE)
- Try: client-server, server-client, client-client (S-O)
- Pick any that connects; verify not half-open
- Reason: client-server hard with NATs, S-O hard on LAN

- Should triggering S-O on LAN be made easier?