Host address availability recommendations

draft-colitti-v6ops-host-addr-availability

Vint Cerf, Stuart Cheshire, Lorenzo Colitti, David Schinazi

IETF 93, Prague

Why?

- IPv6 differs from IPv4 in only a few areas
- One of those areas is addressing, particularly host addressing
- Addressing practices that make sense in IPv4 may not be appropriate in IPv6
  - /64 per link allows “unlimited” host addressing
  - No longer forced to assign one address per host due to address scarcity
  - Many benefits provided by assigning multiple addresses to each host
- This is sufficiently non-obvious to some operators to be worth stating

Common IPv6 deployment model

- IPv6 is designed to support multiple addresses per interface
  - In many deployments:
    - 1 Link-local
    - 1 Stable global (EUI-64, RFC7217)
    - >= 1 Privacy
    - >= 0 DHCPv6

- In most networks, hosts can obtain additional addresses without asking
  - RFC6177, RIRs, 3GPP, BBF, Cablelabs etc. all recommend subscriber be assigned a prefix
    - Most (all?) of these deployments support either SLAAC or DHCPv6 PD

Benefits of multiple addresses

- Privacy addresses
- Virtual machines / Multiple processors inside the host
  - e.g., CPU vs. baseband
- Tethering
- IPv4-over-IPv6 transition mechanisms (e.g., 464XLAT)
- Future applications
  - Identifier-locator addressing
  - Per-application IP addresses, …

- New technologies made possible by multiple addresses:
  - 464XLAT
  - 64-share

What if host does not have enough addresses?

- Features unavailable
- Features available only after explicit request to network
  - High latency
  - No certainty of success
  - Implementation complexity
  - Provisioning load
  - May require human interaction (captive portal, billing, etc.)

Possible host reaction: NAT66

- Hosts can work around limits using NAT66, just like they do in IPv4
- NAT has well-known drawbacks:
  - Application complexity due to NAT traversal
  - Brittleness and support costs due to state maintenance, SPOF, and NAT traversal
  - Battery life impact due to NAT keepalives
    - QUIC uses 15 seconds!
- In IPv4 we have no choice due to address scarcity, but no such pressure in IPv6
- IAB advice: deployment of NAT66 is not desirable

Options for obtaining multiple addresses

	SLAAC	DHCPv6 with multiple IA_NA	DHCPv6 PD	DHCPv4
Extend network	Yes (64-share / ND proxy)	No	Yes (PD / 64-share / ND proxy)	Yes (NAT44)
"Unlimited" endpoints	Yes*	Yes*	No	No
Stateful, request-based	No	Yes	Yes	Yes
Immune to layer 3 on-link resource exhaustion	No	Yes	Yes	Yes

* Subject to network limitations, e.g., ND cache entry size limits.

Semantically most similar to DHCPv4 is DHCPv6 PD

How many addresses?

- Today:
  - Privacy addresses: **7**
  - IPv4-over-IPv6 transition mechanisms: **1**
  - Virtual machines: **5**
  - Multiple processors inside the host: **1 or 2**
  - Tethering: **8 devices, >= 1 per device**
- A host using some but not all of these functions might need ~20 addresses
- How many will future applications need?

Recommendations

- Provide multiple IPv6 addresses from each prefix to general-purpose hosts when they attach to the network
- Don’t impose a hard limit on the size of the address pool assigned to a host
- If the network requires explicit requests, assign a /64 via DHCPv6 PD
  - Using DHCPv6 IA_NA or IA_TA to request a sufficient number of addresses (e.g. 32) would accommodate current clients but sets a limit on the number of addresses available to hosts when they attach and would limit the development of future applications

Host tracking

- Some operators say DHCPv6 is only possible choice due to tracking needs
  - Usually enterprise or university network administrators
  - Frequently cite need to match IP address + timestamp to MAC address
    - “RIAA / MPAA”
- It’s worth mentioning that DHCP by itself doesn’t provide security
  - Only provides security when coupled with L2 security / enforcement features
  - If you have those features, configure them to log via syslog and track without needing DHCP
- Not clear how this will survive DHCP anonymity, MAC address randomization
  - May have to switch to something like 802.1x
- Several large enterprise IPv6 deployments do not use DHCPv6

Address space management

- A /64 per host is equivalent to what we do in IPv4

- Example: 192.168.0.0/16 is 2^{16} endpoints with one IPv4 address each
- In IPv6, 2^{16} endpoints with a /64 each is a /48
  - Easy to get for a small / medium enterprise

- Similarly, 10.0.0.0/8 => 2^{24} endpoints
- 2^{24} endpoints with a /64 each is a /40
  - Easy to get for a large enterprise

Link-layer scaling issues

- Links with lots of addresses consume hardware resources
  - Each IPv6 address is one ND cache entry, perhaps one TCAM entry, etc.
  - More ND cache churn

- Assigning a prefix to each device means only one

Current status

- Version -00 sent to mailing list, received support and feedback
- Uploaded -01 yesterday to address much of the feedback
- Adopt as WG item?

Questions?