How does legacy IPv4 differ from IPv6?

Brian Carpenter

March 2010

Is my title too provocative?

- Recent projections from potaroo.net for unallocated IPv4 address pool exhaustion:
  - IANA 30-Sep-2011
  - RIRs 13-Oct-2012
- Recent ISP survey (30 responses):
  - Majority expect public IPv4 space for customers to run out between 2010 and 2015
  - Predictions when 50% of customers will require IPv6 run from 2011 to 2020
  - Most common target date for IPv6 as standard service is 2011
  - But most predict IPv4-only applications will last at least ten years

Outline

- The Great Disillusionment
- What the IETF has been up to in the last year
  - tunnels++
  - NAT++
  - operations
- Implications outside layers 3 and 4
- Appendix: technical details

Disillusionment: Reality breaks in, as always

- When the IETF first considered deployment scenarios, the idea was that IPv6 would deploy before IPv4 ran out.

FAIL

- This changes the available transition models.
  - Service providers will still need to offer dual stack services, of course.
- More need for interworking than ever expected.
  - The only commercially sane assumption is that v6 clients will need to access v4 services indefinitely.
- This has been driving IETF work for a couple of years.

Tunnels: the SOFTWIRE WG

- “discovery, control and encapsulation methods for connecting IPv4 networks across IPv6 networks and IPv6 networks across IPv4 networks in a way that will encourage multiple, inter-operable implementations.”
- Dual Stack Lite - share IPv4 addresses among customers by combining IPv4-in-IPv6 and NAT. Driven by Comcast broadband model.
- 6rd - blend of 6to4 and ISATAP providing automatic tunneling of IPv6-in-IPv4 to ISP subscribers. Deployed by Freenet.FR
- Older mesh and hub+spoke models also documented, using GRE, IP-in-IP, L2TPv3, IPsec and MPLS.

Ain't misBEHAVing: NAT64

- *Reminder:* the old NAT-PT specification was deprecated, mainly due to irreconcilable differences with DNS.
- However, many operators (especially in the mobile "LTE" world) are convinced that they will soon have millions of IPv6-only subscribers needing access to legacy IPv4-only services.

  {Argument about exactly what the economic incentives are deleted - but just imagine selling mobiles that can't reach PayPal.}
- Therefore, the BEHAVE WG has taken up the NAT64 challenge.

NAT64 only solves one problem

- IPv6-only client (no v4 address, no v4 connectivity) needs to initiate communication with an IPv4-only server.
  - As stated, this requirement cannot be met by the conventional dual stack approach.
  - Whatever *you* may believe, some ISPs believe this is vital, and their suppliers believe they'd like the business.
- NAT64 doesn't tackle any other cases.
- NAT64 comes with a separate DNS64 magic box
  - NAT-PT came with a built-in DNS ALG

Components

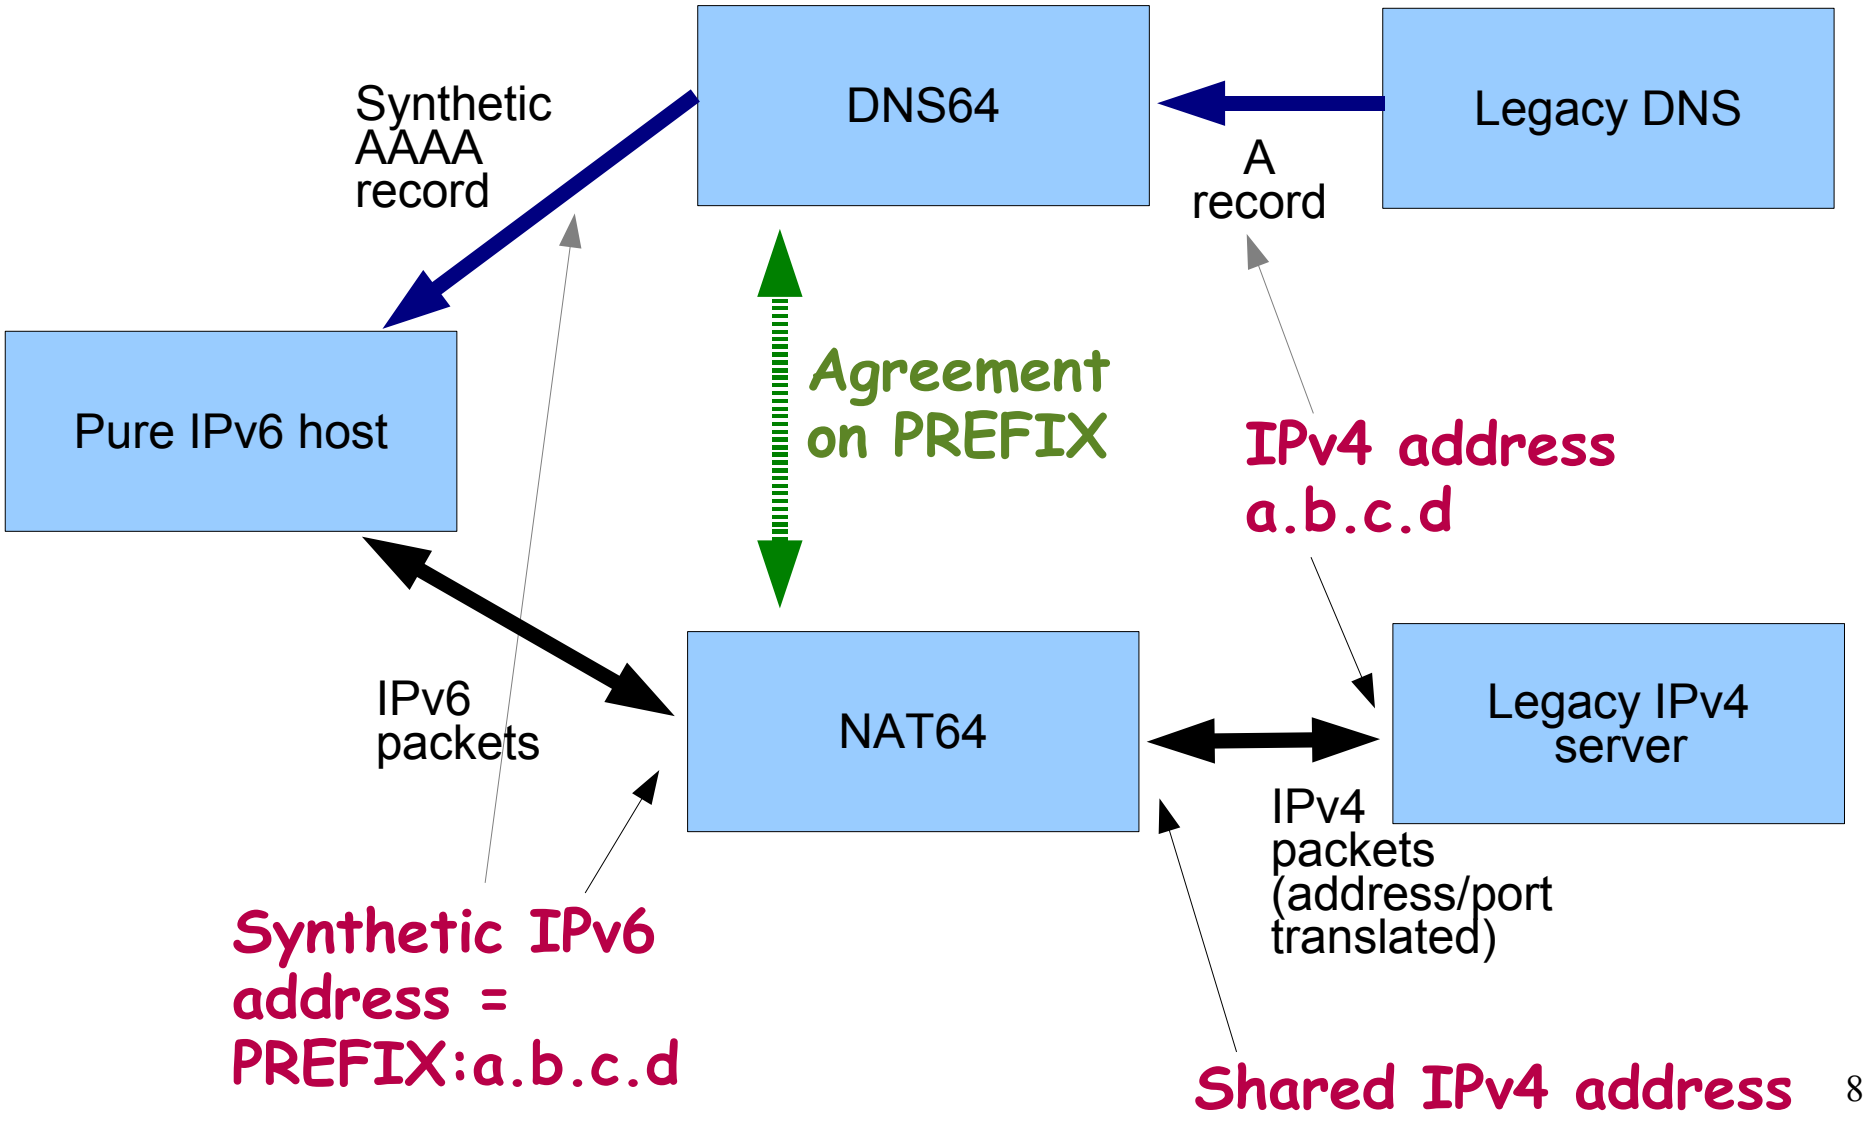
DNS64
Legacy DNS
Synthetic
AAAA
record
A
record
Pure IPv6 host
Agreement
on PREFIX
IPv4 address
a.b.c.d
IPv6
packets
NAT64
Legacy IPv4
server
IPv4
packets
(address/port
translated)
Synthetic IPv6
address =
PREFIX:a.b.c.d
Shared IPv4 address

V6OPS activity

- “The IPv6 Operations Working Group (v6ops) develops guidelines for the operation of a shared IPv4/IPv6 Internet...”
- Current work includes:
  - Requirements for CPE routers
  - IPv6 Deployment in Internet Exchange Points
  - Incremental Carrier-Grade NAT (CGN) for IPv6 Transition
  - ISP scenarios generally.

Other activity

- 6MAN
  - maintaining basics
  - worrying about address selection algorithm
  - defining canonical version of text representation
- SHIM6 - standardising host-based multihoming
- DHCPv6 (polishing)
- Mobile IPv6 (polishing)
- And lots more in various WGs...

Changes outside layers 3 and 4

- Address storage/display/entry
- URI parsing
- AAAA and A records
  - plus ip6.arpa for reverse lookup
- Socket API
  - Handle multiple addresses and address families returned by ~~gethostbyname~~getaddrinfo()
  - There is simplified code in RFC 4038
  - Hurricane Electric page http://owend.corp.he.net/ipv6/
  - itojun's page: http://www.kame.net/newsletter/19980604/
  - Or hide this issue in a library, like the Java JDK

IPv6 Documentation Addresses

- Global unicast address prefix reserved for documentation purposes (see RFC 3849)
  - 2001:db8::/32

Emerging standard for textual representation of IPv6 addresses

- Defining a mandatorycanonical form
  - Fully backwards compatible
  - ~~2001:DEAD:beef:0::1~~
  - 2001:dead:beef::1

Discussion

- What should we be doing to make sure that IPv4-IPv6 coexistence, and later the IPv4 legacy, is handled properly everywhere?

(Some technical details slides follow for reference. For much more, see http://www.ietf.org/proceedings/07jul/slides/ipv6spec-0.pdf*)*

NAT64: Sequence of events

- The IPv6 host uses DNS64 as its regular DNS service to look up servers.
  - For native IPv6 hosts, DNS64 returns normal AAAA records.
  - For hosts with A records only, DNS64 concatenates the agreed PREFIX and the IPv4 address from the A record, and synthesises an AAAA record.
- The IPv6 host just sends normal packets to the synthetic address, which is routed to the NAT64.
  - The NAT64 recognises a new session, extracts the server IPv4 address from the synthetic address, assigns a port on the IPv4 side and other NAT state, and otherwise does its standard NAT thing.
- From an application viewpoint, this looks pretty much like old fashioned NAT44.

NAT64: What is the PREFIX?

- PREFIX is likely to be a /96 (leaving 32 for the IPv4 address)
- Could be locally assigned (one operator controls IPv6 host, NAT64 and DNS64)
- Could be a globally defined WKP (well known prefix)
  - What happens if a synthetic address “escapes” from the scope of the NAT64/DNS64 pair?

Address Architecture

- Not just bigger addresses, a new address architecture… (see RFC 4291)
  - No IPv6 broadcast address, use All-Nodes Multicast address instead
  - There is an IPv6 loopback address
    - 0::1
  - IPv6 addresses with embedded IPv4 addresses
    - Use IPv4-Mapped IPv6 Address, IPv4-compatible address is deprecated
    - 0::FFFF:10.0.0.1

Scoped Unicast Addressing

- IPv6 has concept of scoped unicast addressing
  - Link-local & global [site-local is deprecated]
- IPv6 adds support for unique local addresses (ULAs)
  - Useful for local or private network addressing
  - But deemed to have global scope
  - Not a direct map to RFC1918 addresses
  - ULAs are *[probabilistically but almost certainly]* unique, so a node can be on more than one local network, or local networks can be transparently linked

MTU and Fragmentation

- IPv6 requires lower layers to support an MTU of 1280 bytes (as opposed to 536 in IPv4)
- IPv6 routers do not fragment packets
  - Path MTU discovery is mandatory
  - All fragmentation happens at the source

Not ARP

- IPv6 L2 address resolution is performed using Neighbor Discovery not ARP
- ND is ICMP-based
- Uses Solicited-Node Multicast addresses
- Combines address resolution with host autoconfiguration and reachability detection
- Closely linked to RA (router advertisements)
- IPv6 can configure without DHCPv6