



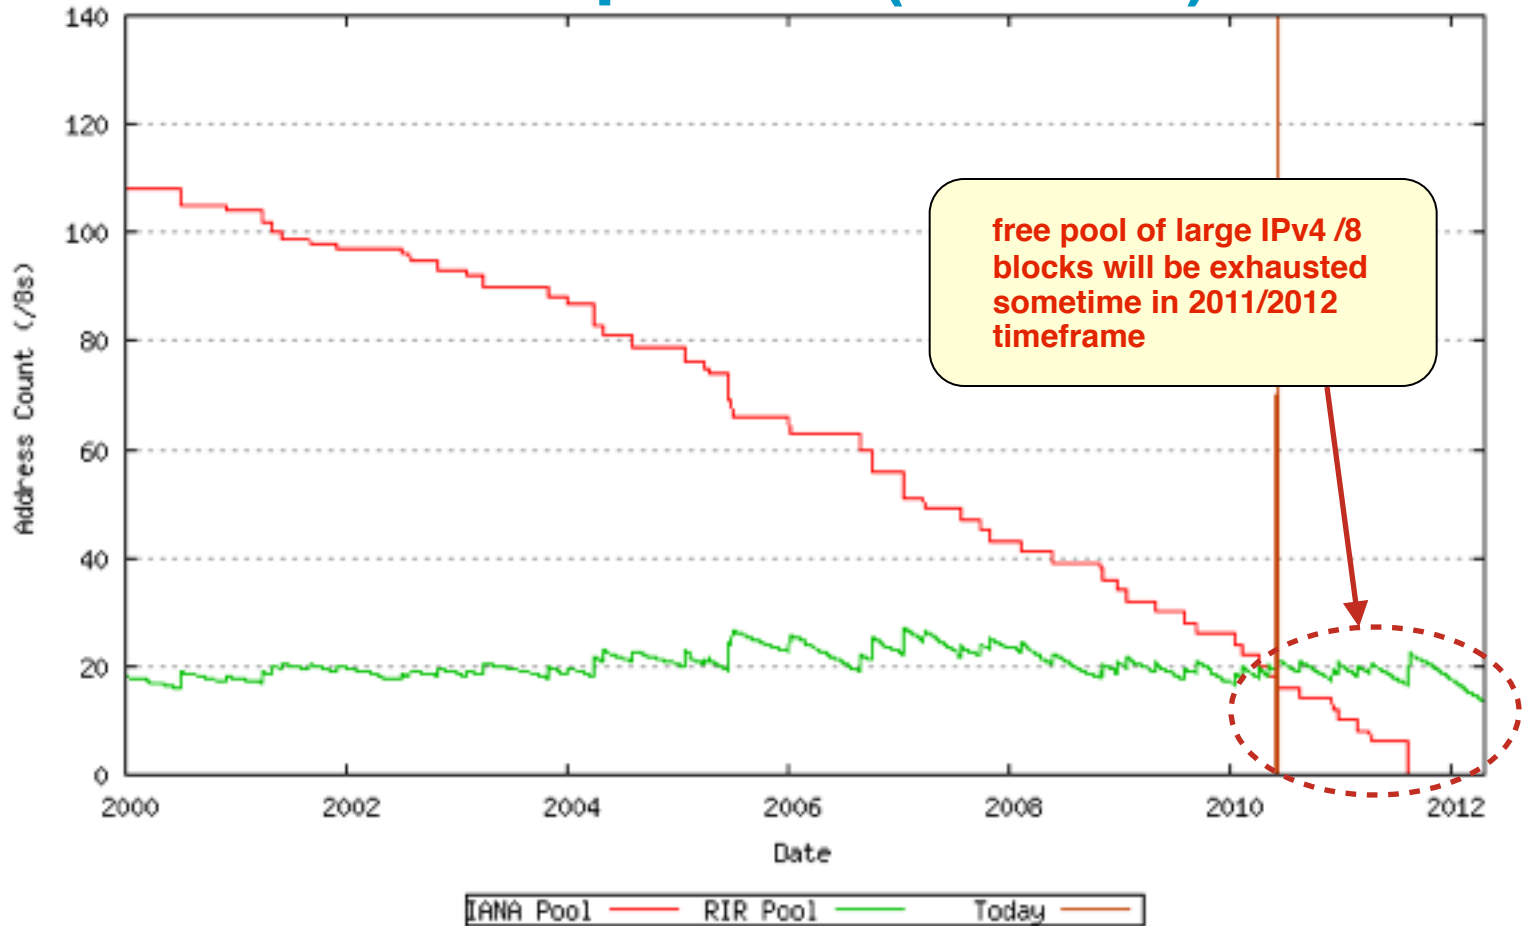
planning your conversion from IPv4 to IPv6



OPASTCO summer convention - 26-july, 2010

steve ulrich (sulrich@cisco.com) - consulting systems engineer


IPv4 address completion (run-out)



	15-dec-2009	08-jun-2010
Projected IANA Unallocated Address Pool Exhaustion:	25-Sep-2011	14-aug-2011
Projected RIR Unallocated Address Pool Exhaustion:	22-Sep-2012	22-apr-2012

see <http://www.potaroo.net/tools/ipv4/index.html>

IPv6 drivers



IPv4 Address space completion

- Public or Private Space
- Limiting network expansion and putting at risk business continuity
- Introducing Operational challenges



National IPv6 Strategies

- Compliance: U.S. Federal Mandate, IPv6 task force
- Next Generation Internet (CNGI) project in China and Japan
- European Commission Recommendation



- IPv6 “on” in Microsoft Vista
- Sensor Networks
- Apple's “Back to My Mac”
- v6 over v4 OTT tunnel providers

IPv6 in Client Software



- Next generation Network architecture require IPv6
- DOCSIS 3.0, Quad Play
- Mobile SP
- Networks in Motion
- Networked Sensors, i.e.: AIRS

Infrastructure Evolution

economics in play

- **basic technical truth about IPv6**

IPv6 in the short term will not solve the IPv4 exhaust problem

over the long term the cost per Gbit/s of an IPv6 system is lower than that of an IPv4 system with NAT.

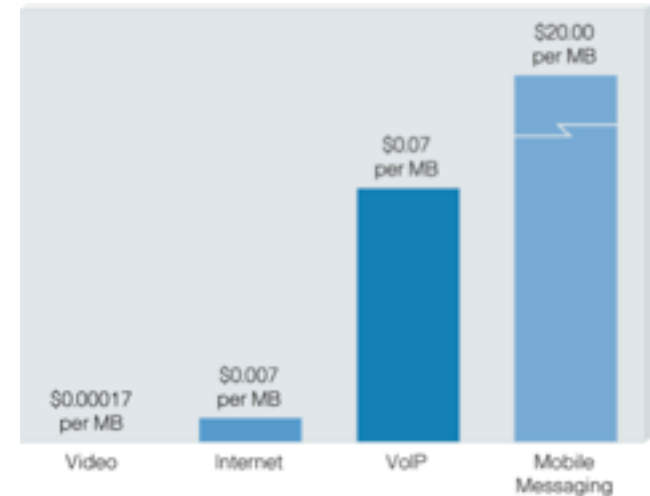
- **the economics of supplying high bit rate services to IP broadband users are key**

IPv6 technology provides a sustainable economical base for internet traffic delivery.


- **The long term economics of video broadband is a business case for IPv6.**

- **IPv6 will eventually be about services.**

to start with it will be about economics that IPv4 can't achieve eg video, mobile broadband end-points.



Price per Megabyte

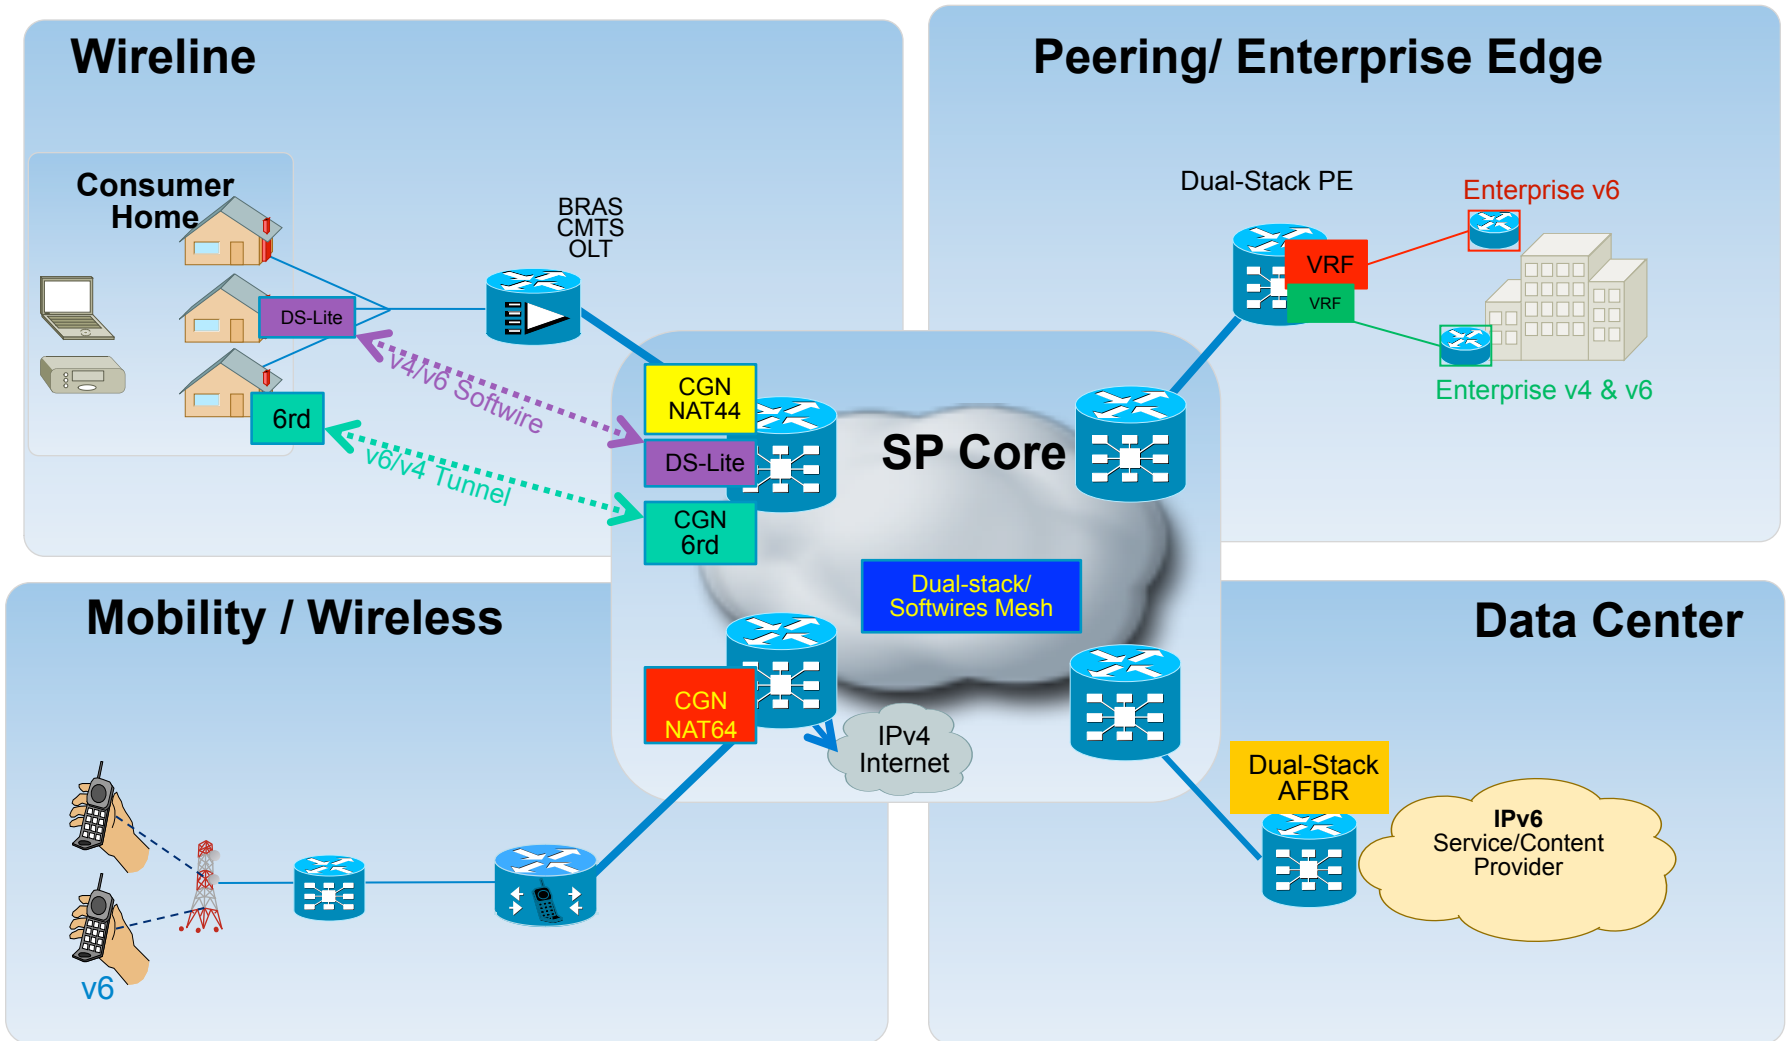


Global Traffic Volume

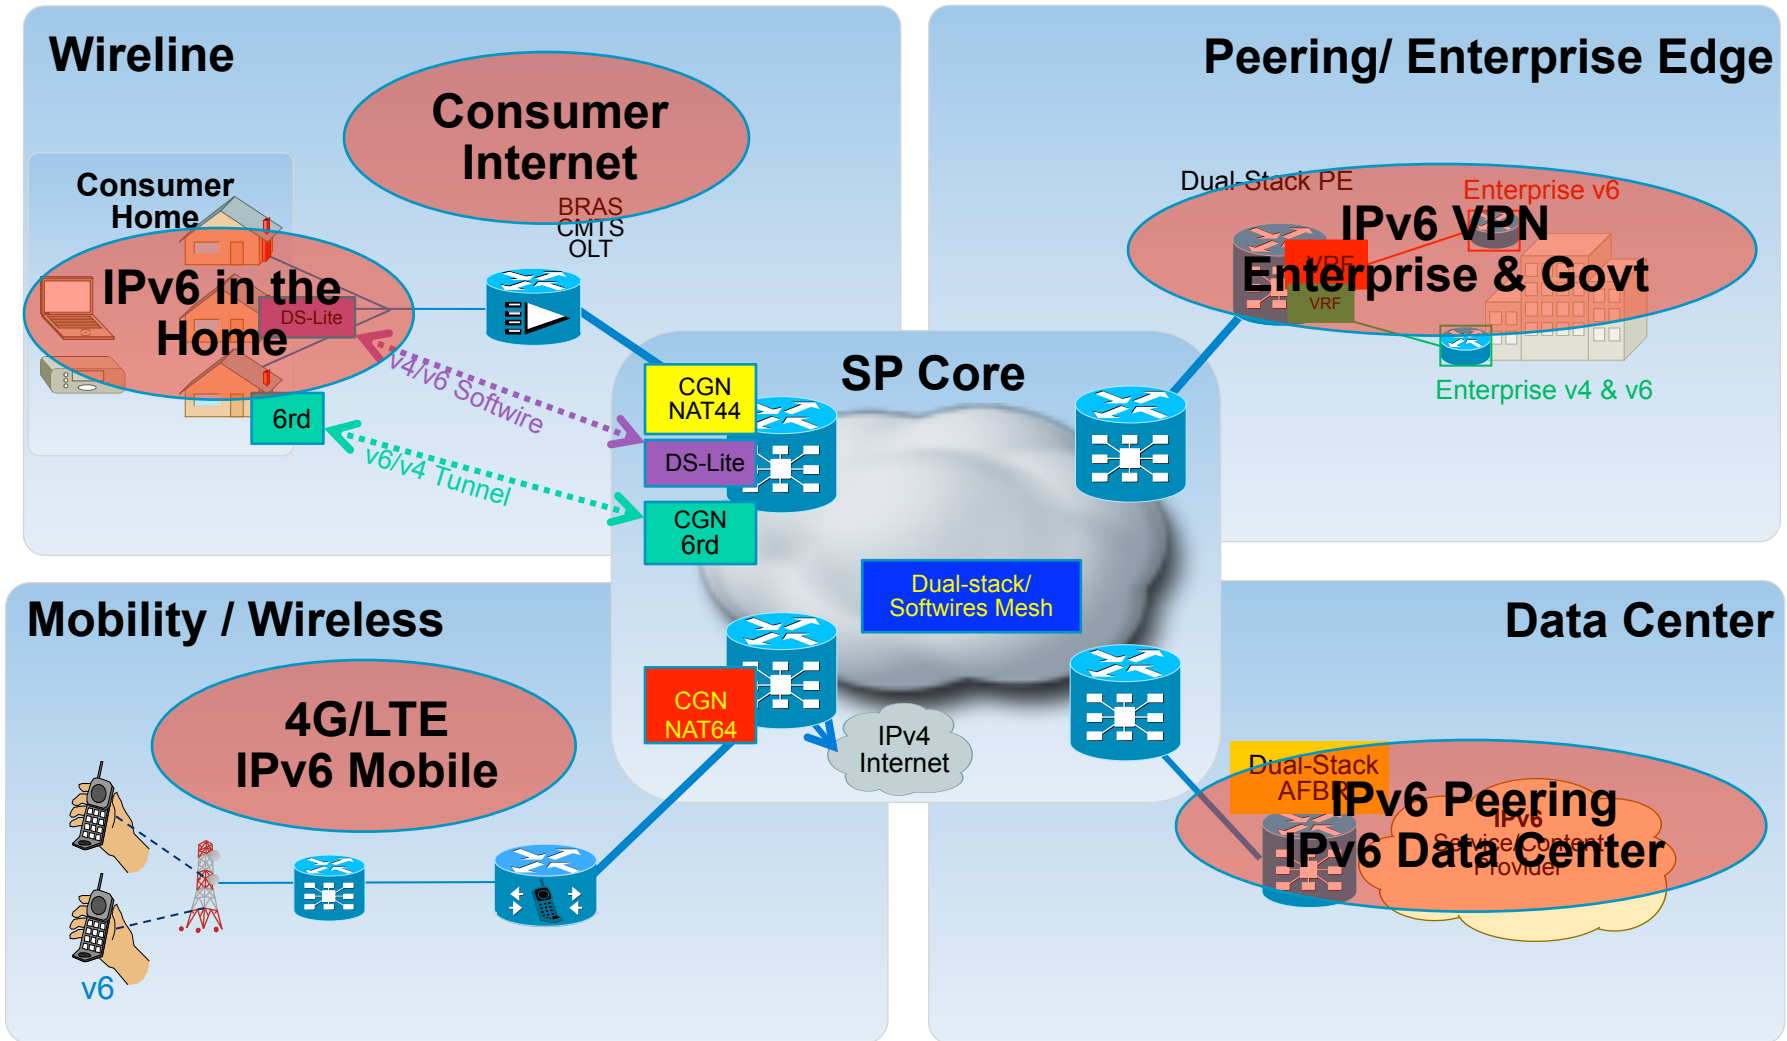
there are services that will run better on IPv6

characteristic	reason	example
Infrequent Use	Maintaining NAT bindings for rare occurrence events is inefficient	<ul style="list-style-type: none"> NTT IPv6 Earthquake Warning service Intelligent Energy Grid
Universal Connectivity	Reachability of devices in the home	<ul style="list-style-type: none"> Microsoft universal connectivity
Energy Efficient Network	Many networked applications per PC results in many keep-alives. Each needs Δ power on many devices.	<ul style="list-style-type: none"> Skype for iPhone drains batteries from application need for keep-alive
Scalable Data Center	Persistent client/server transport connection is needed to keep NAT open (limiting server scale)	<ul style="list-style-type: none"> Web application long polling
High bit Rate services	SP margin per bit for NAT versus a competitor without that cost	<ul style="list-style-type: none"> Netflix On-Demand will be costly. Which SPs extend bandwidth caps?

IPv6 - Backbone-First – 2010-2011

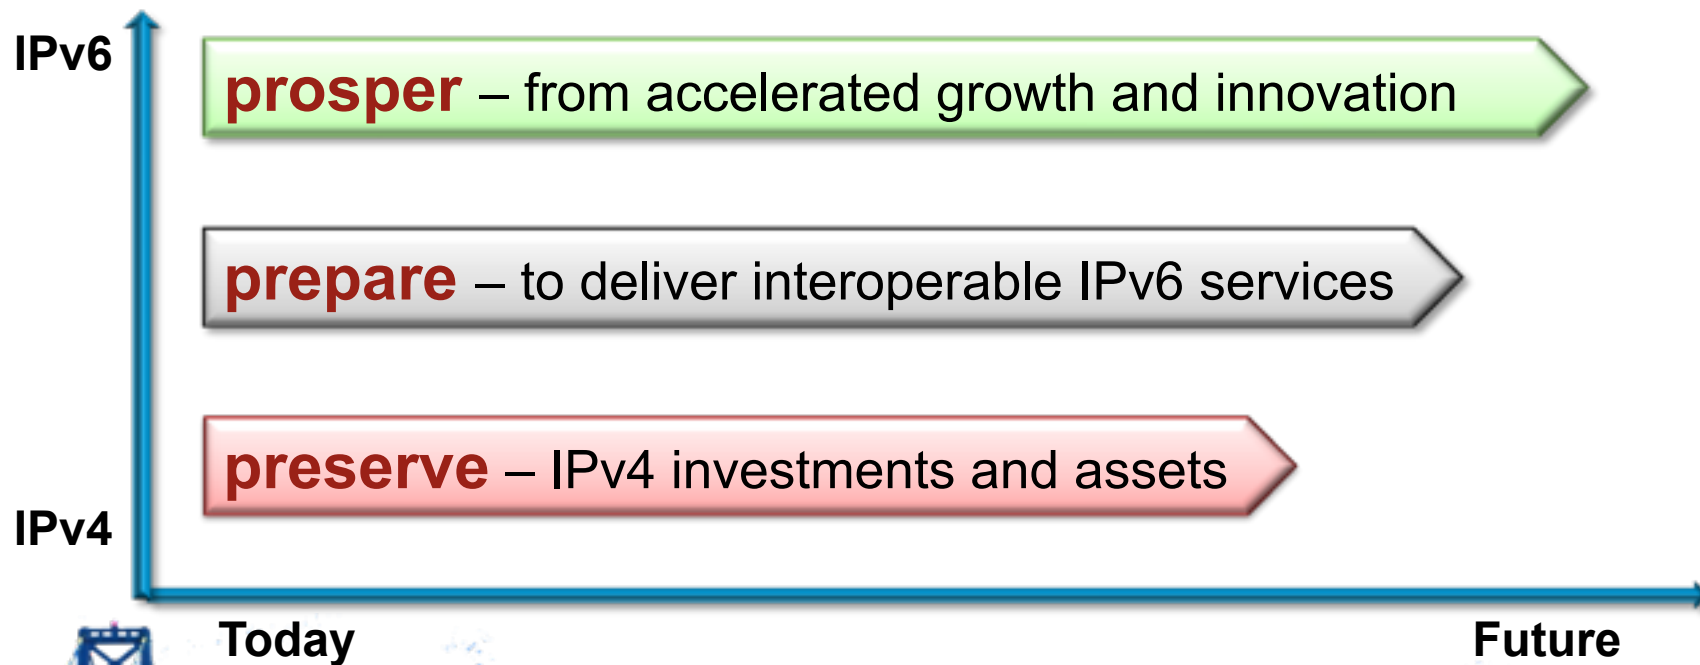


IPv6 Backbone-First – 2010-2011



bridging into the future

Cisco's Carrier-Grade IPv6 Framework – engineered to help service providers:

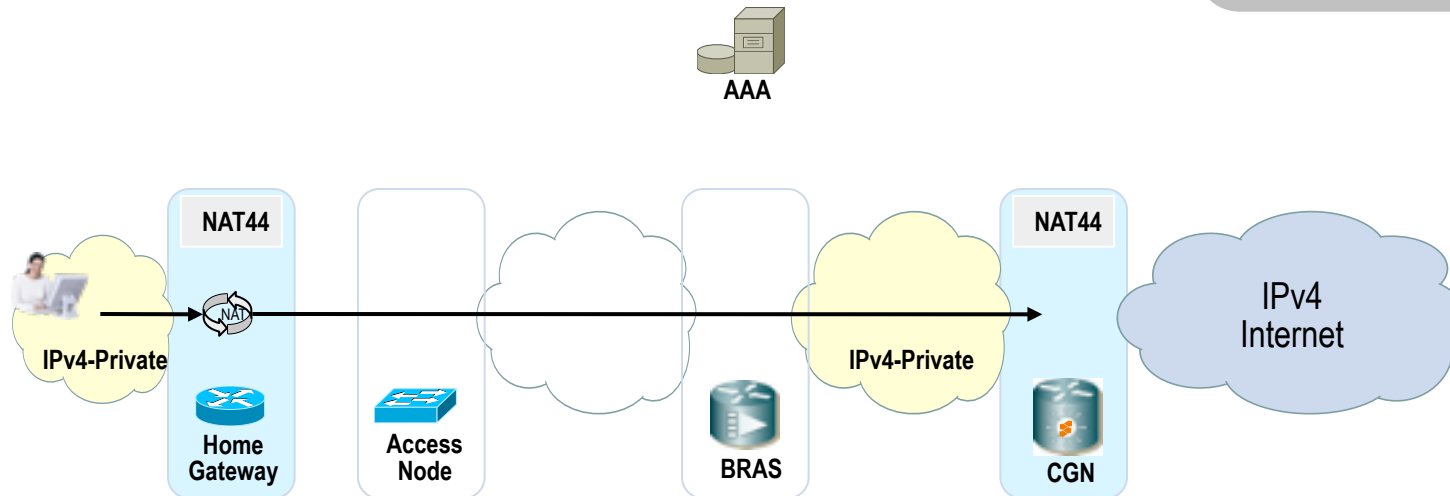


preserve: Carrier Grade NAT44 (CGN44)

- multiple customers multiplexed behind an SP managed NAT device (a Large Scale NAT)
CGN44 multiplexes several customers onto the same public IPv4 address
each customer has unique private IPv4 address

Large Scale NAT44

- O(10G) throughput
- O(10M) bindings
- Transaction logging
- some set of ALGs



- NAT44 can be deployed as centralized or distributed function.
- CPE based NAT44 + CGN44 = NAT444 solution

Carrier Grade NAT44 (CGN44)

SP CGN NAT44

pros

- ISPs can reclaim global IPv4 addresses from customers, replacing with non-routable private addresses and NAT
- addresses immediate IPv4 exhaust problem
- no change to subscriber CPE
- no IPv4 re-addressing in home
- dense utilization of public IP address/port combinations

cons

- SP NAT results in margin & competitive implications
- does not solve address exhaust problem in the long term
- sharing IPv4 addresses could have user behavioral and liability implications
- user control over NAT

preserve: other IPv4 only approaches not involving NAT

- IP-on-demand optimization
 - introduce usage idle-time based mechanism for public IPv4 address assignment. release addresses after idle period. (effectively remove always-on-service for public IPv4)
- IP address trading
 - establish market and regulation registration mechanism
 - enhance anti-IP spoofing/hijacking technology and inter-domain routing



IPv4 subnet trading

- today the cost of getting IPv4 address space is low:
 - Service Provider: RIR membership fee & Registration service fee
 - end-sites usually receive IPv4 address block from SP as part of service
 - many SPs already charge end-site for privilege of public IPv4 address
- when RIRs have no more IPv4 address space to distribute:
 - cost of IPv4 addresses will be higher (today it's close to 0)
 - SPs may “purchase” IPv4 address space from other org's
- current IPv4 address allocation mechanisms were not built to support the dynamic reallocation of subnets
 - facilitating address trading means protecting against address hijacking / false announcements etc.
 - BGP Prefix Validation

IPv4 Subnet Trading

pros


- valuation of IPv4 addresses may hasten IPv6 adoption by encouraging sellers, perhaps more than offsetting costs to move some or all of their network to v6
- receivers of transferred IPv4 address space can prolong their IPv4 networks

cons

- Market may not materialize, so organizations hoping to benefit may not
- depending on region, if RIR doesn't register transfer, there may be no route-ability
- risk to integrity of routing system, as RIRs no longer authoritative for address records. will BGP Prefix Validation be universally deployed in time?
- even more rapid growth of routing system

prepare: IPv6 VPN

enable **enterprise** IPv6 unicast and multicast connectivity across IPv4/MPLS network

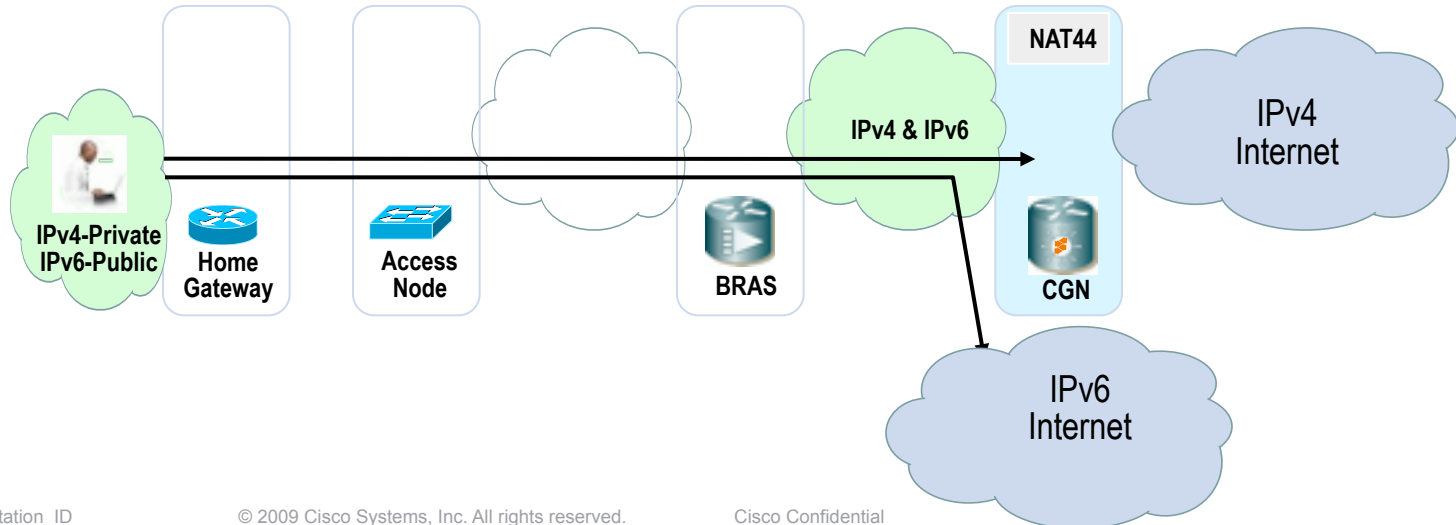


- configure 6vPE routing/forwarding on existing PE routers; core routers stay IPv4/MPLS
- addresses current and future IPv6 VPN route table with business-class SLAs

prepare: native IPv6 and IPv4 dual stack

- classic RFC 4213 solution
 - supported by most end hosts OSes
- in the short term deploying IPv6 in dual stack does not solve IPv4 exhaust ; IPv4 run-out is expected before full deployment

can be easily combined with NAT44 to offer a solution, while allowing IPv6 deployment ramp-up.




native IPv6 and IPv4 dual stack

- **Broadband PPP Access**

dual-stack IPv6 and IPv4 supported over a shared PPP session with v4 and v6 NCPs running as ships in the night.

should not consume extra BRAS session state nor require Access-Node upgrades

note: not all PPPoE clients support IPv6 (eg WinXP)

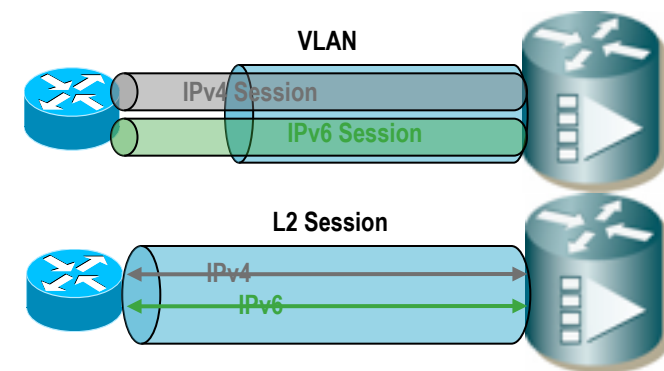


- **Broadband IPoE Access**

form of supporting in “session” form remains to be determined. possibilities include.

- two session model, IPv4 and IPv6 independent sessions.

- L2 session model, IPv4 and IPv6 running on common L2/MAC session



native IPv6 and IPv4 dual stack

pros

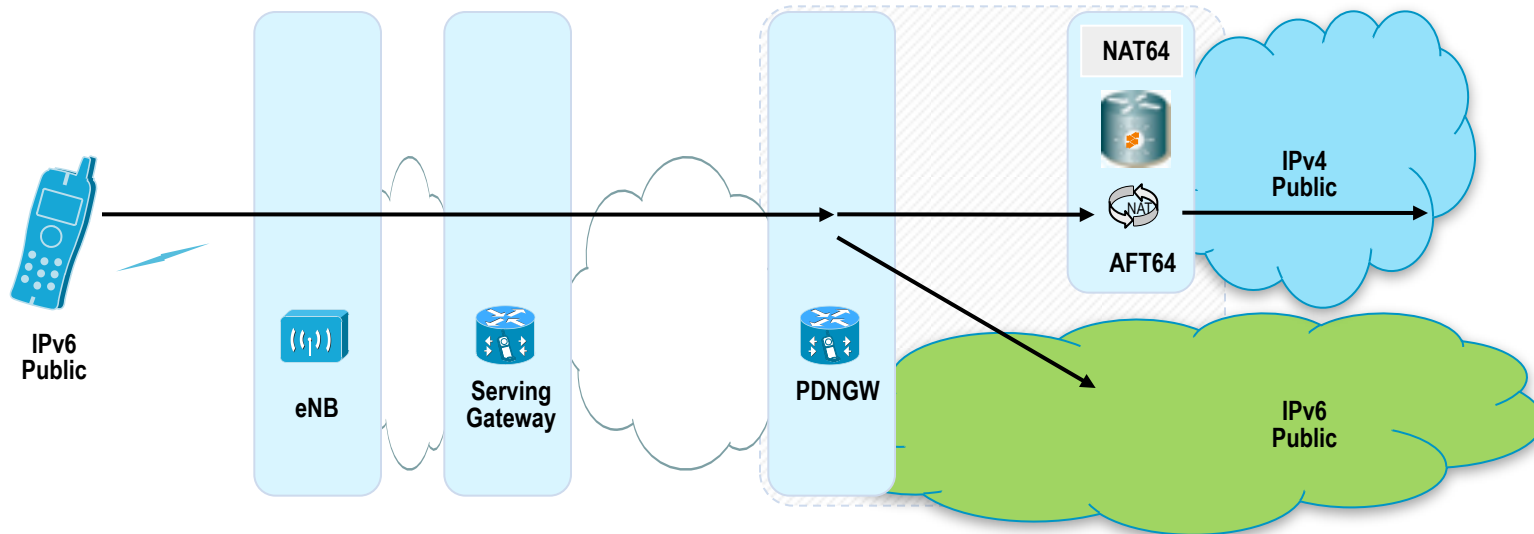
- cost effective long term solution model
- supports legacy (IPv4) applications
- flexible: can be combined with NAT44 deployment for addressing IPv4 exhaustion
- once services are on IPv6, IPv4 can simply be discontinued

cons

- continuing to use public IPv4 doesn't solve IPv4 exhaustion
- IPv6 alongside existing IPv4 infrastructure might cost extra in terms of opex and hardware changes
- in some forms of dual-stack deployments or implementations can lead to decreased network scalability

prepare: native IPv6 and large scale AFT64

- deploy IPv6 only end-points
- use AFT64 to connect IPv6 end-points to Internet and legacy IPv4 user base



- AFT64:= “stateful v6 to v4 translation” or “stateless IV-VI translation a.k.a IVI”

AFT64 translation framework - terminology

- stateful

 - each flow creates state in the translator. supports only IPv6 host initiated communication

 - amount of state based on $O(\# \text{ of translations})$

 - N:1 mappings (like NAPT with NAT44)

- stateless

 - flow DOES NOT create any state in the translator

 - algorithmic operation performed on packet headers

 - 1:1 mappings (one IPv4 address used for each translation to an IPv6 host).

 - for internet access public IPv4 address pool is required for each IPv6 host.

 - supports both IPv6 and IPv4 host initiated communication

native IPv6 and large scale AFT64


pros

- allows IPv6 only clients access to IPv4 content
- IPv6 services and applications offered natively to consumers
- SP network runs IPv6 only, avoiding IPv4 support costs
- stateless technique can be used for IPv4 to IPv6 access

cons

- technical viability of IPv6 only service (IPv6 stack not enabled on all hosts)
- does not address IPv4 customer base
- ALGs required
- DNS infrastructure must be modified to support NAT64
- operations & troubleshooting of transient issues
- stateful NAT has many of the same implications as NAT44

prepare: IP tunneling

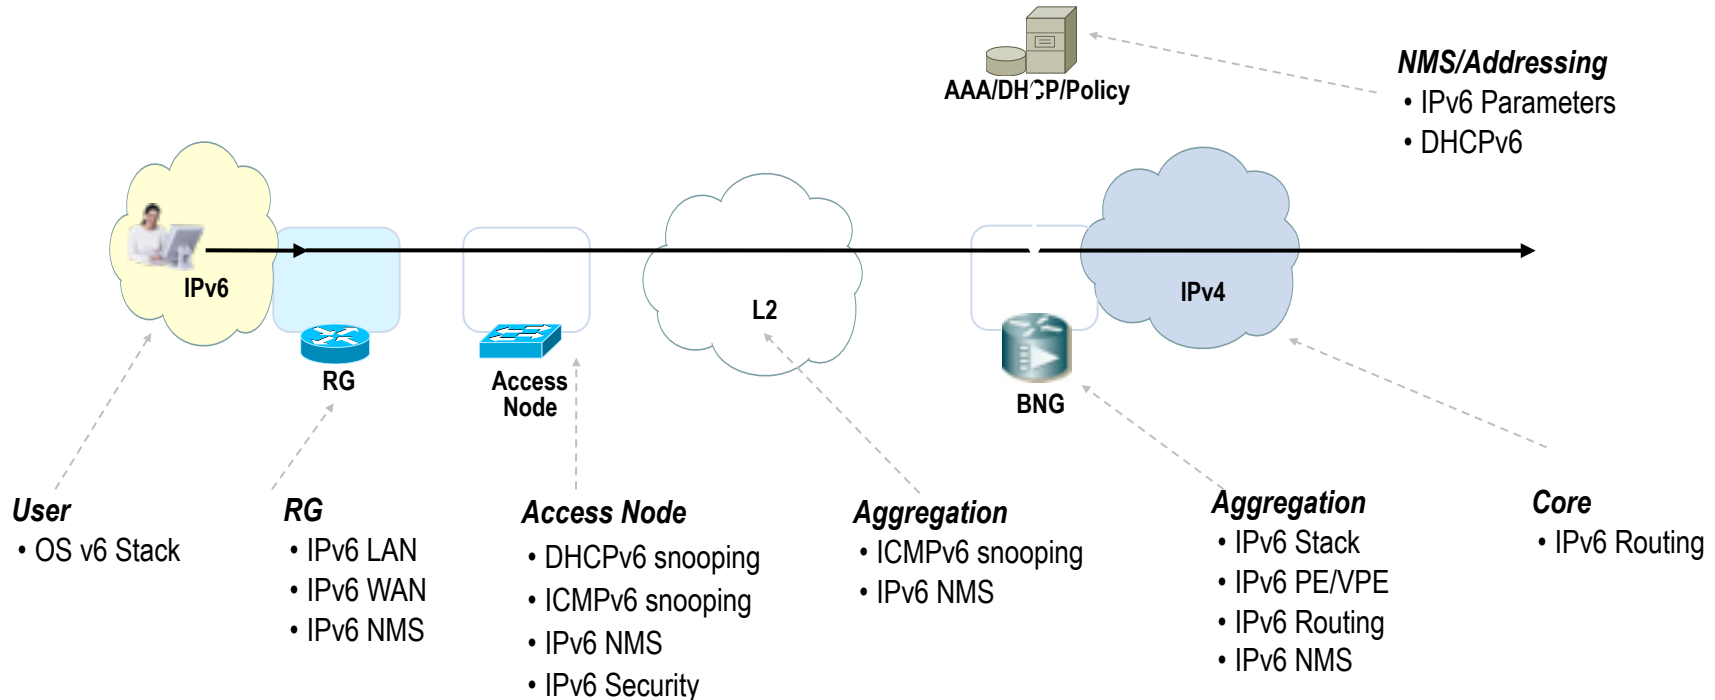


- retains end-end IP semantics
- in theory requires “touching” only tunnel end-points
- in practice, given today’s transition **from IPv4 to IPv6** the different tunneling approaches represent different philosophies:
 - IPv6 Services and IPv6 end-point enablement (IPv6 over IPv4)
 - IPv4 Services and IPv6 as an SP transport (IPv4 over IPv6)

prepare: IP tunneling

IPv6 in IPv4 – why?

- deployment of fully native IPv6 affects numerous system components, aka “touch points”

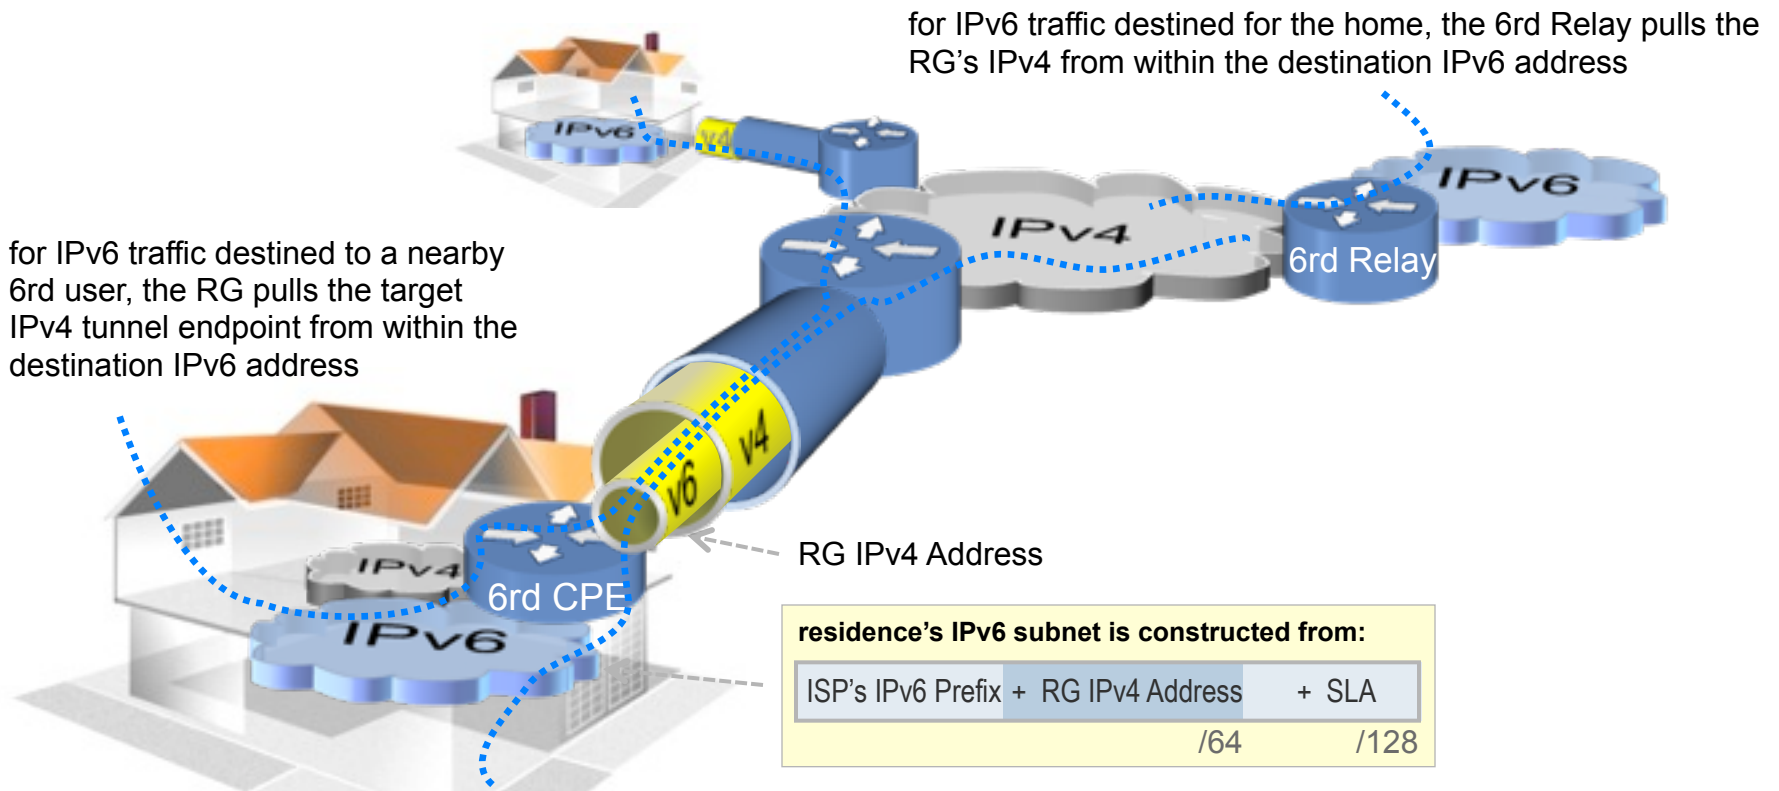


- upgrading some components is more challenging than others:
e.g. IPv6 upgrade of Access Node and/or BNG
- tunneling IPv6 over existing IPv4 infrastructure provides a transition solution which limits the number of “touch points”

Prepare: IP Tunneling

IPv6 Rapid Deployment (6rd) – IPv6 in IPv4

- a form of v6/v4 which traverses v4 aggregation clouds
- <http://tools.ietf.org/html/draft-ietf-softwire-ipv6-6rd-02>



for IPv6 traffic destined to the backbone, the RG uses the destination IPv4 of the 6rd Relay.

tunneling

IPv6 Rapid Deployment (6rd) – IPv6 in IPv4

pros

- enables a v6 service to a routed CPE user
- IPv6 can traverse existing IPv4 infrastructure. no new access CAPEX to enable v6.
- derives IPv6 from IPv4 addresses, eliminating need for much of IPv6 OSS
- efficient local routing of user-user traffic
- stateless = easier to operate
- easily combined with NAT44 to solve IPv4x. decouples NAT44 from IPv6 deployment
- operational models of v4 and v6 are very similar

cons

- continuing to use public IPv4 doesn't solve IPv4 exhaustion. solution may need to be combined with NAT44.
- may require dedicated device in the architecture (BR).

tunneling

dual-stack lite – IPv4 in IPv6

- tunneling IPv4 using IPv6 transport.
- two common options allowed by:

<http://tools.ietf.org/html/draft-ietf-softwire-dual-stack-lite-02>

- dual-stack-lite with NAT44

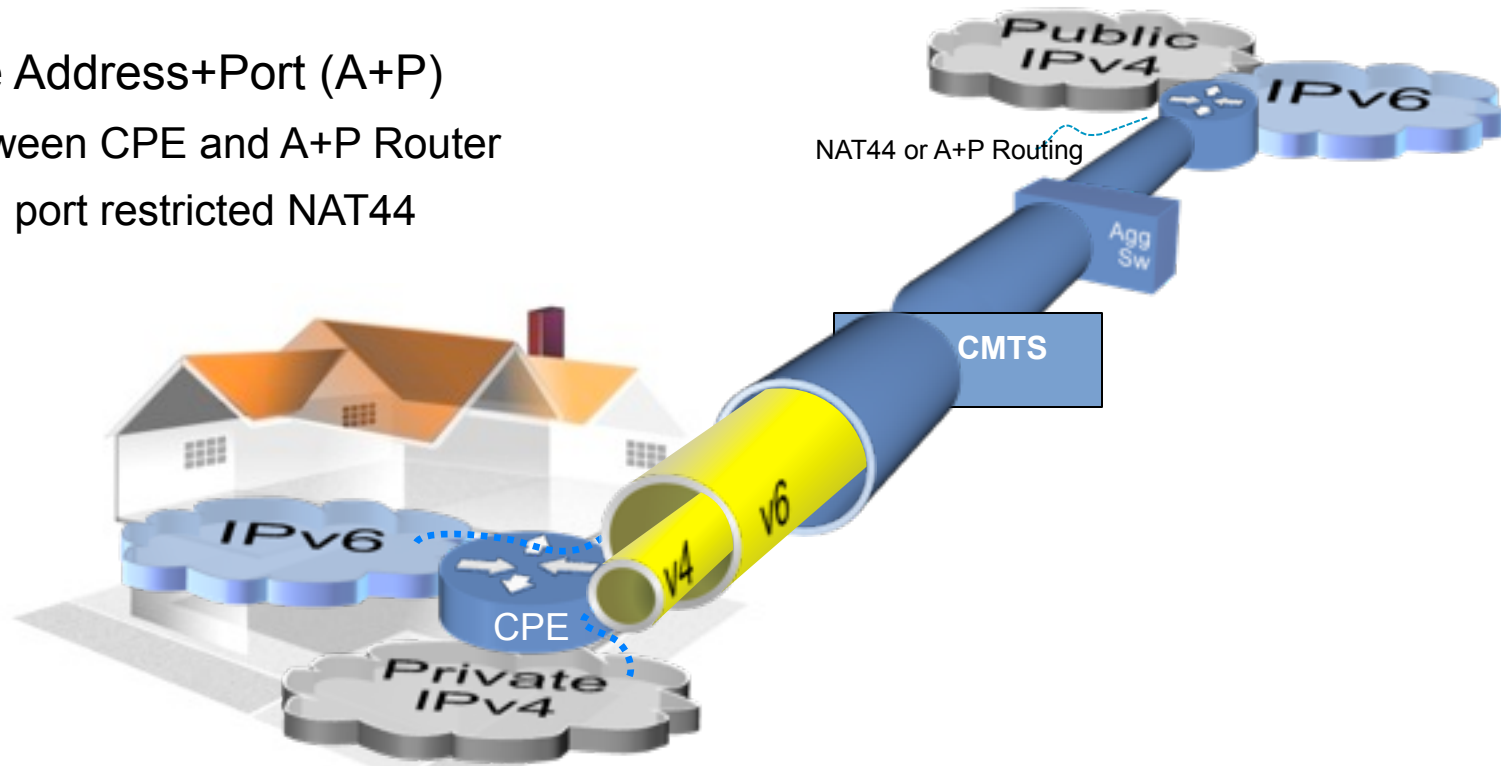
tunnel from CPE is to a LSN NAT44 device.

LSN NAT44 is stateful. No CPE NAT44

- dual-stack lite Address+Port (A+P)

tunnel is between CPE and A+P Router

CPE is doing port restricted NAT44



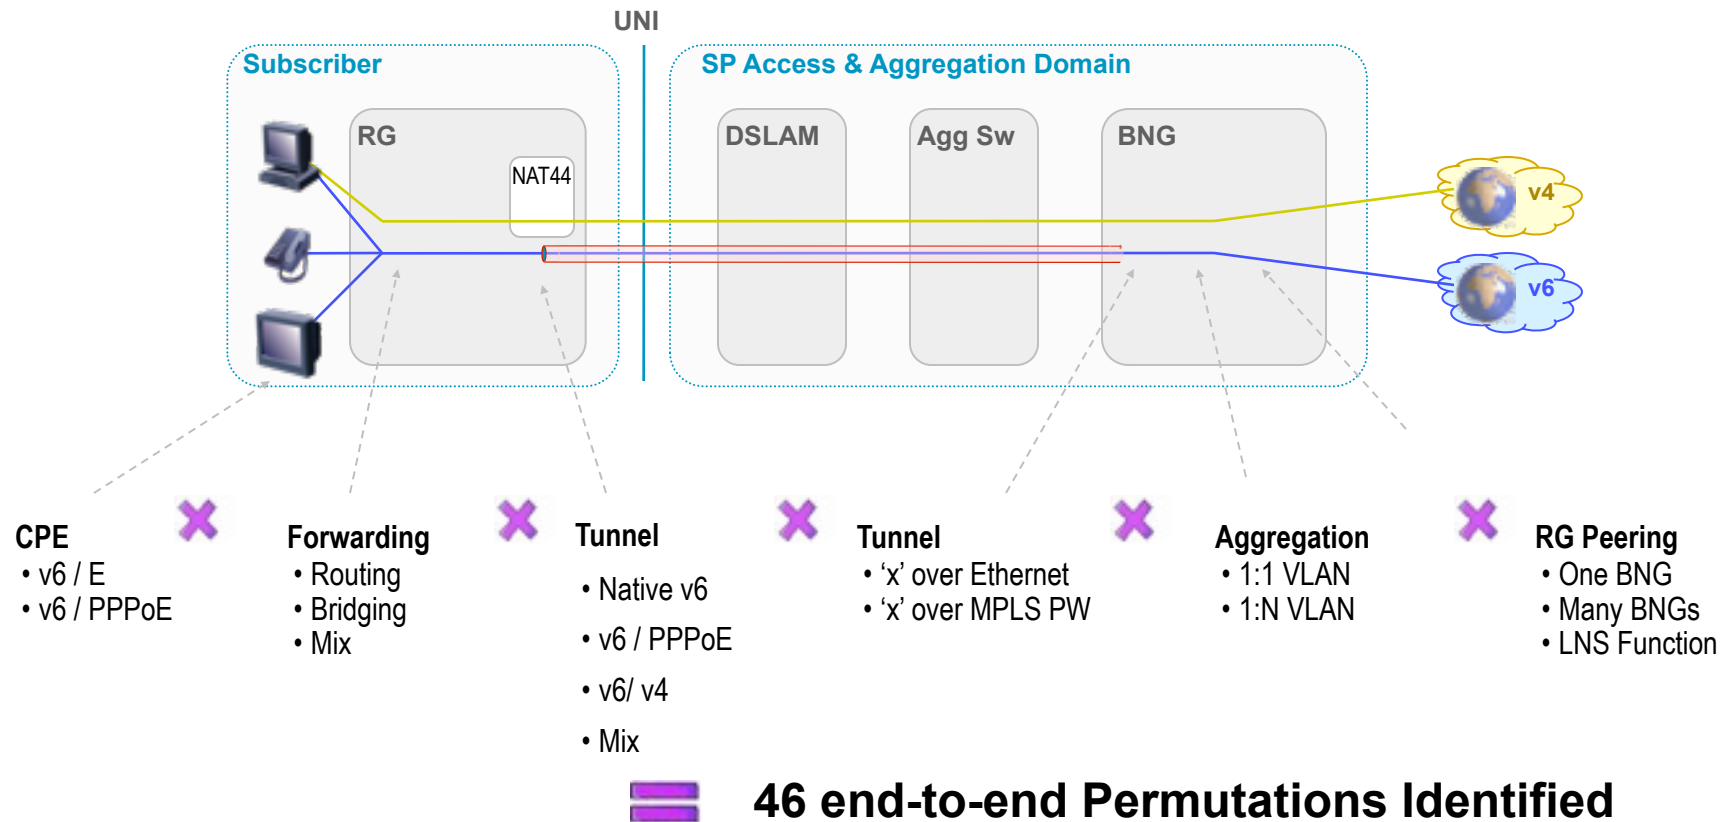


next steps / considerations



v4 + v6 Broadband Residence

many End-to-End Permutations → no convergence yet



broadband direction

- desired end state - fully dual-stack capable network
 - future investment should be oriented around committed dual-stack BNG functionality
 - must support existing BNG feature set as well as support IP_{v4,v6}oE feature sets
 - objective is to minimize non-dual-stack capable platforms within deployed base
- mixture of PPP_{d-s} and IPoE_{d-s} support likely a key requirement given variability in access platforms and CPE

addressing IPv4 only footprint

- ASAP - certify & deploy dual-stack capable CPE w/tunneling capability
- deployment of CGN platforms to terminate v6 tunnels over IPv4 only infrastructure enables incremental rollout over existing IPv4-only infrastructure (use 6RD)
- enables congruency from subscriber policy application
 - i.e.: common QoS and service policy application
- minimizes churn of edge platforms
 - maximizes depreciation cycle

CPE considerations

- functionality timelines aren't necessarily conducive to deployment of D-S CPE - today
- interim requirement to deploy CPE which aren't fully dual-stack capable
 - must insure vendor commitment to necessary IPv6 feature sets to support tunneling and native dual-stack models
 - remote manageability via TR-069 feature sets critical
 - remote code upgrade - critical requirement

access platform considerations

- with the deployment of IPv6 - consideration should be given to the ability of downstream platforms (especially DSLAMs) to support IPv6 features

insure support for IPv6 EtherTypes

insure support for MLD(v2) snooping capabilities if there's a requirement for broadcast video or multicast delivery over IPv6

DHCPv6 L2 relay functionality

IPv4 + IPv6 multicast co-existence - largely unexplored at this point

summary

- long term co-existence (transition) period is very likely
- not doing IPv6 increases the new service adoption risk and existing internet service costs over the long term.
 - IPv6 is a long term must, but not quite a short term solution to IPv4 exhaust.
 - deploying IPv6 incrementally will be key - start now!
- risk of IPv4 exhaust before all services or content move to IPv6 is real, but adoption of IPv6 needs to be done judiciously

start now!

