IPv6 New RFCs

ESNOG/GORE21
Barcelona
April 2018

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 hots 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

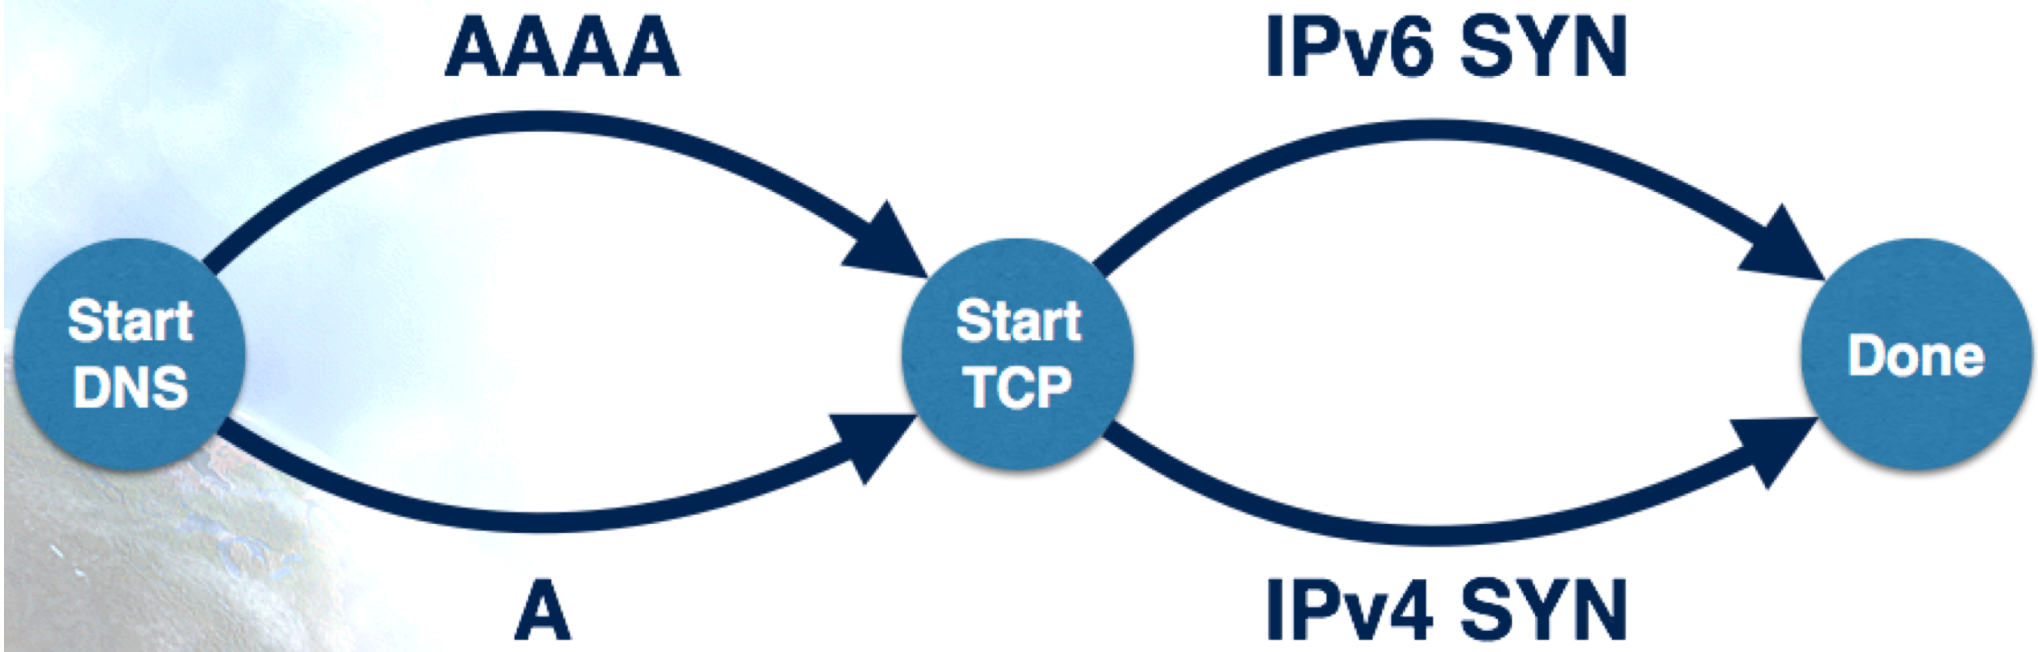

* All figures provided by HEv2 co-authors
David Schinazi, Tommy Pauly
Apple

Happy Eyeballs v1 (HEv1) - 3

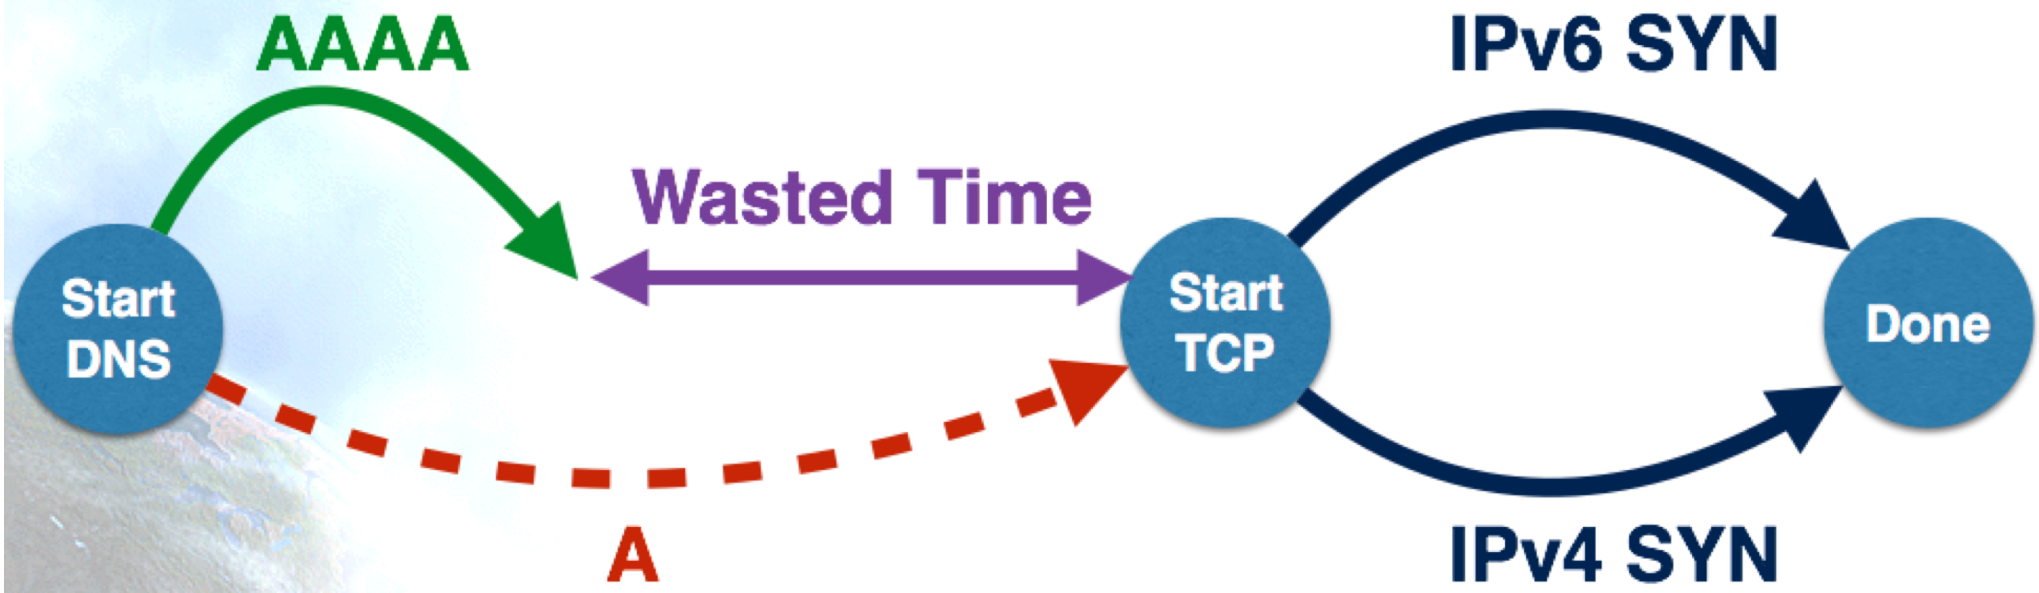

Happy Eyeballs v2 (HEv2) - 1

- RFC8305
  - "Happy Eyeballs Version 2: Better Connectivity Using Concurrency"
- Extends HEv1
- HEv2 is already in production since long time ago in many Apple devices
- Since some years, they did measurements before publishing the RFC
- It accelerates the users experience by "reordering" the address preference, while still trying to keep IPv6 on top

Happy Eyeballs v2 (HEv2) - 2

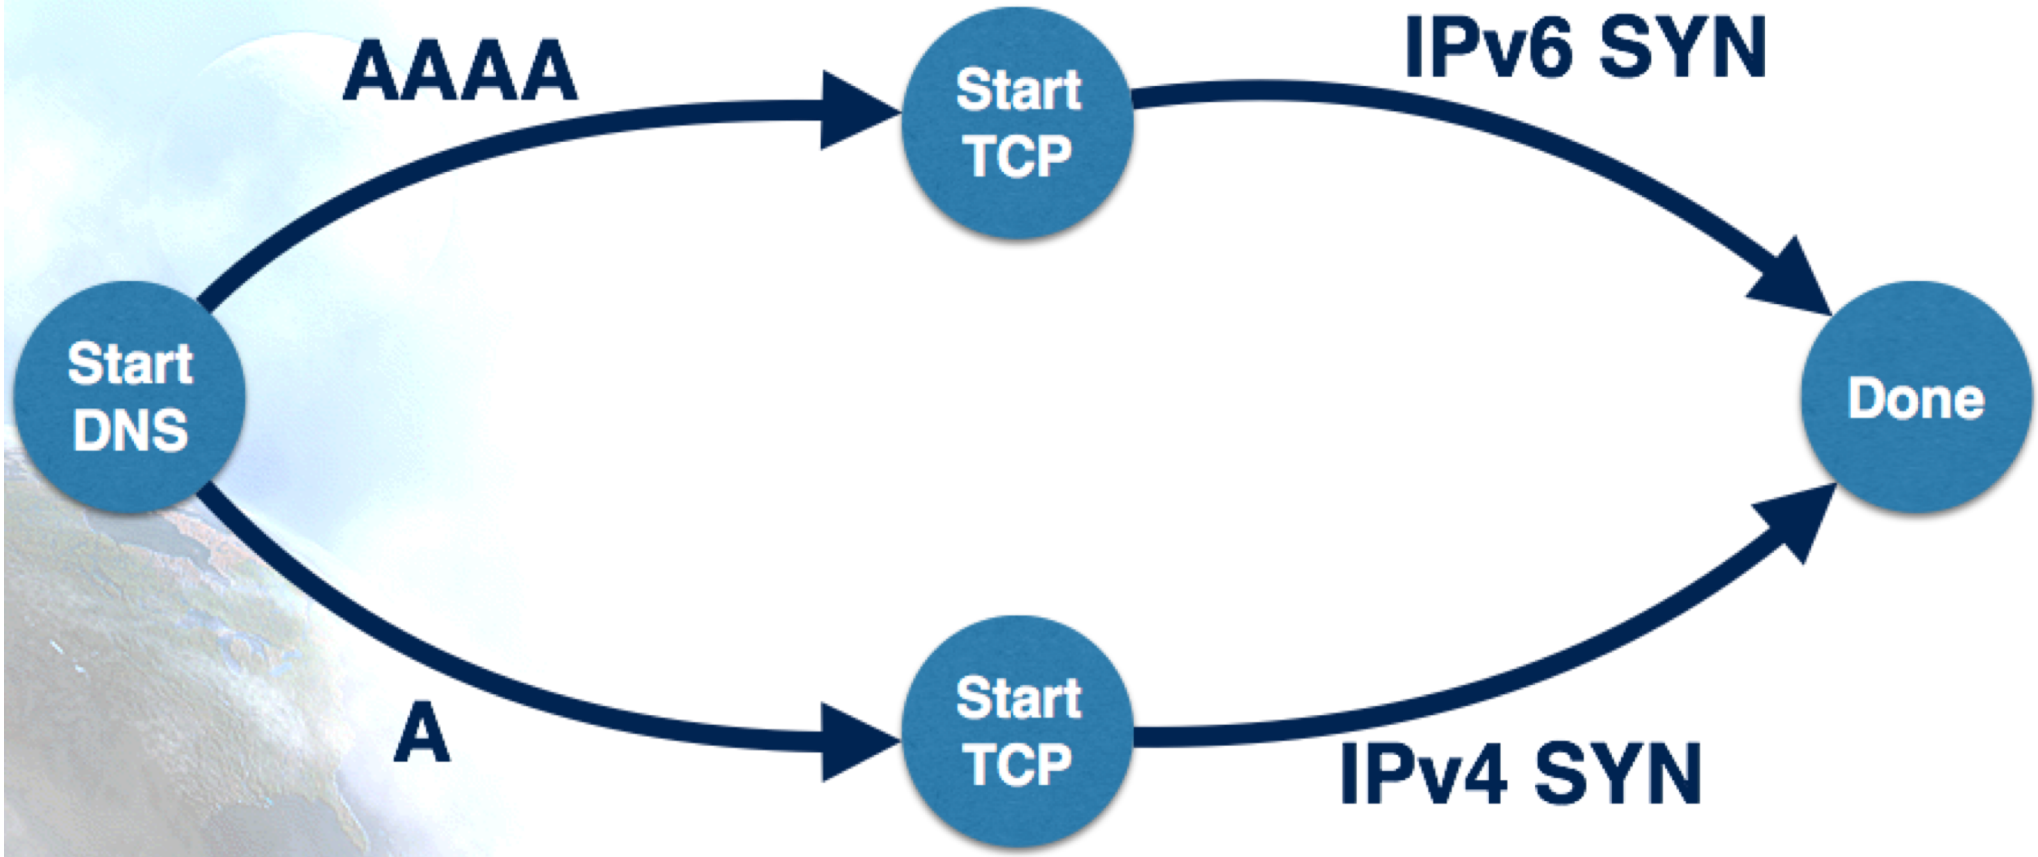

Happy Eyeballs v2 (HEv2) - 3

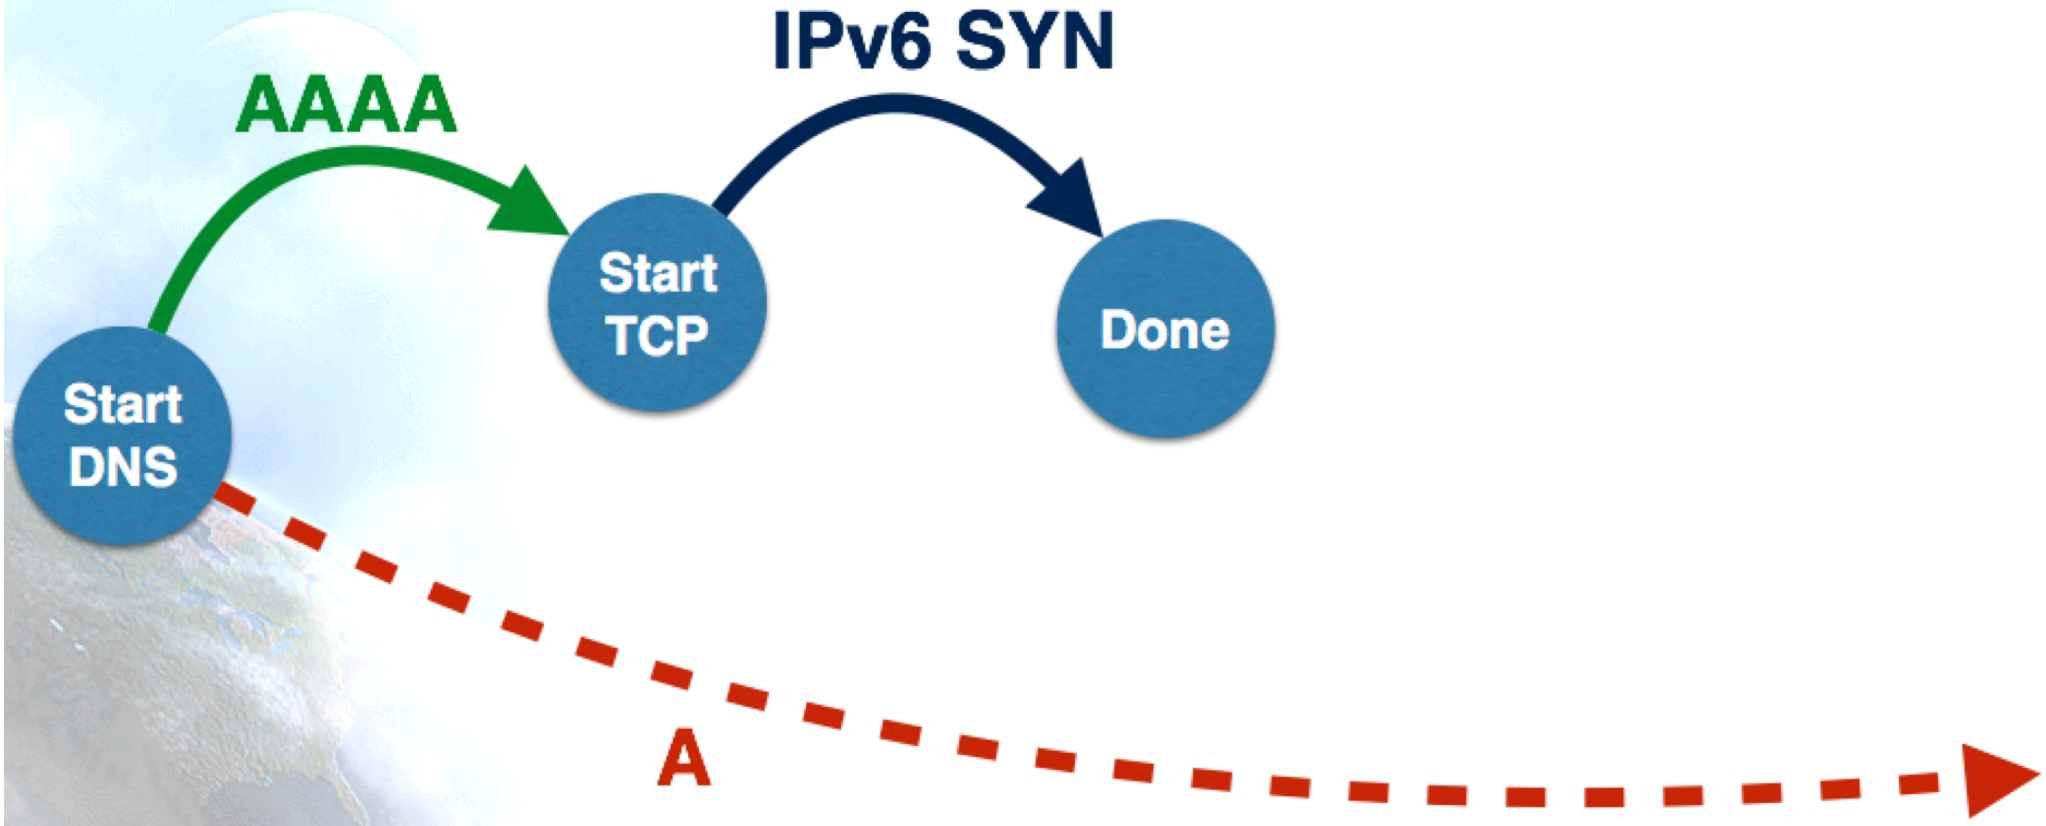

Happy Eyeballs v2 (HEv2) - 4

- RFC6724 (Default Address Selection for IPv6) vs HEv2

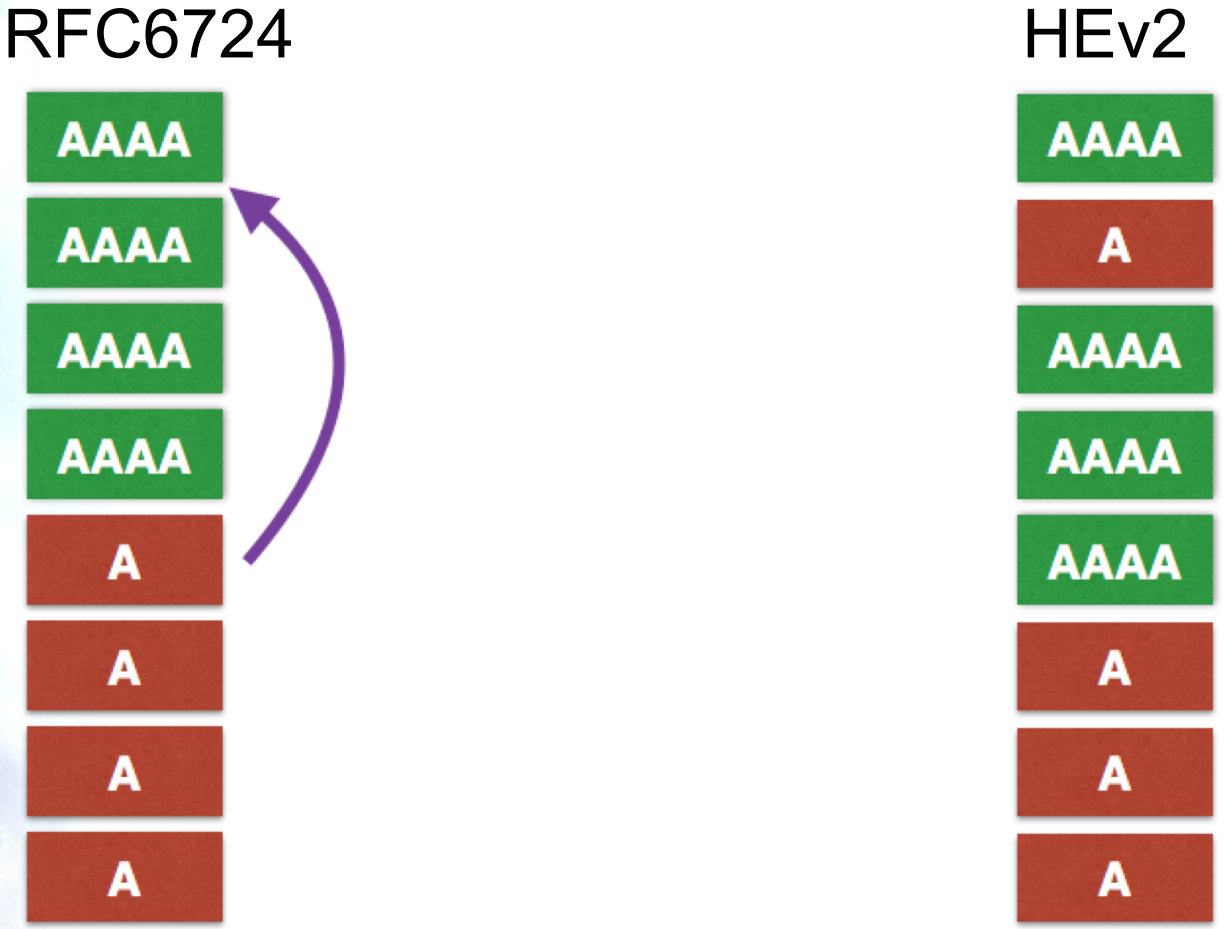

HE good or bad ?

- 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 fall-back 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)

HEv2 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

RFC8273

- RFC8273: “Unique IPv6 Prefix per Host”
- Not a “new” protocol, so already widely supported
  - **Use “existing IPv6 protocols” to allow a unique IPv6 prefix (instead of a unique IPv6 address from a shared IPv6 prefix) to be assigned to a host interface**
- Allows improved host isolation and enhances subscriber management on shared network segments, such as Wireless networks, data centres, among others
- Provides a very simple mechanism for a single host or interface, to be able to run 2^{64} virtual machines, with their own global IPv6 address, not requiring to share a single one

“How To”

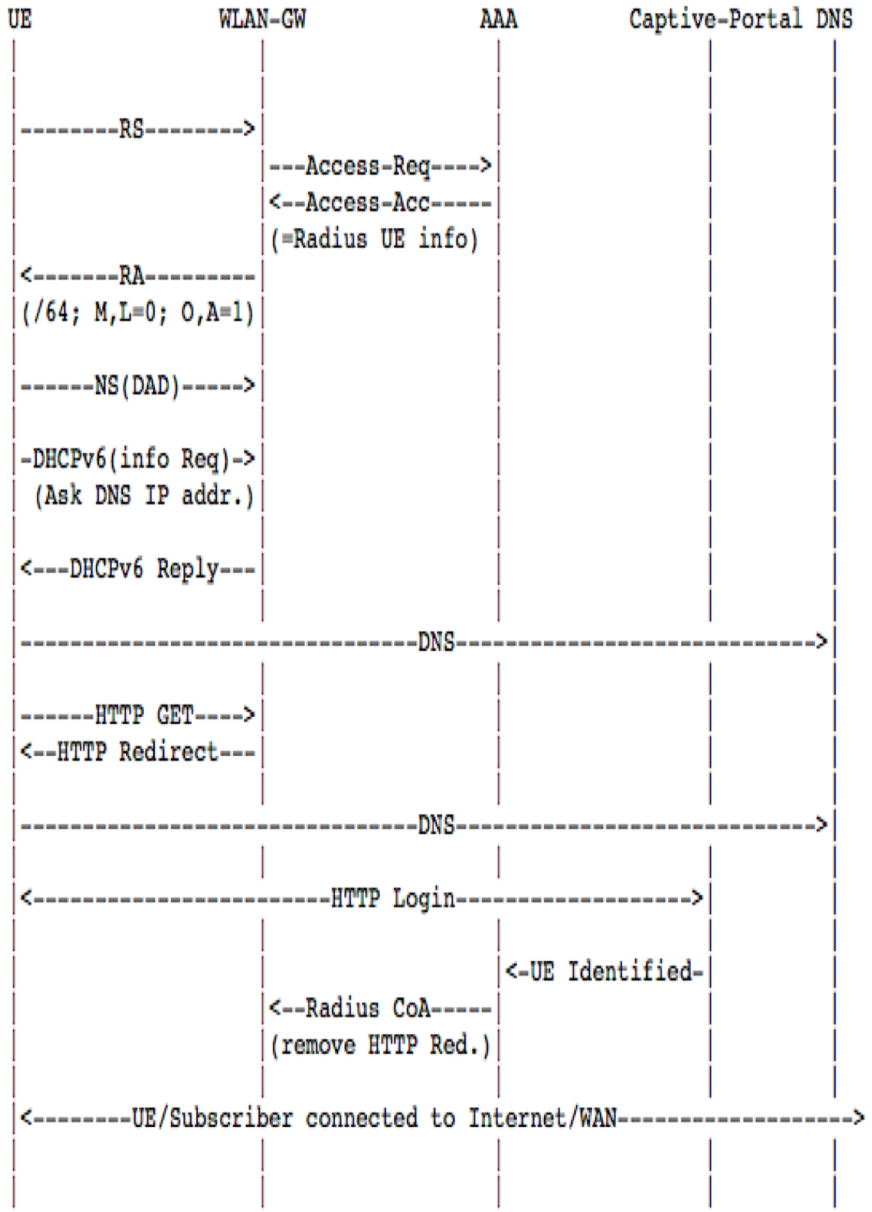

1. First-hop router is a L3 edge router
2. UE connects to the shared-access network and starts IP configuration with SLAAC RS
3. First-hop router sends solicited RA response ONLY to the requesting UE
   - Instead of using the link-layer multicast address (all-nodes group), using the link-layer unicast address of the requesting UE
   - The solicited RA contains the unique prefix (/64) and flags (to indicate if SLAAC and/or DHCPv6 should be used, etc.)
   - Prefix from locally/centrally managed pool, aggregate IPv6 block, …
   - Flags, best practices:
     - M-flag = 0 (address not managed with DHCPv6, 1 for DHCPv6 prefix delegation)
     - O-flag = 1 (DHCPv6 used for other configuration information)
     - A-flag = 1 (UE can configure itself using SLAAC)
     - L-flag = 0 (prefix is not an on-link prefix, everything sent to the gateway)
4. Periodically unsolicited RAs follow same approach

Usage Scenarios

- We are already doing in cellular:
  - /64 per PDP context
  - Prefix sharing with other devices (tethering)
  - Facilitate IPv6-only access (and IPv4-as-a-service)
- Allows extending same concept to other scenarios:
  - Hot-Spot
    - WiFi Calling: Secured Voice over WiFi over “untrusted” connection
      - IPv4 or IPv6 IPsec tunnels to the ePDG (evolved Packet Data Gateway)
  - Corporate networks
  - Data Center
- Allows also IPv6-only access and IPv4-as-a-service
  - Same concept as above for WiFi Calling
    - VPN “on demand” in “own” network for IPv4 services
    - No need for NAT44 (lowers logging costs and fragmentation issues)

Hot-Spot Usage

- WiFi shared-access L2 network
- Provide isolation between user devices either due to legal requirements or to avoid potential abuse
- By using “unique IPv6 prefix per host”, devices only can communicate thru the first-hop router
- Automatically avoids attacks based on link-local ICMPv6:
  - DAD reply spoofing
  - ND cache exhaustion
  - Malicious redirects
  - Rogue RAs
- Better scalability and robustness than DAD proxy, forced forwarding, ND snooping, etc.

Hot-Spot Example

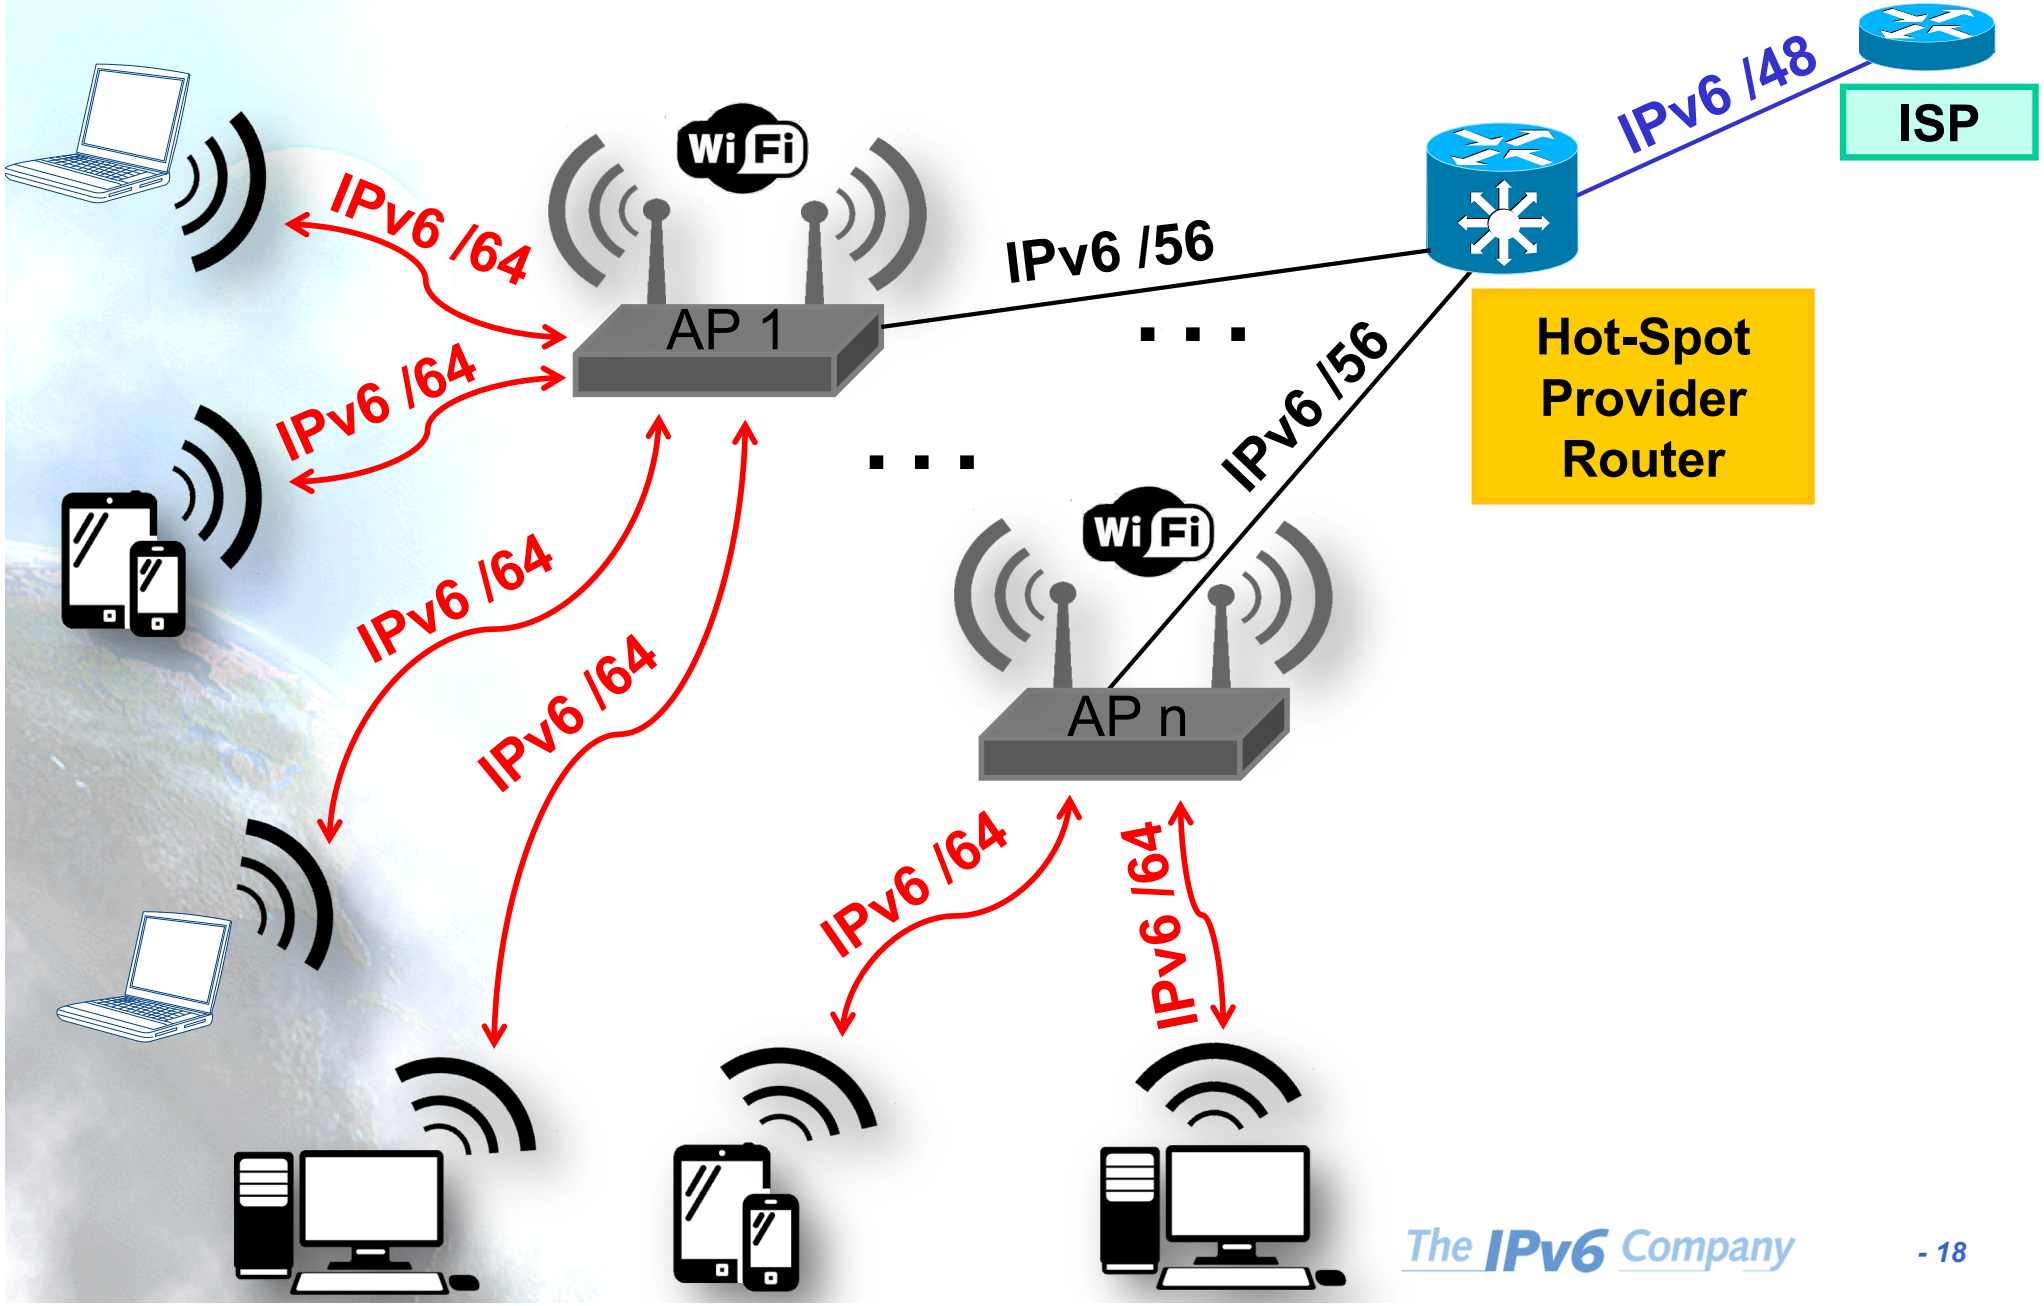

Data Centre Usage

- “How to” same as for the Hot-Spot case
- The UE “server” may need multiple addresses from the same unique IPv6 prefix (VMs, containers), so just need to configure them
- The first-hop router must be able to handle the presence and use of those

Data Center Example

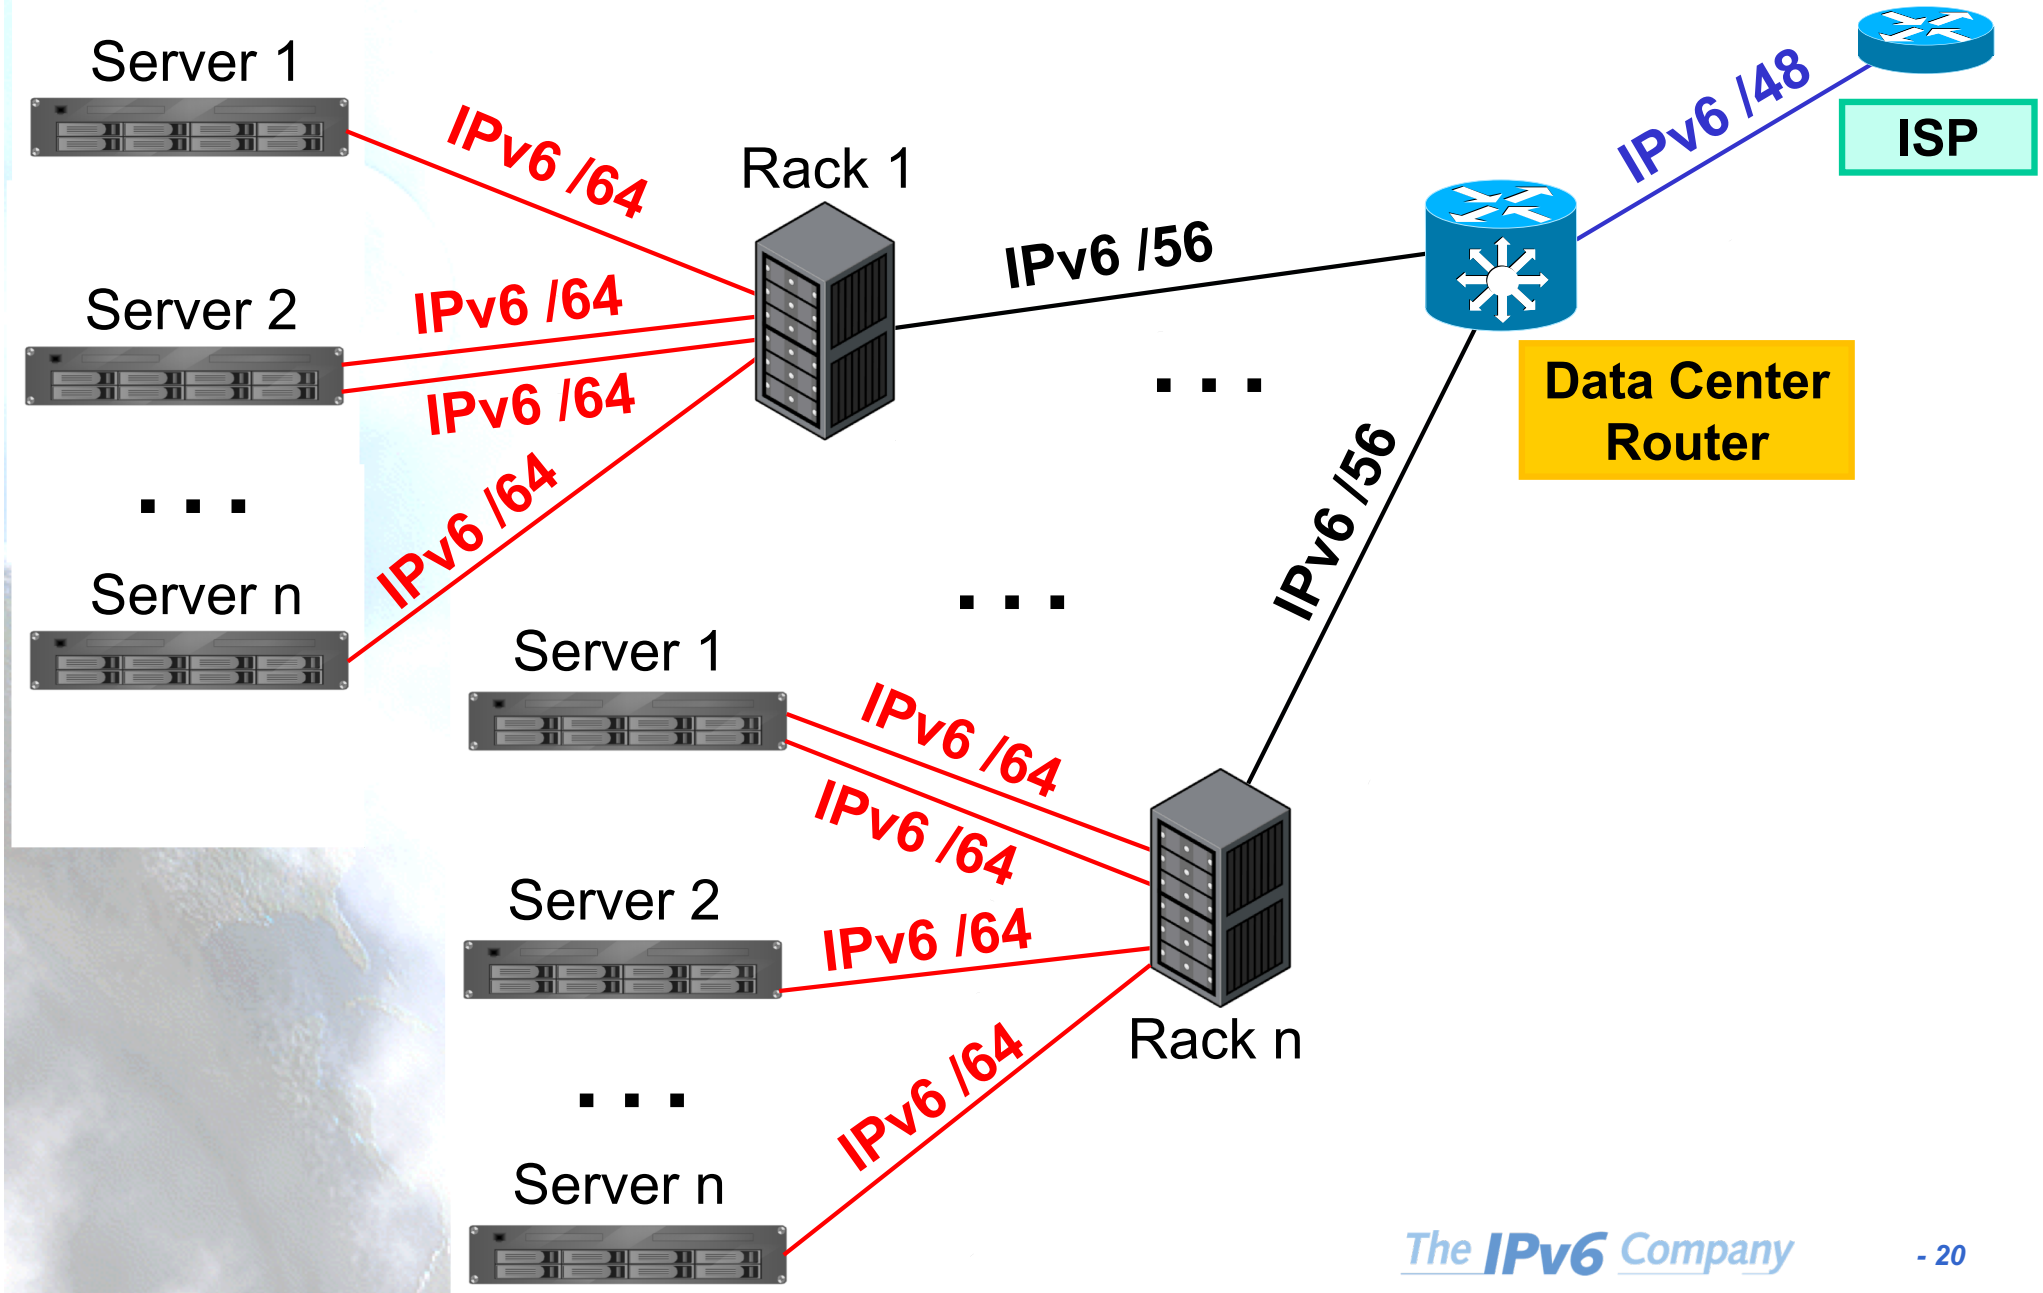

Enterprise Example

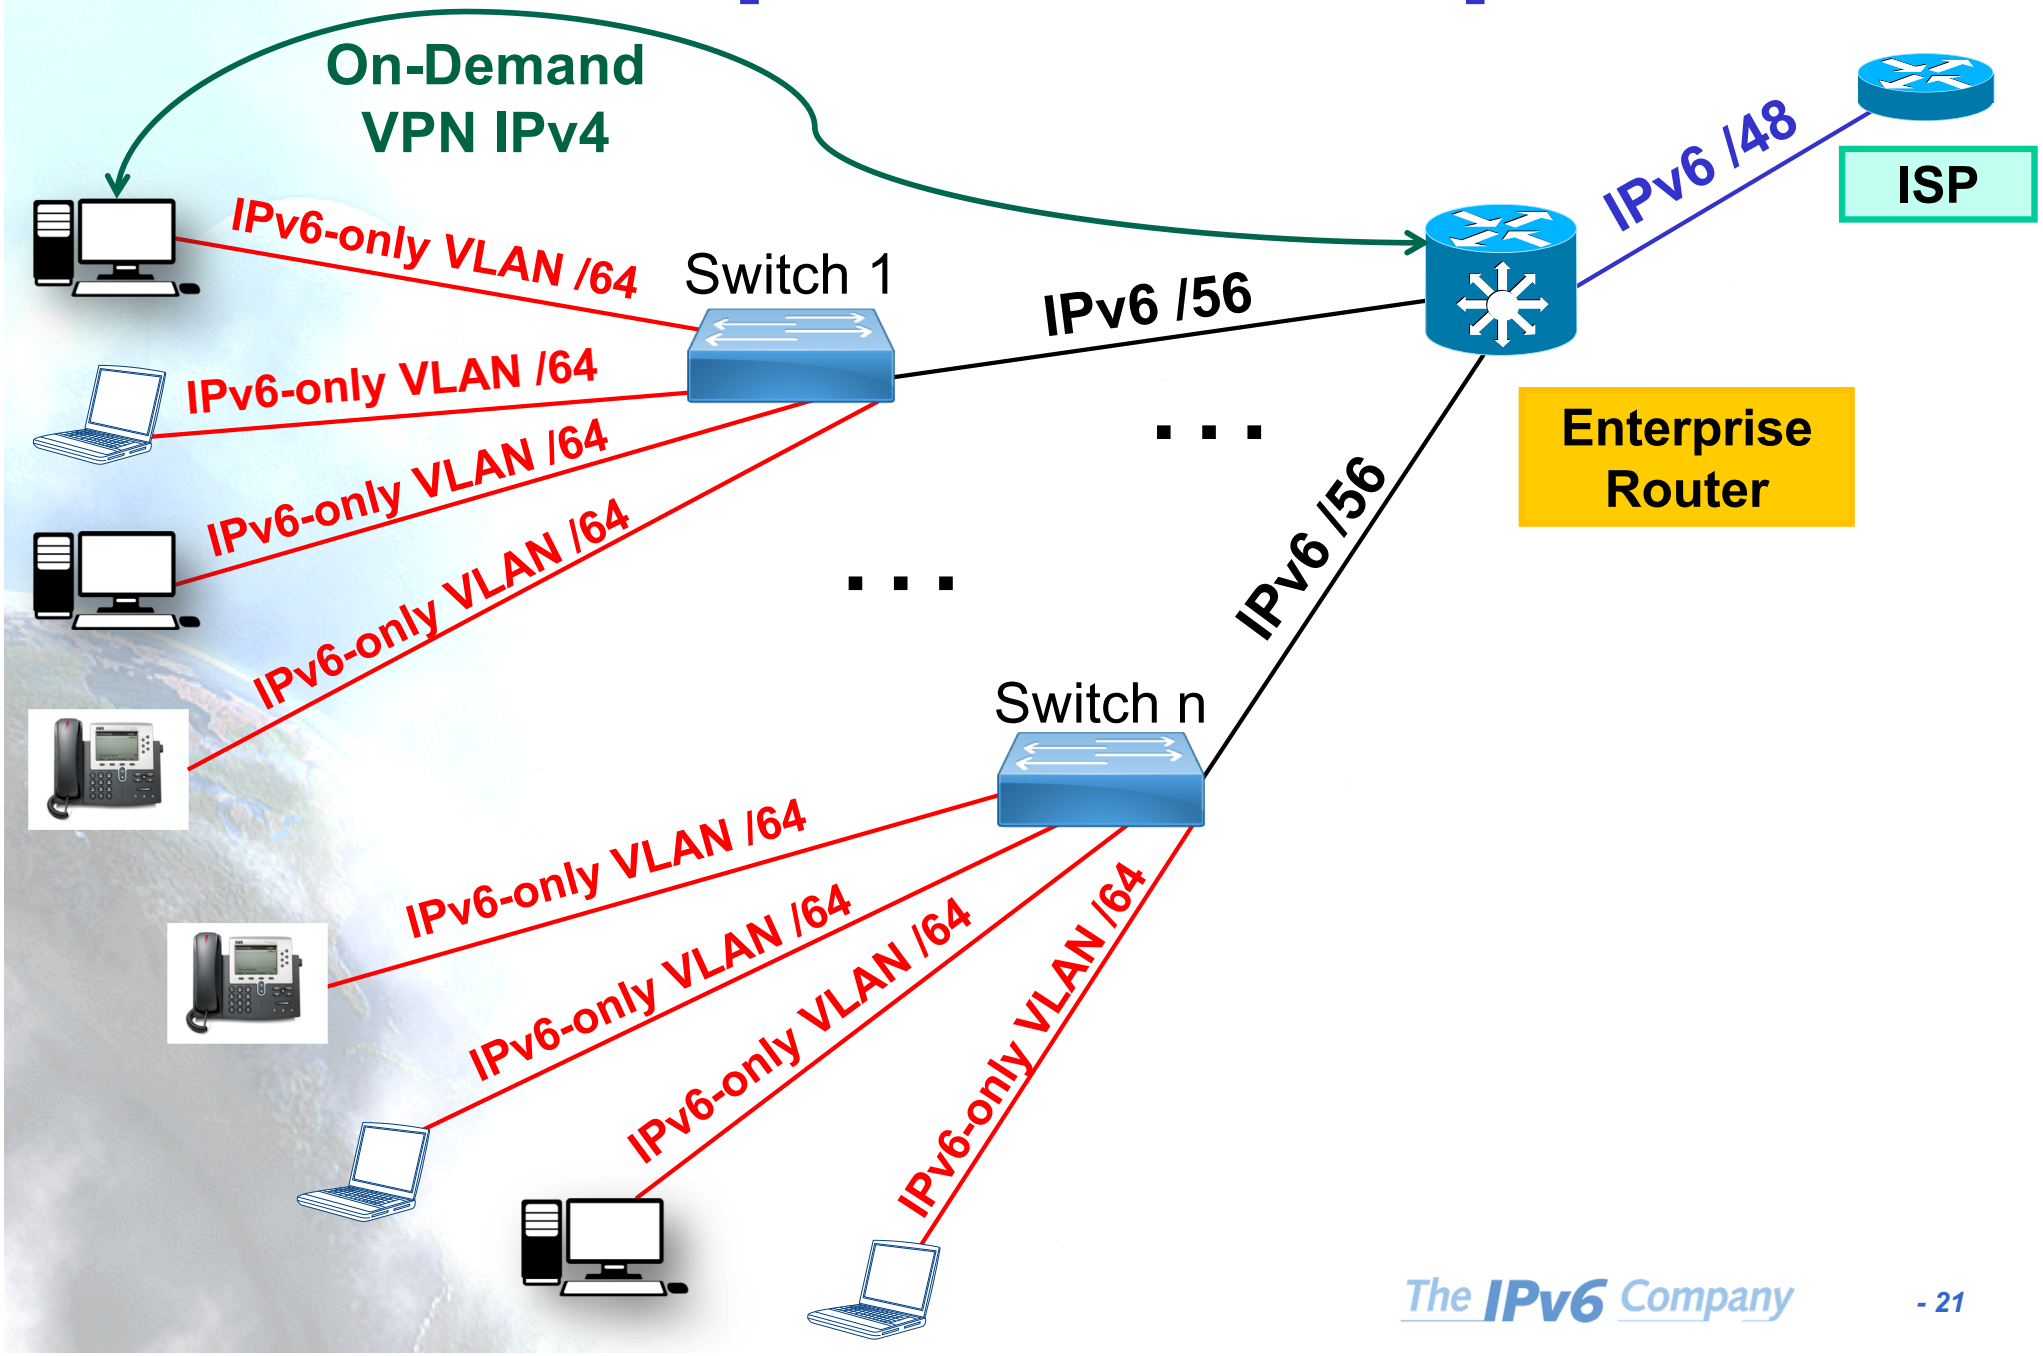

Conclusions RFC8273

- Stable and secure IPv6-only experience
- No performance impact
- Secure host-to-host communication managed by first-hop router
- Each unique IPv6 prefix can function as a control-plane anchor point to ensure that each device receives expected subscriber policy and service levels
  - Throughput
  - QoS
  - Security
  - Parental control
  - Other value-added-services …

Thanks !

Contact:

- Jordi Palet: jordi.palet@theipv6company.com