



Perspectives on IPv6 in the IETF and Broadband Forum




Mark Townsley

Cisco Distinguished Engineer

*With thanks to Fred Baker, Jari Arkko, Dan Wing,
Wojciech Dec, Eliot Lear, and many others...*

Power Throughout the World



NASA "Earth at Night" Visualization Date Oct 23, 2000

Number of Documents at Various Status

Documents about IPv4/IPv6

7 Document Statuses

- Standards

 - Best Current Practice (146)

 - Proposed Standard (1450)

 - Draft Standard (91)

 - Full Standard (77)

- Nonstandards

 - Historic/Obsolete/Just Plain Old (1724)

 - Informational (1510)

 - Experimental (255)

IETF Status	IPv4	IPv6
Informational	933	374
Experimental	151	59
Best Current Practice	86	34
Proposed Standard	772	407
Draft Standard	48	17
Full Standard	48	5

General Areas IETF Is Working on IPv6

- Cross-Registry Information Service
- Addressing
- Dynamic Host Configuration
- Autoconfiguration
- IP over Various Technologies
- Mobility
- Multihoming
- IPv6 Maintenance
- IPv6 Operations

- Translation-Based Transition Technologies
 - IPv4/IPv6 NAT
 - IPv6/IPv6 NAT
- Tunnel-Based Transition Technologies
- Source Address Validation
- Routing
 - Especially Mobile Ad-Hoc Routing
 - Also Global Routing Operations
- Sensor Networks

The Goal...

- In general, the goal is to “Continue the growth of the Internet”:
- For some, that means “**retain simplicity** by extending addressing to more prefixes and more machines.”
- For others, that means “**retain the infrastructure** I am familiar with and have invested heavily in.”
- For us, the goal is:

Continue the growth of the Internet with maximized application options and minimized long-term operational and capital cost.

- That implies:

Deploy IPv6 for more addresses.

IPv4/IPv6 coexistence is required for a turn-up period.

At some point, IPv4 is no longer needed.

At that point, **turn IPv4 off**.

IP Addresses Throughout the World

Internet Map
Connection Density



ChrisHarrison.net

DIMES Data Feb 2007

Changing Conversations

- Used to be:
 - “Do we really need IPv6?”
 - “What’s the Business Case?”
- Now I hear:
 - “I need server load-balancing... better MIB support...”
 - “Should I use ULAs? What firewall policy should I set?”

**“Why should I deploy IPv6?” is slowly being drowned out by
“How do I deploy IPv6”**

IETF Discussions on Transition Plans




Unworkable Approach to Transition: Expect IPv4 and IPv6 to Directly Interwork

■ Problem:

We are running out of IPv4 addresses.

IPv4 and IPv6 are noninteroperable.

If we simply deploy IPv6 networks, they won't be able to talk with IPv4-only hosts.



Preferred Approach to Transition: RFC 4213 Dual-Stack Deployment

■ **Solution:**

Hosts today are IPv4+IPv6:

Windows Vista, Macintosh,
Linux, BSD

Make the network IPv4+IPv6.


When forced to deploy IPv6-only
networks, they will be able to talk
with other hosts.

■ **But...**



We have run out of time for this to
be smooth

In the mean time, we forgot how to
operate multiprotocol networks

IPv4+IPv6 Hosts




IPv4+IPv6
Network



IPv6-only
Hosts or Network


IETF Softwires Working Group (Tunnels)

- Connecting islands of one address family over islands of another
- “Hubs and Spokes” and “Mesh” solutions
- Dual-stack Lite



IETF Behave Working Group (Translators)

- IPv4/IPv6 Translation
- Temporary tool to help coexistence and transition
- IPv4 Addresses
 - May be embedded in an IPv6 prefix in the IPv6 domain
 - Stateless and stateful translation modes
- Connectivity Provided:
 1. An IPv6 network to IPv4 Internet
 2. IPv6 Internet to an IPv4 network
 3. An IPv4 network to IPv6 Internet
 4. IPv4 Internet to an IPv6 network



Work on IPv6 continues...

- IPv6 "Maintenance"

- 6MAN - Fixing bugs in current IPv6 specifications (improvements on address selection - RFC 3484)

- V6OPS - Operational guidance and more... address allocation, firewalls, broadband network deployments, IPv6 home router, etc...

- IPv6 "New features"

- SAVI - "IP source guard" or "First Hop Security" for IPv6 and IPv4

- CSI - Secure Neighbor Discovery (SEND; IPv6-only)

IPv6 Address Independence (6AI) BoF

- IPv6 Network *Address* Translation and Network *Prefix* Translation

Comfortable unscalable algorithms vs. newish scalable algorithms

Renumbering, Multihoming, “Simple Security”, “Topology Hiding”

...

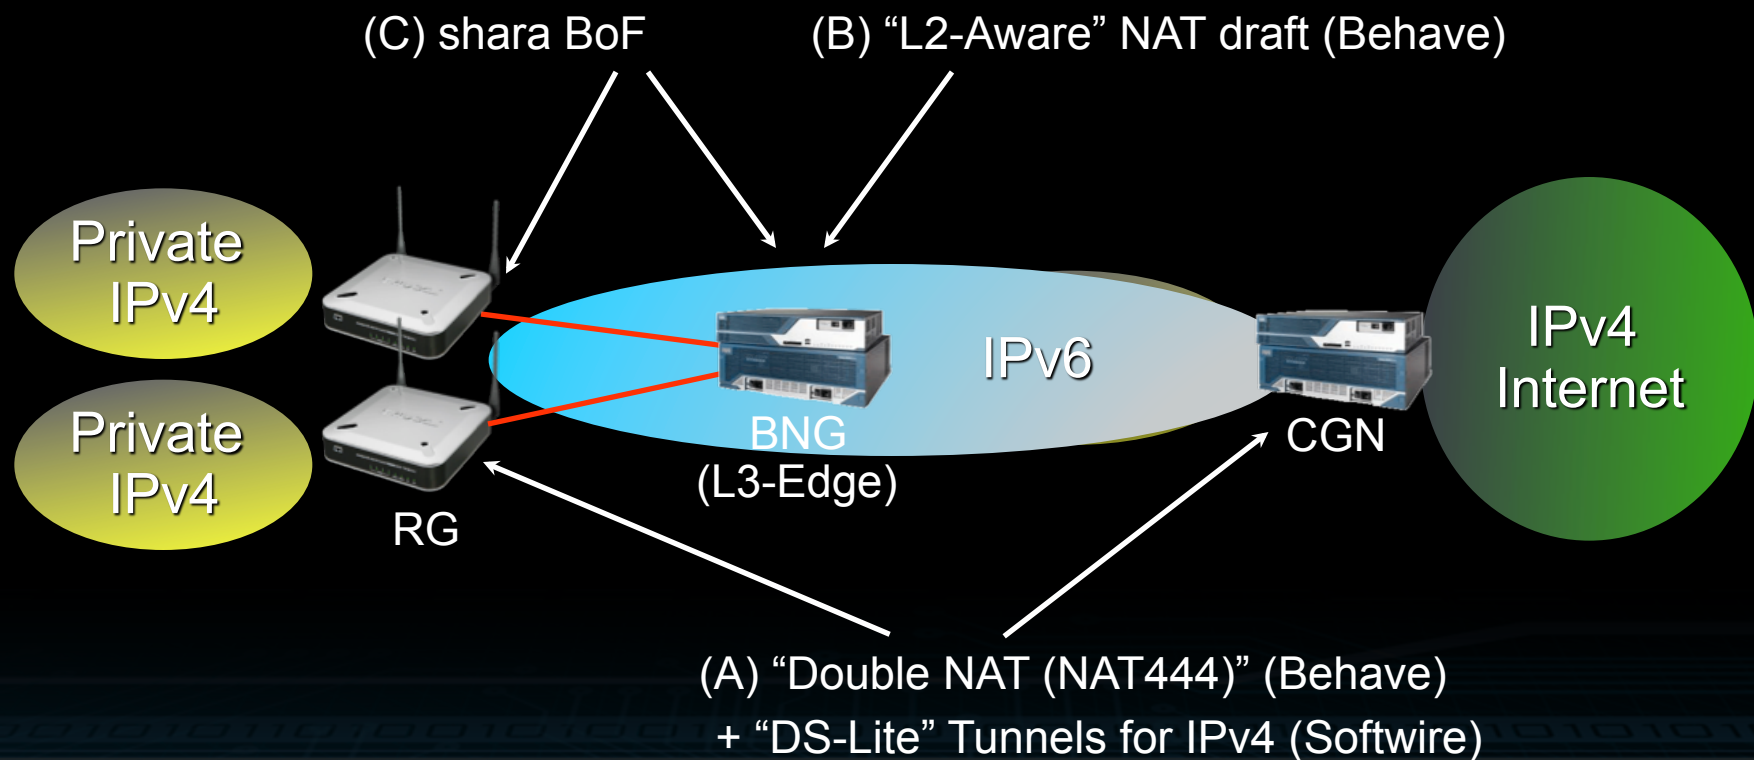
- Issues:

Fear that if these boxes are going to exist, we'd rather have a specification than leave programmers to their creative vices.


Address management, scaling of routing

Business..Business communications

NAT444, Sharing IPv4 Addresses, and a Shared SP Prefix...



Broadband Subscriber Growth and “NAT444”



* source: Infonetics 2008

IP NAT and the Applications and Transports that run on top

- Back in the dialup days, each PC had, even if only temporarily, a global IPv4 address
- When NATs came along, **some applications stopped working** (VPNs, various games, irc, FTP, etc.)

THEY EVOLVED.

The ICE 9-Step Program to Recovery


- Step 1: Allocation
- Step 2: Prioritization
- Step 3: Initiation
- Step 4: Allocation
- Step 5: Information
- Step 6: Verification
- Step 7: Coordination
- Step 8: Communication
- Step 9: Confirmation

ICE Step 6: Verification

- Each agent pairs up its candidates (local) with its peers (remote) to form candidate pairs
- Each agent sends a connectivity check, in pair priority


request from the local to the remote

If the request the generates a response mapped address ie source IP and port request




Signaling Completion

- When controlling agent is done, it inserts a flag into a STUN check
- If passive agent had successfully completed a check in reverse direction, it stops checks for that component of that stream
- Both agents use the pair generated by the check that included the flag
- When 'done' – ring the phone!



ICE Step 1: Allocation

- Before Making a Call, the Client Gathers Candidates
- Each candidate is a potential address for receiving media
- Three different types of candidates



ICE Step 2: Prioritization


$$\text{priority} = (2^{24}) * (\text{type preference}) + (2^8) * (\text{local preference}) + (2^0) * (256 - \text{component ID})$$

Type Preference	Local Preference	Component ID	32 bits
-----------------	------------------	--------------	---------

- Type-Preference: Preference for type (host, server reflexive, relayed)
- Usually 0 for relayed, 126 for host


ICE Step 4: Allocation

- Called party does exactly same processing as caller and obtains its




ICE Step 5: Information

- Caller sends a provisional response containing its SDP with candidates and priorities
- Can also happen in 2xx, but this flow is "best"
- Provisional response is periodically retransmitted
- As with INVITE, no processing by proxies
- Phone has still not rung yet



Peer Reflexive Candidates

- Connectivity checks can produce additional candidates
- Peer reflexive candidates
- Typically happens when there is a symmetric NAT between users
- Peer reflexive candidate will be discovered by both users
- For user A, from the Response
- For user B, from the Request
- Allows direct media even in the presence of symmetric NAT!



Pairing up Candidates

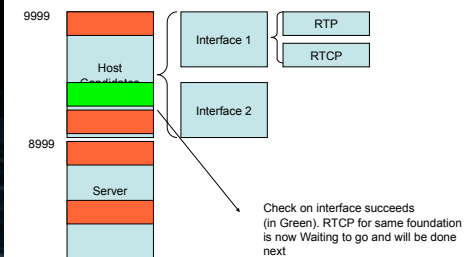
O-P: Offers Priority
A-P: Answers Priority

$$\text{pair priority} = 2^{32} * \text{MIN}(O-P-A-P) + 2 * \text{MAX}(O-P-A-P) + (O-P-A-P \div 1.0)$$

Minimum Priority	Maximum Priority	64 bits
------------------	------------------	---------

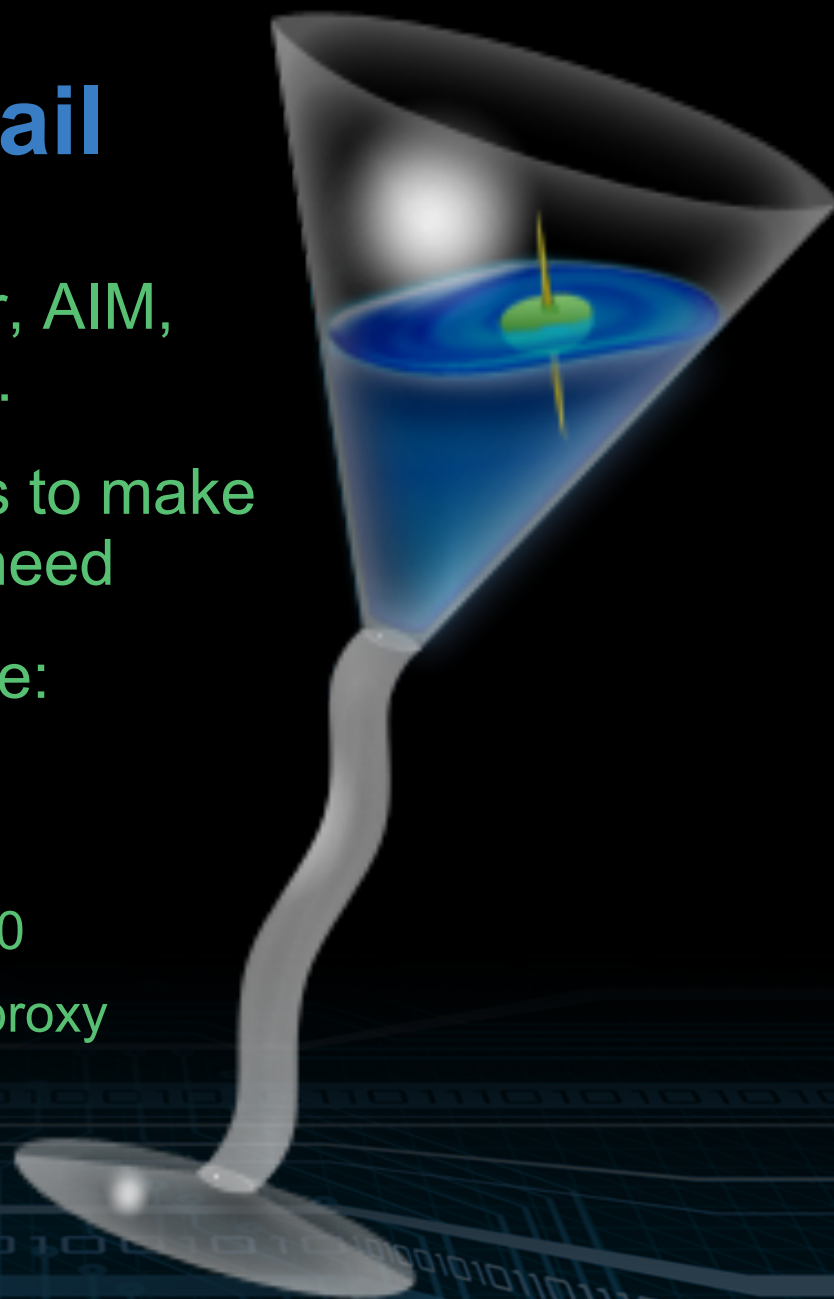
- Pairs are sorted in order of decreasing pair priority
- Each agent will end up with the same list
- Last term serves as a tie breaker
- Min/Max results in highest priority for pair with two host RTP candidates, lowest for pair with two relayed RTP

visualizing Frozen Algorithm



Connectivity Cocktail

- ICE, Skype, Bittorrent, Jabber, AIM, Video Games, Slingbox, etc...
- All use a variety of techniques to make the kind of connections they need
- Active Ingredients may Include:
 - IGD/UPnP, NAT-PMP, RSIP
 - STUN, TURN
 - Tunneling over UDP, TCP, Port 80
 - Proprietary rendezvous servers, proxy servers, etc



When NAT444 Comes...


- Applications that break, will evolve again
- A new cocktail will be formed, with or without IPv6 as an ingredient
- Once its done, its done.
- Break the cycle before it starts...

If you walk away with one message with respect to IPv4 Address Exhaustion:

Do not provide a natted Private IPv4 Address service to your subscribers without also providing a Global IPv6 Address service *at the same time**.

**Or before (this goes without saying....)*

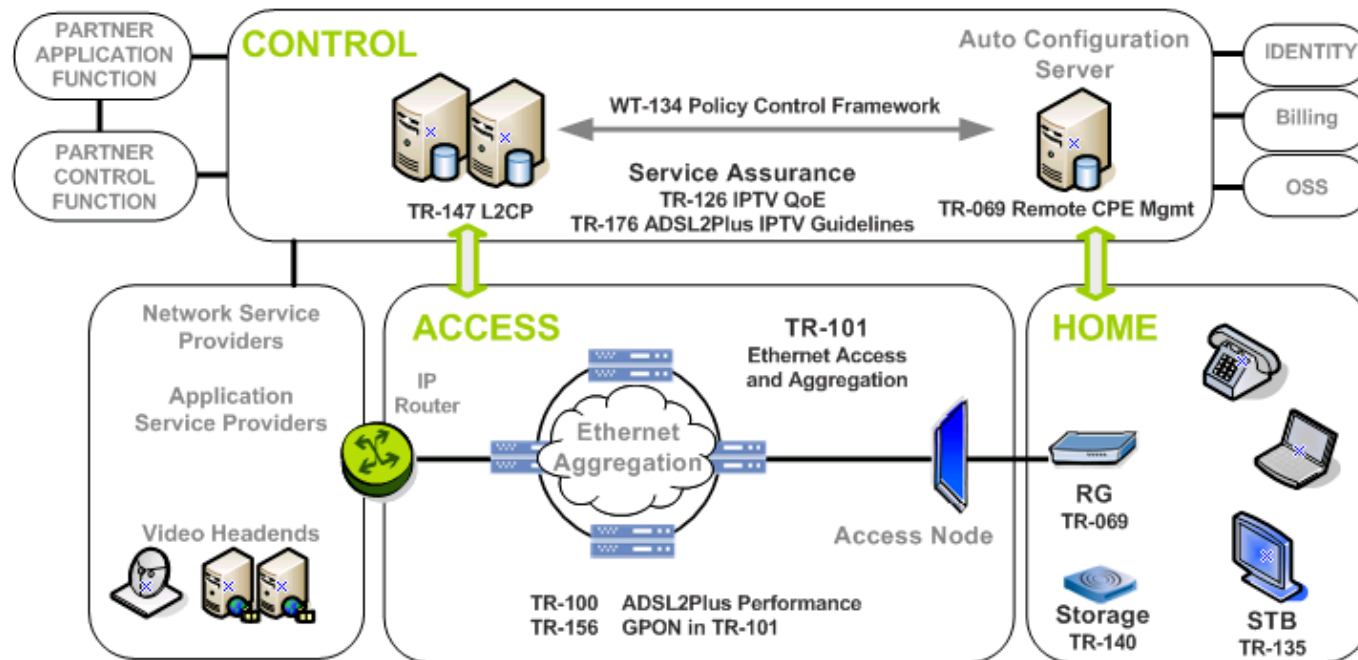
IPv6 @ BBF: what works & what doesn't



Broadband Forum Scope

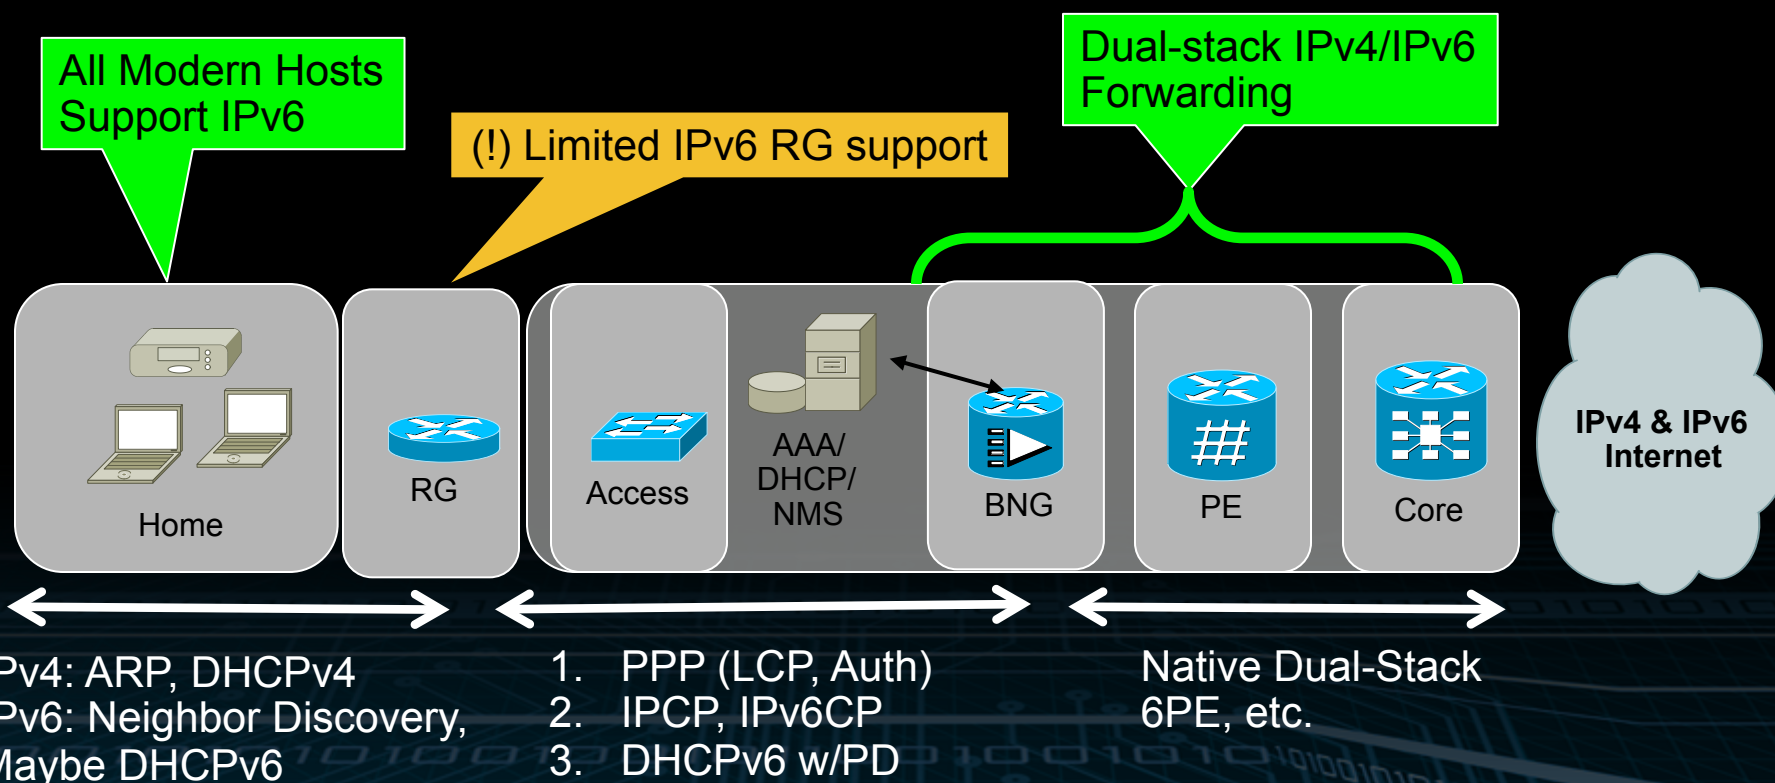
End-to-end Broadband Transport, from Core to Home with associated :

- Control
- Management
- Operational procedures
- Equipment functional requirements
- Interoperability



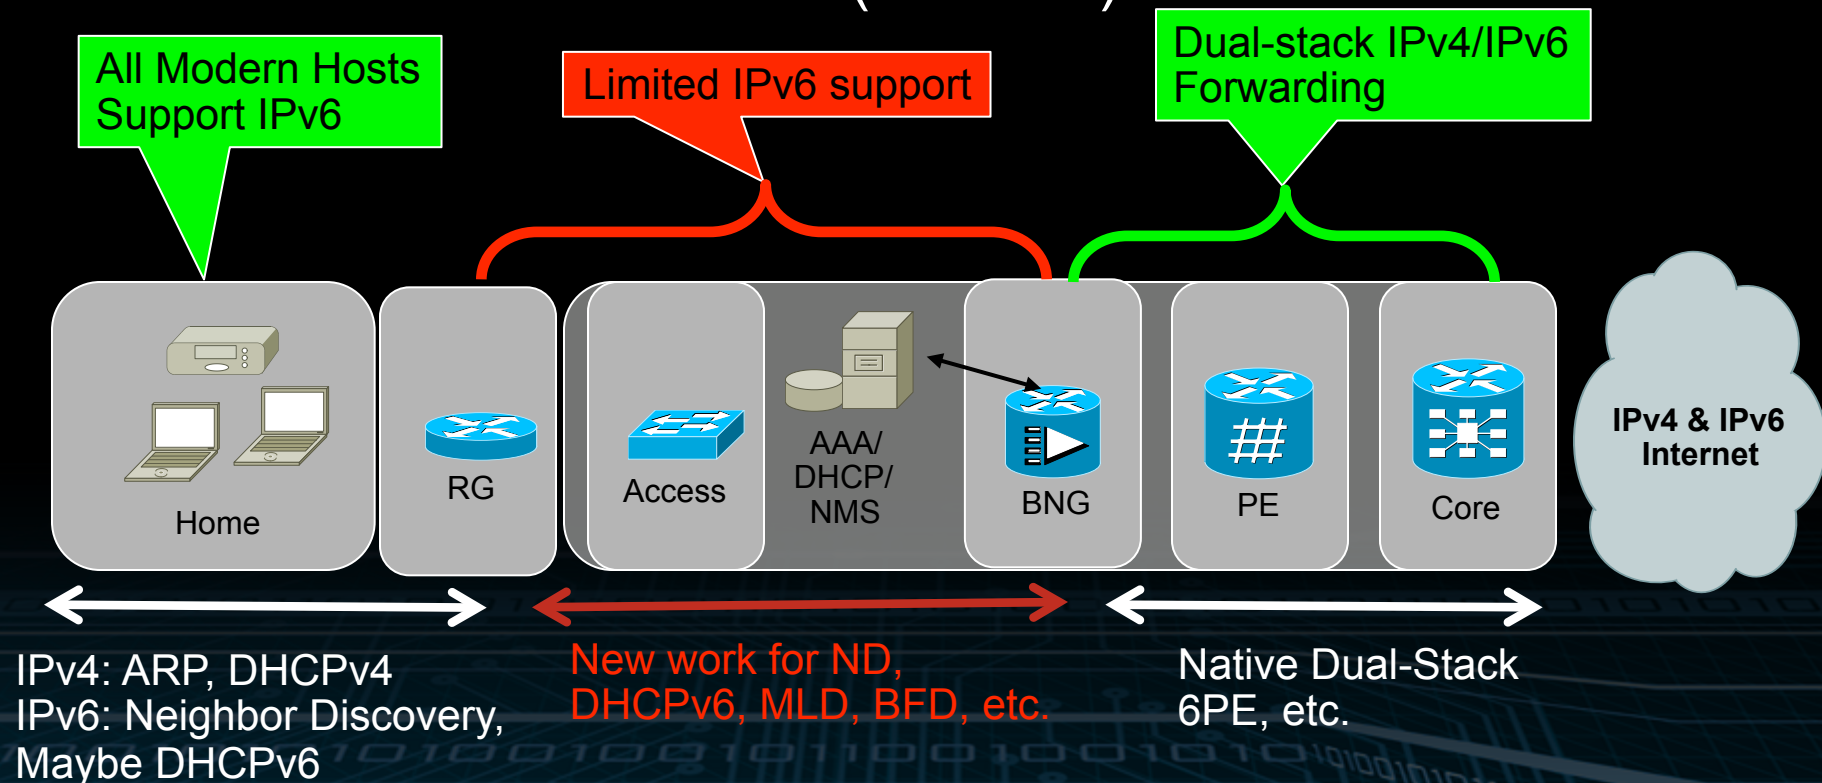
What Works: DSL with PPPoE or PPPoA

- Access network pure L2
- Bridging or Routed RG
- IPv4-centric management, billing, troubleshooting, etc.

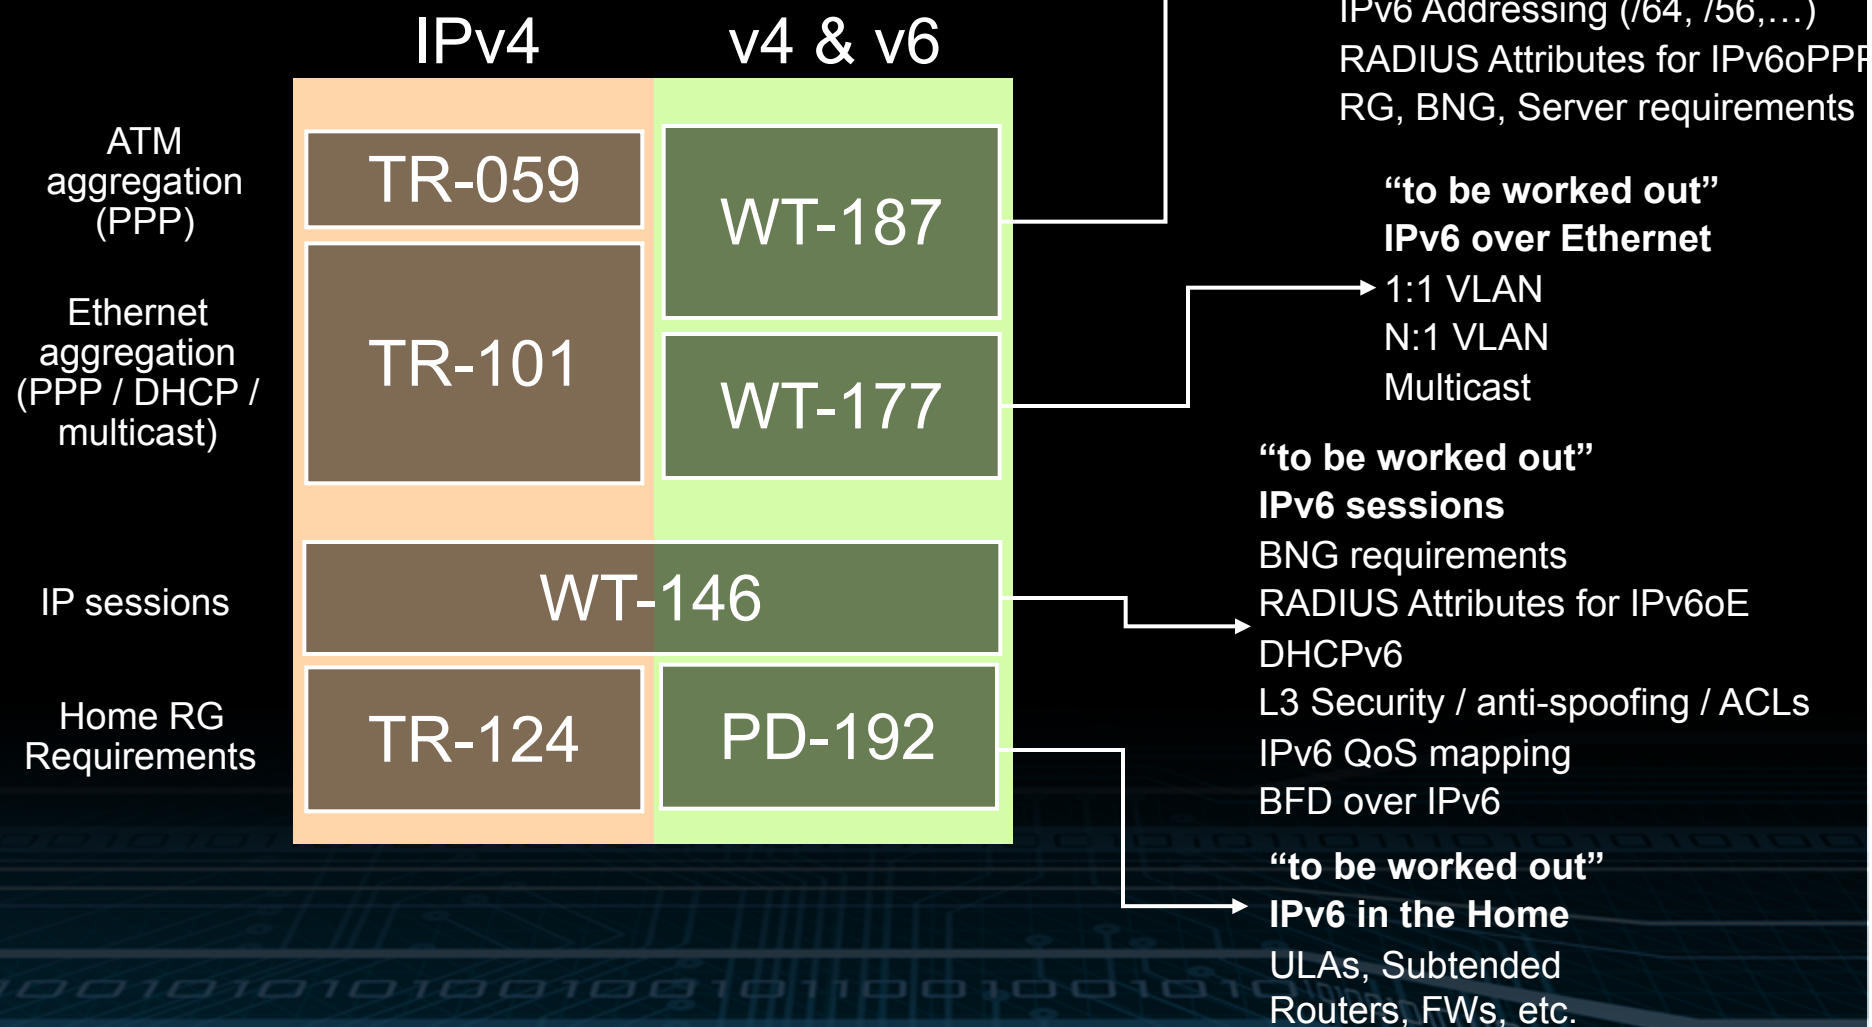


What Doesn't Work So Well...

- DSLAMs with “IP-specific” features (TR-101)
 - DHCP snooping for Line ID and other parameters
 - IGMP snooping, MAC-Forced Forwarding, “MAC-NAT”, etc
- “IP Sessions” at the BNG (WT-146)



BBF IPv6 Work Plan (Launched March 2008)



Recent IPv6 Success Stories...

An inflection point?

- Google over IPv6

- Dec 5, 2007 – Challenged to deploy IPv6 by IETF 73

- Jan 2008 – First production IPv6 router

- Oct 2008 – First “trusted tester” receives AAAA for www.google.com

- Nov 16, 2008 – Challenge met at IETF 73

- Free Telecom

- Nov 7, 2007 – “6rd” presented, decided to deploy


- Dec 11, 2007 - “Opt-in” service made available to 3M subscribers, 250K sign up right away

- March 2008 – Deployed “telesite” IPv6-only service to all 3M subscribers

How did they do it?

1. Turned on what they could
2. Found out what was broken and what was not
3. Filled in the gaps

'Yearly' Graph (1 Day Average)



	Max	Average	Current
In	233.9 Mb/s (2.3%)	72.6 Mb/s (0.7%)	199.7 Mb/s (2.0%)
Out	19.3 Mb/s (0.2%)	5713.0 kb/s (0.1%)	5593.4 kb/s (0.1%)

Conclusion...

- IPv6 is ready for deployment, and there is a great deal of implementation to build upon
- It won't be perfect, system-level gaps exist, but the best way for us to find them at this stage is to look at real deployments
- IPv4 Exhaustion tools are being built too, use them where you must, but not without IPv6 alongside
- Bring your experience back to the IETF, and help us help others
- Thank you, and have a great RIPE week.