Transitioning to IPv6 6RD Demo

Hernán Contreras G, hcontrer@cisco.com

CSE, Cisco Systems Chile

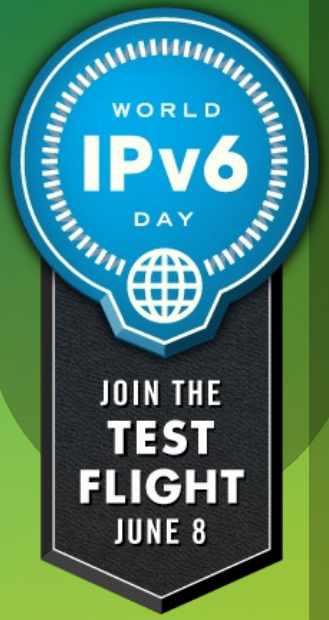

A Rationale for IPv6 Adoption

- IPv4 Run-Out is here now
- Native IPv6 Internet is years and years away
- Entering a period of IPv4/IPv6 Coexistence
- Legacy (IPv4) and new (IPv6) apps and services can only function over an IPv4/IPv6 Coexistence Infrastructure
- Thus we need tools, methods, products and solutions that
  - –Help address IPv4 run-out NOW
  - –Offer incremental means to build out IPv4/IPv6 coexistence infrastructure
- Not one size fits all

Enabling an Orderly, Incremental Transition

Boundless service opportunities with Smart Grid, Connected Cities, Mobile Video, Cloud Computing

Prosper

Prepare

Preserve

= IPv4 = Private IP = IPv6

Smart Grid Opportunity:
110 million households
in US alone

"346": A 3 Tier Transition Framework for Moving from IPv4 to IPv6

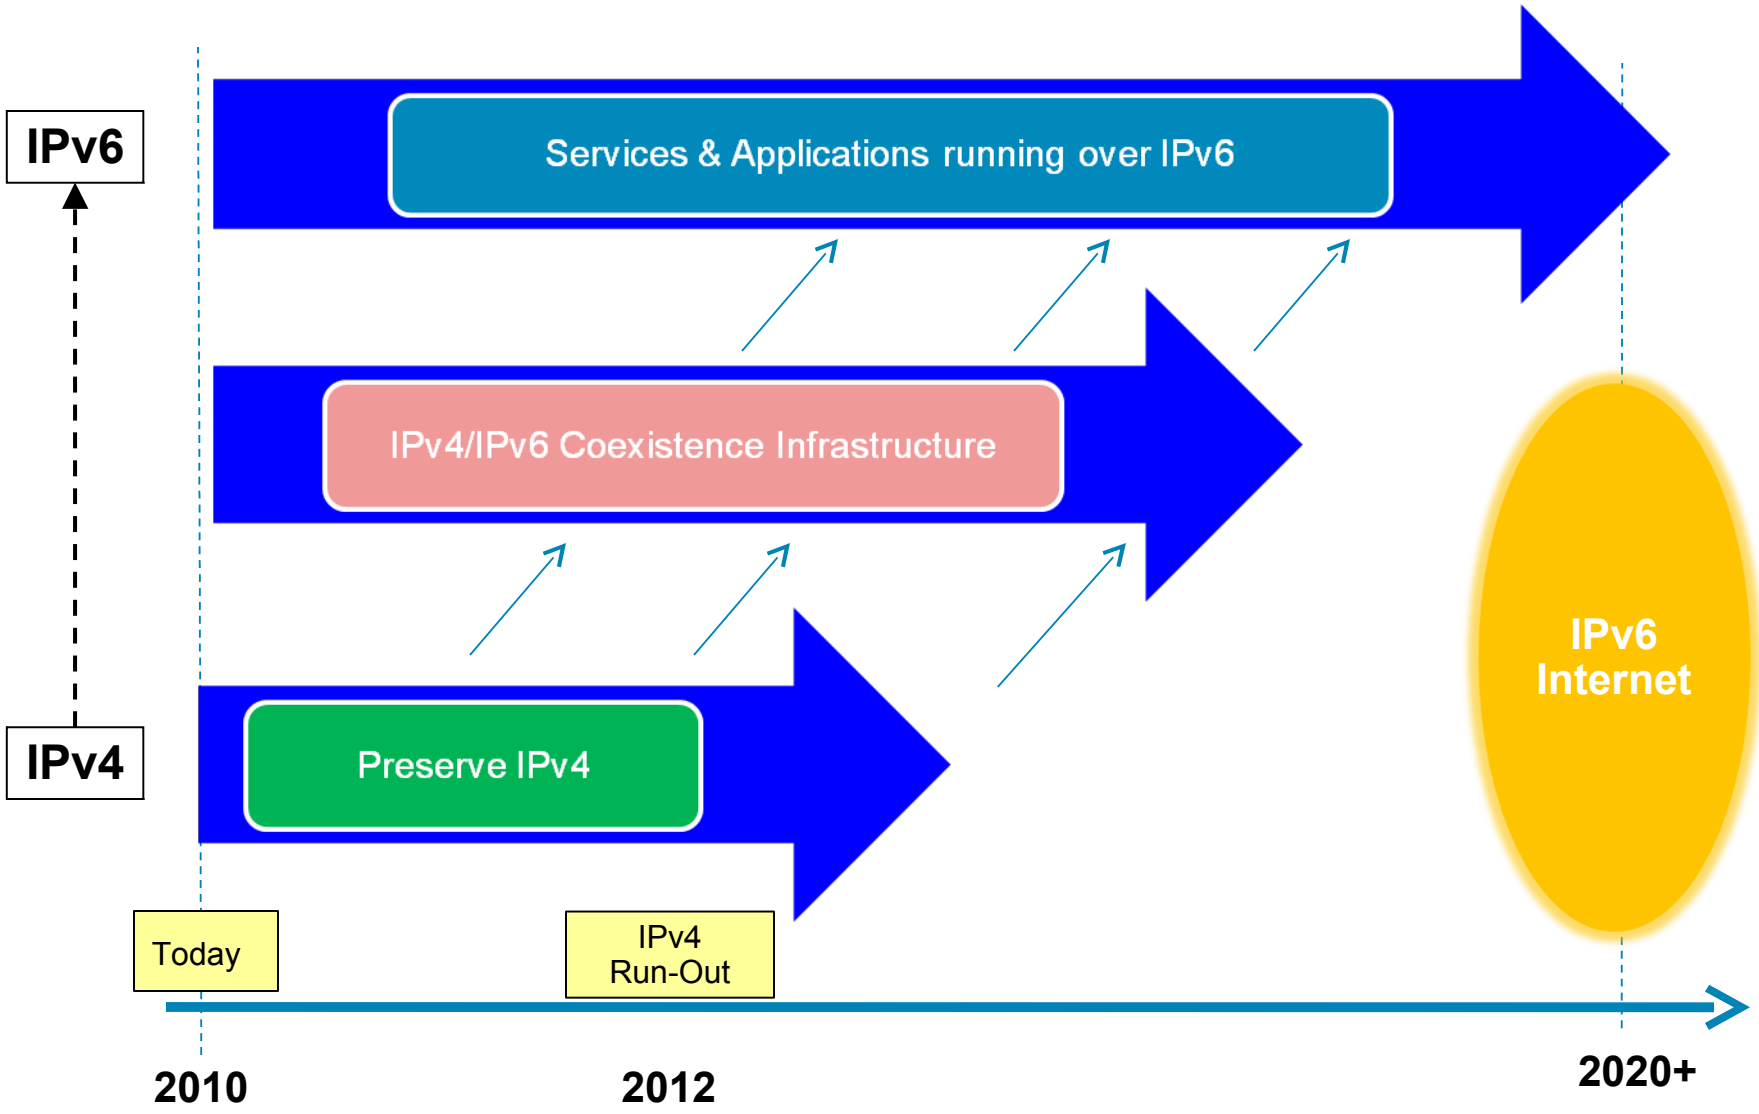

IPv6 Strategy in Broadband Access

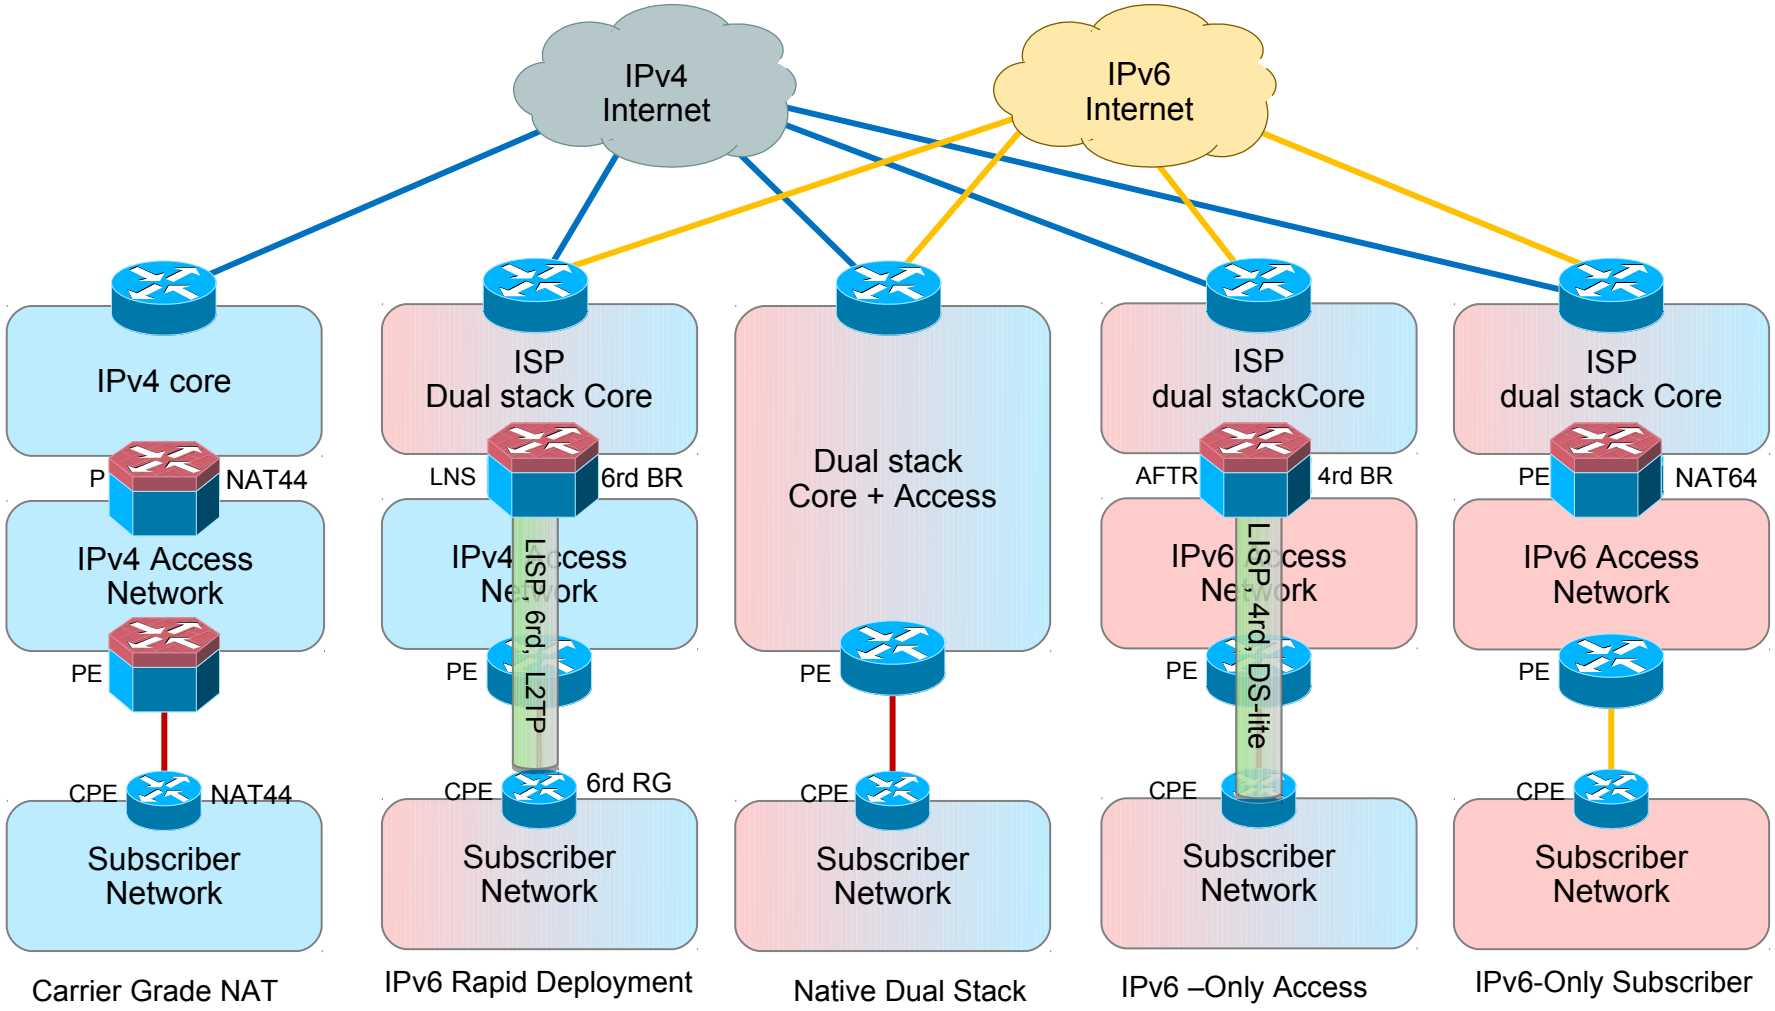

Preserve Prepare Prosper

Prepare, with 6rd (6-over-4)

Subscriber IPv6 traffic is **tunneled** over IPv4 to gateways within the IP-NGN while IPv6 grows

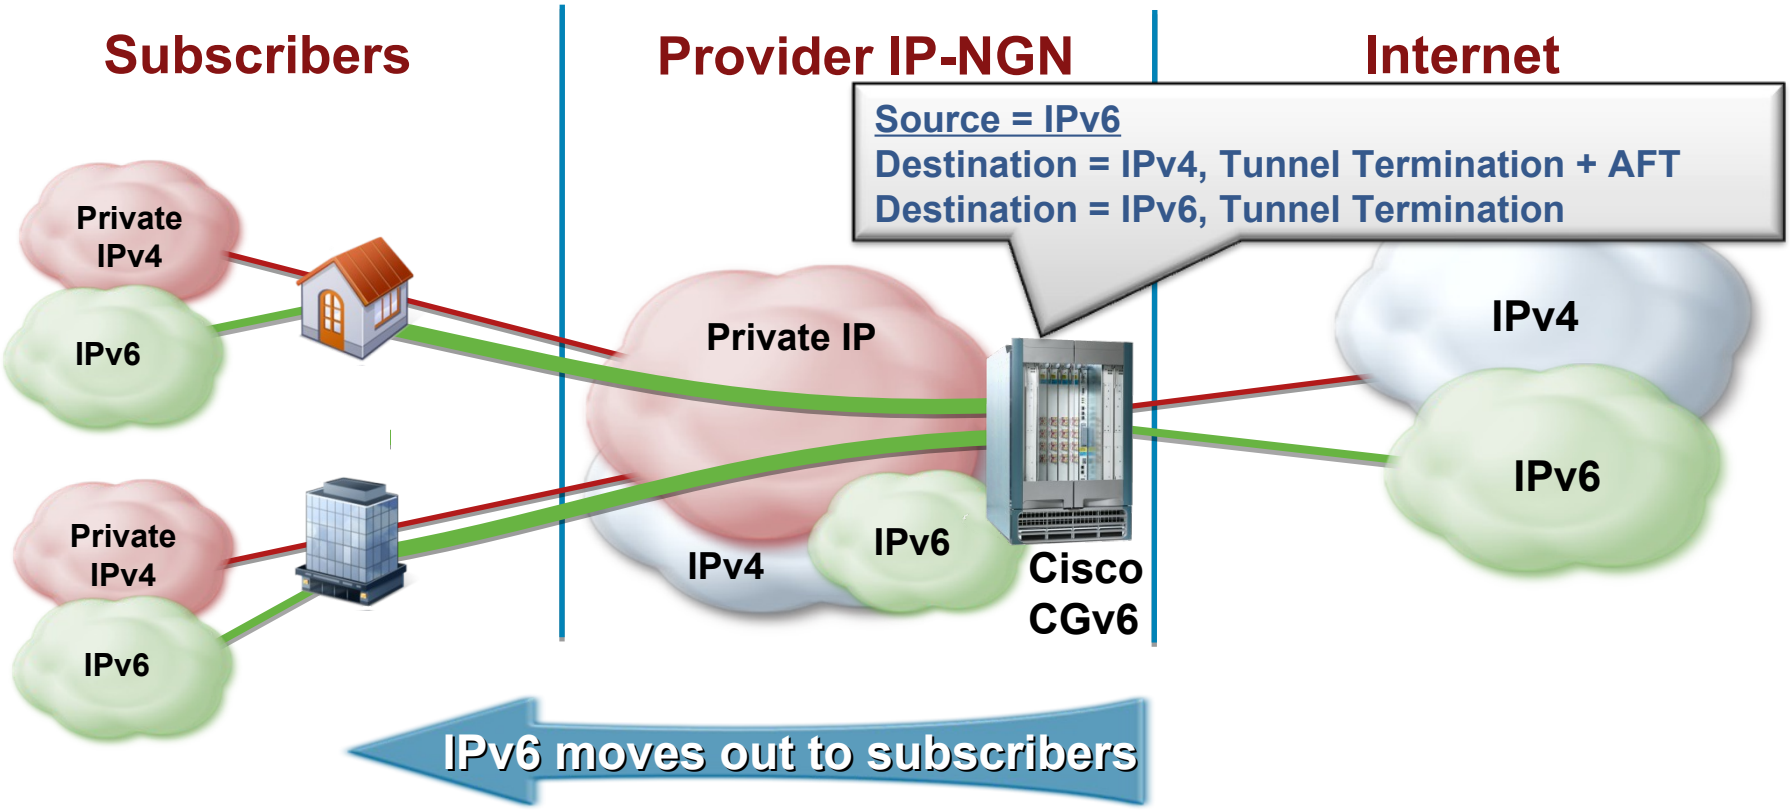

IPv6 Rapid Deployment (6rd) defines such a 6-over-4 model

6rd, IPv6 Rapid Deployment (RFC 5569)

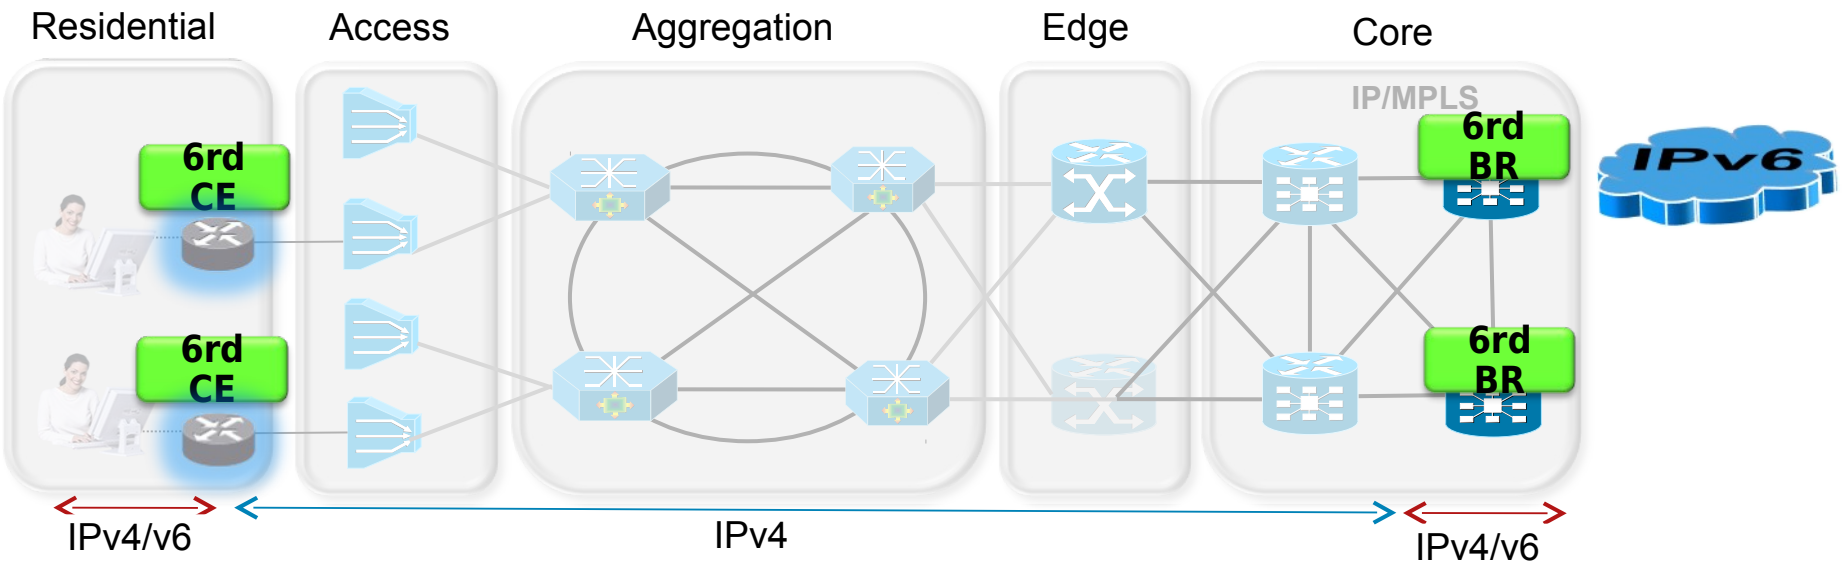

- Introduction of two Components: 6rd CE (Customer Edge) and 6rd BR (Border Relay)
- Reuses IPv4 in the SP (no IPv6 needed at Access/Aggregation network)
- Automatic Prefix Delegation on 6rd CE (no DHCPv6 or IPv6 provisioning system)
- Simple, stateless, automatic IPv6-in-IPv4 encap and decap functions on 6rd (CE & BR)
- Provides native dual stack to subscriber site by using existing IPv4 infrestructure
- 6rd BRs addressed with IPv4 anycast for load-balancing and resiliency
- Limited investment & impact on existing infrastructure

6rd in one slide

6rd specifies a protocol mechanism to deploy IPv6 to sites via a Service Provider's IPv4 network

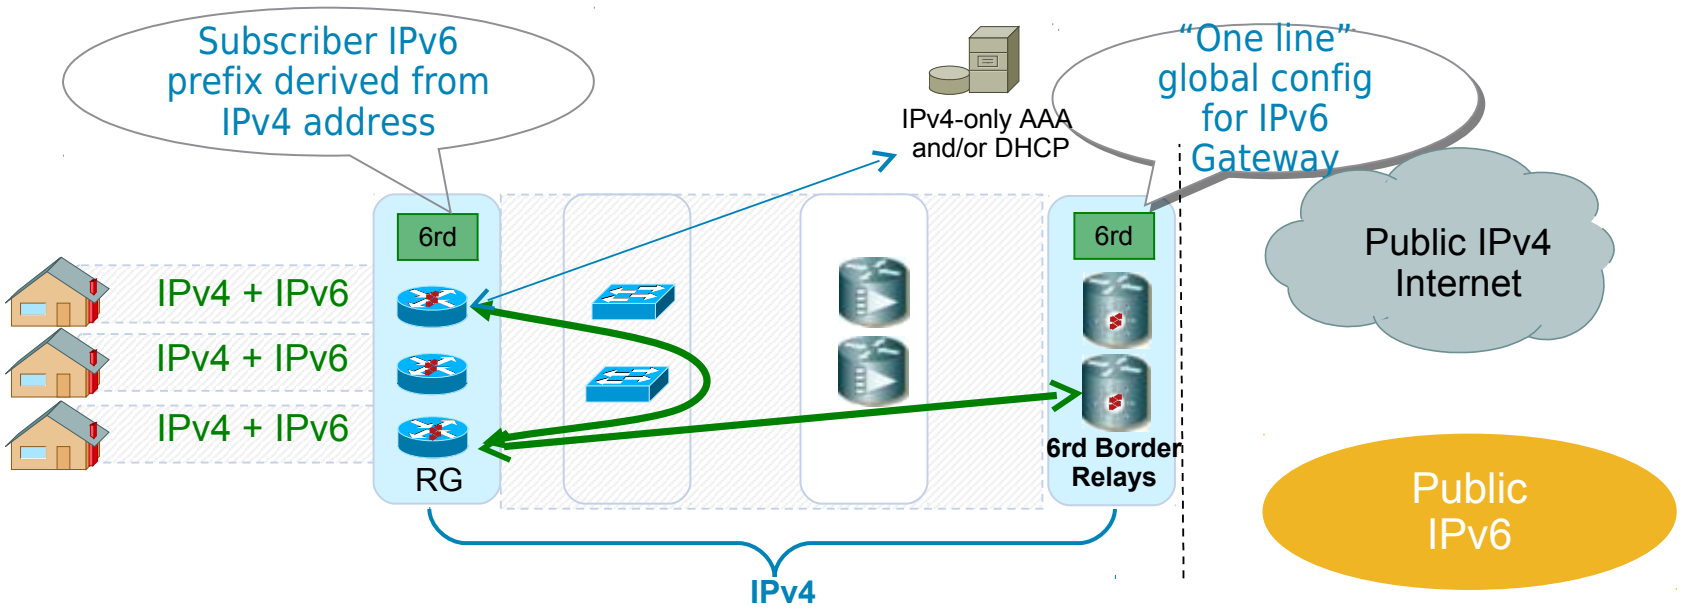

- It builds on 6to4 [RFC3056], with the key differentiator that it utilizes an SP's own IPv6 address prefix rather than 2002::/16
  - IPv6 address in derived from ISP IPv6 prefix and CPE IPv4 address
- Simple, stateless, automatic IPv6-in-IPv4 encap and decap functions
  - RG and BR perform automatic IPv6/IPv4 encap/decap (Protocol value 41)
  - IPv6 traffic automatically follows IPv4 Routing between CPE and BR
- BRs placed at IPv6 edge, addressed via anycast for load-balancing and resiliency

6rd IPv6 Prefix Delegation

The IPv6 prefix used by 6rd-RG to addressing user devices is calculated by combining the 6rd SP Prefix and the RG's IPv4 address obtained via IPv4 configuration methods (public or private)

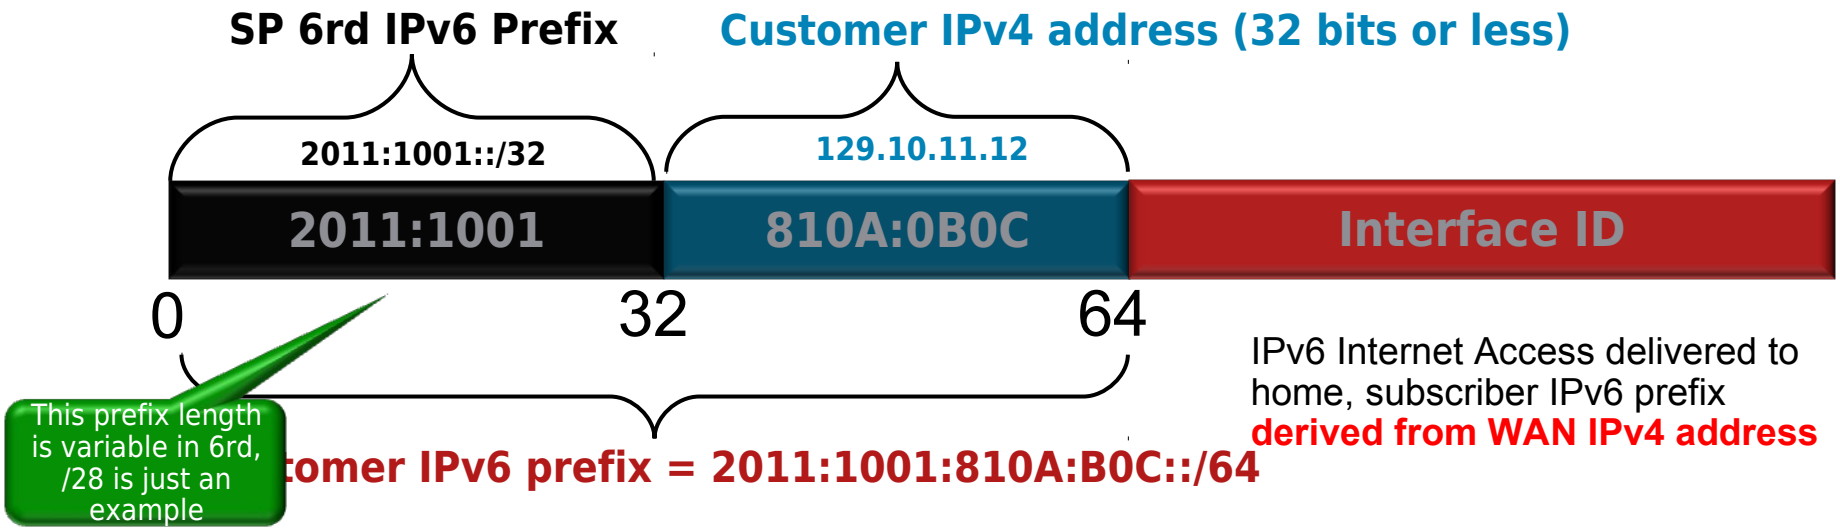

- RG need to get an IPv4 address first, from SP assignment
- RG will generate IPv6 prefix from 6rd prefix and ipv4 address
- End devices configured exactly as for any native IPv6 connectivity to LAN side

  SLACC or DHCPv6
- LAN station use ipv6 prefix to generate ipv6 address.
- Most browsers will prefer to use ipv6 if they can get AAAA record.

6rd Residential Gateway

The 6rd Customer Edge router (6rd CE) plays a critical role in a 6rd deployment.

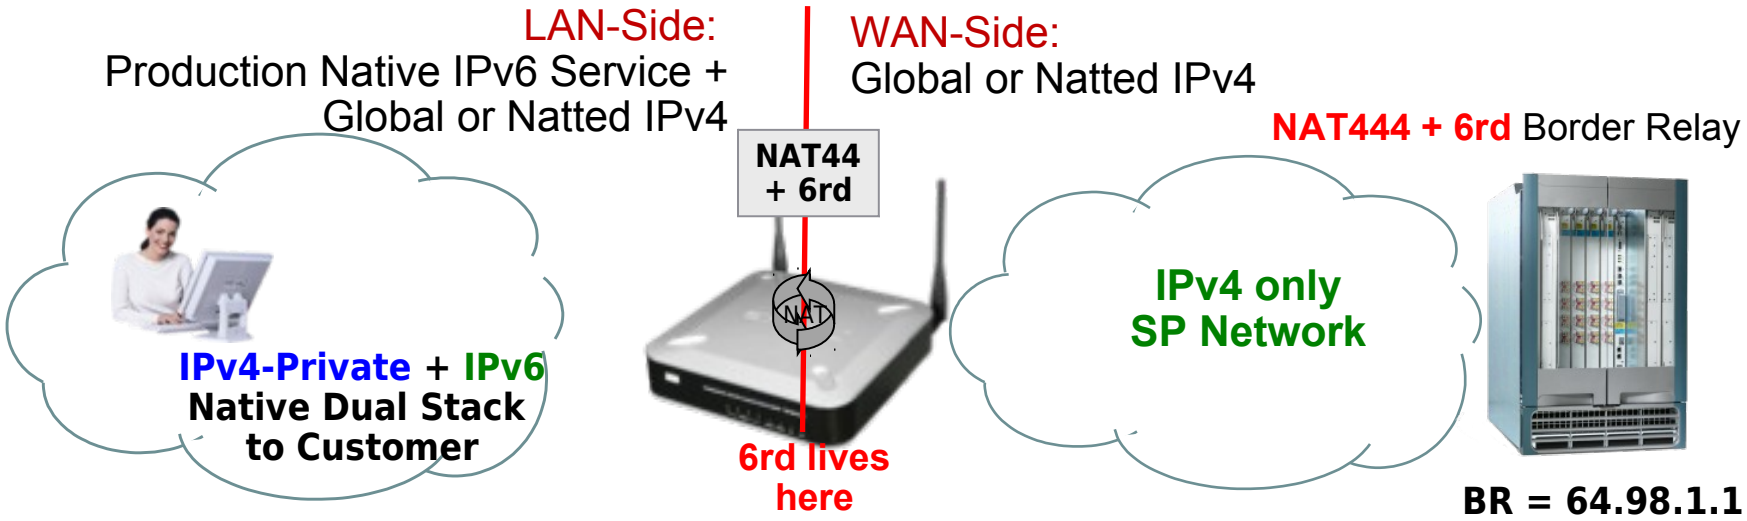

- The 6rd CE router must be configured with:
  1. 6rd IPv6 Prefix and length from SP
  2. Common IPv4 prefix length (v4 suffix length)
  3. 6rd Relay IPv4 address (likely anycast)
- This information can be configured into the device in a variety of ways,
  - including manual configuration, DHCP (option 212) and TR-69
- "Home side" of RG configured exactly as would be for "native" IPv6
  - e.g., same as for a DHCPv6 delegated prefix, using SLAAC
- The 6rd CE router MUST install a default IPv6 route to the relay

Combining NAT44 and 6rd

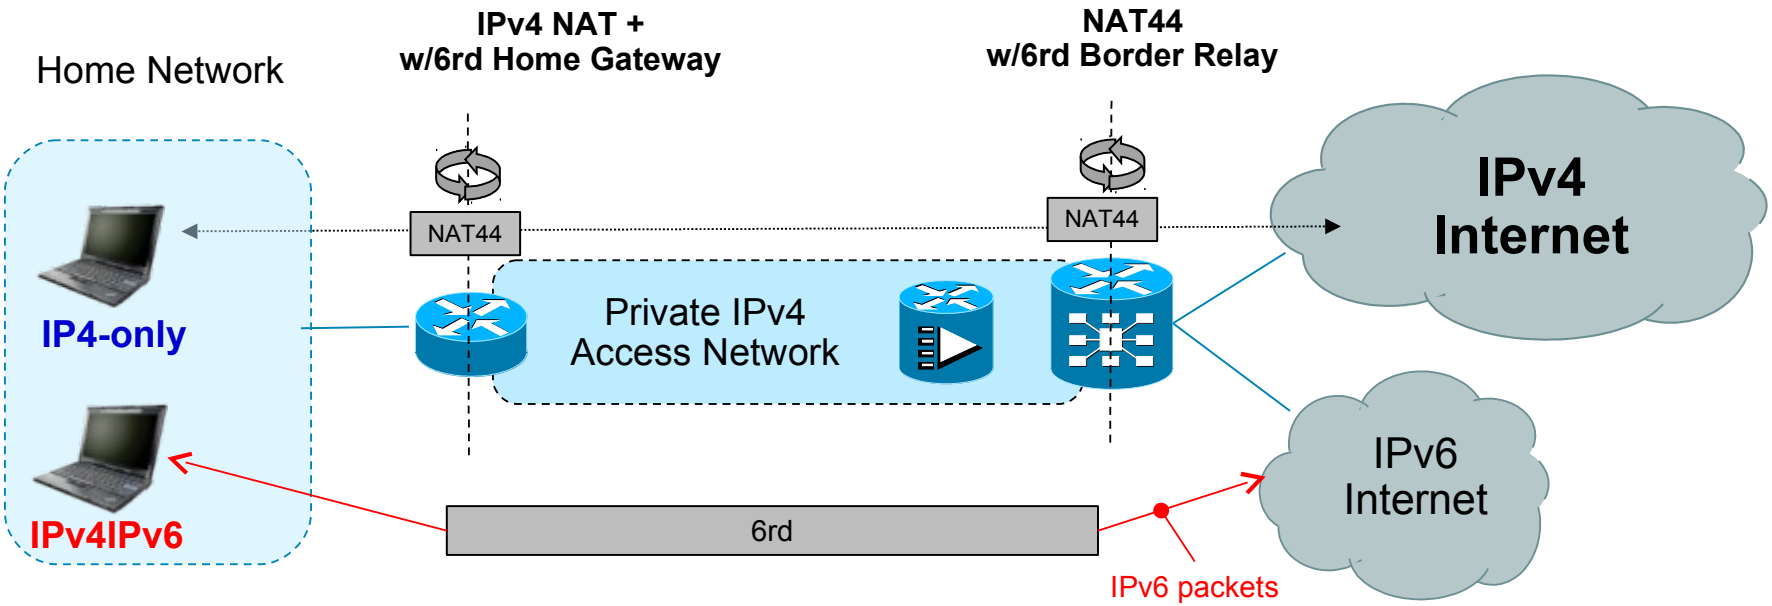

- Addresses IPv4 run-out and enables incremental IPv6 subscriber connectivity over existing IPv4 infrastructure
- 6rd connectivity becomes a NAT44 offload
  - –as more and more IPv4 content becomes IPv6-accessible
- Carrier, Content Provider, and User benefit when traffic runs over IPv6

Packet Flow and Encapsulation

6rd encapsulates IPv6 in IPv4 with a destination IPv4 address which is either encoded within the IPv6 destination address itself, or is the destination address of a preconfigured 6rd Border Relay router that can decapsulate the IPv4 header and route the IPv6 packet outside the SP's IPv4 network.

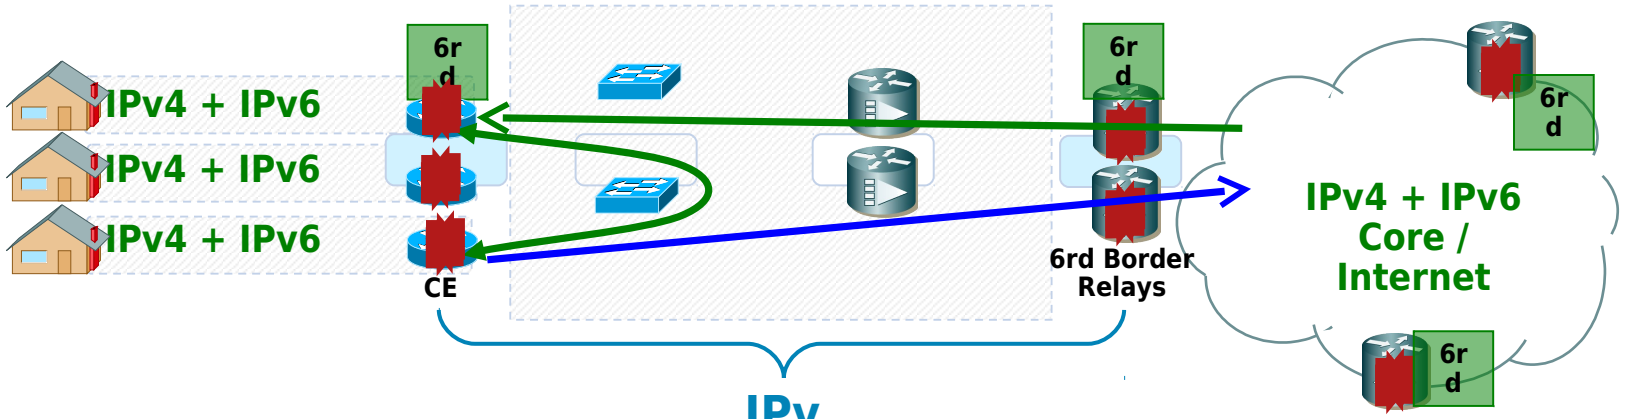

Dest = Inside 6rd Domain

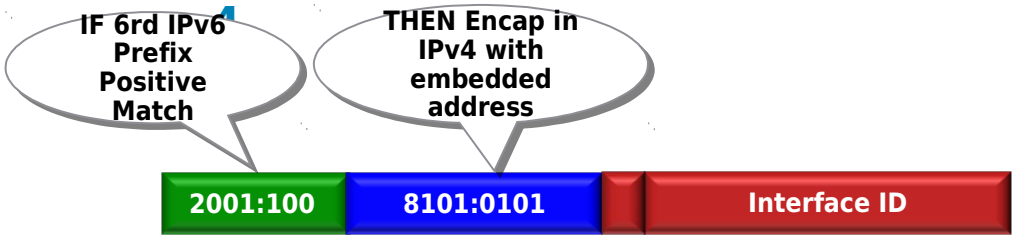

IPv6 Dest = Outside 6rd Domain

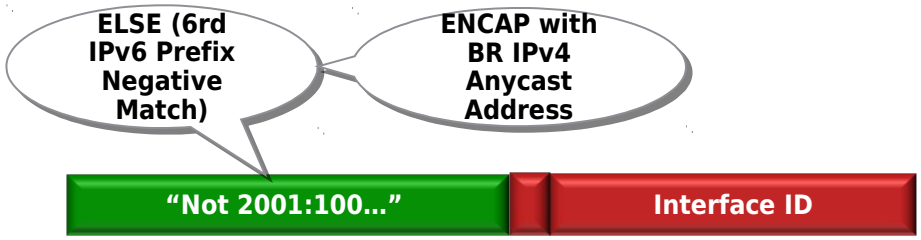

6rd Packet Encapsulation within domain

IPv4 encapsulation automatically determined from each packet's IPv6 destination

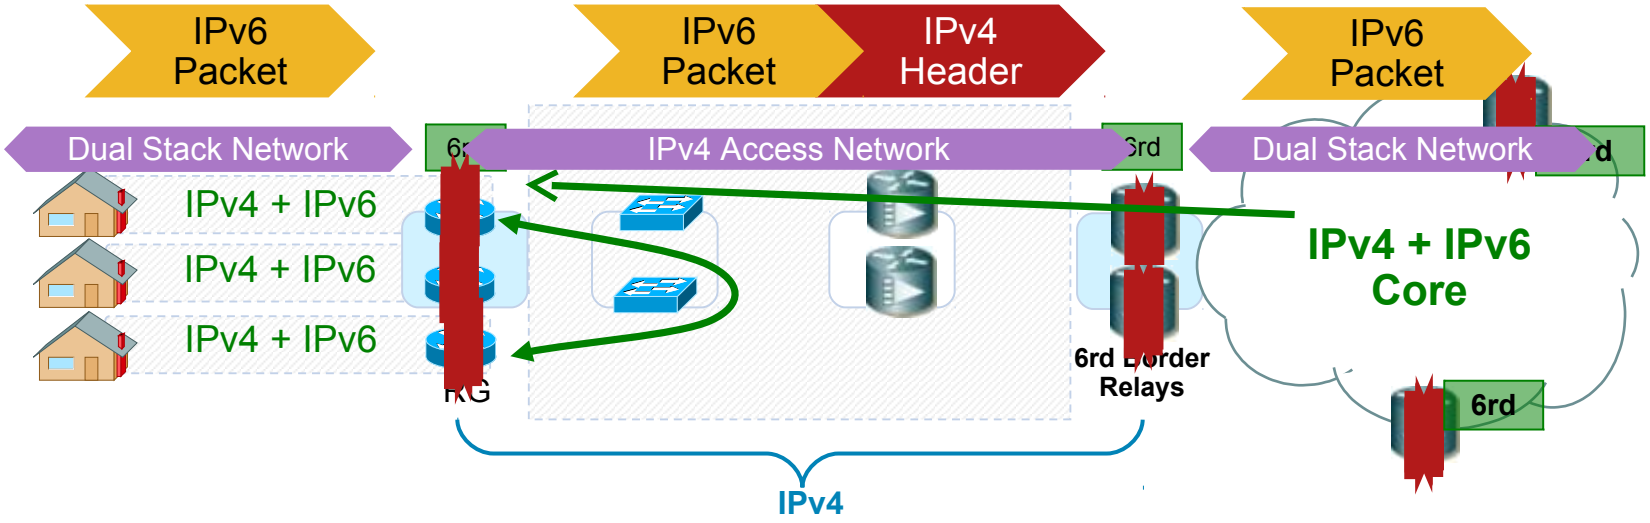

ISP 6rd IPv6 Prefix = 2001:1001:100:/40 IPv4 common bits=8, BR = 10.1.1.1

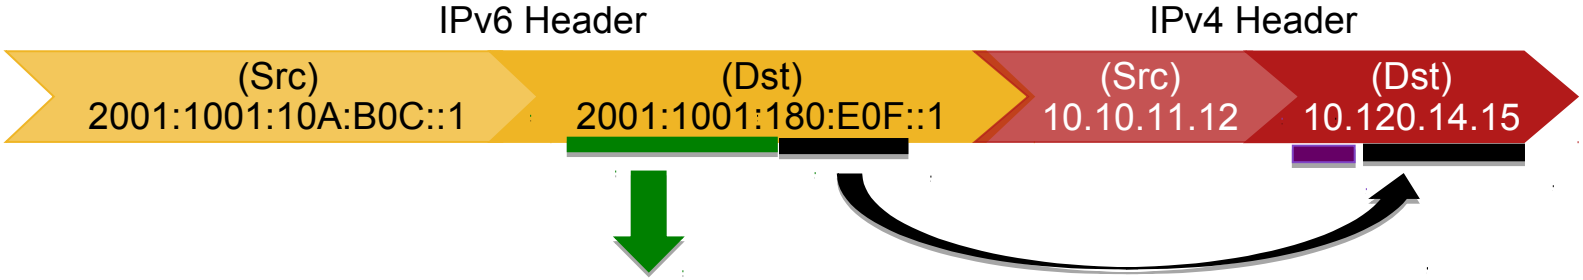

If (dstv6) match ISP 6rd IPv6 Prefix, then (dstv4) derived from (dstv6)

6rd Packet Encapsulation out of domain

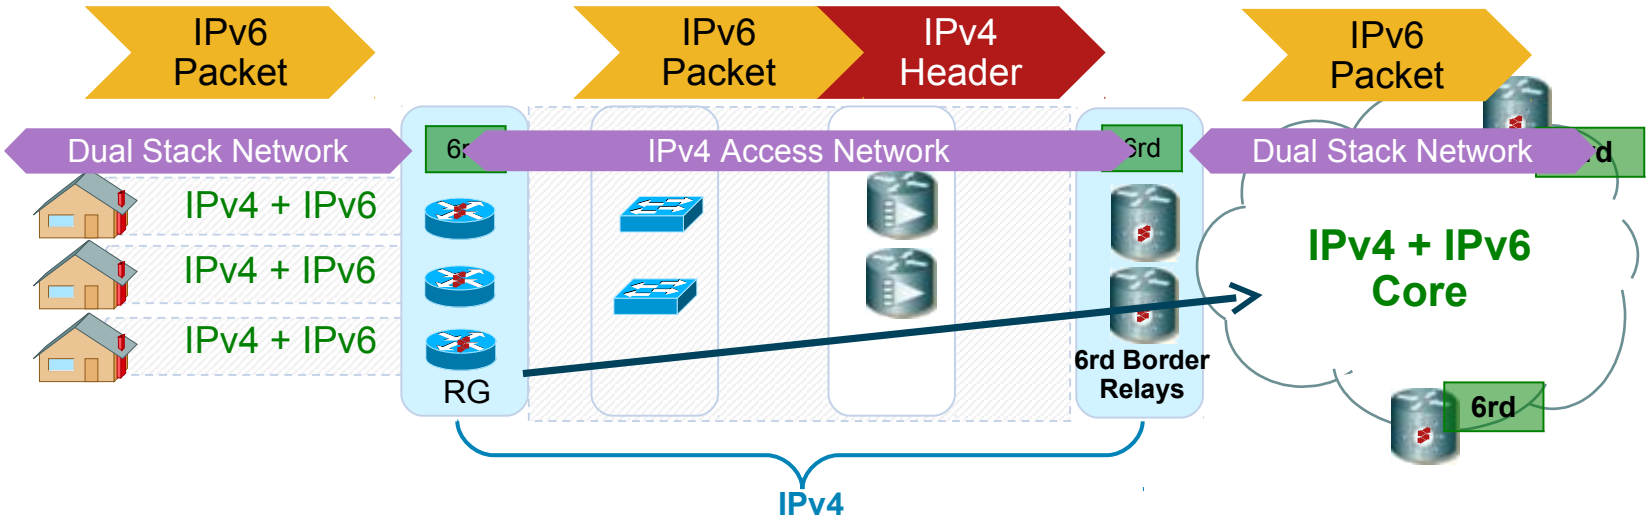

ISP 6rd IPv6 Prefix = 2001:1001:100:/40 IPv4 common bits=8, BR = 10.1.1.1

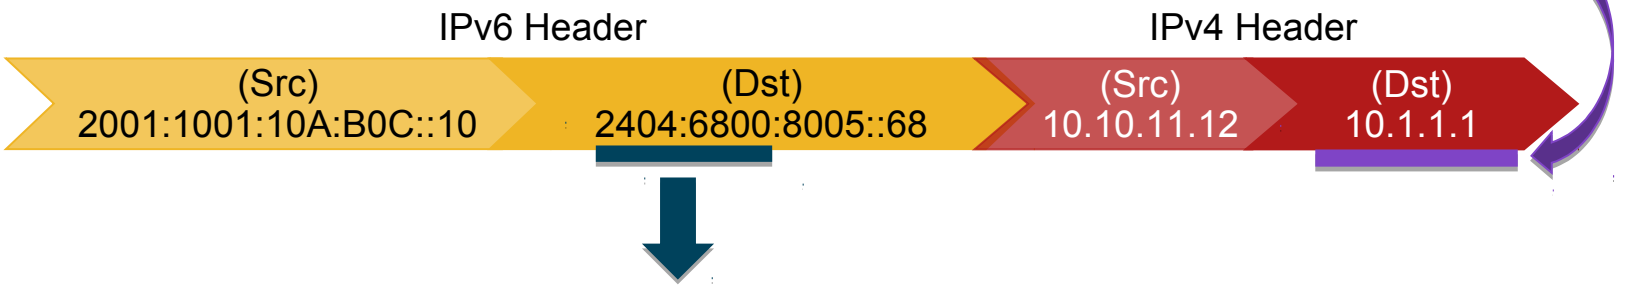

If (dstv6) not match ISP 6rd IPv6 Prefix, then (dstv4) = BR

6rd/NAT44 Lab Trial Diagram

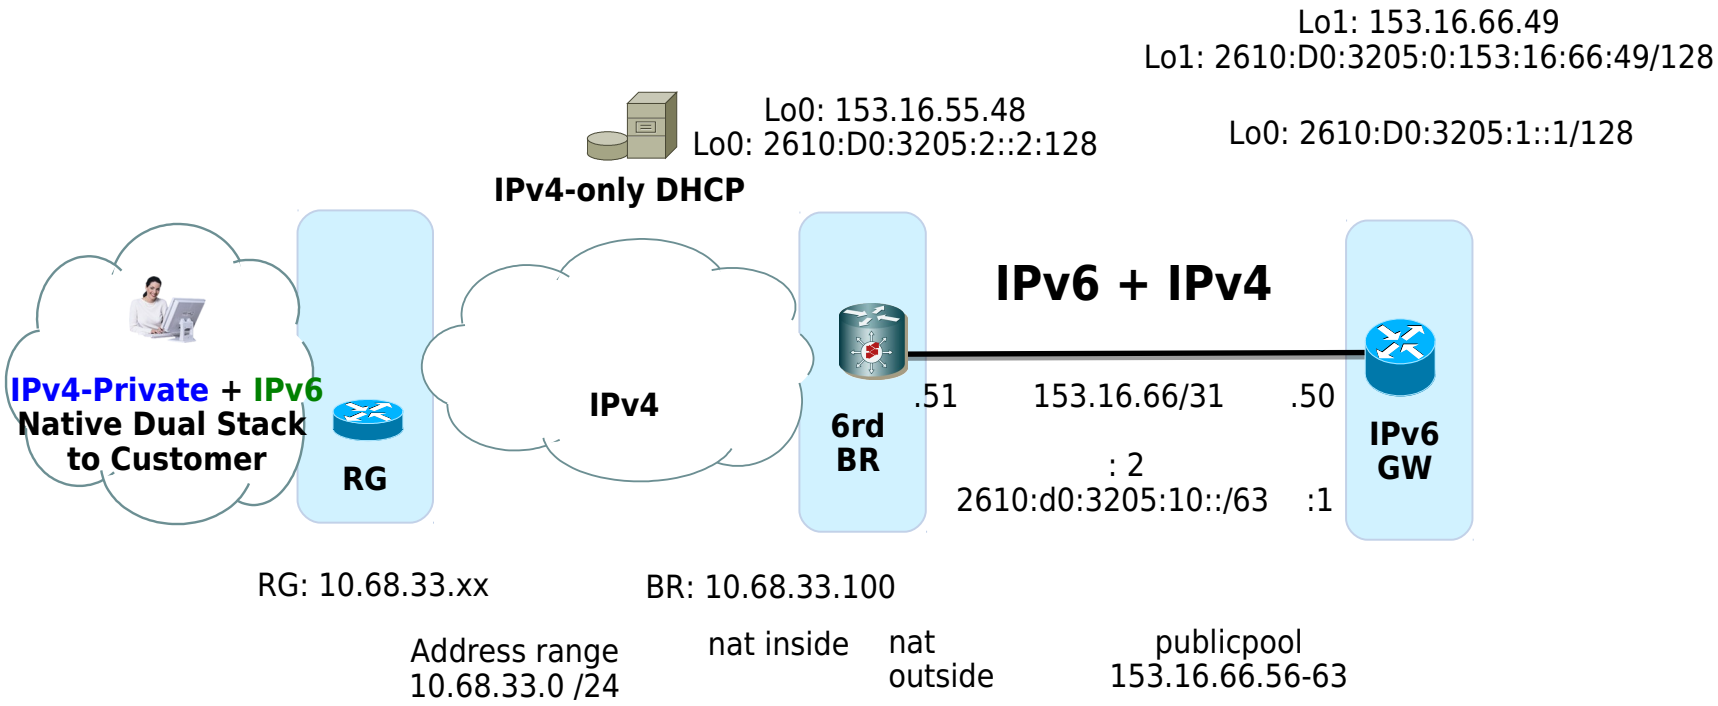

Lab Trial 6rd Automatic Prefix Delegation

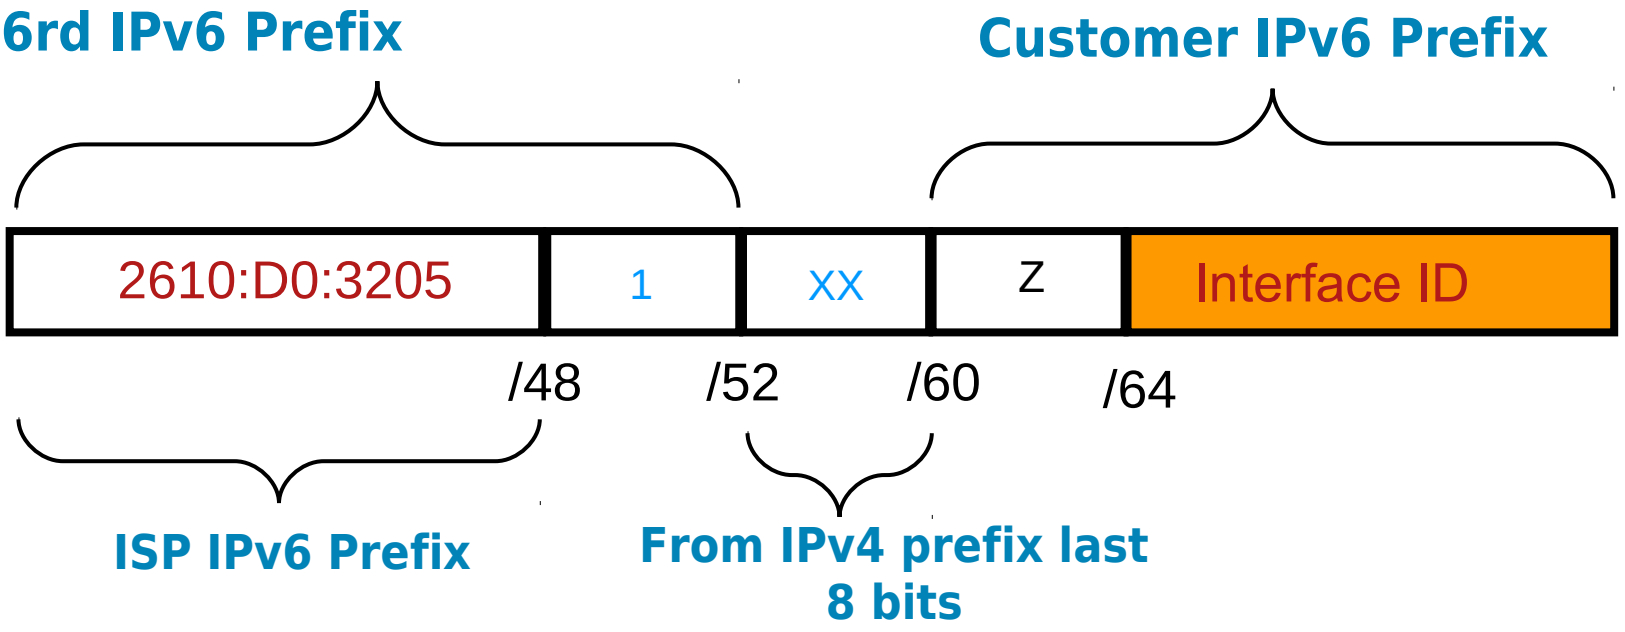

6rd IPv6 Prefix = 2610:D0:3205:1000::/52

IPv4 common bits = 24

BR = 10.68.33.100

6rd CE Linksys Sample

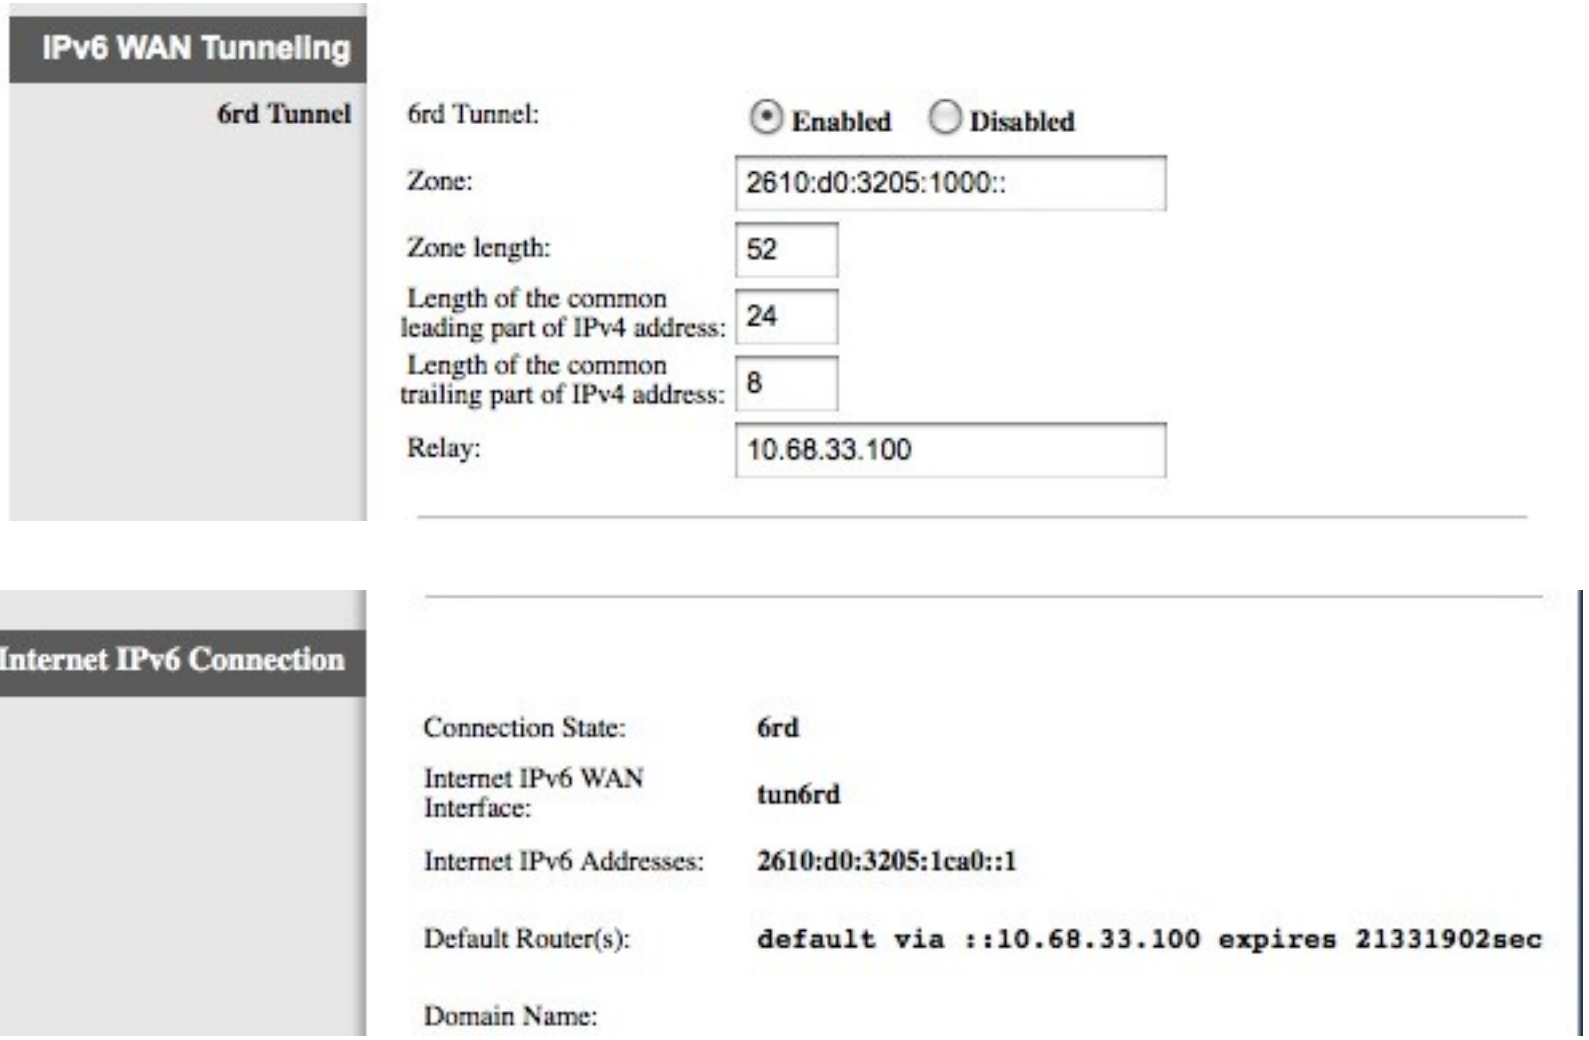

6rd Access IPv6 Internet

Platforms supporting 6rd BR and NAT

Cisco CGSE Module
(Carrier-Grade Services Engine)

Cisco CRS family

Cisco ASR family

Cisco ASR 1002

Cisco ASR 1004

Cisco ASR 1006

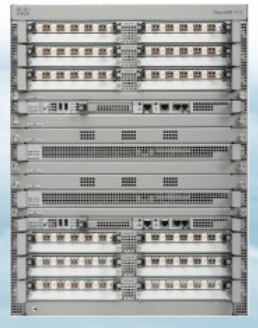

Cisco ASR 1013

Platforms supporting 6rd RG

Home Based

Official Support

Cisco-Linksys WRT610N

Cisco-Linksys E3000

Cisco-Linksys E4200

Open Source Code

(DD-WRT or OpenWRT)

Most Cisco-Linksys Wireless Router

For more detail on Linksys, pls contact Max Koh (makoh@cisco.com)

Enterprise/SMB based

Cisco IOS 15.1(3)T

Cisco 800 Series Routers

Cisco 1700 Series Routers

Cisco 1800 Series Routers

Cisco 2600 Series Routers

Cisco 2800 Series Routers

Cisco 3700 Series Routers

Cisco 3800 Series Routers

Cisco 7200 Series Routers

6rd BR ASR1K Sample

```
interface GigabitEthernet0/1/0
 ip address 10.68.33.100 255.255.255.0
!
interface Tunnel0
 tunnel source GigabitEthernet0/1/0
 tunnel mode ipv6ip 6rd
 tunnel 6rd ipv4 prefix-len 24
 tunnel 6rd prefix 2610:d0:3205:1000::/52
 ipv6 address 2610:d0:3205:1640::/128 anycast
!
ipv6 route 2610:d0:3205:1000::/52 Tunnel0
ipv6 route 2610:d0:3205:1640::/60 Null0
```

Simple and Easy to setup!

6rd BR ASR1K Show

```
#sh tunnel 6rd
Interface Tunnel0:
  Tunnel Source: 10.68.33.100
  6RD: Operational, V6 Prefix: 2610:D0:3205:1000::/52
       V4 Prefix, Length: 24, Value: 10.68.33.0
       V4 Suffix, Length: 0, Value: 0.0.0.0
  General Prefix: 2610:D0:3205:1640::/60
#sh int tu 0
Tunnel0 is up, line protocol is up
  Hardware is Tunnel
  MTU 17920 bytes, BW 100 Kbit/sec, DLY 50000 usec,
Encapsulation TUNNEL, loopback not set
  Keepalive not set
  Tunnel source 10.68.33.100 (GigabitEthernet0/1/0)
   Tunnel Subblocks:
      src-track:
         Tunnel0 source tracking subblock associated with GigabitEthernet0/1/0
          Set of tunnels with source GigabitEthernet0/1/0, 1 member (includes iterators), on
interface <OK>
  Tunnel protocol/transport IPv6 6RD
  Tunnel TTL 255
```

6rd CE Cisco1800 Sample

```
ipv6 general-prefix test 6rd Tunnel0
!
interface FastEthernet0
 ip address dhcp client-id FastEthernet0 hostname 1812
!
interface Tunnel0
 no ip address
 ipv6 address test ::/128
 tunnel source FastEthernet0
 tunnel mode ipv6ip 6rd
 tunnel 6rd ipv4 prefix-len 24
 tunnel 6rd prefix 2610:D0:3205:1000::/52
 tunnel 6rd br 10.68.33.100
!
interface Vlan1
 ipv6 address test ::1:0:0:0:1/64
 ipv6 nd autoconfig prefix
!
ipv6 route ::/0 Tunnel0 2610:d0:3205:1640::
```

6rd CE Cisco1800 Show

```
#show tunnel 6rd
Interface Tunnel0:
  Tunnel Source: 10.68.33.199
  6RD: Operational, V6 Prefix: 2610:D0:3205:1000::/52
       V4 Prefix, Length: 24, Value: 10.68.33.0
       V4 Suffix, Length: 0, Value: 0.0.0.0
       Border Relay address: 10.68.33.100
  General Prefix: 2610:D0:3205:1C70::/60
#show ipv6 general-prefix
IPv6 Prefix test, acquired via 6rd
  2610:D0:3205:1C70::/60 Valid lifetime infinite, preferred lifetime infinite
   Vlan1 (Address command)
   Tunnel0 (Address command)
ipv6-1812# show ipv6 route
S   ::/0 [1/0]
     via 2610:D0:3205:1640::, Tunnel0
LC  2610:D0:3205:1C70::/128 [0/0]
     via Tunnel0, receive
C   2610:D0:3205:1C71::/64 [0/0]
     via Vlan1, directly connected
```

NAT44 ASR1K Config

```
interface GigabitEthernet0/1/0
 ip address 10.68.33.100 255.255.255.0
 ip nat inside
!
interface GigabitEthernet0/1/1
 ip address 153.16.66.51 255.255.255.252
 ip nat outside
!
router ospf 4
 redistribute static subnets
!
ip nat pool publicpool 153.16.66.56 153.16.66.63 prefix-length 29
ip nat inside source list 10 pool publicpool
!
ip route 153.16.66.56 255.255.255.248 Null0
!
access-list 10 permit 10.68.33.0 0.0.0.255
```

NAT44 ASR1K Show

```
#sh ip nat tr
Pro  Inside global         Inside local          Outside local         Outside global
---  153.16.66.56          10.68.33.100          ---                   ---
tcp  153.16.66.56:445      10.68.33.100:445      93.123.60.32:3993     93.123.60.32:3993
tcp  153.16.66.56:445      10.68.33.100:445      122.116.97.160:4668   122.116.97.160:4668
!
#sh ip nat st
Total active translations: 3 (0 static, 3 dynamic; 2 extended)
Outside interfaces:
  GigabitEthernet0/1/1
Inside interfaces:
  GigabitEthernet0/1/0
Dynamic mappings:
-- Inside Source
[Id: 1] access-list 10 pool publicpool refcount 3
 pool publicpool: netmask 255.255.255.248
          start 153.16.66.56 end 153.16.66.63
          type generic, total addresses 8, allocated 1 (12%), misses 0
```

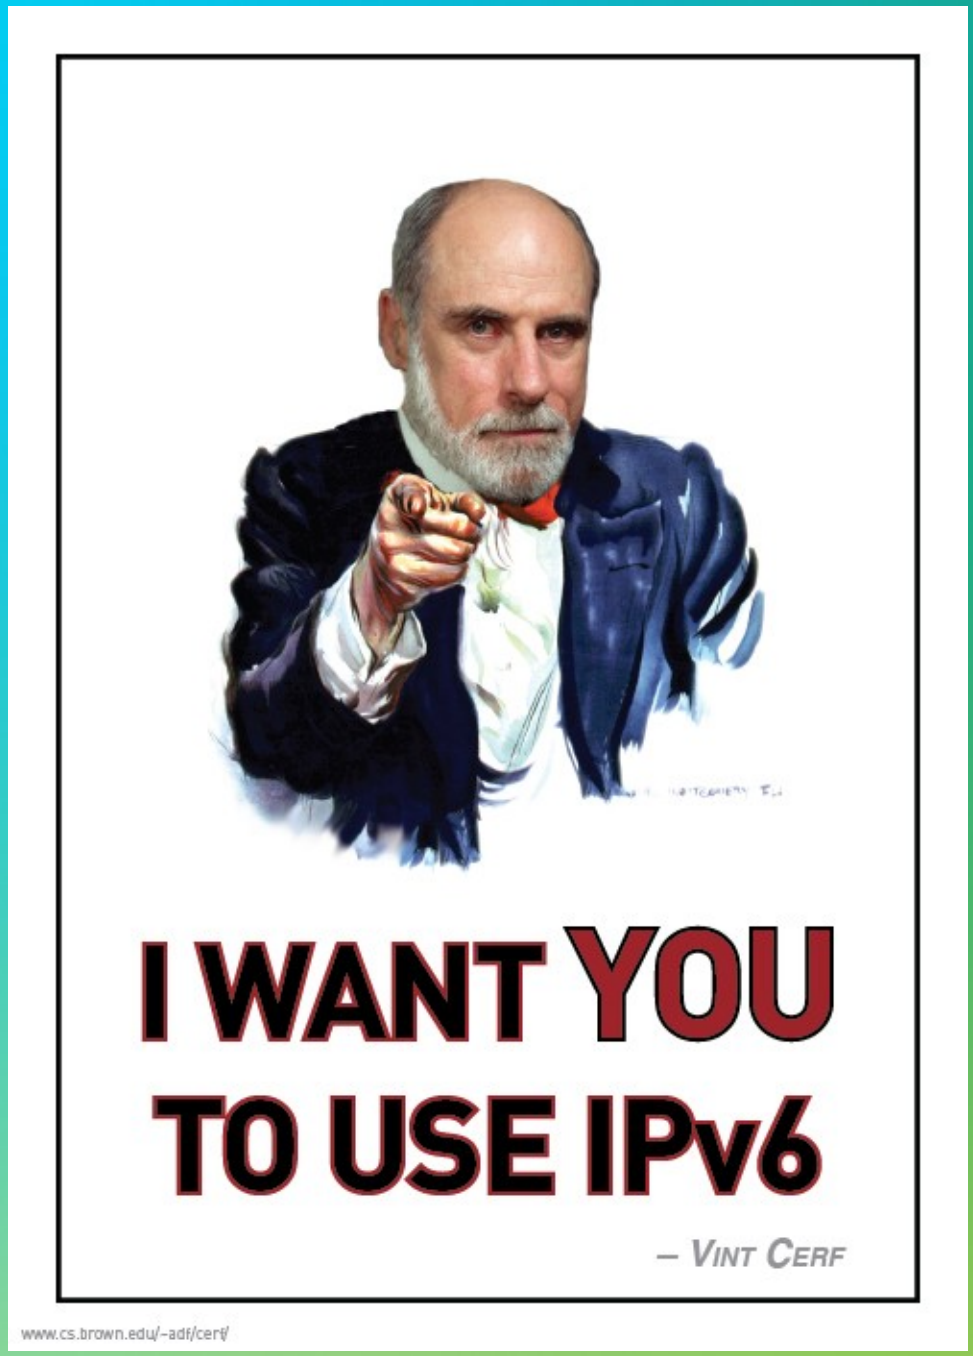

CISCO