

IPv6 and IPv4: Twins or Distant Relatives

Infohlox

Infoblox .

Infoblox 🗧

Infoblox :

Paul Ebersman, IPv6 Evangelist NANOG54, San Diego (5-8 Feb 2012)



What you'll see immediately

- More addresses
 - 340 undecillion
- Bigger, beefier addresses
 - 2001:db8:dead:beef::1
- Lots more addresses per interface
- More "magic"

7 Layer View (Sorta)



Routing Efficiencies

- Fixed header size
- Extension header chain
- Flow labels in header
- No intermediate fragmentation
- PMTUD



- No broadcast
- Multicast
- NS/Solicited Node, no ARP
- ICMPv6





- Unicast
- Multicast
- Anycast

6



- Link Local
- Global Unicast
- Unique Local
- Transition
- Misc (Site Local, Reserved, Special)



- Local (broadcast) Domain
- fe80::/64
- Similar to APIPA (169.254.0.0/16)
- Reusable on all interfaces



- Globally routable
- Unique
- "Public"
- Use 'em everywhere!



- Not globally routable
- Not unique (but registerable...)
- Replacement for RFC1918 (if you must)

More addresses per interface!

- Use 'em; we'll make more
- Multiple default routes
- Quiescence
- How do it know? (RFC 3484)

Subnet Planning

- You can get enough IPv6 space
 - Do the architecture you want, not the one you're stuck with
 - Use GUA space everywhere, make NAT a choice
 - Map your subnets to your process/provisioning or business model
 - Do a scheme that aggregates and makes ACLs sane

Sample /32 Plan by Geography

- 2001:db8:abcd::/36

- City: 4 bits = 16 possible locations
- 2001:db8:abcd::/40
 - **Hub**: 4 bits = 16 possible hubs per city
- 2001:db8:abcd::/48
 - Floor: 8 bits = 256 floors per hub.
- 2001:db8:abcd:12xx::/56
 - Switch: 8 bits = 256 Switchs per floor.
- 2001:db8:abcd:1234::/64
 - VLAN: 8 bits = 256 VLANs per switch.

- /48 is minimum routable chunk
- /64 for all non-p2p subnets
- 127 for p2p links (RFC 6164)
- 128 for loopbacks
- Use /64 each for p2p/lb, pair for each routing domain

- If you qualified in v4, you still do
- If PI space would have been useful in v4, it still is
- If you didn't understand it in v4, v6 won't help you...

SLAAC vs DHCP



- SLAAC == StateLess Address AutoConfiguration
- Uses Router Advertisement (RA) messages
- Network policy moved to the edge

- EUI-64 uses the mac address and an algorithm to generate interface ID
- Windows7/Vista randomly generates interface ID by default
- Servers and LINUX/UNIX mostly use EUI-64

MAC-Address to Interface ID



20

Client configures link-local address

- Generates 64 bit host ID (EUID from MAC, random)
- Uses link local prefix and EUID to generate tentative address (such as fe80::028c:f5ff:fe05:4235)
- Does DAD (Duplicate Address Detection)
 - Sends a multicast Neighbor Solicitation message containing its new tentative address to the solicited node address
 - If no other node responds with a Neighbor Advertisement using that address, the host configures itself with that address

SLAAC Sequence cont.

- Host now looks for Router Advertisement (RA) Messages
 - Sends multicast Router Solicitation message
 - Listens for RA messages
 - Configures itself based on contents of RA message, including doing DHCPv6

RA Message Contents

- Local prefix(es), including A (autonomous address configuration) flag
- Router info
 - Router's link-level address
 - Lifetime of route
 - Router priority
- Flags: M (ManagedAddress) flag and O (OtherConfiguration) flag
- Maximum Transmission Unit (MTU) of upstream link

Not in RA Messages...

- RDNS server
- NTP or "other" configuration
- RFC 6106 for RDNS in RA
 - Lack of client support…



- Must run both RA and DHCPv6 for most sites...
 - No DHCPv6 without an RA message with M or O flag on
 - Many options not available to clients without DHCPv6
 - No default gateway in DHCPv6
 - Must configure DHCP and edges

DHCPv6 features

- Not BOOTP! ③
- No broadcast!
- New ports (546, 547)
- Vendor options in TLV tuples
- Reconfigure now secure

DHCPv6 vs DHCPv4 messages

DHCv6 Message type	DHCPv4 message type
SOLICIT (1)	DHCPDISCOVER
ADVERTISE (2)	DHCPOFFER
REQUEST (3), RENEW (5), REBIND (6)	DHCPREQUEST
REPLY (7)	DHCPACK/DHCPNAK
RELEASE (8)	DHCPRELEASE
INFORMATION-REQUEST (11)	DHCPINFORM
DECLINE (9)	DHCPDECLINE
CONFIRM (4)	
RECONFIGURE (10)	DHCPFORCERENEW
RELAY-FORW (12), RELAY-REPLY (13)	



- IPv4 address shortages made pool size precious
- IPv6 has plenty
- Protect from brute force scans
- Do pay attention, though...

Failover is so 5 minutes ago

- Drivers in IPv4 for failover:
 - Changing client IP address was user noticeable
 - Lack of large enough pools
- If you use up 18 quintillion addrs...
- Fragility of keeping state
- High availability more useful



- Client and server must support and be configured for it
- Now has security
- With quiescence and reconfigure, renumbering is easy (mostly)



• Mac address as ID is flawed:

- Not always unique
- Can be altered
- Multi-interface hosts confuse things
- But it's what most of the eyeballs on the Internet are ID'ed by currently
- DUID (DHCP Unique Identifier) is the replacement in IPv6

DUID issues

- Yes, mac addresses sucks.
- DUIDs suck differently.
 - Can't correlate v4 and v6 addrs to same host
 - Can't get mac address from DUID
 - Persistent storage of DUID may cause surprises

- One DUID per DHCP server or client
- One Identity Association (IA) per network interface on a host
- A host can DHCP for all interfaces via DUID/ IA as unique key

Identity Associations

Types:

- **IA_TA**: temporary address(es), i.e. privacy addrs
- IA_NA: non-temporary address(es), i.e. not privacy addrs
- **IA_PD**: prefix delegation

- Delegate a prefix to a device
- Device can delegate longer prefixes to its own clients
- Likely scenario is home/CPE routers
- Lots of potential but not lots of gear available now

ICMPv6

36



- Required for:
 - DAD
 - Finding routers (RA/SLAAC)
 - Finding servers (DHCP)
 - PMTUD
 - Connectivity (echo request/response)
 - Network errors

- Filter it all and you don't have a useful network
- ICMPv6 much more detailed/precise in types and functions
- RFC 4890 has excellent filtering practices

Security

39



- Most issues much the same as IPv4
- Misconfiguration more likely than malice
- Untested code and lack of experience
- Security vendor claims must be validated

- Subnet size makes brute force scanning pointless (if you really use it...)
- Privacy addresses
- IPSec

Security realities

- Bad host numbering schemes
- IPSEC:
 - Good news: just like IPv4
 - Bad news: just like IPv4
 - Exception: Microsoft DirectAccess...

- Test all your firewall and security appliances for IPv6
 - ACLs for IPv6
 - Detect various tunneling (ISATAP, Teredo, 6in4, 6to4, etc)
- Make sure all your NMS and logging deal with IPv6, both for transport and data

Q & A