

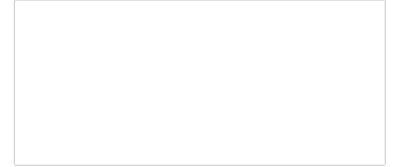
# IPv6: Are we there yet?

**NANOG 58 – New Orleans**  
**Paul Ebersman – IPv6 Evangelist**  
**@Paul\_IPv6, pebersman@infoblox.com**



So where are we?

# The more things change...



- **How wide is deployment?**
- **SLAAC vs DHCP**
- **Identifying users/machines**
- **Interface “magic”**
- **Org/political challenges**

# The more things keep changing...

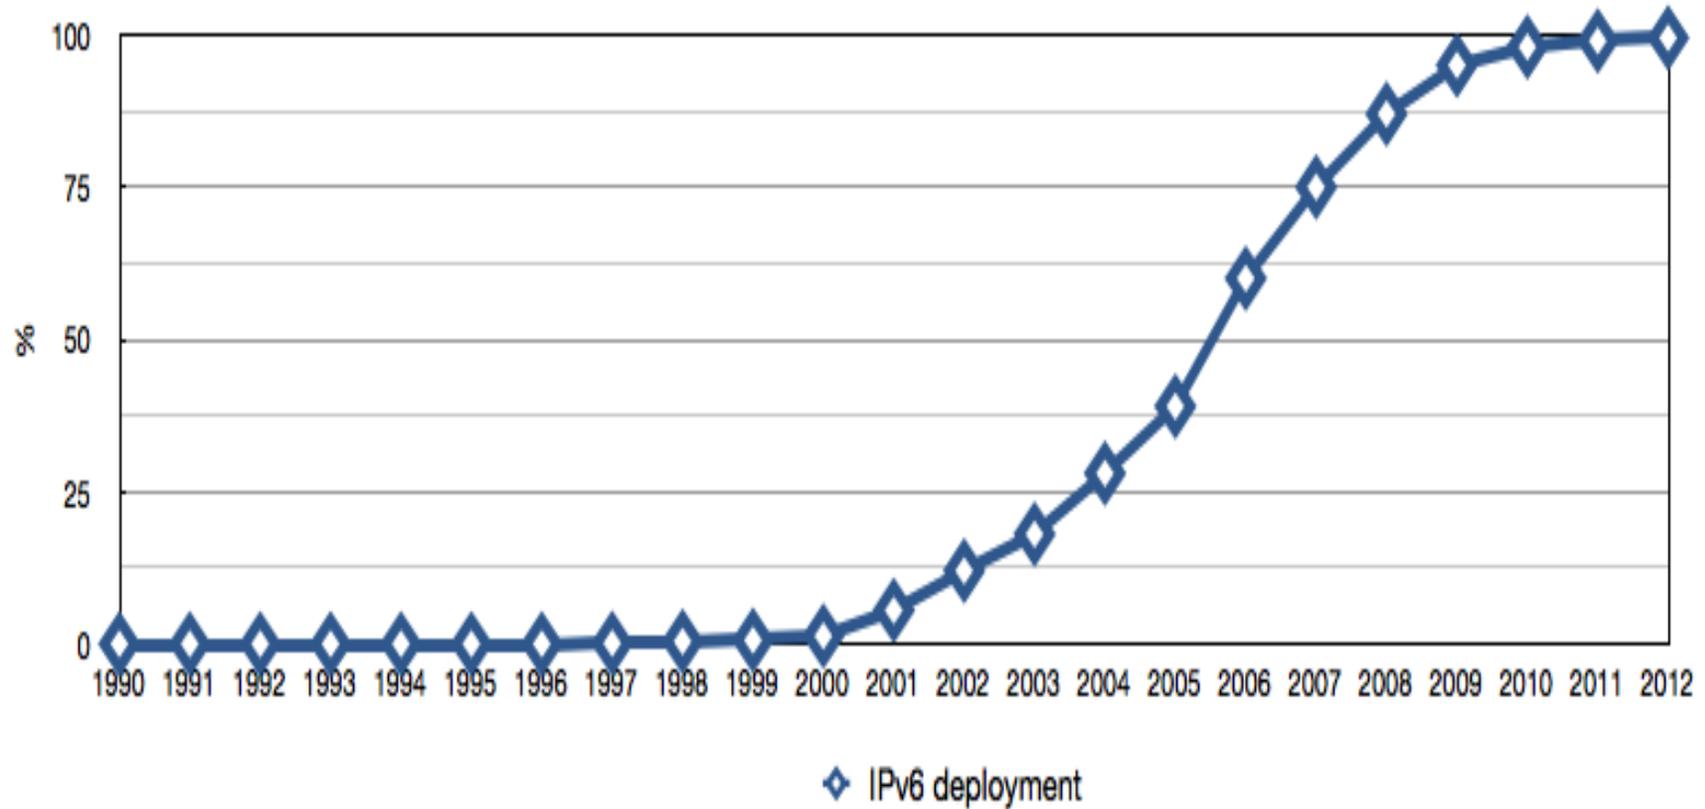
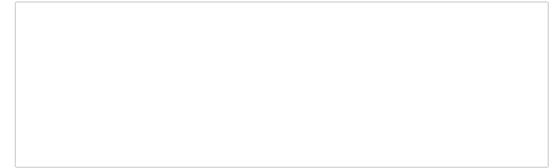
---

- **App changes (esp. browsers)**
- **Policy changes (PTR)**
- **Security and “broadcast domain” changes**
- **IPSEC**
- **IETF Process**

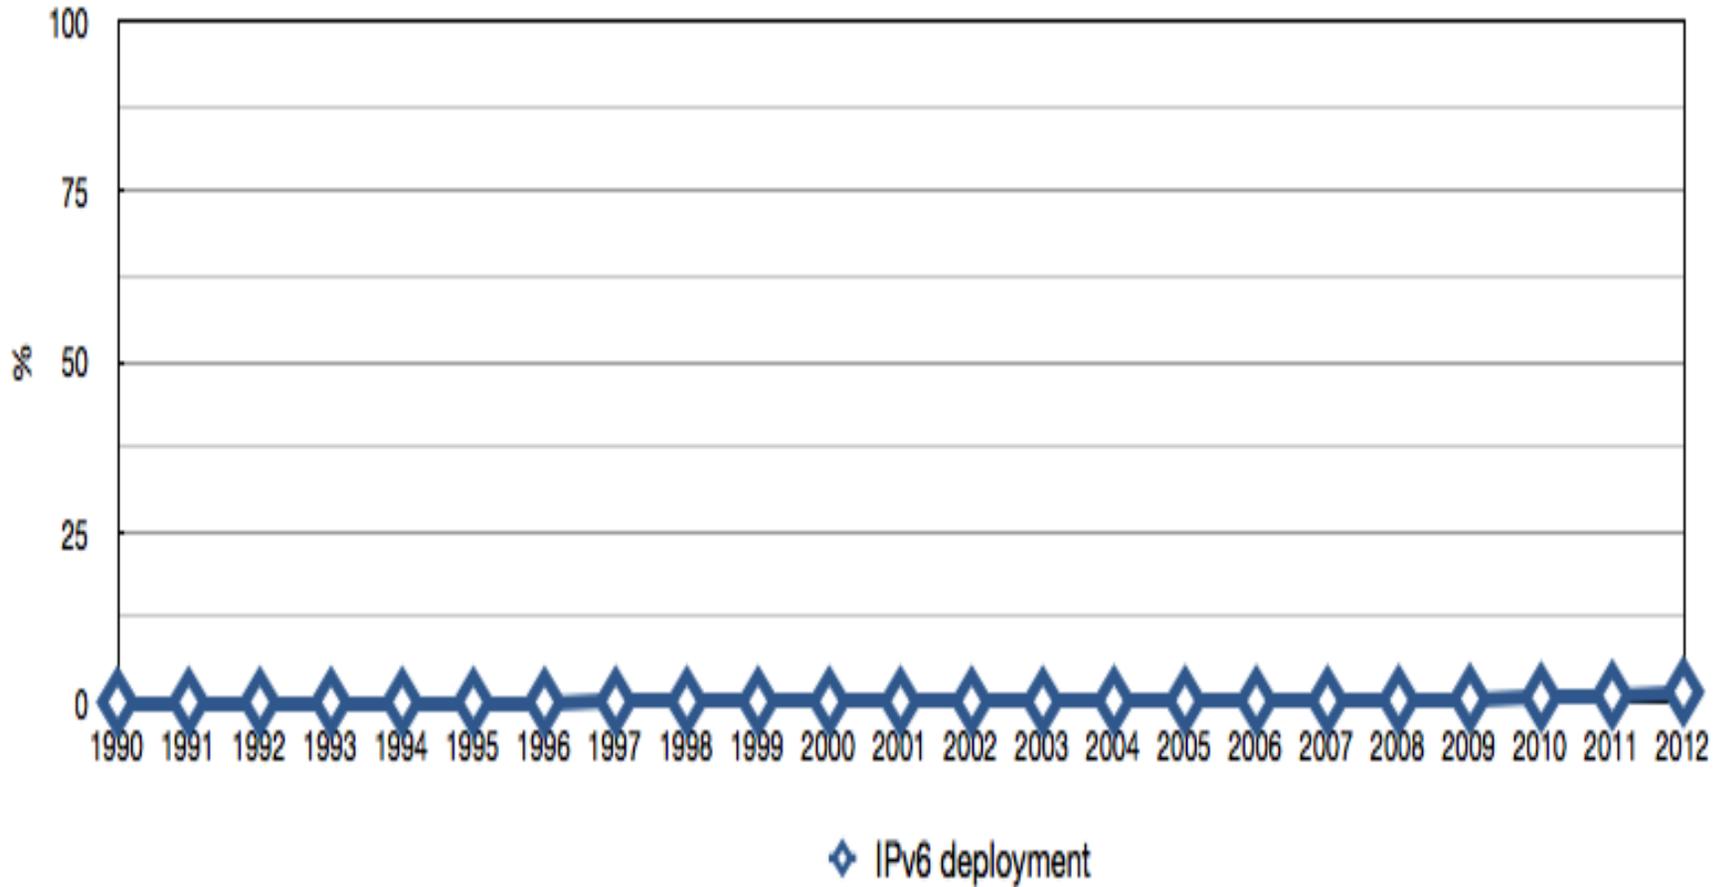
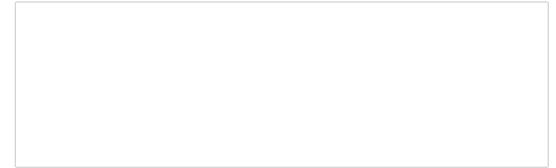


**100% by 2012...**

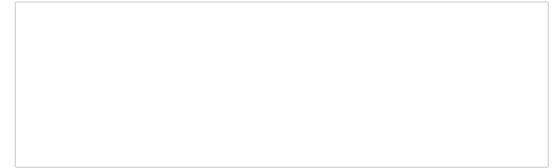
# The dream...



# And the reality

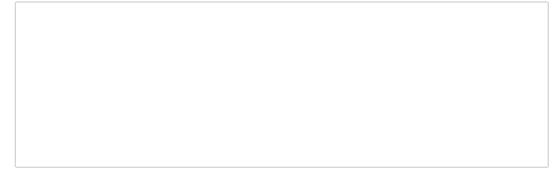


# Some better than others



Index	ISO-3166 Code	V6 Use ratio	V6 Users (Est)	Country
1	RO	10.28%	889683	Romania
2	EU	9.29%	0	European Union
3	LU	7.50%	35259	Luxembourg
4	CH	7.38%	476755	Switzerland
5	FR	5.50%	2761915	France
6	JP	4.18%	4210674	Japan
7	BE	3.58%	304461	Belgium
8	DE	3.06%	2078627	Germany
9	US	2.65%	6619195	USA
10	PE	2.16%	227924	Peru

# Some better than others

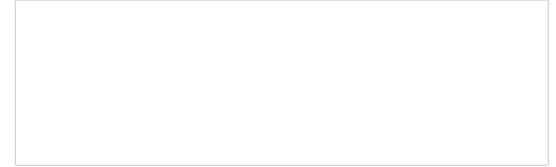


- <http://www.worldipv6launch.org/measurements/>
  - **VZW at 30%**
  - **France, Romania & Switzerland**
  - **12% of Alexa 1,000**
  - **IX traffic in EU approaching double digits**



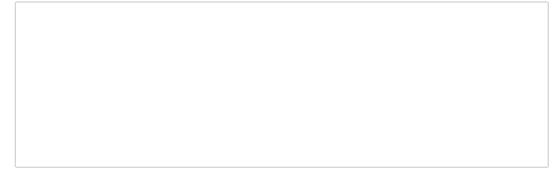
But v4 to v6 is **\*HARD\***...

# DUID > Mac address



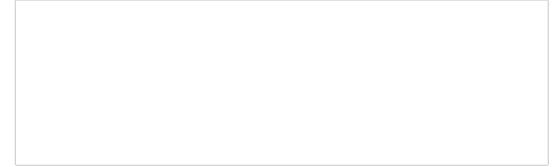
- **Mac address as ID is flawed:**
  - Not always unique
  - Can be altered
  - Multi-interface hosts confuse things
- **But it's what most of the eyeballs on the Internet are ID'ed by currently**
- **DUID (DHCP Unique Identifier) is the replacement in IPv6**

# What DUIDs do right



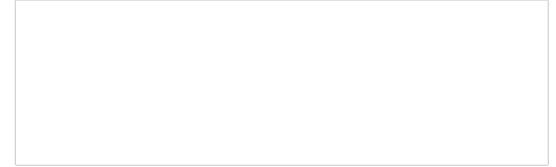
- **One DUID per DHCP server or client**
- **One Identity Association (IA) per network interface on a host**
- **A host can DHCP for all interfaces via DUID/IA as unique key**

# 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**

# But I do dual stack...

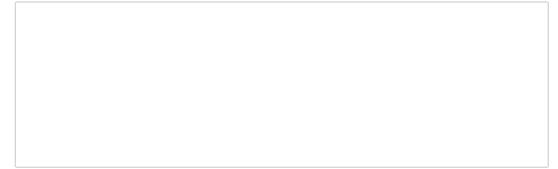


- **How to correlate all addrs to same client:**
  - **draft in ietf: draft-ietf-dhc-dhcpv6-client-link-layer-addr-opt (headed to IESG)**
  - **circuit-id/remote-id work as with DHCPv4**



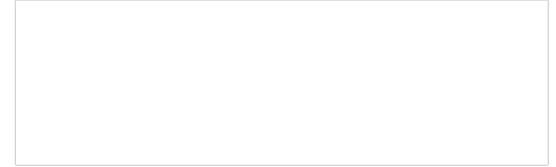
# RA Season! Wabbit Season!

# IPv4 routing



- **Static default route**
- **DHCP server gives default route**
- **Changing network might miss changing DHCP default route**

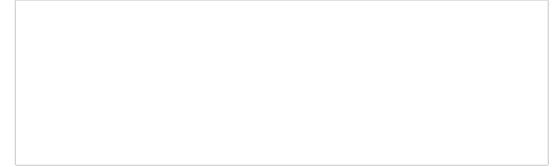
# IPv6 routing



- **Static default route (link local). Ick.**
- **DHCP server can't give default route...**
- **Folks changing routers probably own RA configs**

# Layer 9 (political)

---

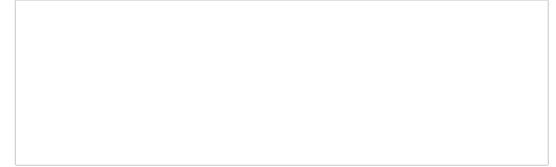


- **Different groups for DNS, DHCP, routers, RAs, IP addr assignment?**
- **Can't just change DHCPv6 or RA, need to coordinate with systems, network, maybe security**



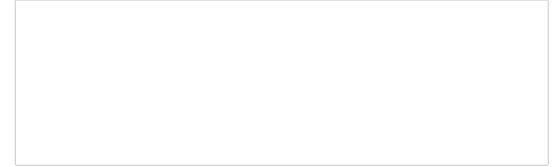
# Happy Eyeballs

# IPv6. Yes. Have some.



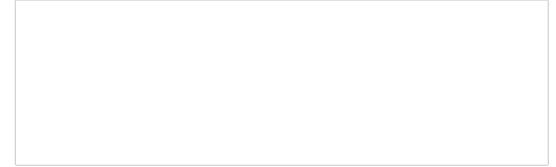
- **Original plan: Always use IPv6/AAAA if available**
- **Result: poor user experience (long timeouts, use of slower links, etc.)**

# Err... We meant Happy...



- **Next attempt was to specify draft/RFC**
- **“But that doubles DNS traffic”...**
- **And OS and browser folks both dived on it**

# Hence “Hampering Eyeballs”

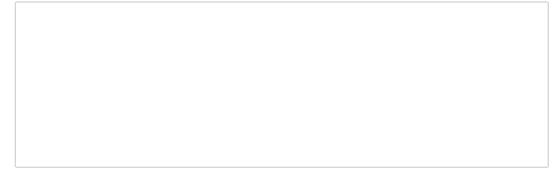


- **Testing by Geoff Huston**
- **Problems with browsers**
- **Lots of problems with OS X**
- **Windows trying to fix at network layer...**

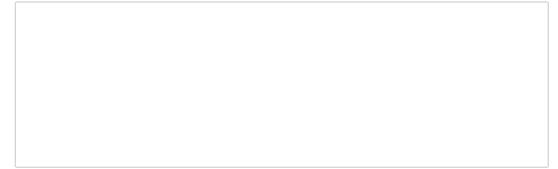


How do it know?

# Source/Destination Address

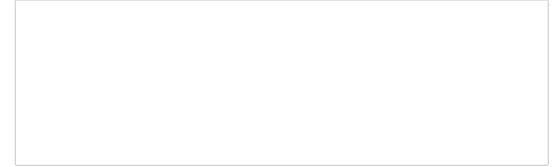


- **Multiple interfaces w/ multiple addrs**
- **Multiple prefixes**
- **Dual stack...**
- **How to choose...**
- **RFC 6724 (formerly RFC 3484)**



- **Types of addrs:**
  - IPv6: GUA, ULA, Link Local, privacy
  - IPv4: public, APIPA, 1918
  
- **Some better than others**
  - Consider scope, type, prefix length
  - Avoid deprecated
  
- **Allow local policy overrides**

# Debugging will be fun

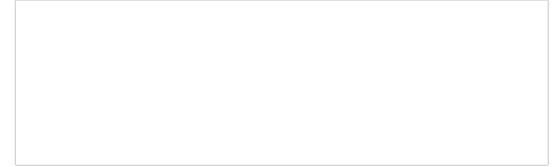


- **Decisions time/context sensitive**
- **How to train staff and users**
- **Local tools to dump all info**
- **Packet sniffers?**



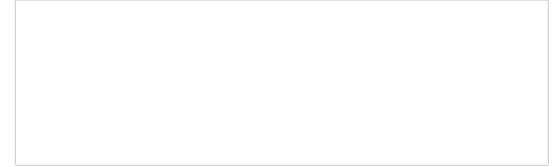
And what don't  
we know yet...

# Default route



- **Multiple default routes from RAs**
- **No more HSRP/VRRP! Maybe...**
- **But does this actually work?**
- **Not all Oss did the right thing (Fedora, ???)**

# What else will we find...

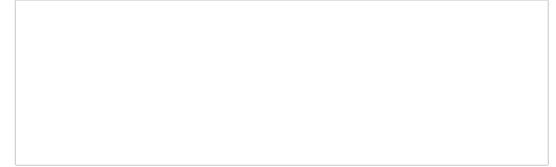


- **AIX makes multiple AAAA/ip6.arpa queries with no working IPv6 stack**
- **draft-liu-bonica-dhcpv6-slaac-problem**
- **Making apps IPv6 ready:**
  - **<https://ripe66.ripe.net/archives/video/1194>**
- **And there will be more...**



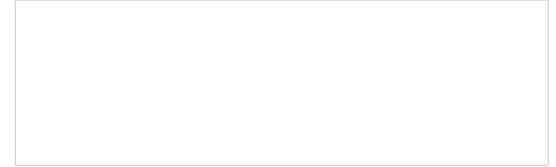
# CPEs & PD

# I can save \$5/customer and...



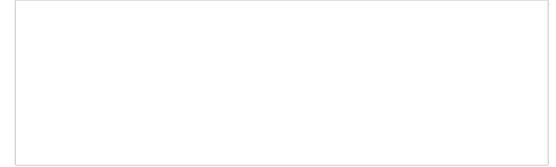
- ... make my support burden a nightmare
- RIPE 554: <http://www.ripe.net/ripe/docs/ripe-554>
- CableLabs and UNHIOL

# I can't NAT?



- **No...**
- **And /64 for house isn't enough**
- **But how big?**
- **And how to route?**

# Sorry, Mom. Use OSPF

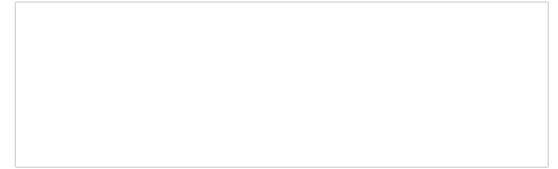


- **OSPF or IS/IS in homes? Really?**
- **HOMENET WG**
- **HIPNET**



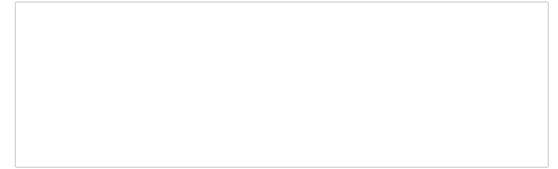
# Reverse/PTR goo

# How did this all start?



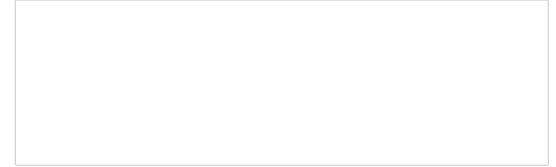
- ftp (<ftp.uu.net>, <ftp.wustl.edu>)
- SMTP
- Security devices
- Silly web things

# How did we do it IPv4



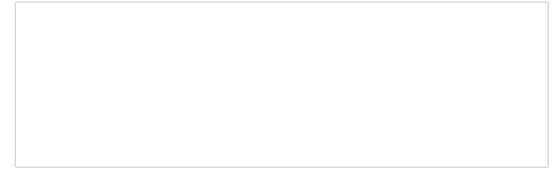
- **By hand (ow)**
- **Scripts**
- **\$GENERATE**
- **IPAM**

# 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?**

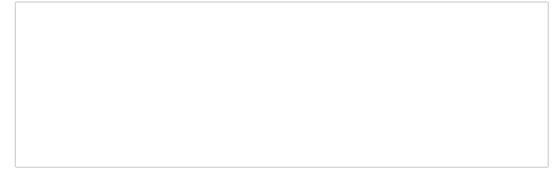
# So what are we left with?



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

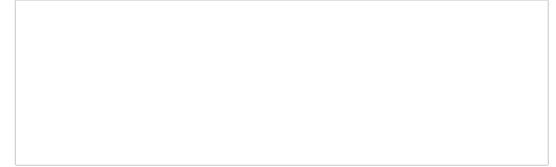


# ICMPv6



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

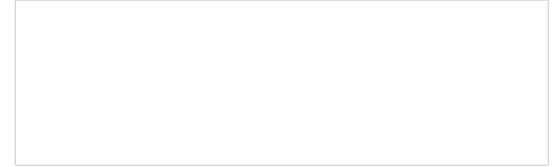
# ICMPv6 Filtering



- **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**

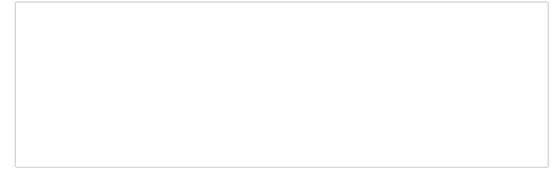


# IPSEC Myth vs Reality



**IPSEC in IPv6 is better than IPv4 because it was designed in and mandated.**

# And the reality

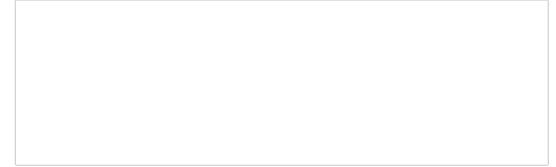


- **RFCs said “MUST” support IPSEC (but softening to “SHOULD”...)**
- **Didn’t define “support”, let vendors do it**
- **Vendors shipped, didn’t enable**
- **No PKI...**



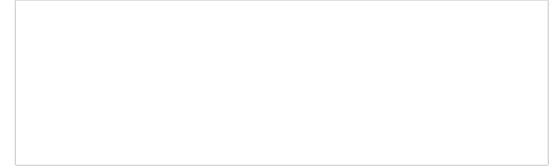
# IETF Blue Light Special

# The more things change...



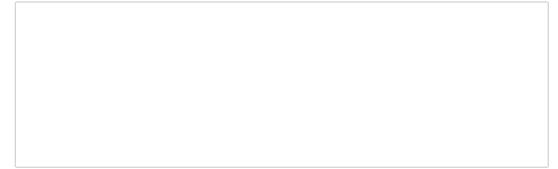
- ... the more they keep changing
- **DHC: 19 drafts, 73 RFCs**
- **IPv6: 12 drafts, 52 RFCs**
- **More every IETF meeting**

# And some RFCs are old...



- **RFC 3315 needs rewrite (liu/bonica)**
- **/etc/resolv.conf & RFC 6731**
- **Or problems are new (MIF)**
- **PKI and key mgmt (DNSSEC/DANE/RPKI)**

# What to do?



- **Join the WG mailing lists**
- **Come to IETF if you can**
- **Coordinate with other operators (IPBCOP.org)**
- **Beat on vendors**



# Q & A



Thank you!