Issues with IP Address Sharing
draft-ford-shared-addressing-issues-02

M. Ford (Ed.), P. Roberts (Internet Society)
M. Boucadair, P. Lévis (France Telecom)
A. Durand (Comcast)

Purpose of the document

- Lots of documents specifying address sharing solutions
  - NAT444, NAT64, DS-Lite, etc.
- Capture the issues that address sharing (in any form) creates, document them in one place
- Not intended to get into very detailed solution-specific discussions

Main changes

- Addressed all comments received during IETF76 and over email
- Extraneous text removed/moved to Annex
- Added basic analysis of issues as they relate to first and third parties
- Re-organised to bring more significant issues to the top of the list
- Added text on ports in TIME-WAIT state, TCP control block sharing, rDNS, load balancing, impact on battery life for mobile handsets, ICMP attacks

Issue	1st party	3rd parties
Overly restrictive allocations of outgoing ports will impact performance for end users	x	
Incoming port negotiation mechanisms may fail	x	
Incoming connections to Well-Known Ports will not work	x	
Some applications will fail to operate	x	x
TCP control block sharing will be affected	x	x
Reverse DNS will be affected	x	x
Inbound ICMP will fail in many cases	x	x
Amplification of security issues	x	x
Fragmentation will require special handling	x	

New text (1)

- TIME-WAIT state
  - Ports enter this state for ~ 4 minutes after a connection has concluded
  - Port consumption measurements must count ports in this state as used
- TCP control block sharing
  - CPE NAT already creates issues for this technique today
  - Large-scale address sharing will make the issue more severe and widespread

New text (2)

- Reverse DNS
  - Reverse DNS strings no longer sufficient to identify a discrete subscriber
- Load balancing
  - Deterministic algorithms based on IP addresses may see sudden imbalances in load as address sharing is enabled
  - Growth of address sharing will require re-evaluation of load balancing algorithm designs

New text (3)

- Battery life for mobile hosts
  - Maintaining NAT state requires hosts to send frequent keep-alive messages
  - Sending these keep-alives may significantly reduce the battery life for mobile hosts
- ICMP attacks
  - Malicious user could send Packet Too Big reducing the MTU down to 68 octets
  - Value will be cached by server for all subscribers sharing the IP of the malicious user
  - Could lead to a DoS condition for the server and the NAT

Concluding

- Is this suitable for adoption as an intarea WG work item?
- Is there support for adopting it?