Stateless Dual Stack Lite*

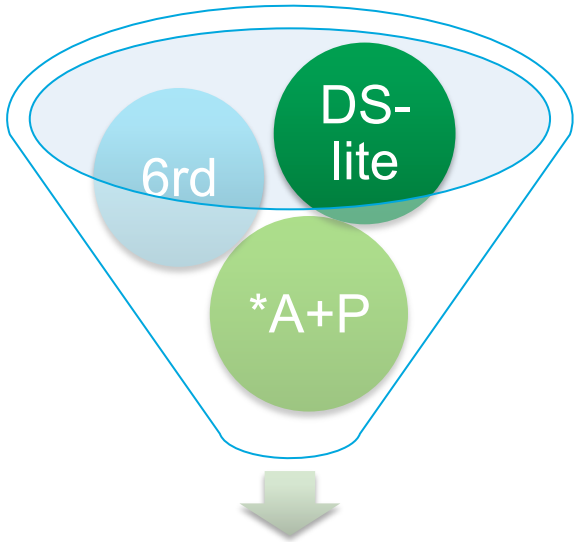

Mark Townsley, Cisco Fellow

Slovenia Go6 Summit, Nov 2011

Dual Stack Lite (DS-Lite)

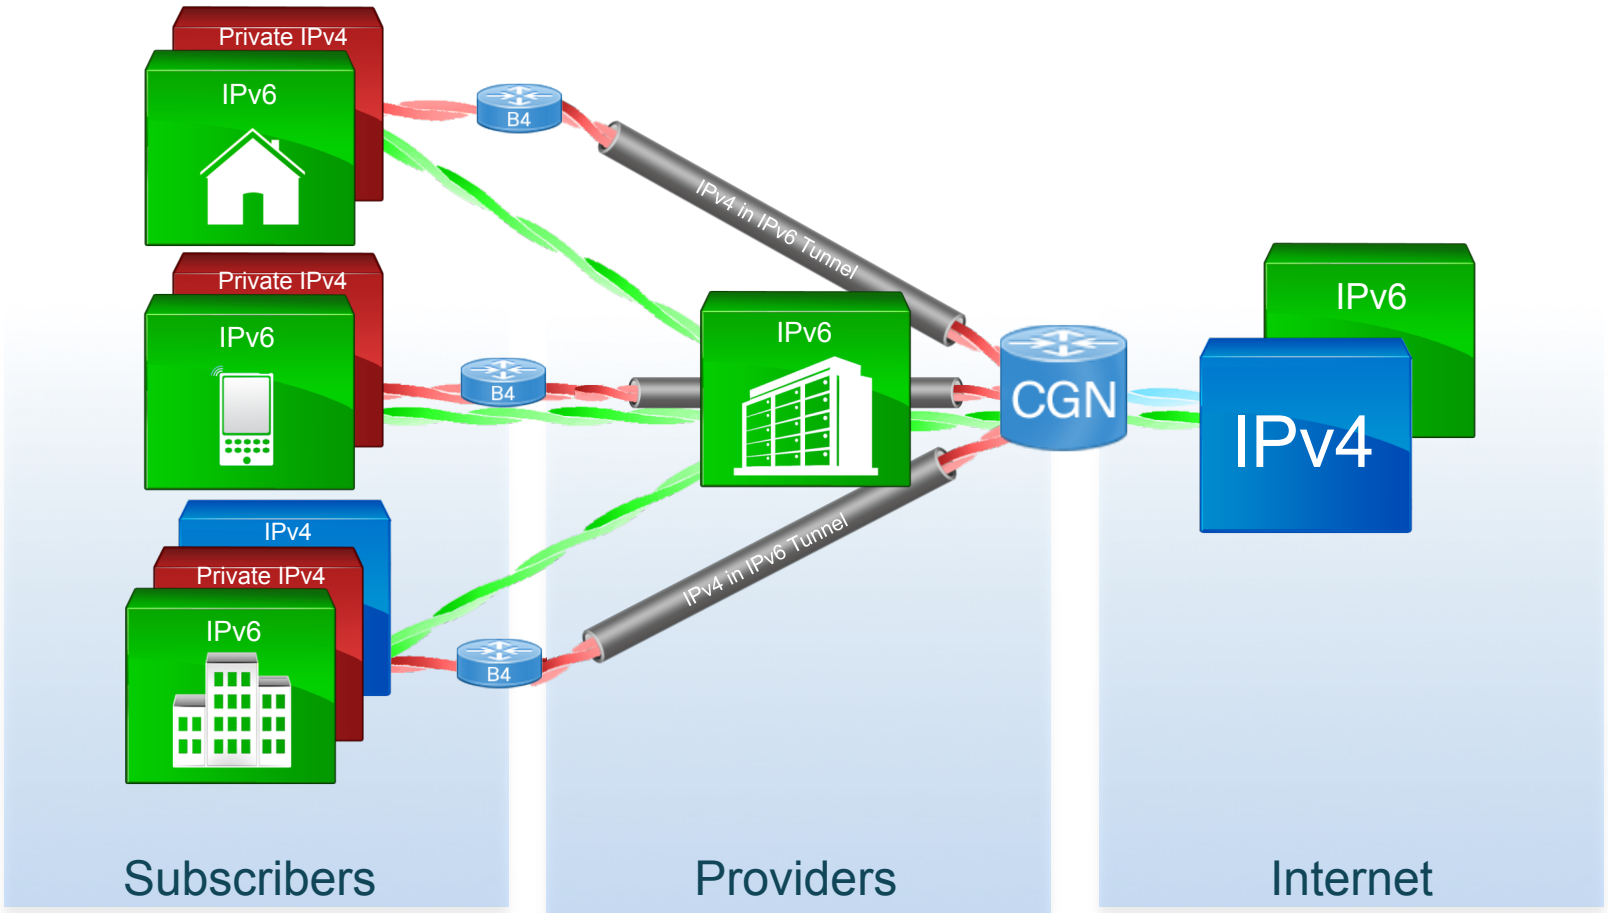

IPv6 Residual Deployment (4rd)

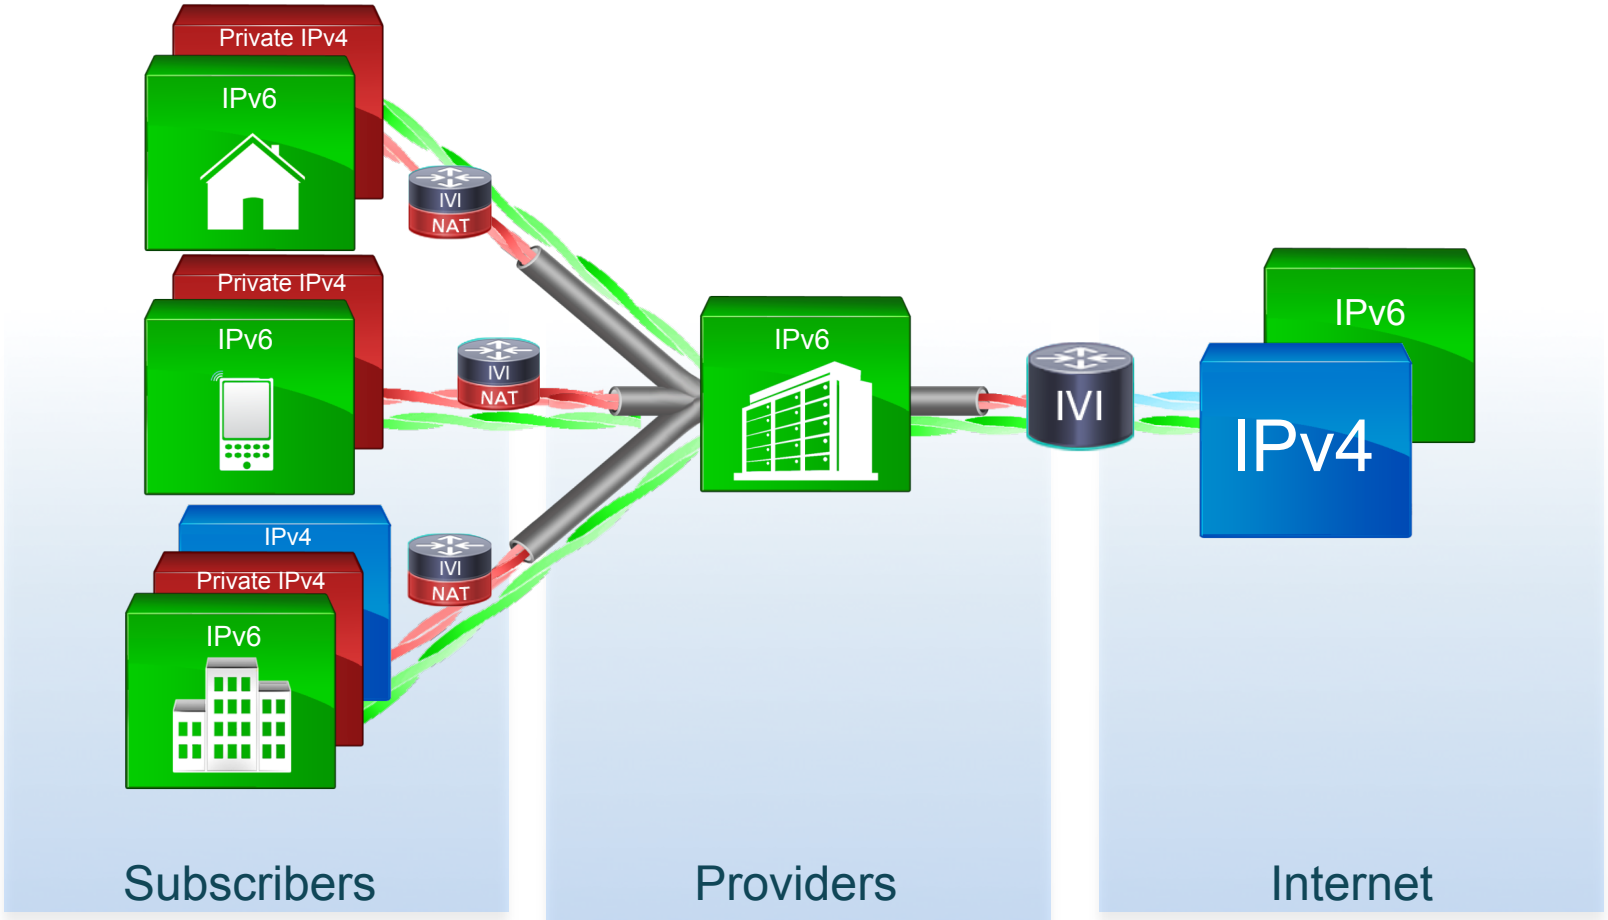

Protocol Details:
Three essential parts of 4rd:

① Global IPv4 Address and Port Set derived from IPv6 configuration

- 4rd programs an IPv4 NAPT that stays at the Edge of the network

② Algorithmic mapping and encap of IPv4 over IPv6

- IPv6 destination determined from each packet's IPv4 destination address and port via a stateless algorithm
- No per-subscriber tunnel state or provisioning

③ Stateless Border Relays

- No new centralized NAT State
- Border Relays reachable via anycast

① IPv6 → IPv4 + Port Mapping

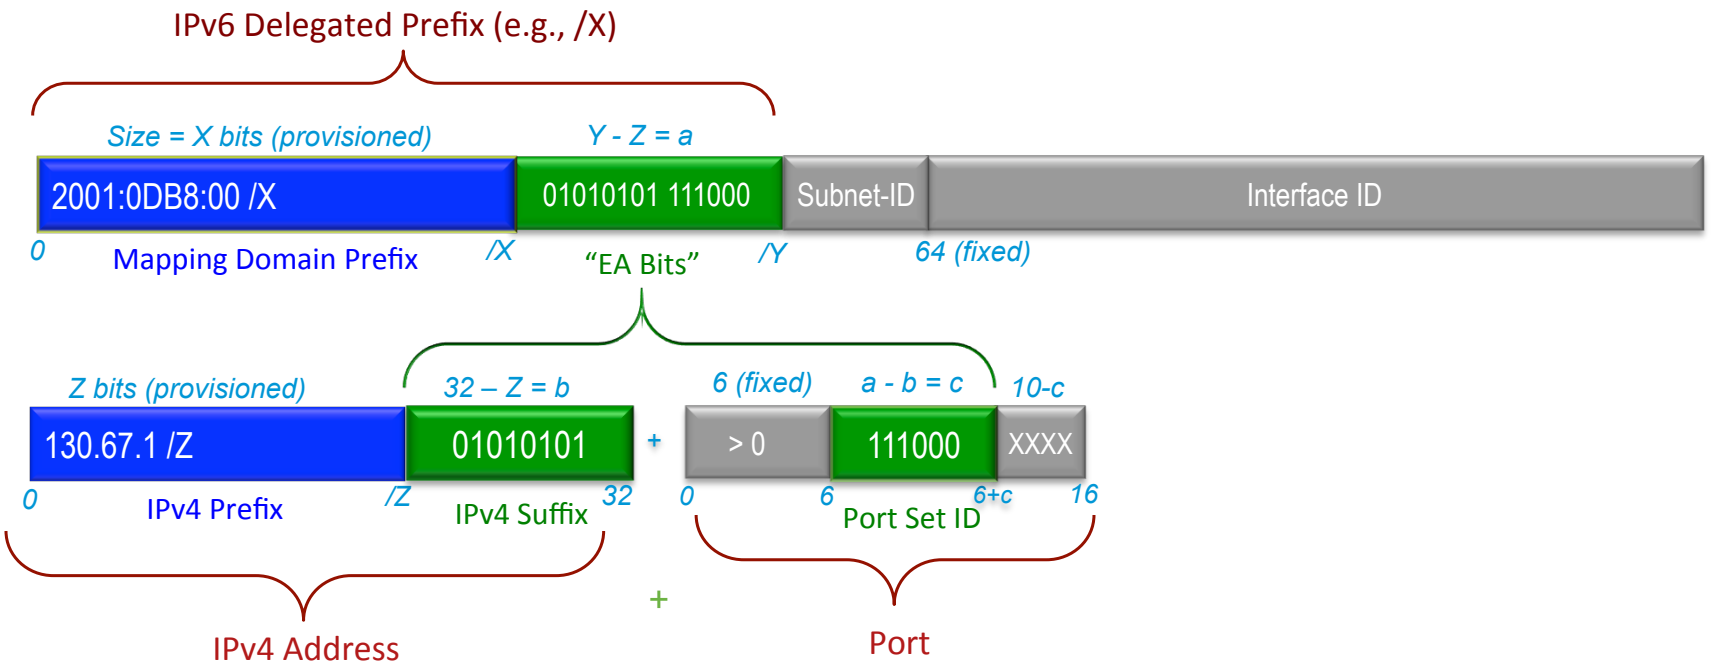

① IPv6 → IPv4 + Port Mapping

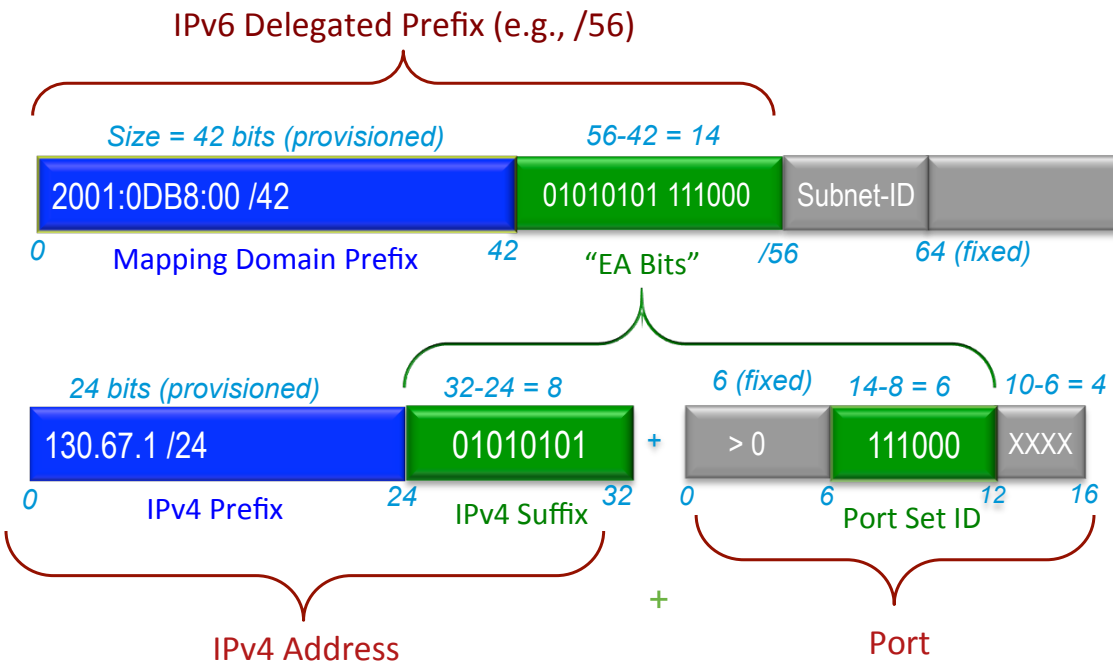

For this Example...

2^6=64 port sets per IPv4 Address

Ports 0-1023 skipped, each CPE gets $2^{16}/2^6 - 2^4$ = 1008 ports

One IPv4 /24 serves $2^{(6+8)} \approx 16{,}384$ (vs.≈256) subscribers

② Packet Flow and Encapsulation

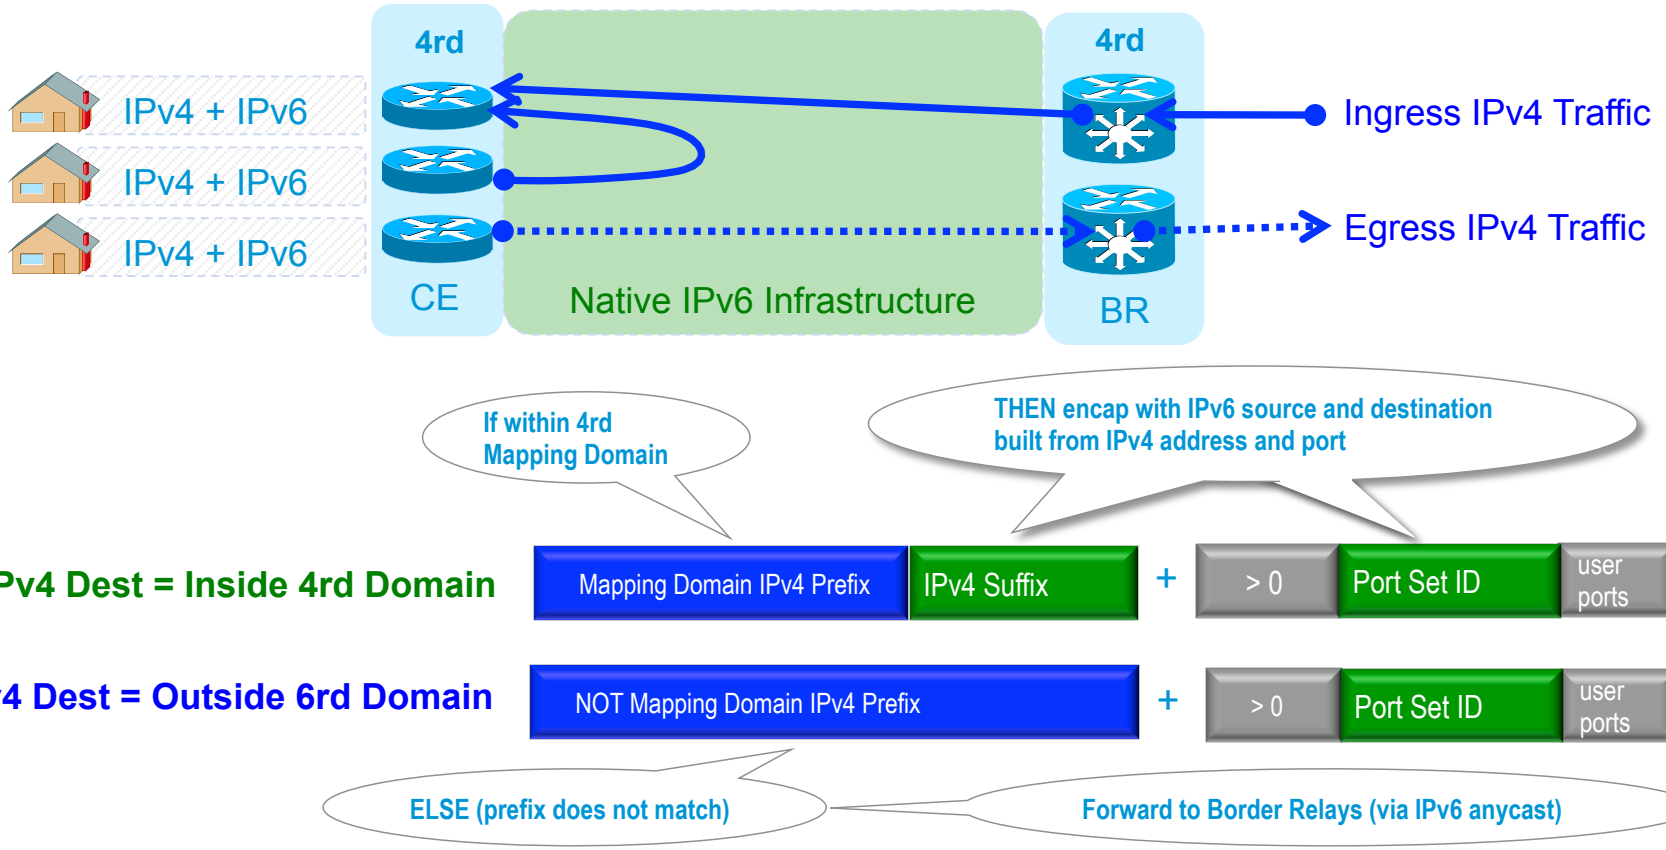

③Stateless Border Relays

- Handle traffic to/from the 4rd domain
- Simple algorithmic mapping and encap/decap function
- No per-packet state, so may be reached via anycast
- BR placement is solely a function of IPv4 traffic, not the number of subscribers

Residential Gateway (4rd CE) Implementation

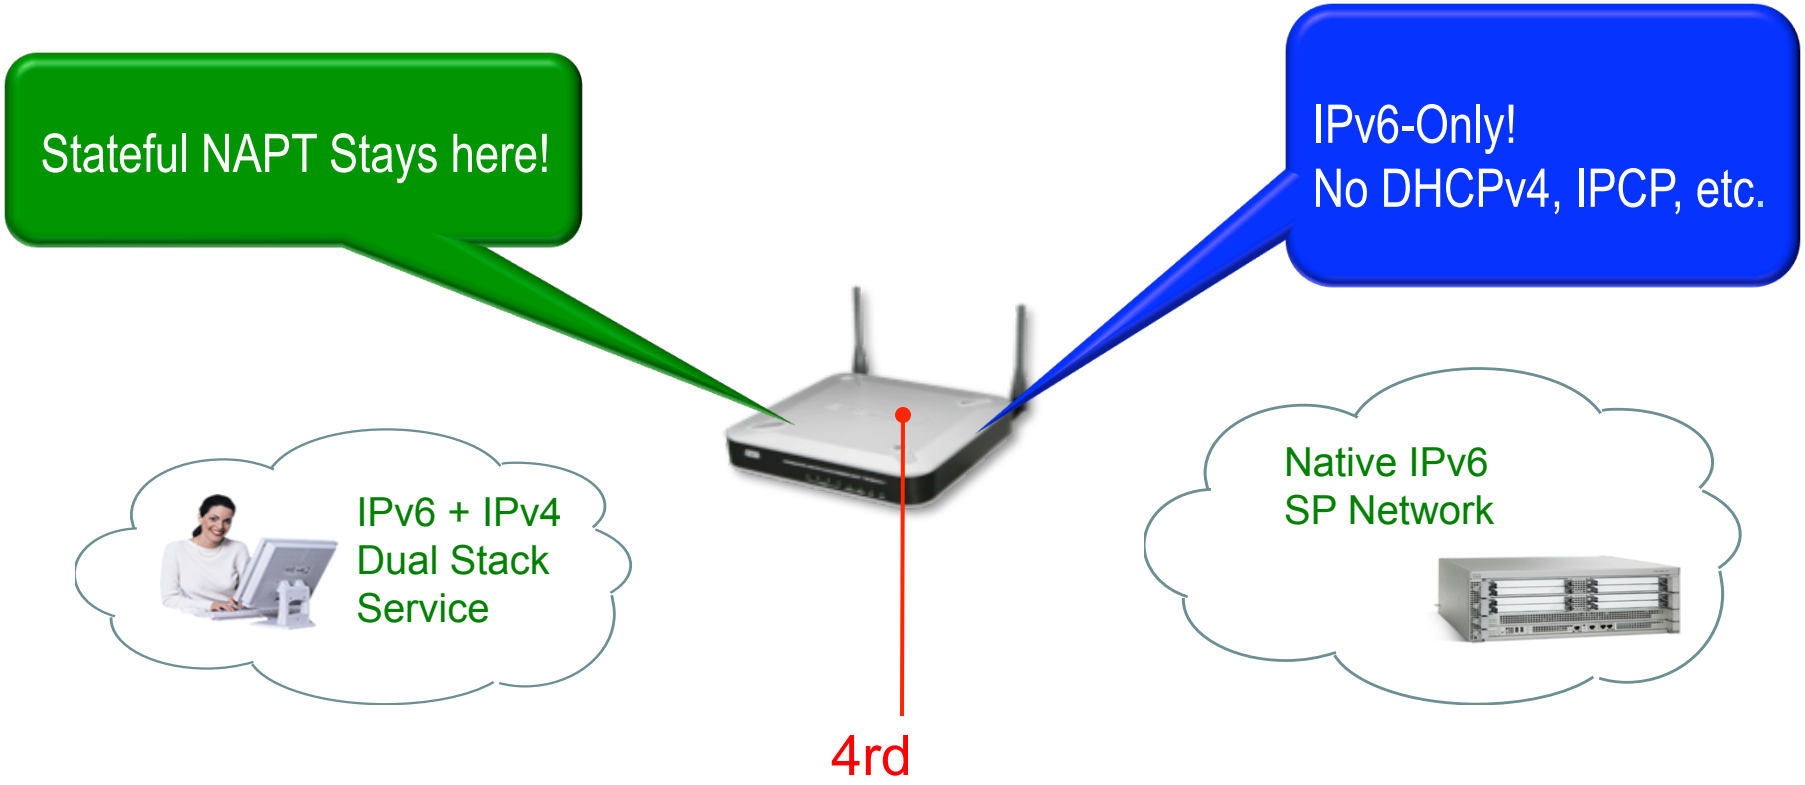

Residential Gateway Configuration

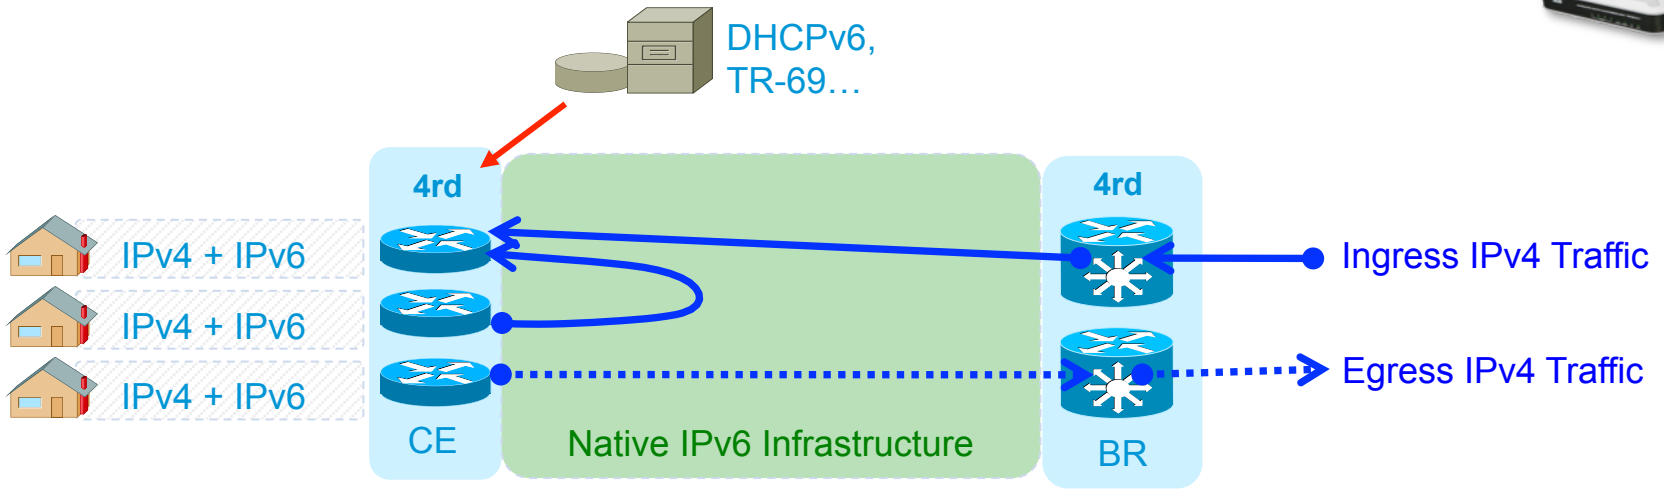

- Home is configured for dual-stack, 4rd config elements identical for all CEs in a Domain
  1) ISP IPv6 Mapping Domain Prefix and Length
  2) Global IPv4 Prefix and Length
  3) 4rd Border Relay IPv6 address or prefix (anycast)

Border Relay Implementation

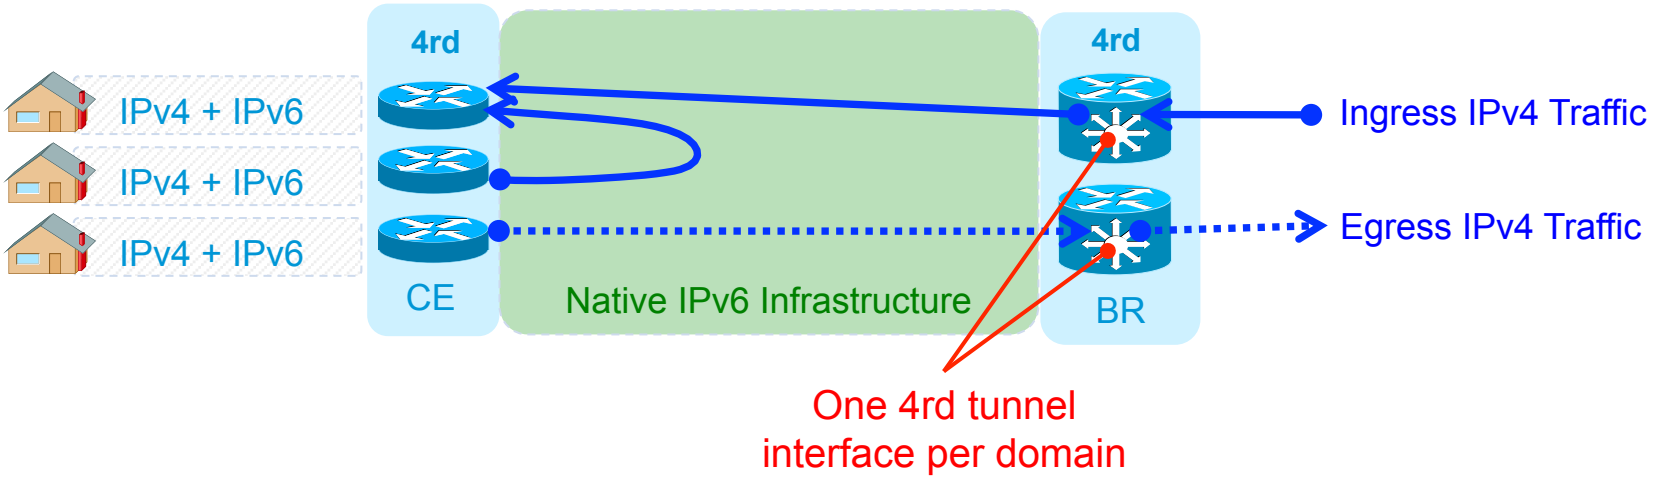

- Single multipoint tunnel interface in Border Relay
- No per-user state or CGN, serves all users in 4rd Domain

Stateless DS-Lite Forwarding Architecture

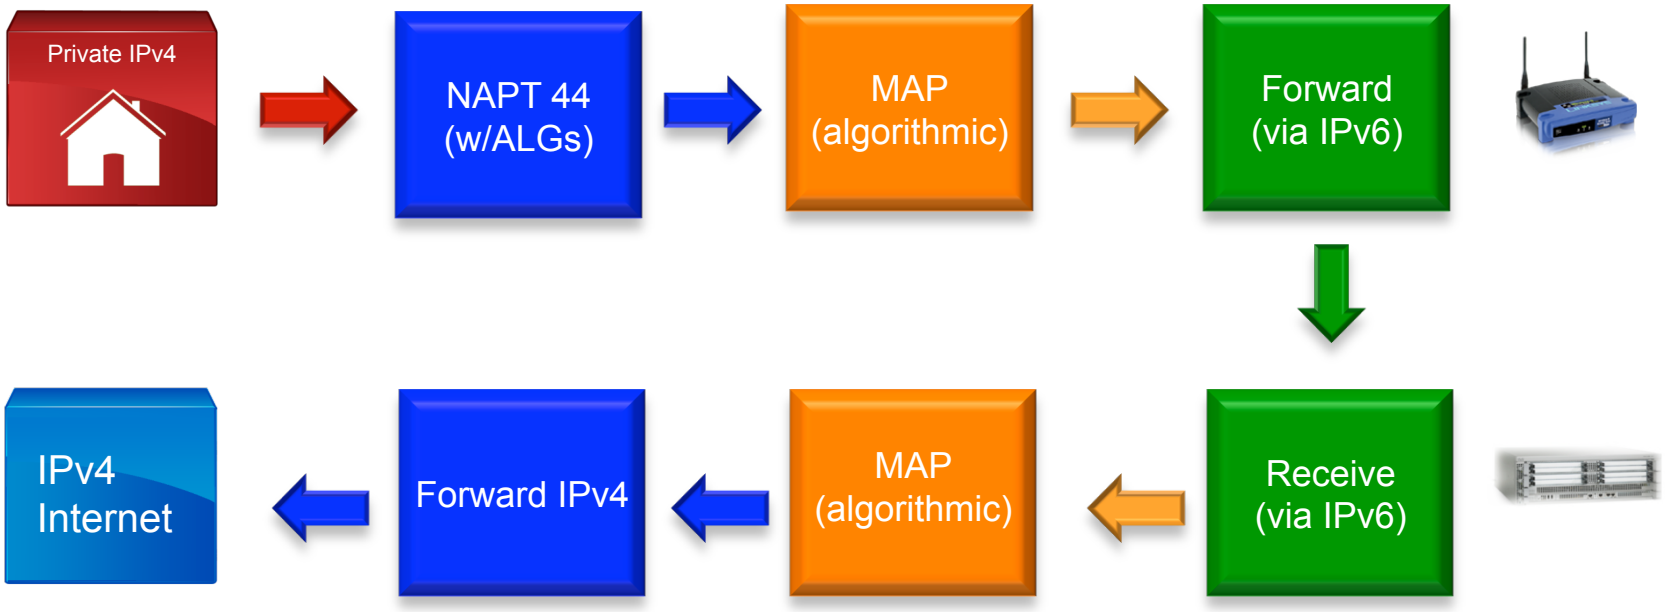

Encapsulation Forwarding Architecture

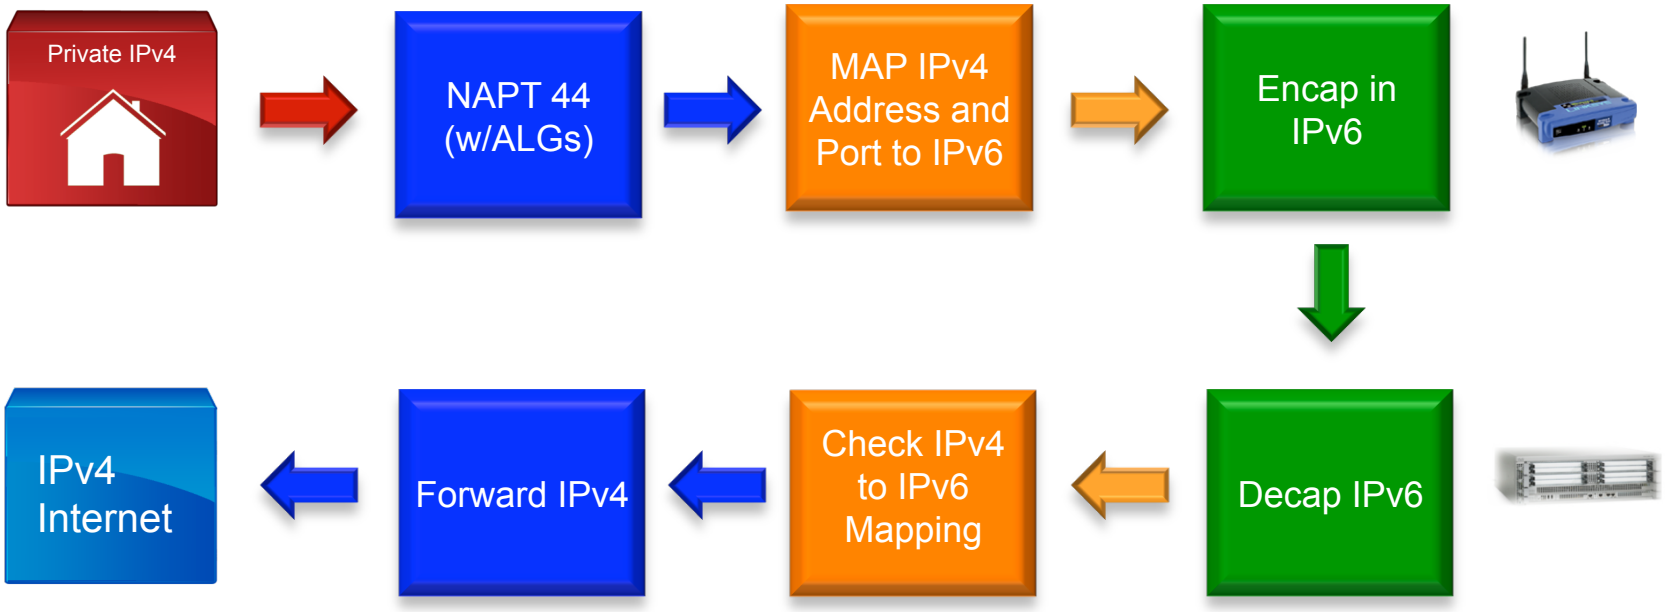

Translation Forwarding Architecture

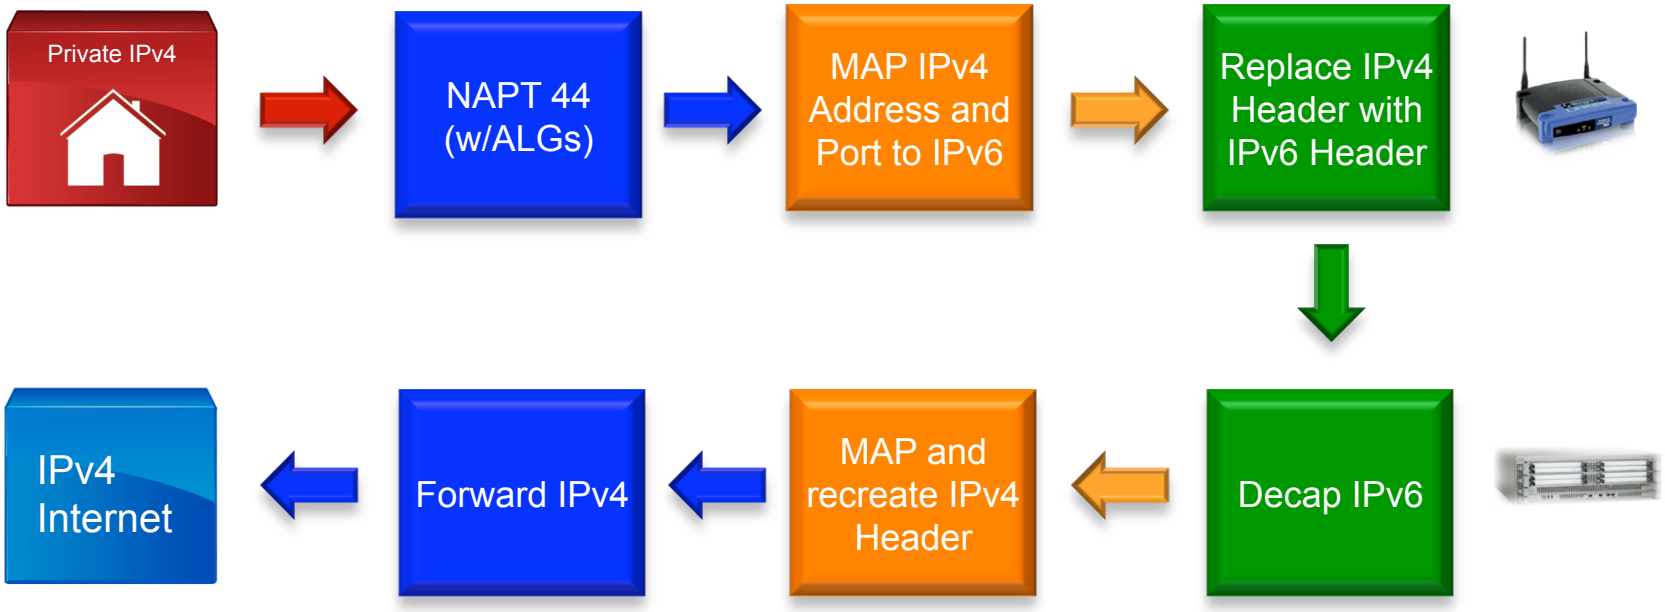

Encapsulation or Translation – Boils down to 20 bytes

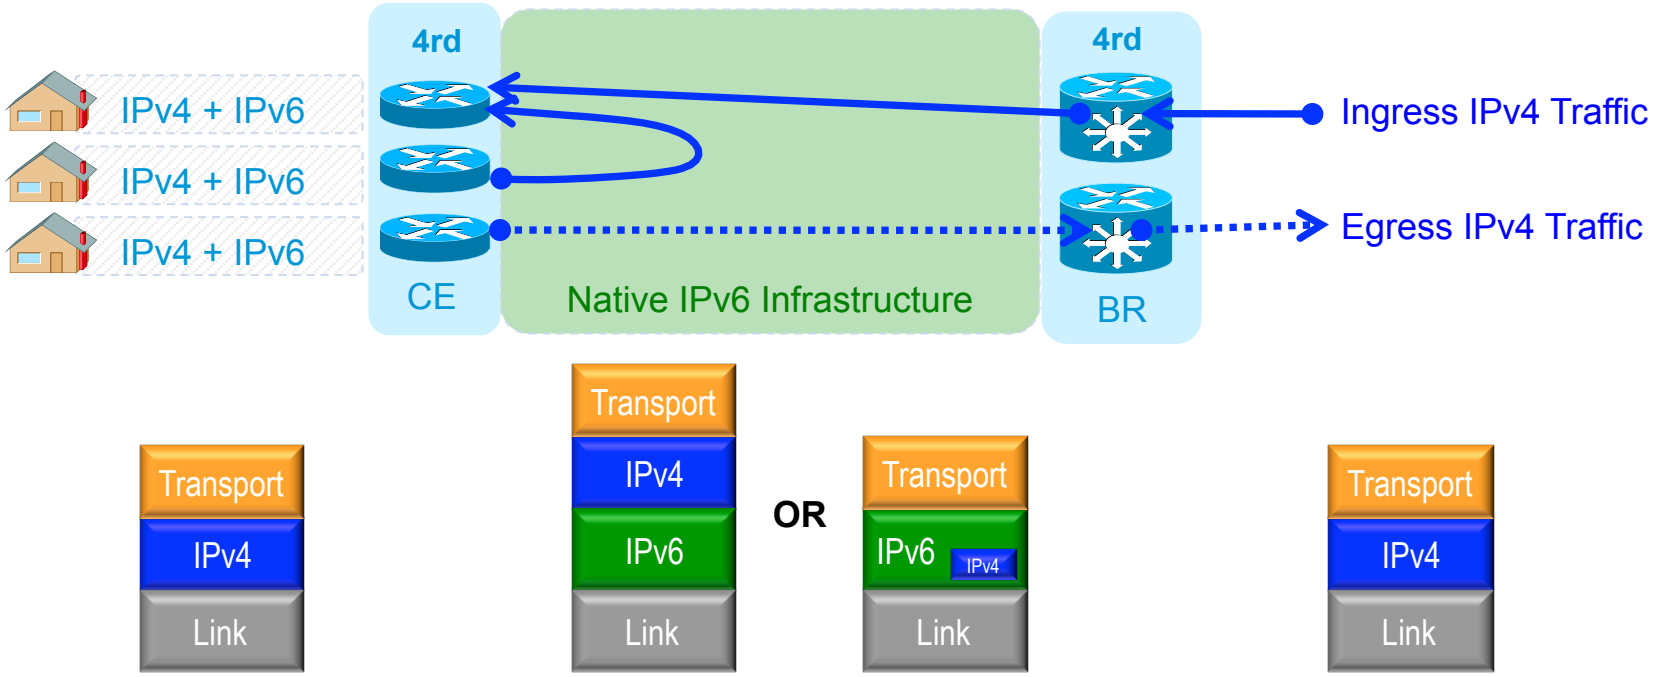

Encapsulation or Translation – Boils down to 20 bytes

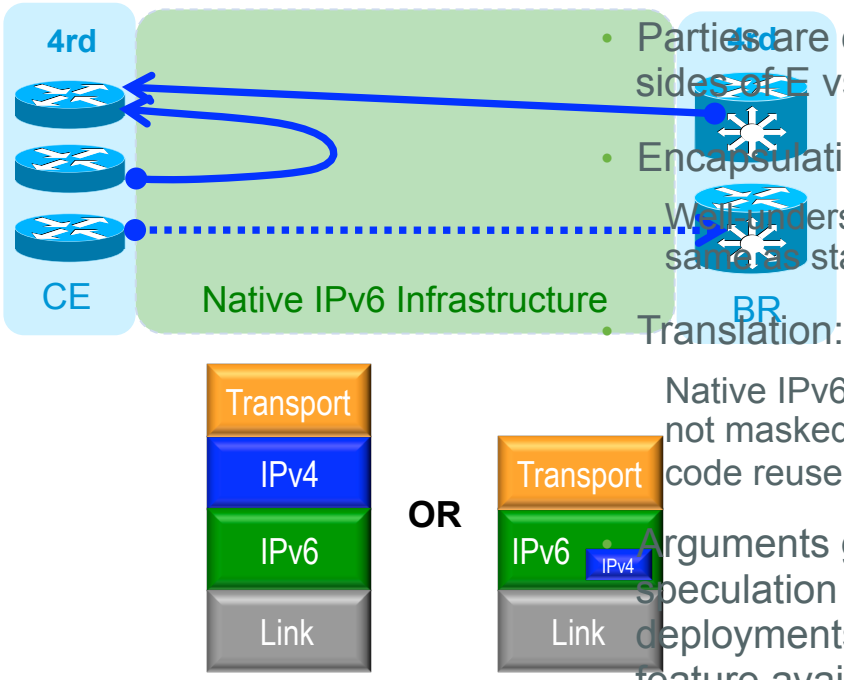

- Parties are entrenched on both sides of E vs T (vs U!) debate
- Encapsulation:
  - Well understood, simple, transparent, same as stateful dual-stack lite
- Translation:
  - Native IPv6 ACLs and DPI functionality not masked by IPv4 header. NAT64 code reuse. Looks like “Real IPv6.”
  - Arguments gravitate towards speculation about what future IPv6 deployments will require and what feature availability will be

So Where Are We Headed With All This?

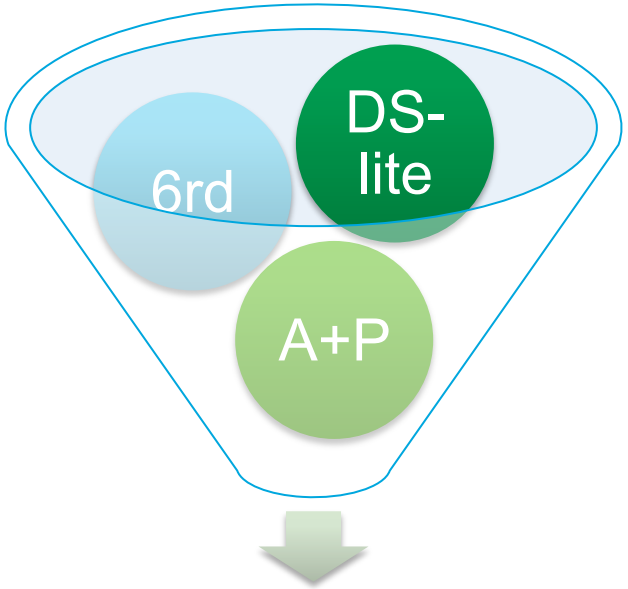

History

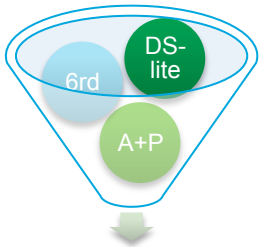

- IPv4 over IPv6 mechanisms date as far back as DSTM, proposed in 2000
- Softwires created in 2006
- 6rd and DS-Lite emerge in ~2008 (RFC5969 & RFC 6333)
- IETF Shara BOF focusing on “A+P” in 2009
  - Resulted in RFC6269 on Shared IP address issues (applies to CGN, DS-lite, any A+P solution)
  - No consensus on the need for A+P based solutions
  - Experimental RFC 6346 published in 2011
- A myriad of solutions appear with reckless abandon, searching for a home in the IETF
  - 4rd, 4over6, dIVI, ...
- “Solution Motivation” document published in 2011
  - Authors include France Telecom, Softbank, Comcast, DT, Portugal Telecom, China Mobile…

More Recently…

- Consensus to agree upon a single mapping algorithm, coupled with an encapsulation and translation option
  - Softwires Interim Meeting in China (Sept 2011)
- Mapping Design Team formed
  - Latest result (Oct 31): draft-mdt-softwire-mapping-address-and-port-01
- Since China, Remi Despres has brought forth a compromise between Encapsulation and Translation “4rd-U” in hopes of merging the two. No consensus yet.
- A great deal of activity around this topic, demanding strong leadership in the IETF

Stateless A+P Technology – Bottom Line

- Seems to work, and comes with very compelling scalability properties
- Could become a very strong alternative to stateful DS-Lite
- Standardization effort has been troublesome until recent months
- Cautiously optimistic the IETF will produce a standard (or two) in 1H2012, and choices among interoperable products across the industry thereafter.

Thank you.