The background of the slide is a photograph of Earth as seen from space, showing the curvature of the planet and the atmosphere. The text is overlaid on this image.

Happy Eyeballs v2 & Reporting IPv6 Failures


RIPE 75
Dubai
October 2017

Jordi Palet (jordi.palet@theipv6company.com)

Happy Eyeballs v1 (HEv1) - 1

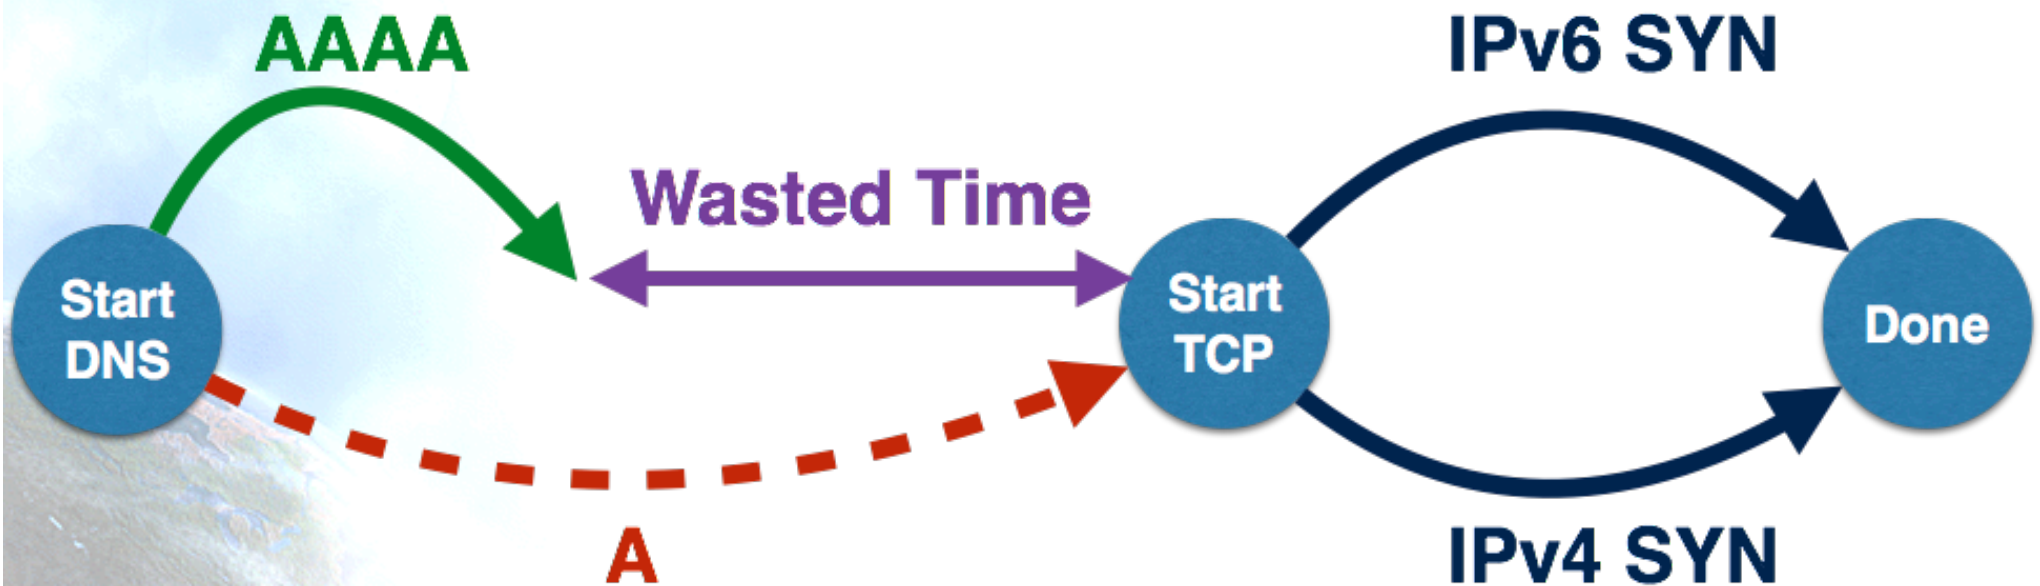
- Transition is based in preferring IPv6
- RFC6555 (April 2012)
 - Happy Eyeballs: Success with Dual-Stack Hosts
- In dual-stack hosts if IPv6 fails apps in the client present delays, compared with IPv4, which can be so high that may ruin the user experience
 - Up to 21 seconds in every web object
- HE sorts it out
 - Querying for both A y AAAA
 - Sending TCP SYN to both (IPv4 & IPv6)
 - Using the faster one, unless difference is small, so still giving preference to IPv6

Happy Eyeballs v1 (HEv1) - 2



* All figures provided by HEv2 co-authors
David Schinazi, Tommy Pauly
Apple

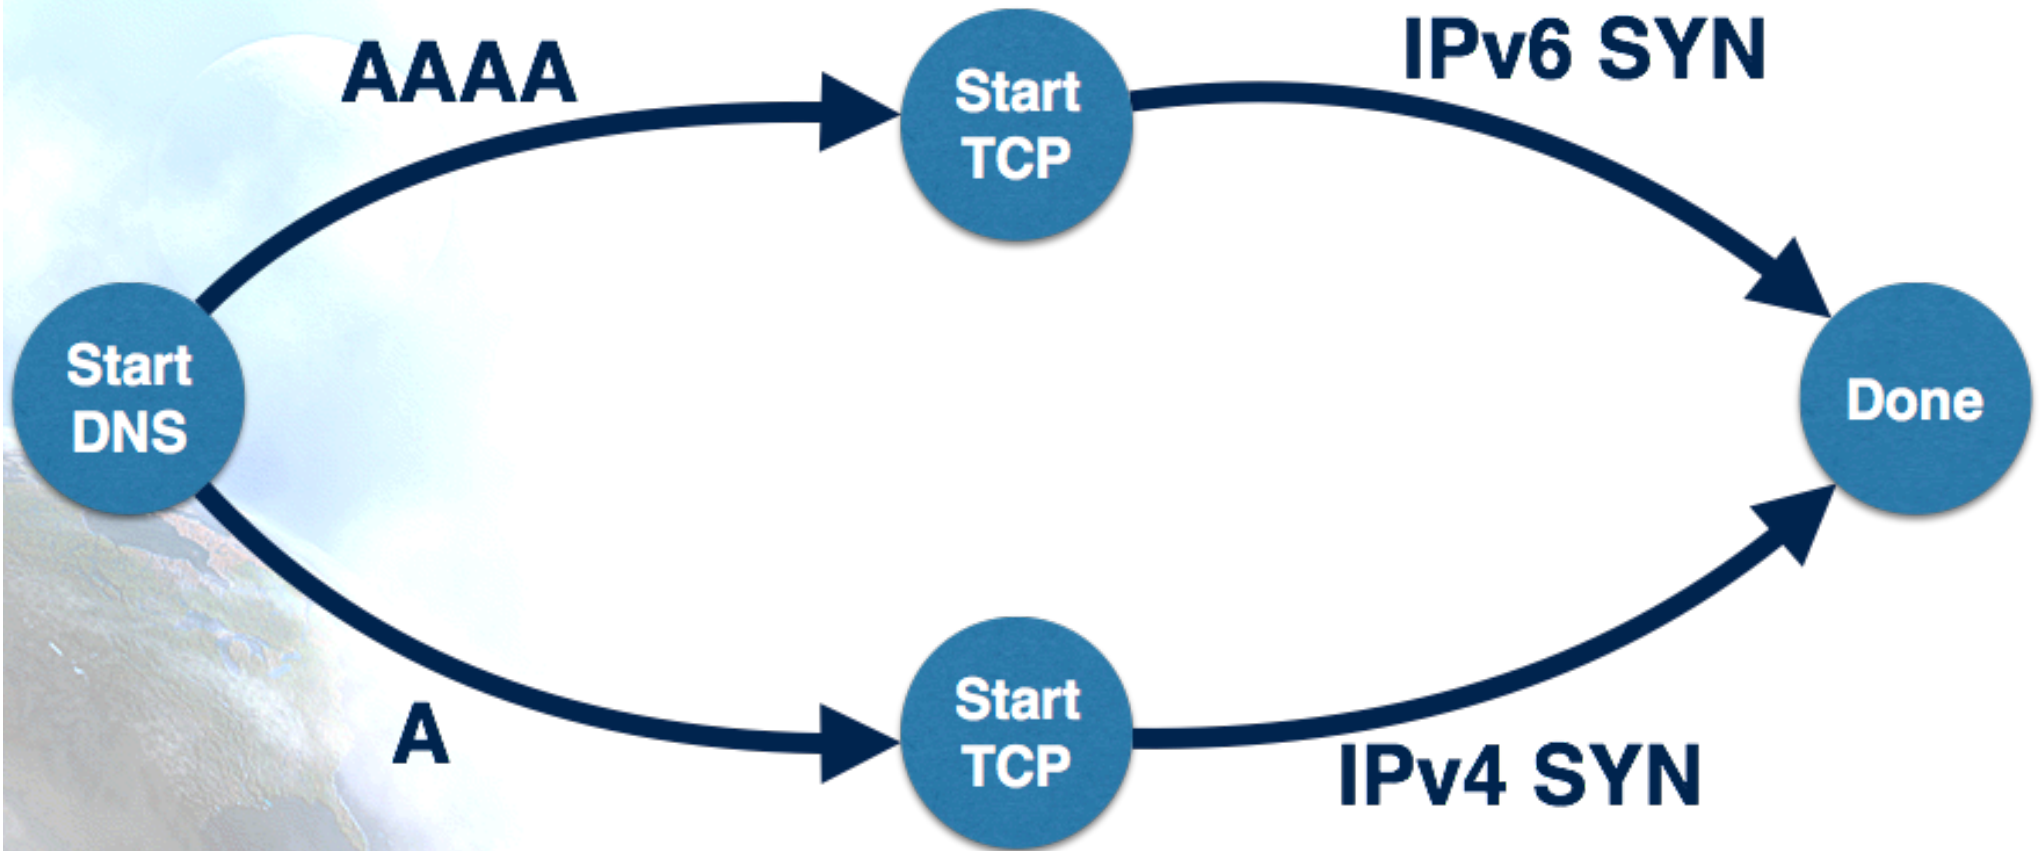
Happy Eyeballs v1 (HEv1) - 3



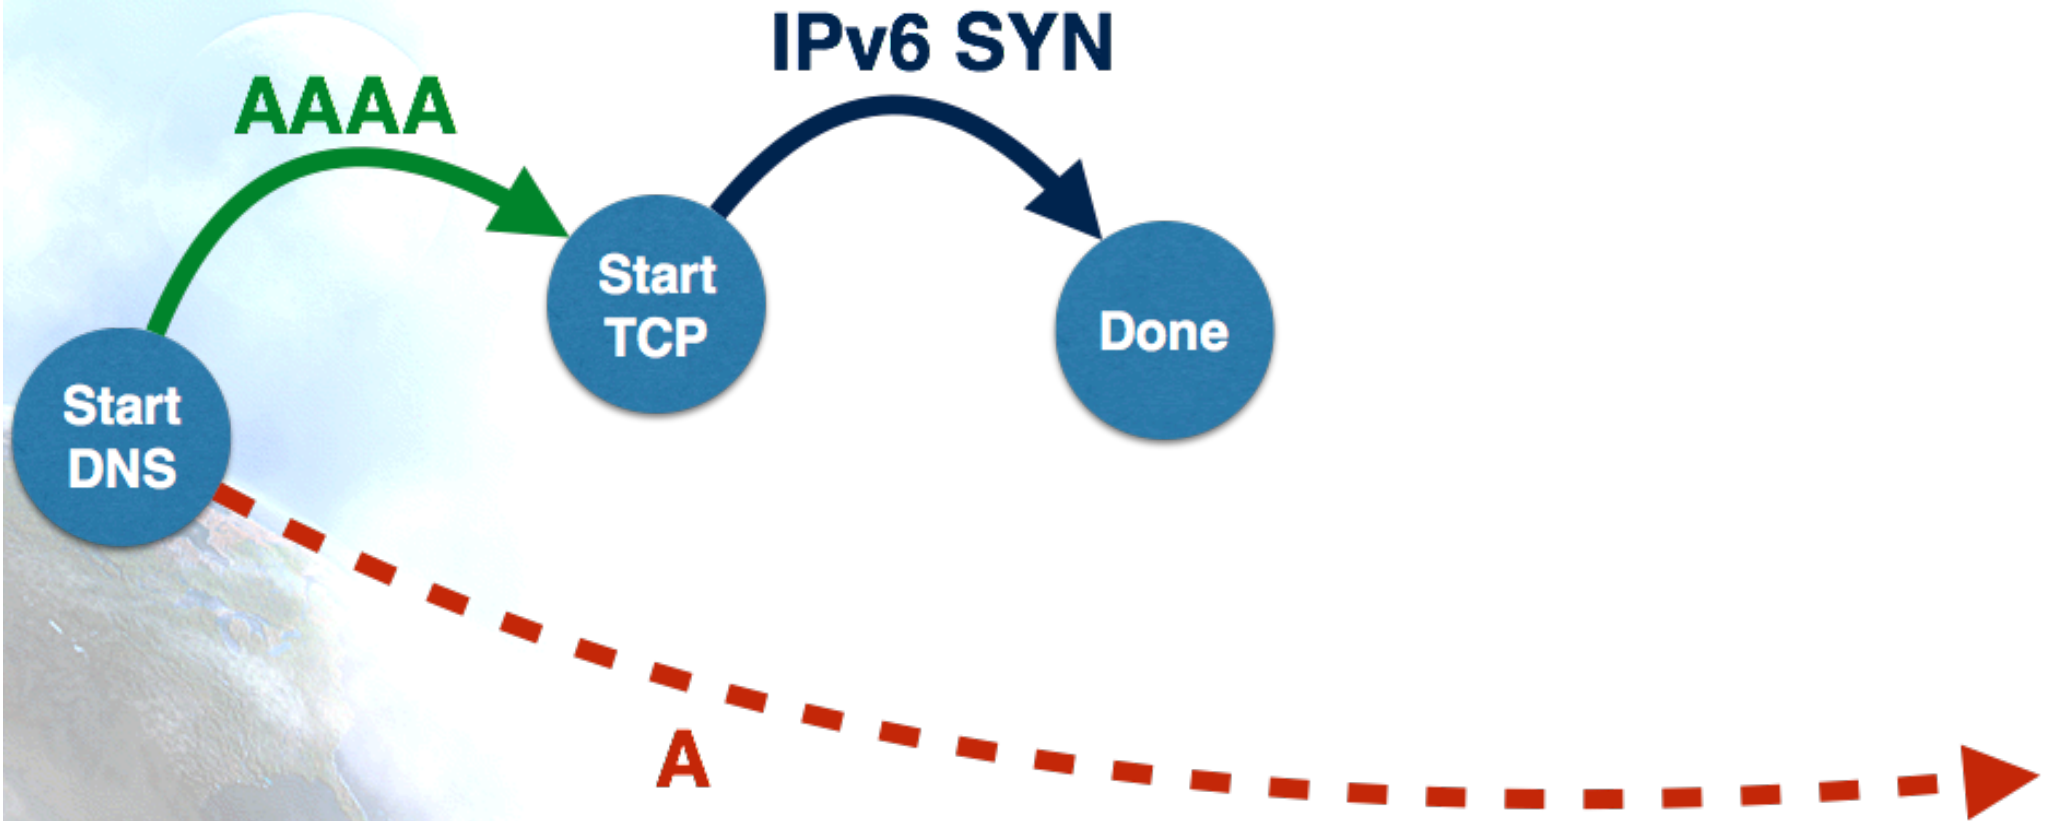
Happy Eyeballs v2 (HEv2) - 1

- draft-ietf-v6ops-rfc6555bis
- Extends HEv1
- HEv2 is already in production since long time ago in many Apple devices
- Since some years, they did measurements before publishing the ID (RFC in weeks?)
- It accelerates the users experience by “reordering” the address preference, while still trying to keep IPv6 on top

Happy Eyeballs v2 (HEv2) - 2




Happy Eyeballs v2 (HEv2) - 3



Happy Eyeballs v2 (HEv2) - 4

- RFC6724 (Default Address Selection for IPv6) vs HEv2

RFC6724



HEv2



HE good or bad ?

- Happy Eyeballs is good for the users
- However, “hides” IPv6 failures, so is bad for operators if they don’t have appropriate ways to monitor their correct IPv6 deployment
 - Big content providers often block IPv6 (by hiding AAAA records) for operators with “bad” IPv6 quality
 - Consequently, IPv6 traffic will not grow in those networks, which is the main goal
 - Badly performed IPv6 deployments are counterproductive and may bring bad technical and business decisions

Common IPv6 Failures

- IPv6 deployment, is unfortunately, many times, done in a “broken” way because not “unlearning” IPv4, so it creates troubles which reduce the users perceived “QoS”
 1. ICMPv6 filtering
 - Breaks PMTUD and the destination becomes non-reachable
 2. IPv6 path doesn’t work or has higher delay
 - Fallback to IPv4

Reporting of HEv2 Failures

- draft-palet-ietf-v6ops-he-reporting
- This document describes a HE (v1 & v2) extension, to do an automated failure reporting when the client fallback to IPv4
- ¿How?
 - KISS: Reusing existing and commonly available protocols
 - syslog, only UDP port 514 (RFC5424/26)
 - Very common in many networks
 - No need to ask the operators to install anything “new” or “different”

Automating the Reporting

- Syslog sorts-out the operator network side
- We also want “zero-config” in clients
- ¿How?
 - Reusing again ...
 - This only makes sense if the ISP already has IPv6 to customers
 - The ISP uses a NSP (Network Specific Prefix)
 - HE discovers that prefix by means of RFC7050 (Discovery of the IPv6 Prefix Used for IPv6 Address Synthesis)
 - Add to it a well known and no longer used IPv4 (192.88.99.0/24, it was 6to4 anycast, deprecated by RFC7526)
 - So we have an IPv6 GUA (or /96 for HA) for clients to report to:
 - Network-Specific Prefix::192.88.99.1 (example 2001:db8::192.88.99.1)

Conclusions

- HEv1/v2 DO NOT solve PMTUD failures
 - Operators need to avoid breaking ICMPv6
- If “draft-palet-ietf-v6ops-he-reporting” becomes an RFC, is NOT a “solution”, but
 - Having data for error allows sorting them out
 - In your network or tell to third parties
 - Monitoring your network it is will very important:
 - Same issues than IPv4, consider longer-term for IPv6
 - Traffic quality
 - Quantity
 - Stability
 - Prefix visibility
 - ...
- RIPE ATLAS can help to that
 - Also paid services available

Thanks !

Contact:

– Jordi Palet:

jordi.palet@theipv6company.com