IPv6 via IPv4
Service Provider Networks – “6rd”

Technical & Operational Overview

Mark Townsley

Distinguished Engineer, Cisco Systems

October 2009

IPv4 Address Exhaustion

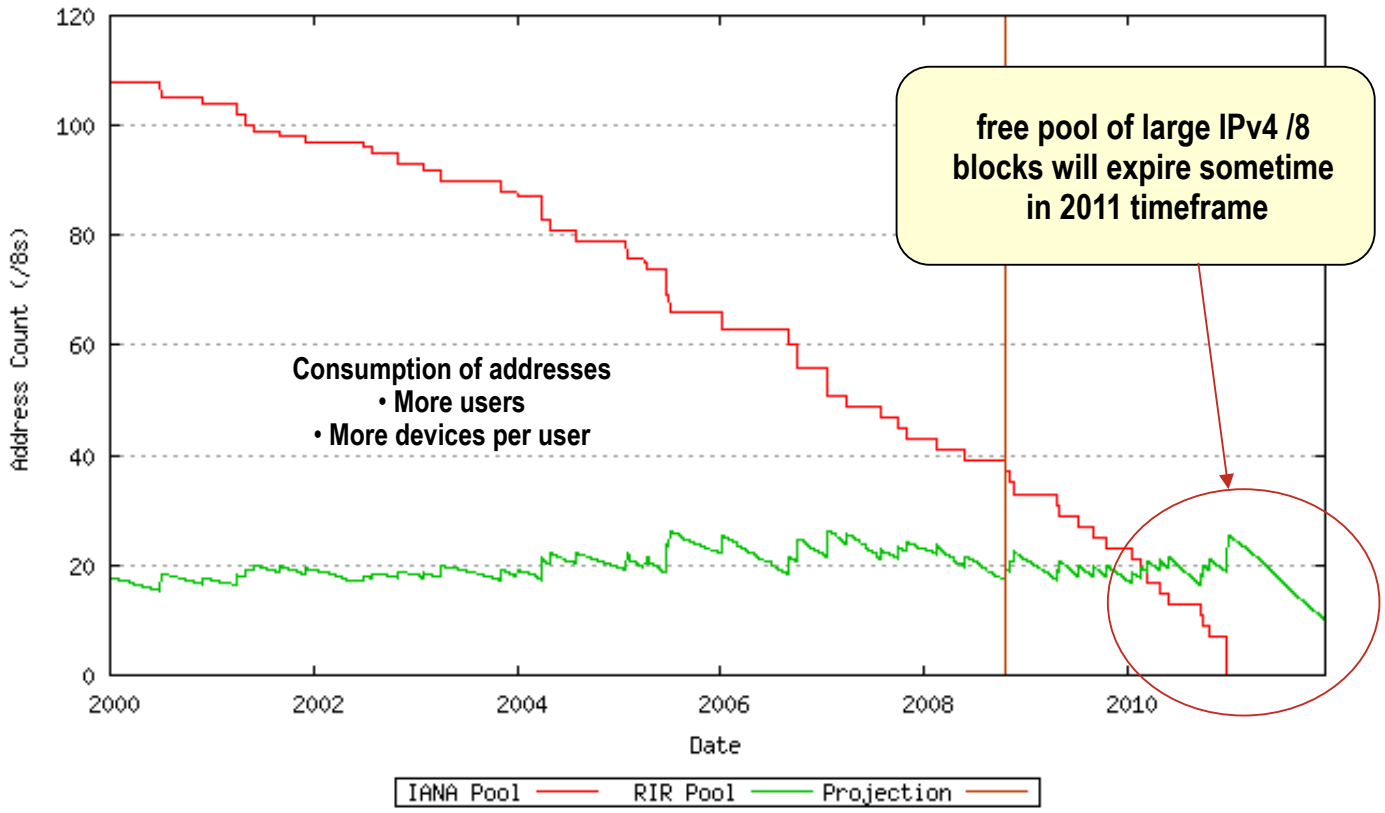

- **see http://www.potaroo.net/tools/ipv4/index.html for more details**

Problem: Gap in IPv6 Availabilty

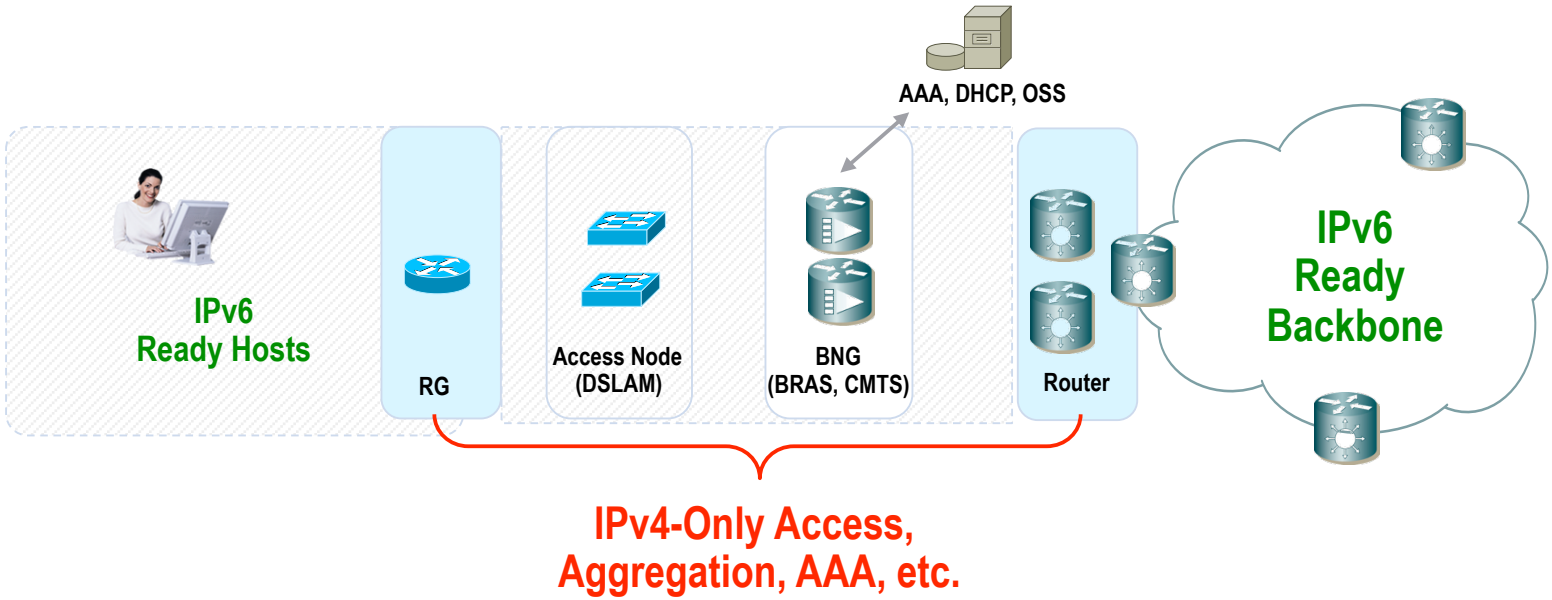

IPv6 via IPv4 SP Networks using “6rd”

- Incremental method for deploying IPv6
- Not an IPv6 “trial” service. IPv6 to subscriber is production-quality, native IPv6 + IPv4 dual-stack
- Reuses IPv4 in the SP - No v6 support needed in Access and Aggregation infrastructure, no DHCPv6 servers, no Neighbor Discovery, etc.
- Similar to 6PE in that it provides a native dual-stack service to a subscrbier site by leveraging existing infrastructure, operations, etc.

6rd in One Slide

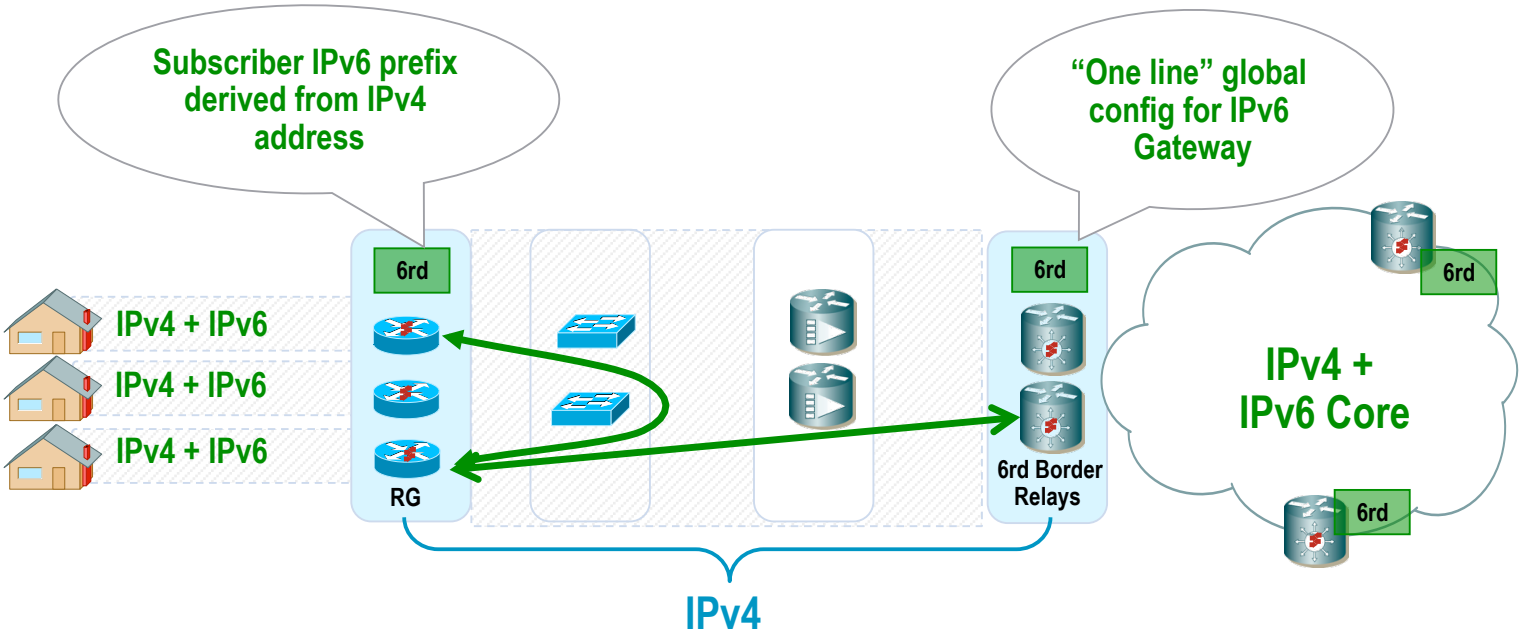

- Native dual-stack IP service to the Subscriber
- Simple, stateless, automatic IPv6-in-IPv4 encap and decap functions
- IPv6 traffic automatically follows IPv4 Routing
- BRs placed at IPv6 edge, addressed via anycast for load-balancing and resiliency
- Defined in `draft-ietf-softwire-ipv6-6rd`

Residential Gateway Implementation

LAN-Side:
Production Native IPv6 Service + Global or Natted IPv4

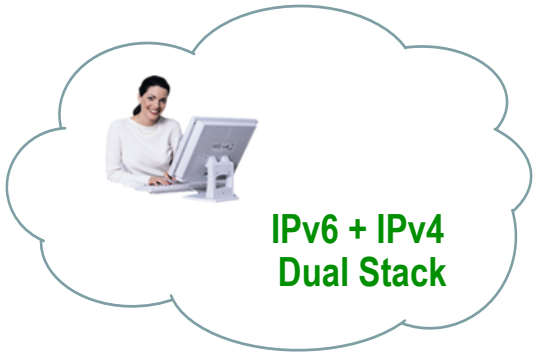

WAN-Side:
IPv6 via IPv4 Global or Natted IPv4

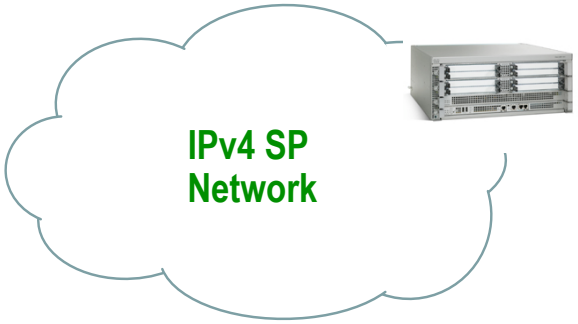

IPv6 Internet Access delivered to home, allowing IPv6 enabled applications and content to remain unaffected by IPv4 Exhaustion

6rd lives here

IPv6 in SP Network evolves at its own pace, with its own balance of costs and incentives

RG 6rd Setup & Provisioning

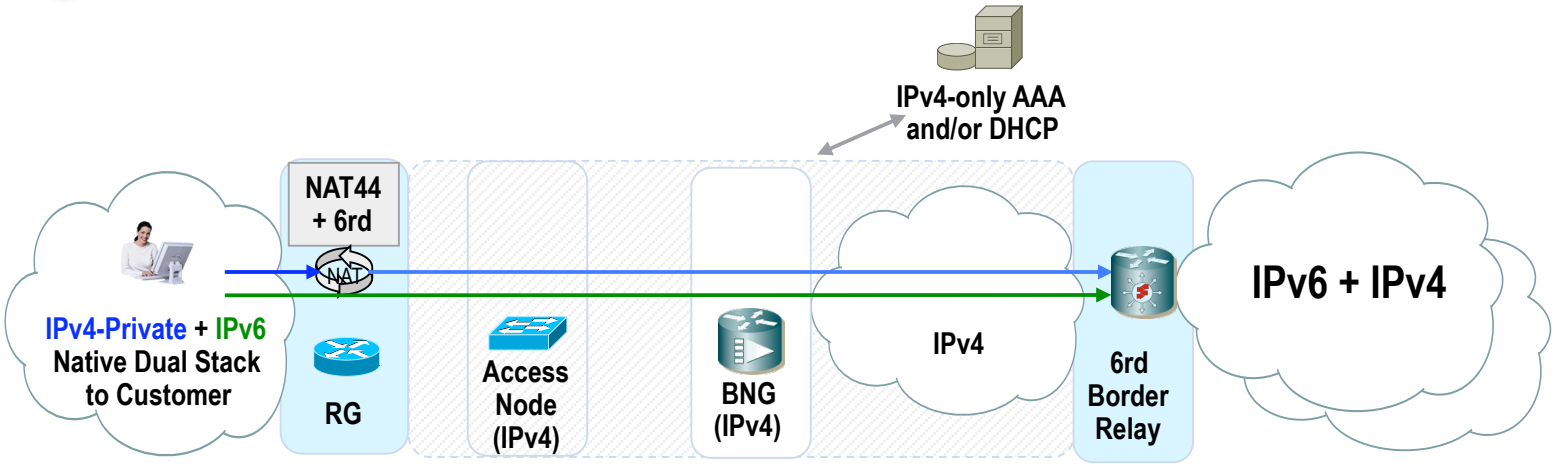

- RG configuration, same for all subscribers (via TR-69, DHCP, etc)
  1. ISP 6rd IPv6 Prefix and length
  2. Common IPv4 bits suffix length
  3. 6rd Relay IPv4 address (likely anycast)
- "Home side" of RG configured exactly as would be for "native" IPv6, e.g., same as for a DHCPv6 delegated prefix

Border Relay Implementation

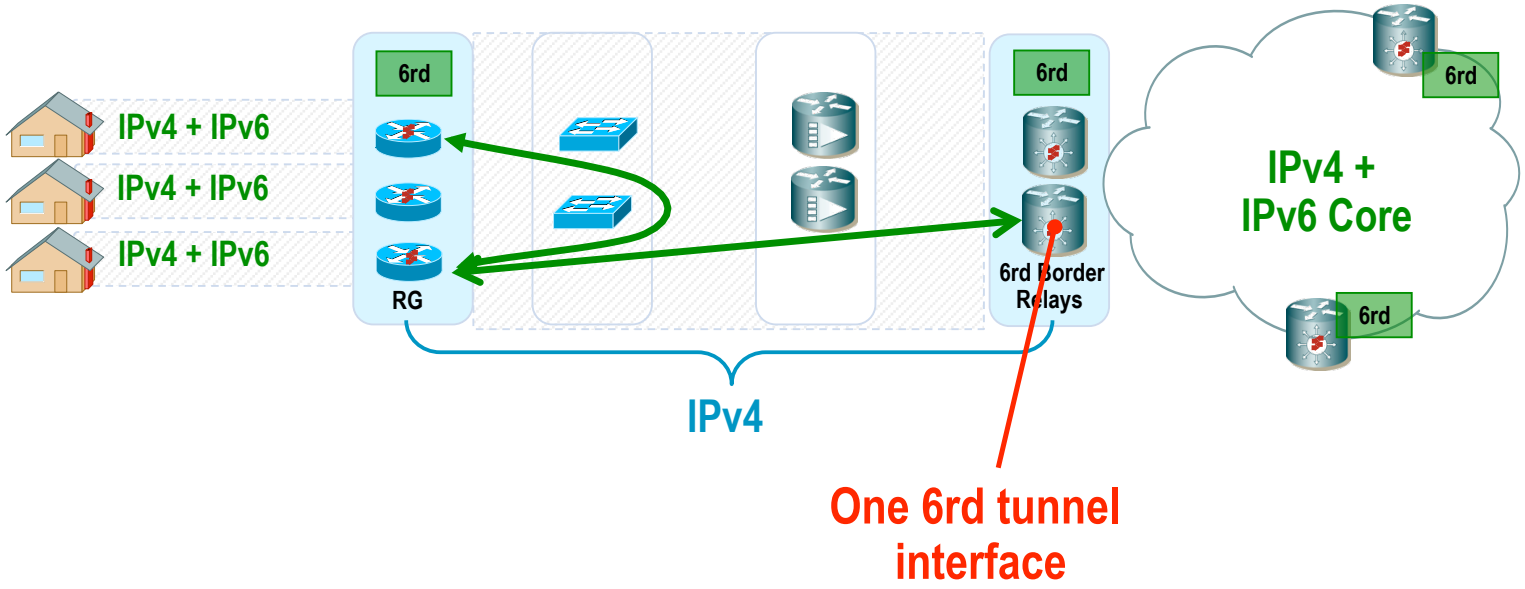

Single multipoint tunnel interface in Border Relay

No per-user state, serves ALL users in 6rd Domain

Border Relay Provisioning

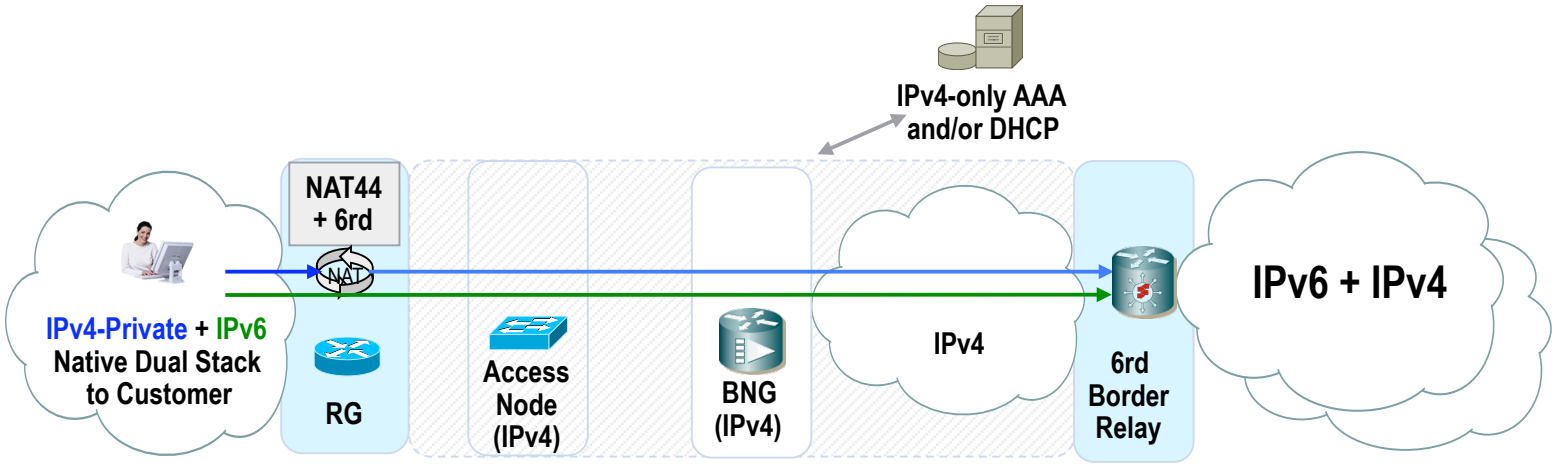

BR must have IPv6 reachability (Native, 6PE, GRE Tunnel, etc).

1. ISP 6rd IPv6 Prefix and length
2. Common IPv4 bits suffix length
3. 6rd Relay IPv4 address (likely anycast)

BR Implementation in CGSE

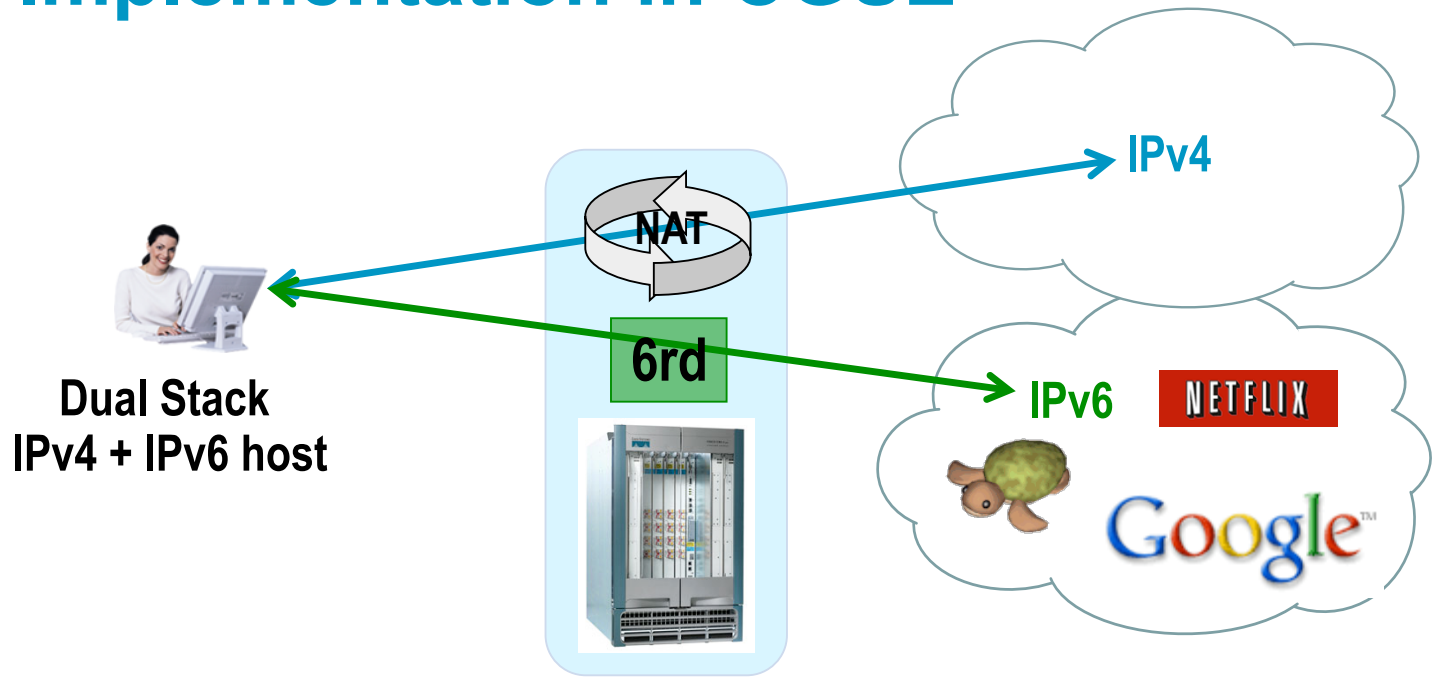

- **NAT444 + 6rd** Border Relay allows IPv6 capable content to flow without taking IPv4 NAT Resources
  - e.g., A MacOS or Vista user surfing to Google maps or Netflix would automatically run via 6rd, not requiring IPv4 NAT state in the CGSE or RG

Gory Details:
Three parts of the “6rd” Mechanism

- IPv6 Prefix Delegation derived from IPv4
  - Global IPv4 or Natted IPv4 in same deployment
- Stateless mapping and Encapsulation of IPv6 over IPv4 (RFC 4213)
  - IPv4 encapsulation automatically determined from each packet’s IPv6 destination
  - No per-subscriber tunnel state or provisioning
- IPv4 Anycast to reach Border Routers

6rd Automatic Prefix Delegation
(From a Global IPv4 Prefix)

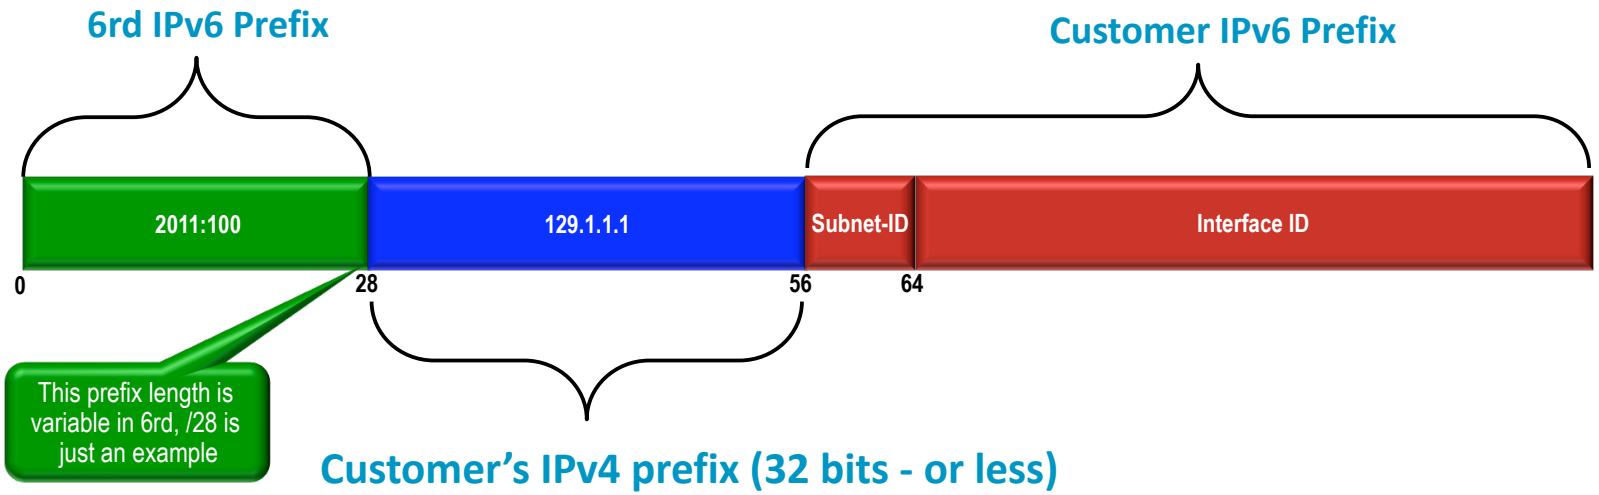

6rd Automatic Prefix Delegation
(From a Private IPv4 Prefix)

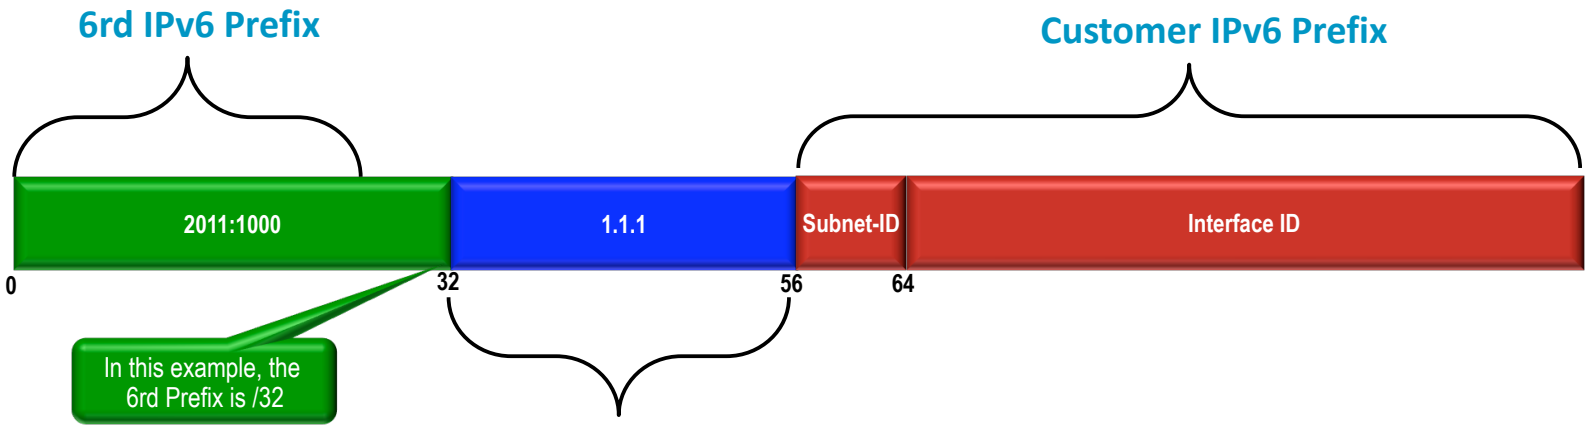

Any number of bits may be masked off, as long as they are common for the entire domain. This is very convienent when deploying with a CGSE , but is equally applicable to aggregated global IPv4 space.

Packet Flow and Encapsulation

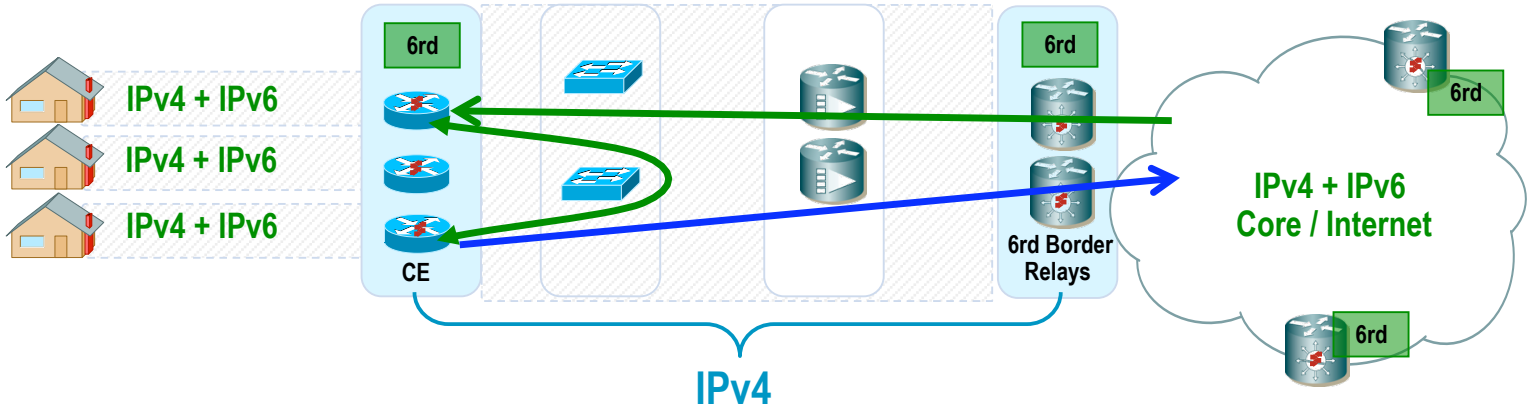

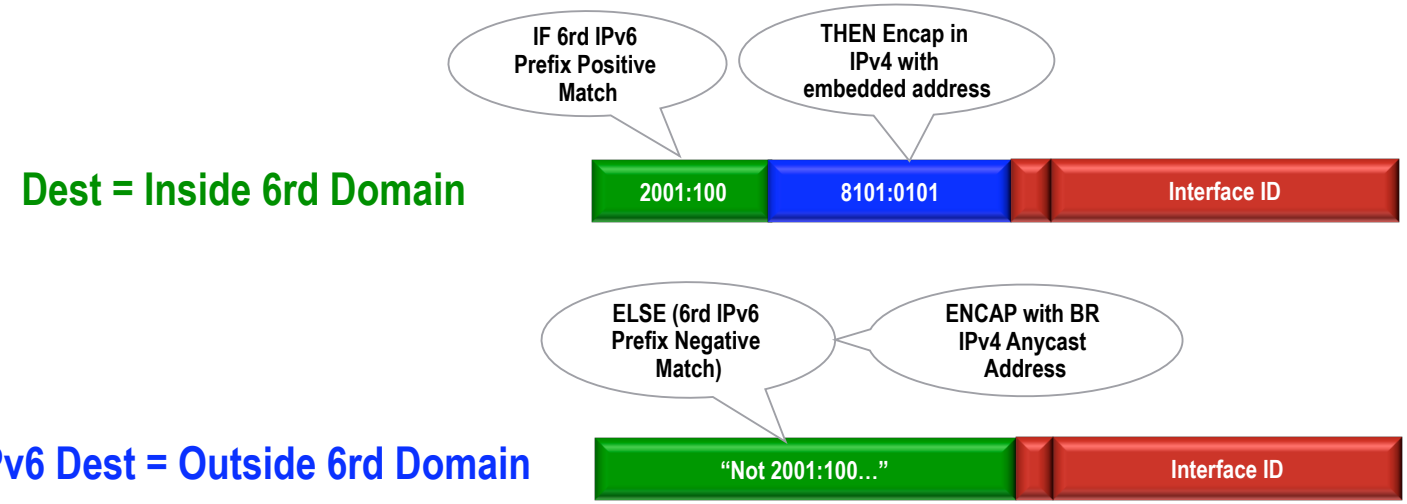

Border Relay via Anycast

- 6rd is stateless, so no need for packets within a flow to traverse the same Relay
- Allows use of IPv4 routing for load-balancing, resiliency and redundancy
- Border Relays are installed only in strategic locations where native IPv6 is available:
  - IPv6 Internet uplinks
  - Edge of internal IPv6-enabled network
  - BR placement is a function only of IPv6 traffic, not the number of sites

Standardization Status

- Was defined in draft-townsley-ipv6-6rd-01.txt
- Now, **draft-ietf-softwire-ipv6-6rd-00.txt** as this now an IETF Standards Track WG document
- Idea has been circulating in the IETF since 2007 when Free Telecom first deployed it based on the invention of Remi Despres (RFC 5569 to be published shortly describing this)
- On track in the Broadband Forum to be part of their IPv6 Technical Recommendations (PD-192 RG Specification) – Strong support from AT&T, Verizon, Swisscom, Fastweb, etc. at meeting

6rd Deployment

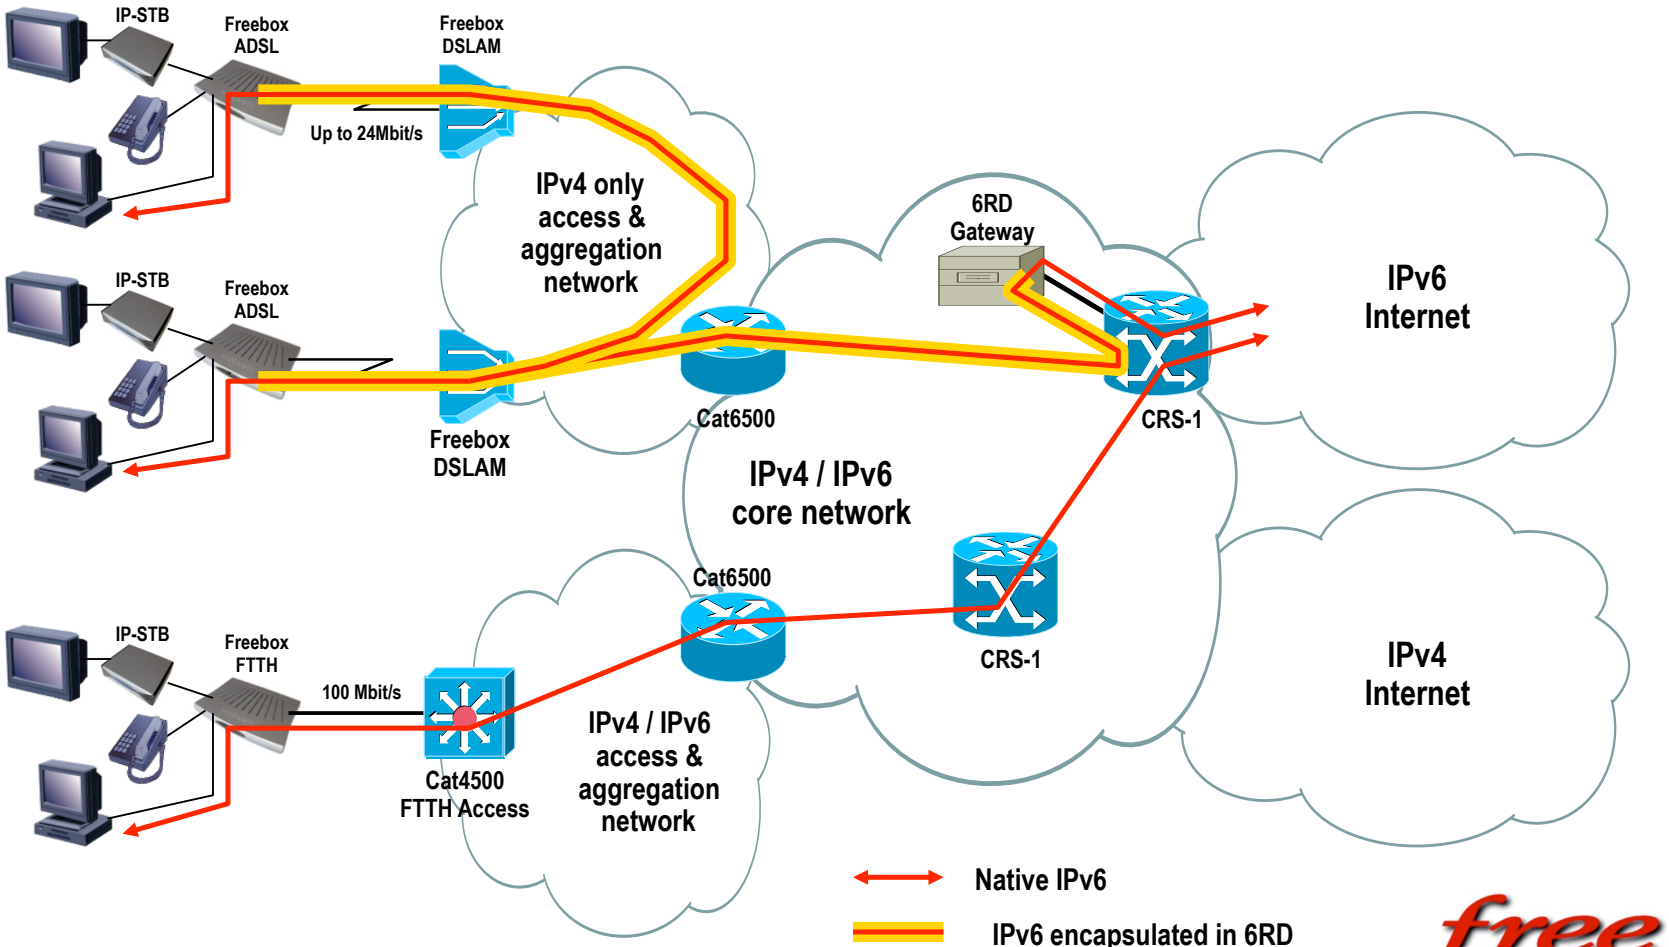

Some Stats 2/3

- 6rd-gw1 Yearly Traffic (1Day AVG) :

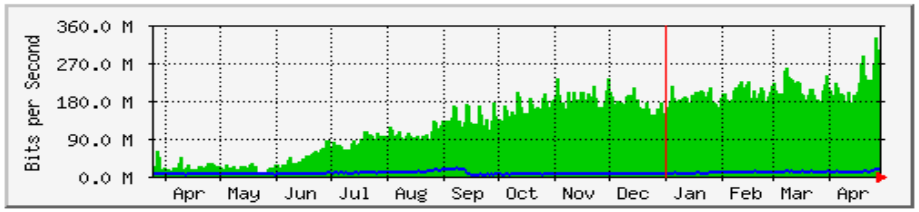

- 6rd-gw2 Yearly Traffic (1Day AVG) :

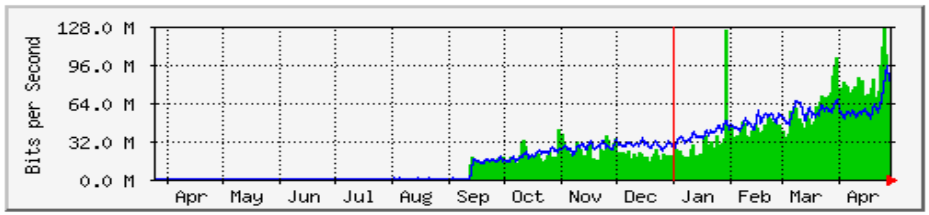

- Customers : 310K
- Global Daily traffic (5min AVG) :

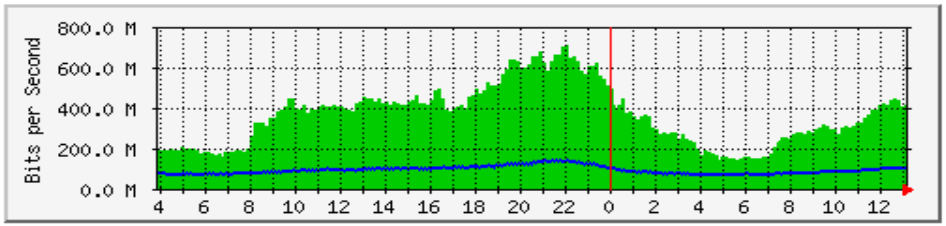

	Max	Average	Current
In	706.3 Mb/s (3.7%)	335.9 Mb/s (1.8%)	414.4 Mb/s (2.2%)
Out	138.9 Mb/s (0.7%)	89.9 Mb/s (0.5%)	101.6 Mb/s (0.5%)

Summary for 6rd

- Like 6PE, delivers Production-Quality IPv6 by only touching edge points around your network
- Capitalizes on what access networks do well, provisioning and transport of IPv4, adapted for carrying IPv6
- Stateless operation, easy to provision, low overhead
- Proven deployment, standardization underway
- Coming soon as part of our “IPv6 Transition Services” in IOS, ASR1K, CGSE, etc…