IPv6 Unique Prefixes

Draft Update and generic reflections of IETF progress

Gunter Van de Velde

Product Line Manager @ Nokia

Unique IPv6 Prefix Per Host

- draft-ietf-v6ops-unique-ipv6-prefix-per-host-02
  - Last version update: 13 March 2017
  - Past WGLC on 11 April 2017
  - IESG Telechat requested after 28 April 2017
  - Intended status: BCP
- What is it about?
  - IPv6 prefix recommendation for shared public LAN's (wired and wireless)
  - Traditional address planning
    - LAN: /64
    - Host: 128 bit host address from the LAN /64
  - New suggested address planning
    - LAN: non-defined
    - Each host on LAN gets a /64 IPv6 Prefix
- Why?
  - Improved subscriber management and improved subscriber security

Unique IPv6 Prefix Per Host

- Applicability Scope:
  - Stable and secure IPv6 only experience
  - Zero performance impact
  - Supported by largest set of hosts
  - Allow innovation
  - Secure host-to-host communication managed by first hop router
  - BCP guidelines

IPv6 Wi-Fi Subscriber Onboarding Procedures (1)

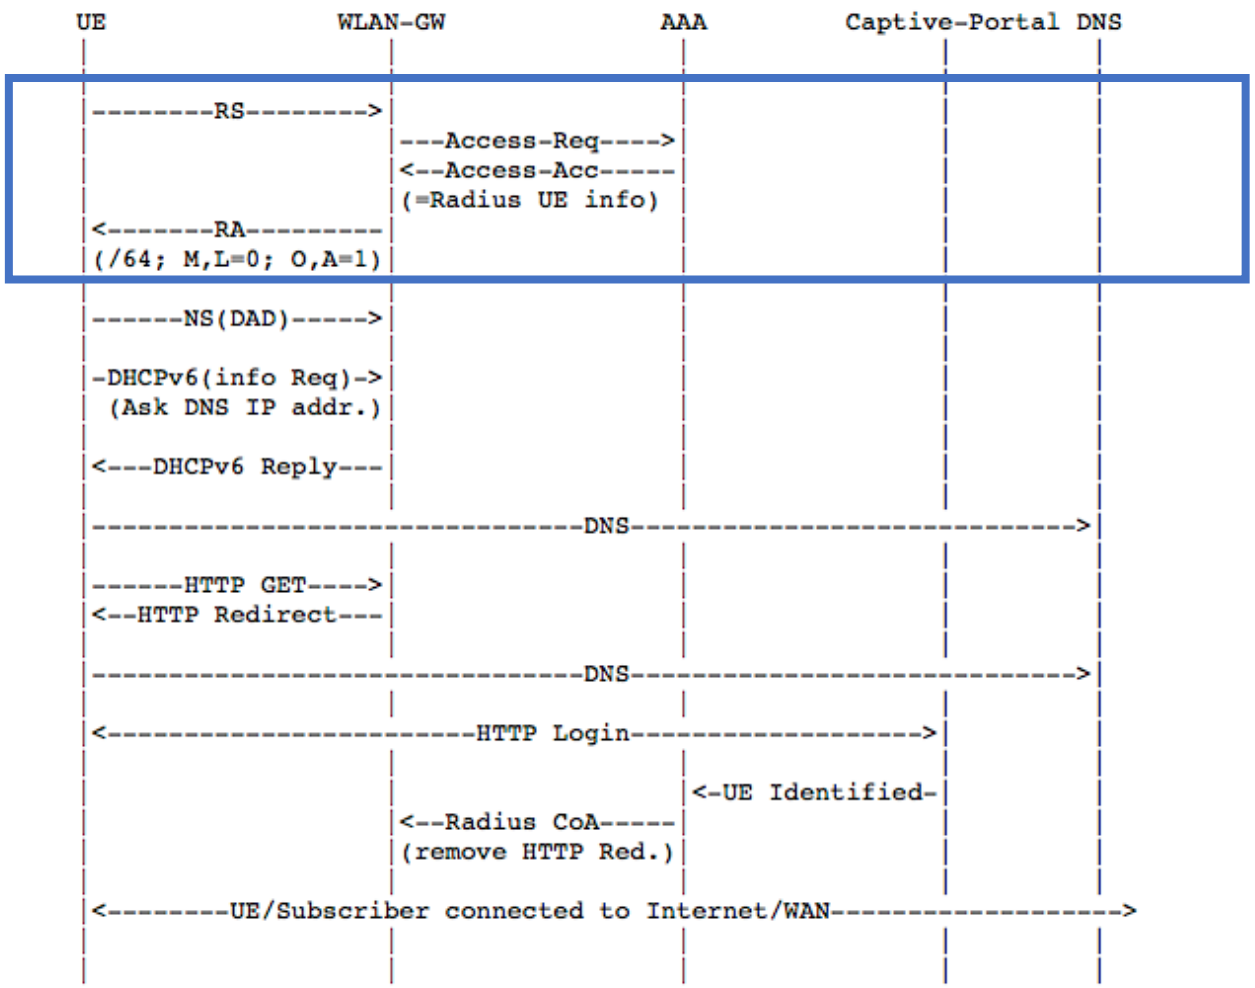

- When UE connects it sends a RS to learn
  - IPv6 Gateway, Prefix information, DNS, remaining info for global routing
  - RS send from UE via the AP-bridge to the WLAN-GW
  - Due to split-horizon for BUM traffic the RS is not seen by other UE's connected to the same AP
- First time UE connects it is not Authorized and WLAN-GW queries AAA server
- AAA server checks policy DB and returns /64 together with http-redirect to Captive portal via Radius-acknowledge message

IPv6 Wi-Fi Subscriber Onboarding Procedures (2)

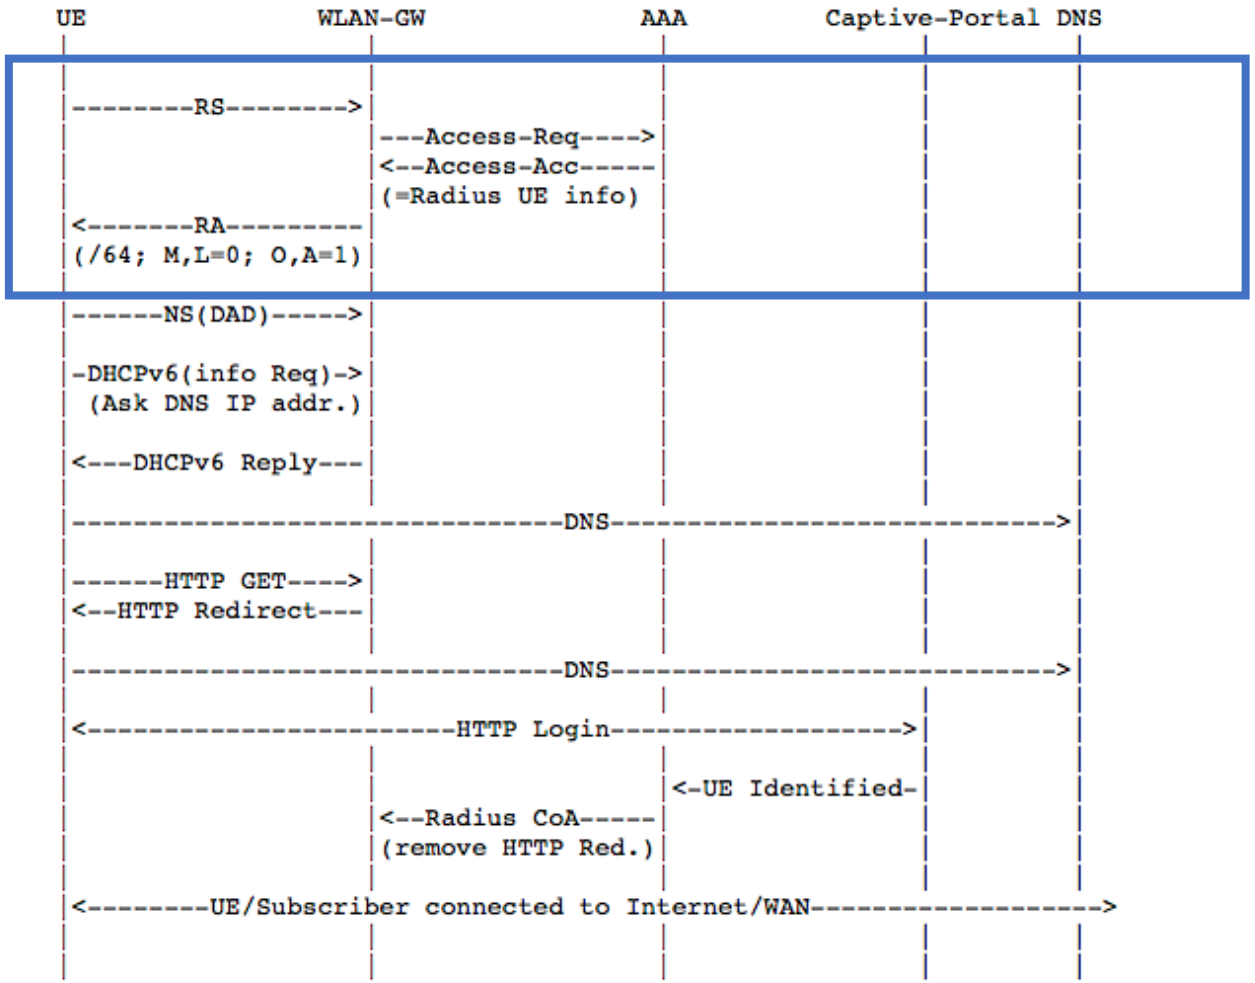

- WLAN-GW uses received Radius info to compose the “RA” response to the UE originated “RS” message
- RA contains few important bits of information
  - A IPv6 /64 prefix
  - Some flags
- (1) IPv6 /64 prefix
  - Locally managed pool on WLAN-GW
  - Pool signaled through Radius
- (2) Some flags
  - Indicate to use SLAAC and/or DHCPv6
  - Prefix is on/off-link
  - Is there need to request ‘Other’ information (e.g DNS)?

IPv6 Wi-Fi Subscriber Onboarding Procedures (3)

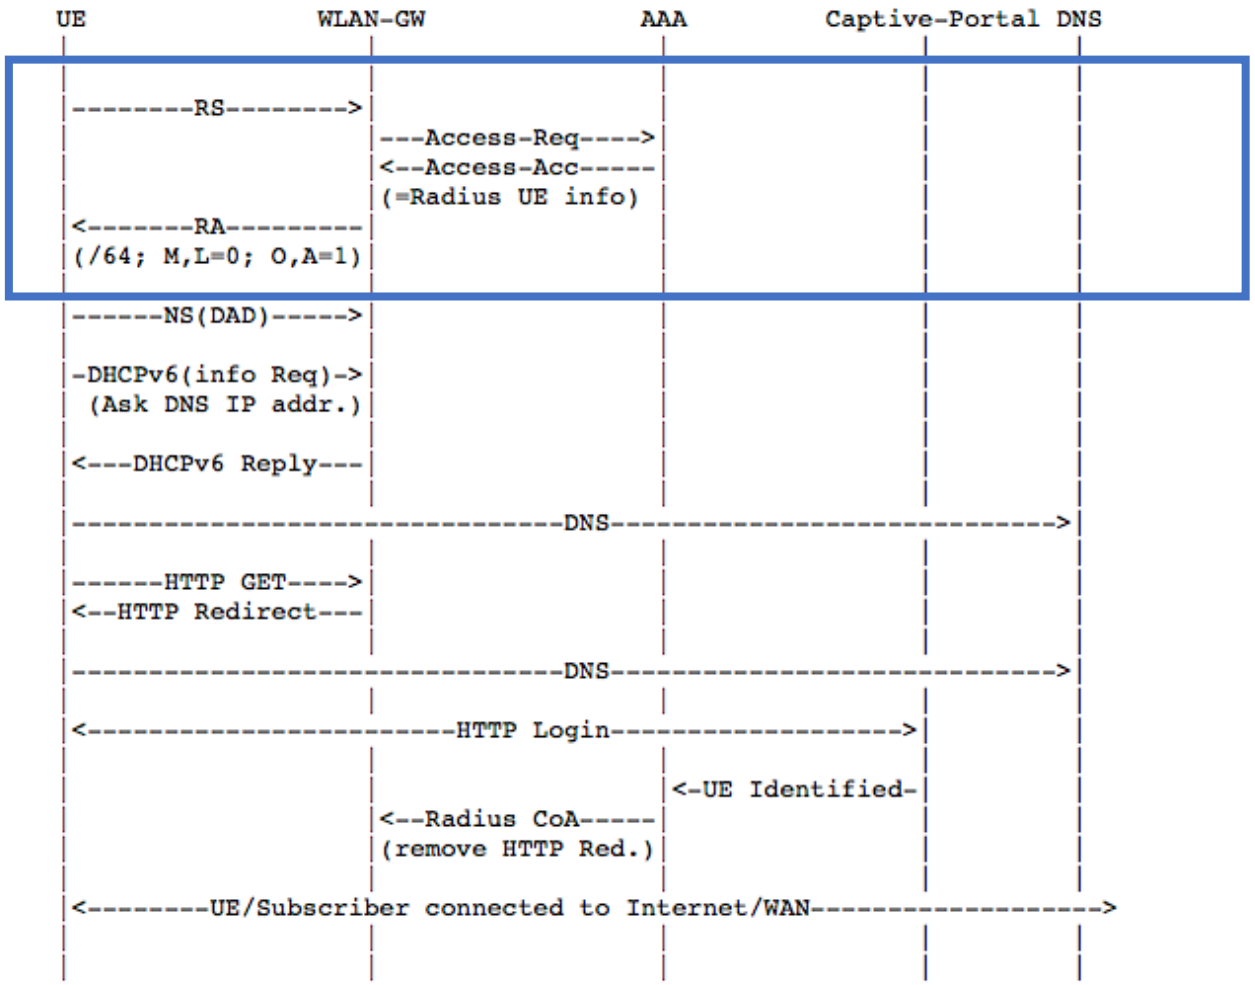

- IPv6 RA flags for best common practice
  - M-flag = 0 (UE/subscriber address is not managed through DHCPv6), this flag may be set to 1 in the future if/when DHCPv6 prefix delegation support over Wi-Fi is desired)
  - O-flag = 1 (DHCPv6 is used to request configuration information i.e. DNS, NTP information, not for IPv6 addressing)
  - A-flag = 1 (The UE/subscriber can configure itself using SLAAC)
  - L-flag = 0 (The UE/subscriber is off-link, which means that the UE/subscriber will send packets ALWAYS to his default gateway, even if the destination is within the range of the /64 prefix)
  - Timers:
    - IPv6 Router Advertisement Interval = 300s
    - IPv6 Router LifeTime = 3600s
    - Reachable time = 30s
    - IPv6 Valid Lifetime = 3600s
    - IPv6 Preferred Lifetime = 1800s
    - Retransmit timer = 0s

IPv6 Flow Label Saga

- General feeling at IETF:
  - Do not touch the flow-label (non-permutable field)
  - Flow label is unpredictable and ill defined
  - In general it has been useless 20 bits in IPv6 header (and it is not expected to change)
- However SPs desire to make use of IPv6 Flowlabel
  - https://tools.ietf.org/html/rfc6294

Alternate marking
by Telekom Italia

- Draft: draft-ietf-ippm-alt-mark
- What is so cool about Alternate Marking
  - Metrics harvested using REAL traffic
  - No injection of probes or other artificial metrics
  - Method has standard effort happening on IPv4, MPLS, BIER
  - IPv6 has been left out… until now…
- How does it work?
  - 5 minute window to mark (5-minute mark ‘0’, 5-minute mark ‘1’, 5-minute mark ‘0’, etc)
  - ACLs are used to count packets egress/ingress of each interface
  - Counters are harvested each 5 minutes and sent to management system
  - Analytics are done on the management system

Alternate marking
by Telekom Italia

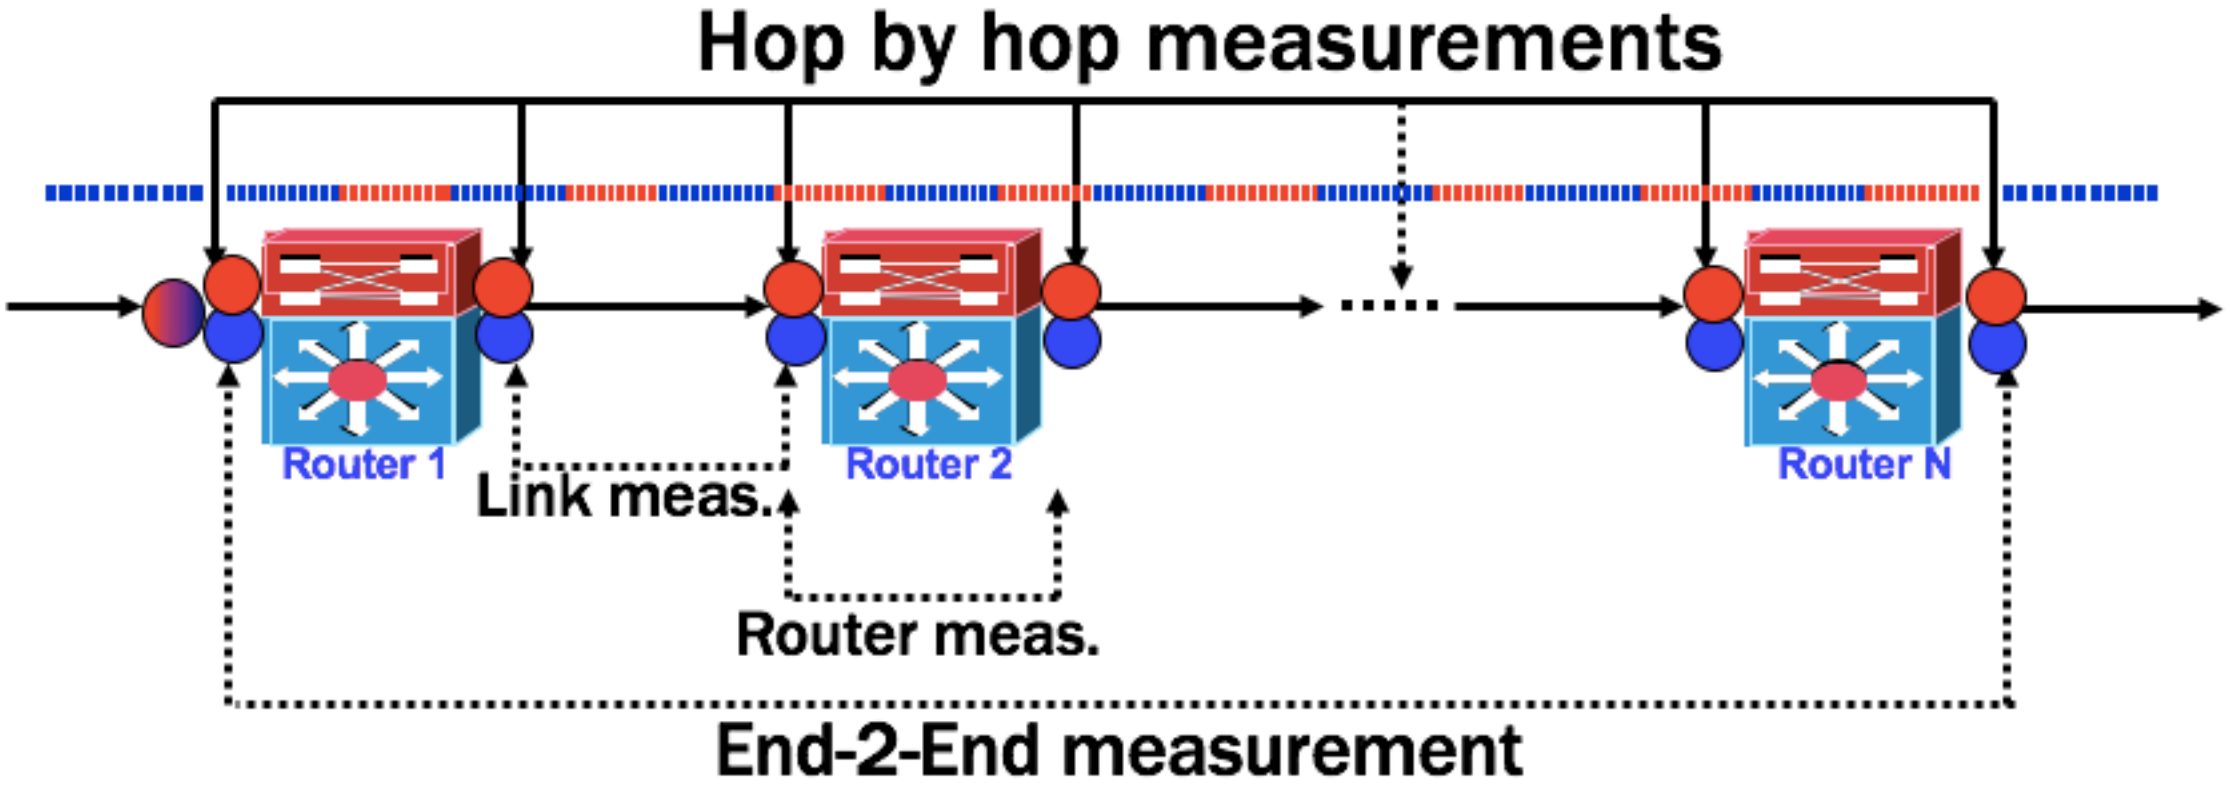

Ref. Telekom Italia (Giuseppe Fioccola)

Alternate marking
IPv6 Flow Labels for Telemetry (by Telekom Italia)

- IPv6 based Alternate marking
- Note
  - only in managed and controlled domain
  - Original flow-label reconstructed when leaving SP controlled domain
- Technology can be used together with SRv6
- SRv6 policy and other IPv6 tunnel encapsulations
  - IPv6 SRv6 Encap
    - Outer SRv6 header uses 1 or 2 bits from 20 bit flow-label field
    - Outer SRv6 header removed when exiting the SP domain
    - Original Flow-label restored at egress
  - IPv6 SRv6 EH insert
    - SRv6 EH insert proposes insertion of IPv6 Extension Headers (however in RFC2460bis it was ruled against insertion of EHs, making SRv6 header insert future uncertain)
    - Original flow-label can be carried as Opaque data TLV in SRv6 headers, and by egress device used for original header construction

Segment Routing for IPv6 - two proposals
(IPv6 Extension Headers or Unified Routing instructions)

- **draft-ietf-6man-segment-routing-header**
  - IPv6 EH to drive Segment Routing
  - Clean IPv6 purist approach
  - However
    - HW unfriendly (not easy to push multiple 128 bit addresses with HW support)
      - (just look at MPLS MSD of your favorite vendor or merchant silicon)
    - Lots of IPv6 header overhead
    - Iffy IPv6 Extension Header insertion proposal to reduce overhead
      - base IPv6 Protocol needs to be modified to even support this
      - IETF 6man (30 March 2017): **proposed RFC2460bis EXPLICITELY forbid intermediate nodes to insert EHs**

- **draft-xu-mpls-unified-source-routing-instruction**
  - **New and presented at IETF98 – well received by SPRING Working Group**
  - Native IPv6 dataplane using UDP
  - No changes needed to IPv6 at all
  - Size matters: Per segment only 20 bits are needed (SRv6 EH needs 128 bits)
  - Virtualization and entropy friendly
  - already supported on any device supporting MPLS control plane and IPv6 data-plane
    - MPLS encapsulated in a IPv6 UDP header, and each segment is identified by 20 bit's
    - Works seamless across NATIVE IPv4, IPv6 and MPLS

What else IPv6 is happening at IETF? Not a full view… just short overview

- draft-ietf-opsec-v6 is in WGLC right now
  - Check with Eric Vyncke because he is editor of this work
- 6MAN Last Call Summary
  - https://www.ietf.org/proceedings/98/slides/slides-98-6man-ietf-last-call-summary-00.pdf
  - *Updates: rfc2460bis, rfc4291bis, rfc1981bis*
  - *Elevation (but no change): rfc4443, rfc3595*
- *Requirements docs*
  - IPv6 Node Requirements (draft-clw-rfc6434-bis)
  - Requirements for IPv6 Router (draft-ali-ipv6rtr-reqs)
  - Basic Requirements for IPv6 Customer Edge Routers (draft-palet-v6ops-rfc7084-bis)

What else IPv6 is happening at IETF? Not a full view… just short overview

- v6ops
  - The eternal discussion on “making RDNSS a MUST”
    - Two mechanisms for same purpose (-> send config info to host)
    - Only one tends to be used, so developing 2 is kind of silly
    - Very religious discussion
    - No outcome yet… virtual hum ongoing…
  - WGLC for draft-ietf-v6ops-v4v6-xlat-prefix
    - Start: 11 April – end: 25 April
    - Updates RFC6890 (= Special-Purpose IP Address Registries)
    - IPv6 prefix 64:ff9b:1::/48 for local use within domains that enable IPv4/IPv6 translation mechanisms
  - 7084bis (Basic Requirements for IPv6 Customer Edge Routers)
    - HNCP is proposed as a MUST by some, but strongly debated (google is pushing for inclusion)
    - LISP is proposed as MUST, but again strongly debated
    - Prefix delegation more specified (/48-/64 for CPE routers)
    - Proposes multiple transition technologies, and right now it is discussed if that is really the “best”?