

# IPv6-only at the IETF

How to retire IPv4 from our networks



# The *sunset4* working group

“ In order to fully transition the Internet to IPv6, individual applications, hosts, and networks that have enabled IPv6 must also be able to operate fully in the absence of IPv4. The Working Group will point out specific areas of concern, provide recommendations, and standardize protocols that facilitate the graceful "sunsetting" of the IPv4 Internet in areas where IPv6 has been deployed. This includes the act of shutting down IPv4 itself, as well as the ability of IPv6-only portions of the Internet to continue to connect with portions of the Internet that remain IPv4-only. ”

# The problems to be solved

- Remotely disabling IPv4
- Altering client connection establishment behaviours
- Disabling IPv4 in operating systems and applications
- On-Demand provisioning of IPv4 addresses

# Remotely disabling IPv4

- The absence of a DHCPv4 response is a failure
  - There is currently no way to signal that there should be no response
- Home router devices often provide a default IPv4 route
  - Regardless of whether the WAN port has an IPv4 address
- Auto-configured IPv4 addresses can enable IPv4 services
  - Monitoring and logging of IPv4 will not be in place
- “ARP for everything” can cause problems
  - Time-outs, large amounts of ARP traffic, etc

# Client connection establishment

- Happy Eyeballs implementations may prefer NAT44
  - NAT44 may have insignificant delays
  - Some implementations do not give IPv6 an advantage
  
- Operating systems send DNS queries for A records
  - Availability of IPv4 does not alter this behaviour
  - Causes many useless A record queries to be made

# Disabling IPv4 in software

- Libraries and applications depend on IPv4 O/S support
  - struct sockaddr\_in
  - AF\_INET
- Processes and applications depend on 127.0.0.1
  - Hard-coded dependencies on the IPv4 loopback interface
- Maintaining IPv4 support adds cost and security risks

# On-demand provisioning

- IPv4 could be provisioned on-demand to end users
  - CPE device can monitor for IPv4 traffic before requesting an address
- Happy Eyeballs implementations generate IPv4 traffic
  - Even if IPv6 works well
- Operating systems perform IPv4-based connectivity checks

# Potential solutions

- Approaches to indicate the network is IPv6 only
  - Using DHCPv6 (or RAs) with a new DHCP OPTION
  - Convey DHCPv4 messages over DHCPv6 transport
  - Transport DHCPv4 over IPv6
  - Transport DHCPv4 over SoftWire IPv4 tunnel on top of IPv6
- DNS A record filtering
  - Experiment run by WIDE
  - Use NAT64 and remove all A record responses from DNS
- Update RFC3927 s3.3 to fix “ARP for everything”
  - There is no agreement yet on how to update this
  - Only needs a single vendor change to start working



# Additional expertise is helpful

- These problems need addressing before IPv6 is common
- Contributing to the IETF will help deliver good solutions
- Comments, reviews of documents, and new ideas all welcome

<http://tools.ietf.org/wg/sunset4>