Infoblox

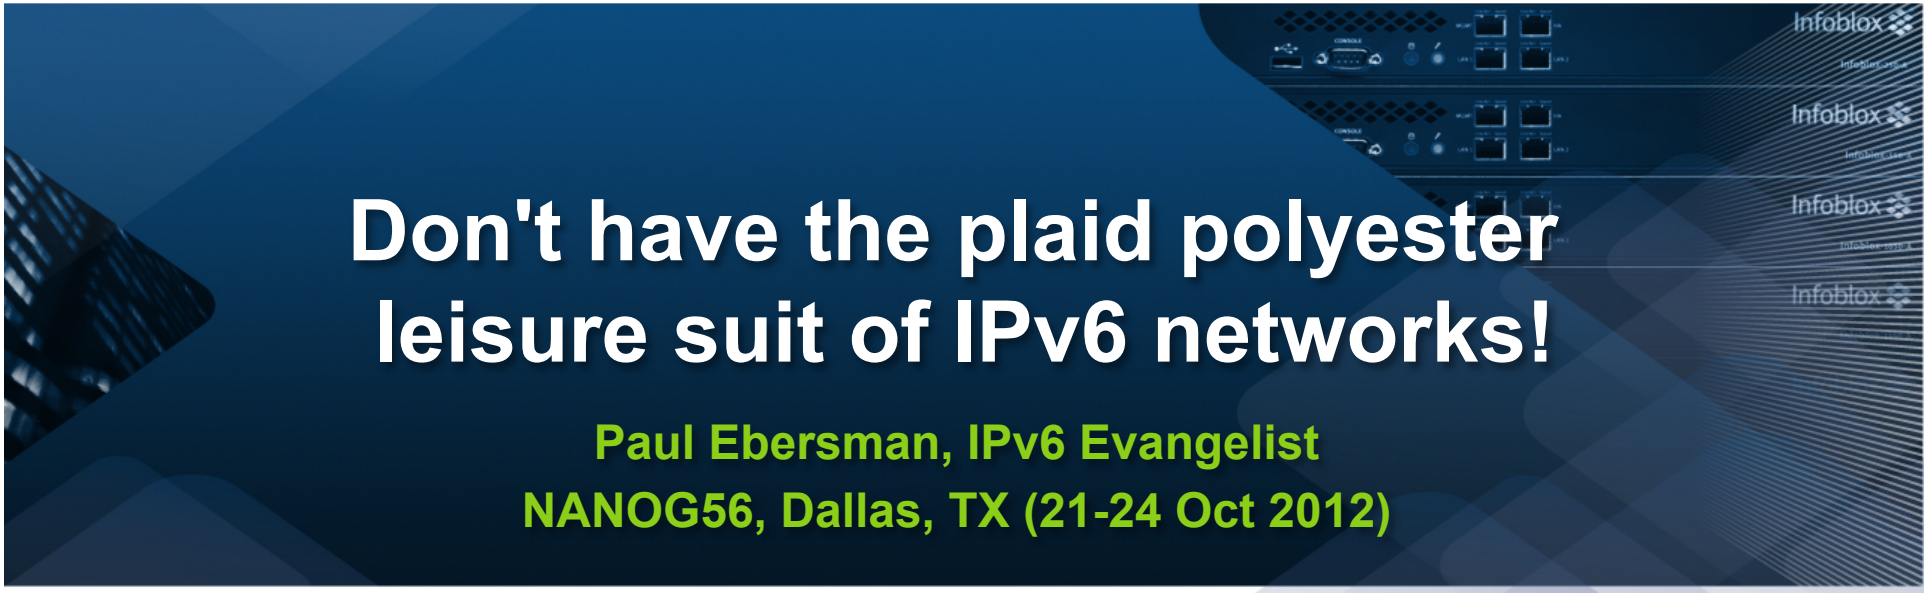

Don't have the plaid polyester leisure suit of IPv6 networks!

Paul Ebersman, IPv6 Evangelist

NANOG56, Dallas, TX (21-24 Oct 2012)

Lots of Changes

Change is good

- **Well, change is inevitable…**
- **Many constraints from IPv4 now gone**
- **Classful vs CIDR**

Routing Efficiencies

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

Network Efficiencies

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

Be an architect

Sane Subnetting

- 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

Prefix Lengths

- /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

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 Switches per floor.
- **2001:db8:abcd:1234::/64**
  - **VLAN**: 8 bits = 256 VLANs per switch.

Subnets, not hosts

18 quintillion…

- Addresses > L2 capacity
- RIR/ISP allocations based on subnets
- Enjoy your nibbles while you may

Use the whole /64!

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

1918/NAT.
Die die die.

How did it ever make sense?

- Shortage of IPv4 for consumers
- IPv6 not widely available
- Desperation
- Mushrooms?

Why is it still around?

- **Still not enough IPv4**
- **The “It’s more secure” myth**
- **Have bent/twisted apps (Skype)**

How naked is the emperor?

- NAT != security
- Debugging/logging hard
- Breaks end to end

But we *like* to suffer

- No NAT66. Yet…
- Stateful FW also painful
- ULA ~= RFC 1918

I'm a Mac

DUID > Mac address

- **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**

What DUIDs do right

- 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

Where DUIDs don’t work…

- Anyone using mac address for identification or filtering
- Anyone trying to correlate IPv4 and IPv6 to the same machine/user
- Persistent storage of DUID may cause surprises

Interface ID generation

- 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

But I do dual stack…

- How to correlate all addrs to same client:
  - hwaddr draft in ietf
  - circuit-id/remote-id

DHCP. Or not.

The good old days

- With IPv4, only two methods:
  - Static
  - DHCPv4

More choices!

- Classic: static
- StateLess Address Auto Configuration (SLAAC)
- Stateless DHCPv6
- Stateful (full DCHPv6)

SLAAC

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

Not in RA Messages…

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

DHCPv6

- “public” or “private” (temporary) addresses
- RDNS server, NTP, TFTP, Vendor options
- Update DNS with A/PTR
- But no default route!

Differences

- DHCPv6
  - Filter/control access
  - Update IP address management system
  - Update A/PTR records in DNS
  - Further from client, more centralized
  - Handles more complex configs, phones, printers, etc.

Differences

- SLAAC
  - Local/fast
  - Light weight
  - Decentralized
  - No logging, A/PTR updates or IPAM updates

Your priorities

- Do you have auditing or logging requirements?
- Centralized or distributed management
- Technical level of support staff
- Range of different gear?

Centralized model

- Need auditing
- Need access control
- Senior technical staff not everywhere
- DHCPv6 is your friend

Coffee House

- Baristas are not networking folks
- Customers just need it to work
- No logging, lease churn would be burden
- Small range of client machines
- SLAAC!

DHCPv4-like DHCPv6

- Send RA messages with A=0, O/M=1
- DHCP for all configurations except default route
- DHCP server does A/PTR and IPAM updates

Coffee House Setup

- Send RA messages with A/O=1, M=0
- Send RDNS in RA messages
- DHCP server does no leases, just gives DNS for clients that can't do RFC 6106

PD

Prefix Delegation

- Dynamic Heirarchical Networks
- DHCPv6 reconfigure and your network
- Vendor support…
- Potentially cool

ICMP

ICMPv6

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

ICMPv6 Filtering

- 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

Reverse/PTR goo

How did this all start?

- ftp (ftp.uu.net, ftp.wustl.edu)
- SMTP
- Security devices
- Silly web things

How did we do it IPv4

- By hand (ow)
- Scripts
- $GENERATE
- IPAM

How would that work for IPv6

- A single subnet is a /64
- A /64 has 18 quintillion (4 bil x 4 bil) addrs
- A PTR record has 34 labels in IPv6
- Anyone got a computer with enough disk or RAM to hold one /64 zone file?

So what are we left with?

- Admit that PTRs are pointless
- Pre-populate (assuming FTL travel…)
- Pre-populate statics for routers & big servers
- As previous plus DHCP server adding clients
- Lie on the fly (if not doing DNSSEC)

The nice thing about standards…

We're not done yet

- Over 200 RFCs relating to IPv6
- But over 200 drafts in active revision too…
- More drafts added every IETF (3 meetings/year)

What can we do?

- **Participate!**
- **Make sure your vendors participate and implement the new standards**
- **Pick your battles**

Q&A