

Issues with address sharing

Mat Ford
ISOC Standards & Technology

Pierre Levis
France Telecom

Address Sharing

- Current practice: give a unique IPv4 public address to each customer
 - this address can be shared into the Home LAN through a NAT device (in the CPE)
- With IPv4 free-pool allocation completion this is no longer possible
 - Scalability of RFC1918 space also creating problems
- A possible solution: allocate the same IPv4 public address to several customers at the same time

Address sharing is the common objective of all IPv4 shortage solutions (CGN, DS-lite, Double NAT, Layer2-Aware NAT, Port Range, ...)

Some principles

- End-to-end principle* may be under pressure but disregard it at your peril
 - *“The end-to-end principle is arguably the fundamental principle of the Internet architecture. In a sense the Internet is the embodiment of the principle. By allowing either tacit or explicit erosion of the principle as we apply our understanding to the construction and operation of the global network, we allow the dismantling of the utility itself.”*
- IPv6 is the goal

*See RFC1958, 3724

Criteria for judging ISP responses

- How easy is it for the end user to control the impact of the address sharing solution on the end-to-end communication?
 - No helpdesk calls to get incoming ports allocated, please
- Extent to which solution offers a natural progression to widespread deployment of IPv6
 - Tunnel over IPv6, NAT464, etc.

Issues

- Broken apps
- Port distribution, port reservation, port negotiation
- Connections to WKPs
- UPnP
- Security and Subscriber identification
- Performance/Resilience

What breaks?

- Apps that
 - Establish inbound communications
 - Carry port info in the payload
 - Carry address info in the payload
 - Use Well-Known Ports
 - Do not use ports (ICMP)
 - Assume uniqueness of IP addresses
 - Explicitly prohibit multiple simultaneous connections from the same IP address
- Fragmentation

Port distribution

- Outgoing connections
 - Source port number usually irrelevant
- Incoming connections
 - Specific numbers matter
 - External referrals
 - Stable over time – for how long?
- How many ports is enough?
 - How far off do you need to push IPv6 deployment?
 - Active subscribers use >80 - >160 ports at a time on average^{*}
 - Distribution heavy-tailed
 - How many of your customers is it OK to annoy?

^{*}http://www.wand.net.nz/~salcock/someisp/flow_counting/result_page.html

Port distribution

- If port-ranges uniformly sized then how many ports is enough?
- If ports dynamically allocated without upper limit then heavy-users, or infected hosts can exhaust the shared resource
 - DoS against a single address would effect multiple subscribers.
- Ports must be allocated based on assumptions about the average number of ports per active-user and the typical number of simultaneous average users.

Well-known ports

- Which port is your webserver on today?
- Proposals for application service location protocols haven't gained much traction, historically

UPnP

- UPnP monotonically increases port number until it finds something usable, or gives up
- Can't be redirected to use a valid port
- So it might work, if you get lucky
- UPnP IGD 2.0 will probably fix this for new/upgraded devices

Security and Subscriber identification

- Logging IP addresses no longer sufficient
- If you see hundreds of connections from multiple ports, you have to log them all as may be multiple users - increased OPEX
- Penalty boxes
 - No way of knowing whether multiple connection attempts from a single IP are a result of shared addressing or abuse
- Port randomisation has reduced effectiveness in mitigating blind attacks
 - Randomisation on OS may be defeated by non-implementation on shared-addressing CPE

Subscriber management for SPs

- Customer identification no longer possible based solely on IP address
- OSS will require updates for
 - activation of services
 - management of customer profiles
 - LI
 - Traceability and mandated logging
 - Considerably more onerous where CGN is present

Performance/Resilience

- Additional latency due to having to request a new port prior to every DNS query?
- Reduced network resilience
- Malicious users sharing my address can now impact my connectivity

Aren't we here already?

- To some extent modem pools and NAT already cause these problems
- Widespread adoption of shared-addressing mechanisms will make them much more prevalent and much more severe
 - Broken apps are more annoying when there's nothing you can do about it
 - Users must be able to determine their representation on the network for some semblance of end-to-end to work

Conclusions

- Shared addressing will mean
 - Degraded network experience for many users
 - Reduced security
 - Higher costs for service providers
 - As yet unclear, but potentially significant, impacts on content providers
- The potential for all Internet users to be service providers is fundamental to the utility of the network
- Are there circumstances under which this could be worth it?

