IPv6 ADOPTION
APPROACHES TO LSNAT AND IPv6

SANOG

Alastair JOHNSON
alastair.johnson@alcatel-lucent.com

September 2011

Alcatel·Lucent

LARGE SCALE NAT
DISCLAIMER

- This is not a pro-LSNAT talk!
- NAT and Large Scale NAT are not the very best of ideas
  - It hurts you the operator
  - It hurts your customers
  - It hurts the Internet as a whole

- But it is an unfortunate necessary evil for some environments
- It should **never** be seen as a replacement for IPv6 deployment

IPv6 DEPLOYMENT SUMMARY

SANOG

Deployment **has** happened

- Many (but not all) backbones are dual stacked
- Wireline access networks remain a challenge
  - Architecture issues
  - Product support issues
- Wireless access networks are evolving (2G => 3G => 4G), and with it is coming IPv6
- Content providers are out there with IPv6 today
- However, traffic remains low – why?

IPv6 DEPLOYMENT CHALLENGES

SANOG

- The typical challenge for IPv6 adoption is around the access networks
  - CPE is a huge - and expensive – challenge
  - Architectural challenges
    - DSL – N:1 VLANs vs. 1:1 VLANs
    - DOCSIS – no support in non-standard environments/pre-DOCSIS-2.0
    - Etc…
- The other major challenge is in the home network
  - CPE again
  - Home devices – PCs, game consoles, cellphones, etc
- However, **customer bases and traffic continues to grow**

IPv6 DEPLOYMENT
IPv4 EXHAUSTION

- IPv4 and IPv6 are not interoperable
- Networks/hosts are not upgraded to IPv6 simultaneously – IPv4 and IPv6 endpoints will be required to communicate
- Requirement for mechanisms to support communication between IPv4 and IPv6 endpoints
- As networks continue to grow, customers are going to continue to need IPv4 resources
- As hard as we've tried, IPv6-only operation is still a while away

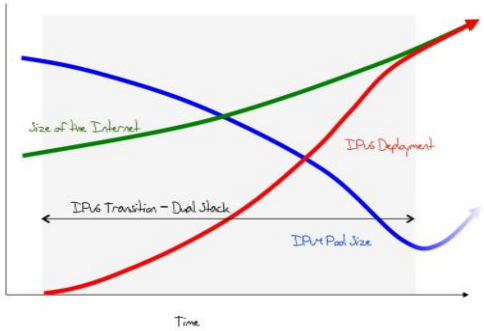

The plan

The reality

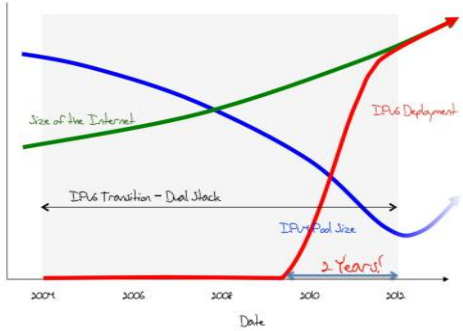

Graphs from Geoff Huston
http://potaroo.net

METHODS OF TRANSITION

SANOG

Home device	Access network	Destination	Solutions
IPv4	IPv4	IPv4 Internet	Large Scale NAT
IPv4	IPv6	IPv4 Internet	Dual-Stack Lite NAT464
IPv6	IPv6	IPv4 Internet	NAT64 Stateful NAT64 Stateless/IVI
IPv6	IPv4	IPv6 Internet	6to4 6RD
IPv6	IPv6	IPv6 Internet	Dual-Stack

SCENARIO

DUAL-STACK THROUGH THE WIRELINE NETWORK

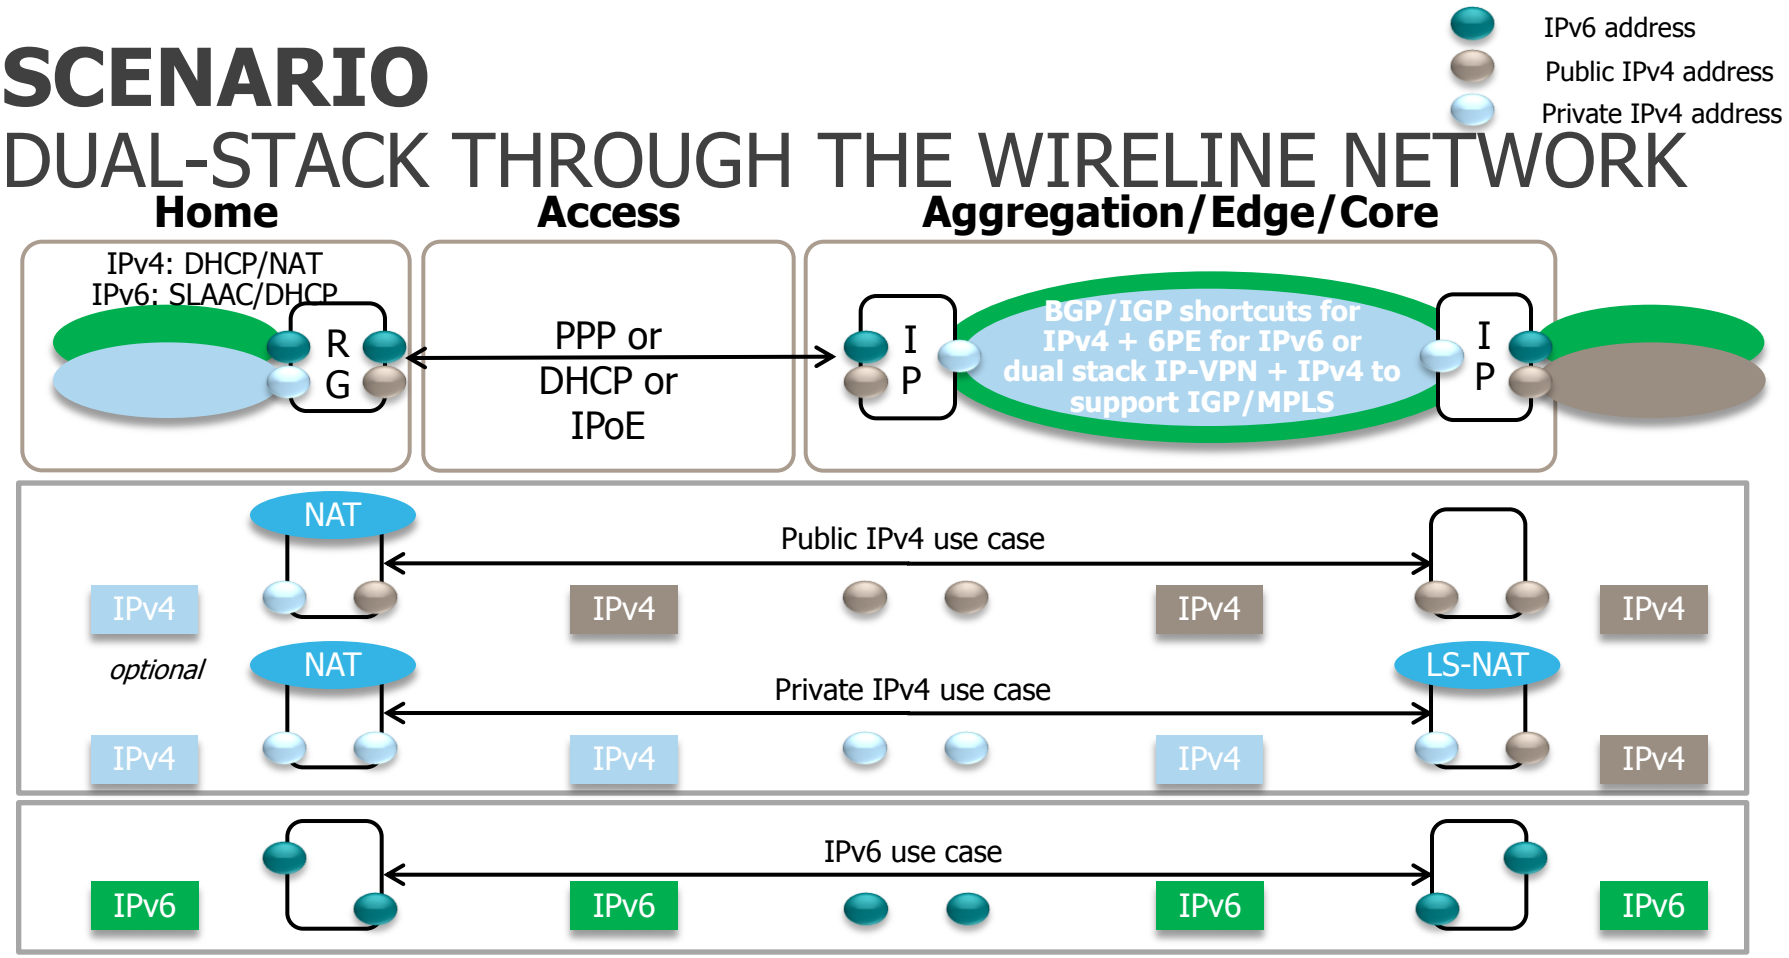

SCENARIO
IPv6-ONLY OPERATION

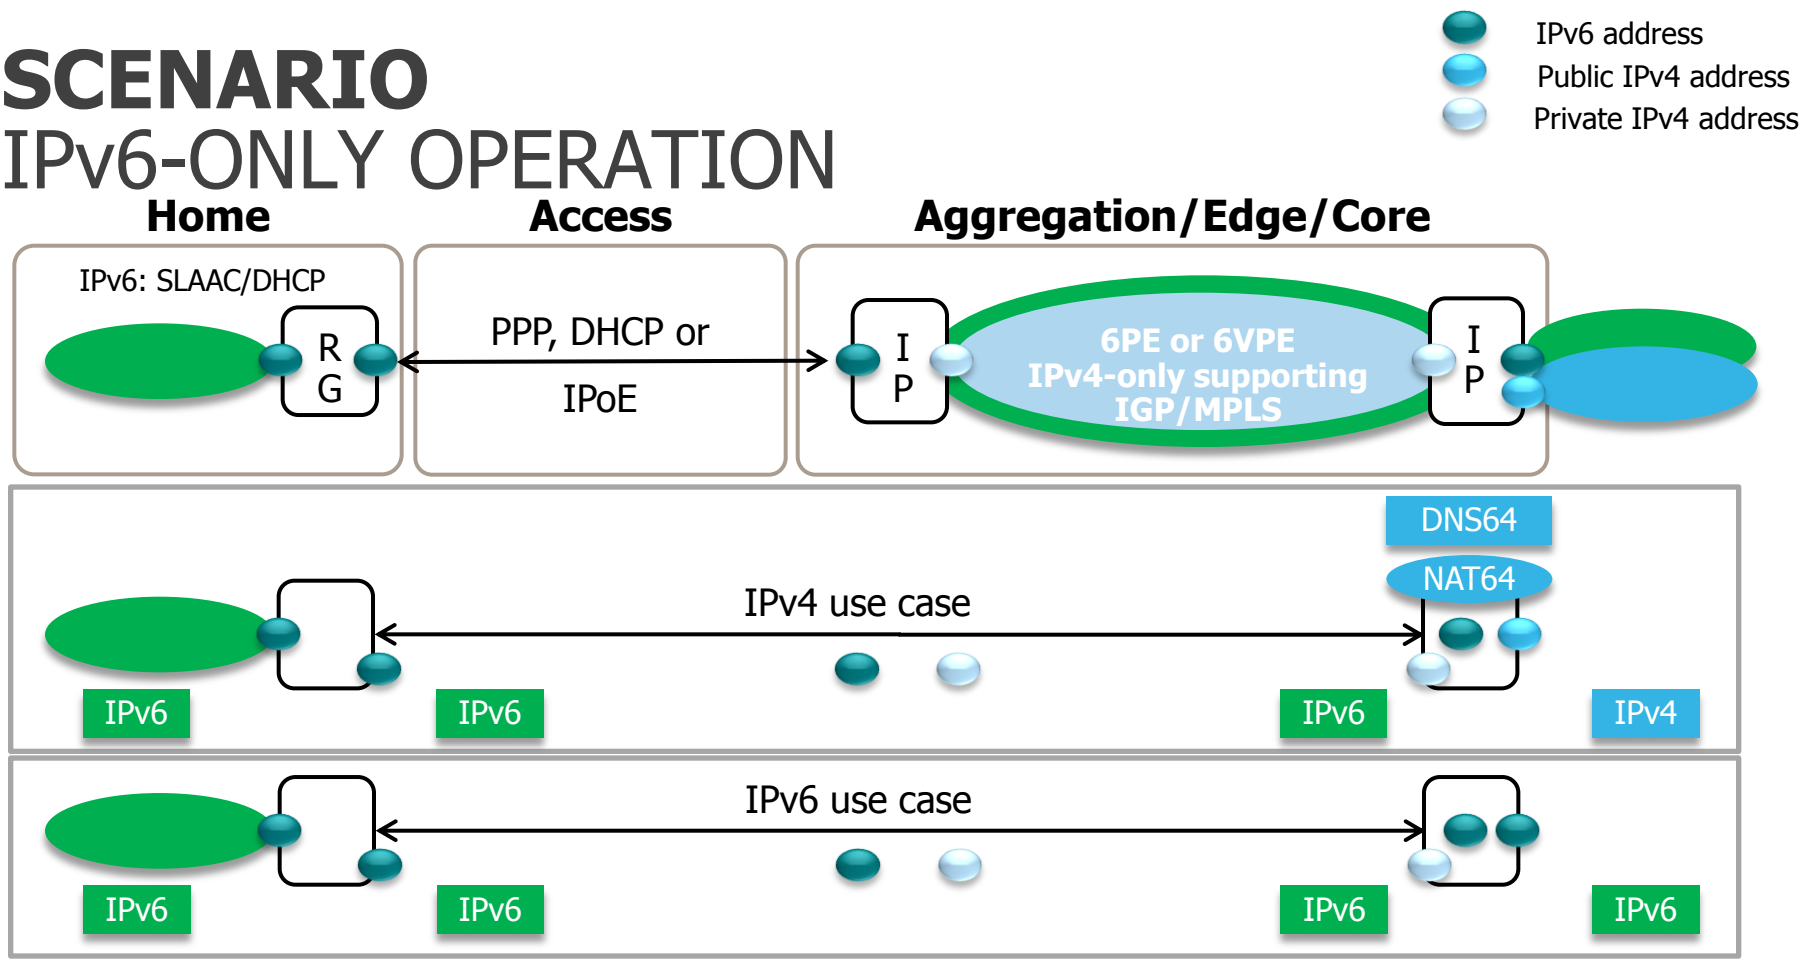

IPv6 DEPLOYMENT OBSERVATIONS

- IPv6 deployment is accelerating
  - February 2011: IANA run-out
  - June 2011: World IPv6 day
  - The chicken-egg scenario is being resolved
    - Content is available!
- Most operators are favoring native dual-stack where possible
  - Tunnels are still used in difficult network domains
  - Tunnels are still often (☹) used for transit connectivity
- Interim transition and rapid deployment protocols are being deployed, or at least evaluated
  - 6rd
  - DS-Lite

LARGE SCALE NAT
ABUSE MANAGEMENT

- LSNAT makes it hard to identify your subscribers for abuse tracking
  - Port Ranges / Port Blocks are a way to manage this
    - Subscriber_1234 => 192.0.2.1:[2000-2999]
  - Persistent external address mapping is used to map a subscriber to the same external address
- Collateral damage from abuse
  - Persistent external address mapping helps, but still means N subscribers may be blocked (/32 level).
  - Use a sane subscriber:IP mapping. 1000:1 is not a great idea!
  - Trying to put as much identification info into the flows as possible (ALGs?)
- Source-port logging
  - Anyone who depends on log data for abuse management needs to include source port data
  - RFC6302
- Lawful Intercept

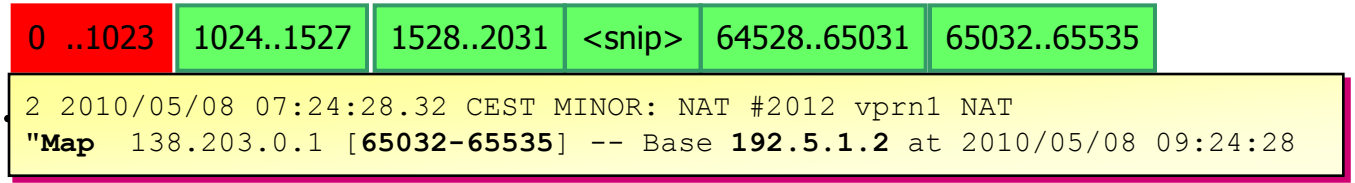

LARGE SCALE NAT
PERFORMANCE AND PORT ALLOCATIONS

SANOG

- Each host is guaranteed some number of external ports
  - Multiple port blocks can be assigned to a host, preferring to stick to the same external IP
- When the ports are exhausted, return ICMP type 3 code 13
- Allow for “burst ports” or “reserved ports”
  - Prevent subscribers overwhelming their port allocation and DNS breaking
  - Based on IP criteria, DSCP, etc…
- ‘Google maps will break’
  - We haven’t observed this, particularly if you handle port exhaustion gracefully (ICMP unreachables) **and** use sane port allocations

LARGE SCALE NAT CONNECTIVITY

- ALGs
  - You'll need ALGs for all your favorite protocols, SIP/FTP/PPTP/IPSEC
  - Potential to reinvent the home NAT brokenness but hopefully we've learned, or simply do not NAT these customers
- Port forwarding
  - Some apps will require either static port-forwards configured in the NAT
  - Non-scalable
- Port Control Protocol (PCP)
  - Provides an interface for hosts to request how a NAT or firewall translates and forwards incoming IPv4 or IPv6 packets

LARGE SCALE NAT
SOME DEPLOYMENT EXPERIENCE

SANOG

UDP	IMAP	Rapidshare	Break Videos	SSH	MSN pc to pc voice	Skype Phone to PC	RDP
DNS	STUN	Hotfile	Pandora Tv	SSL	MSN Messenger Webcam	Skype PC to PC AV Chat	Nintendo Wii Web Browsing
ICMP	Myspace	rutube	Deezer	Gmail	MSN Messenger Remote Assistance	Skype Generic	Nintendo Wii Control
NTP	Windows Live	Justin.tv	Baidu	Hotmail	Google Talk	Skype PC to PC	Nintendo Wii Data
Pop3	Facebook	Metacafe	Godtube	yahoo mail	Jabber	PC: Valve's Steam Service	PSP
SMTP	YouTube Web	YouTube HD	Shockwave Flash	µtorrent	Google Talk Data	PC: World Of Warcraft	XP PPTP
FTP	Twitter	Yahoo Video	Google Video	Vuze (bitttorent)	Google Talk Voice	PC: Counter-Strike	<proprietary> VPN client
ftp passive	Mega Upload	Baidu Hi	Ipsec	eDonkey	IRC file transfer (client)	Windows Update	
ftp active	Flickr	Zshare	isakmp	MSN Messenger	IRC	Windows Activation	
Flash Video	Dailymotion	Bliptv	SSHD	MSN Messenger file transfer	IRC file transfer (server)	AVG Update	

LARGE SCALE NAT
SOME DEPLOYMENT EXPERIENCE

- Collateral damage
  - Somewhat unavoidable
  - Wireless environments have already had to deal with this
  - Non-transparent proxies have also caused this for a number of years
- Some of the apps on the previous page do suffer collateral damage issues
- Both content operators and the network operators will have to work on how to manage this
  - 1IP is not necessarily 1User
- All of these problems applies to NAT64 as well – so even a pure IPv6 network can suffer in the short term

LARGE SCALE NAT
THE PAIN/VALUE BARRIER

SANOG

- In the short term, to continue growing your IPv4 customer base, CGN does make sense => value!
- In the long term, LSNAT does have a pain threshold
  - How willing are you to support customer applications that break?
  - How willing are you to overload your addresses? 3:1? 4:1? 10:1? 100:1?
- Abuse management and lawful interception do become major operational headaches
- Scaling LSNAT as IPv4 continues to grow will be expensive

Continuing IPv6 deployment and overcoming the challenges is the only viable, and cost-effective, long term plan

IPv6 DEPLOYMENT AND LSNAT CONCLUSION

SANOG

- IPv6 deployment is the only long term solution for the network growth problem
- LSNAT gives us an interim crutch to get through the pain
- But in the long term, LSNAT will be very painful for operators and end users
  - Including even NAT64

NAT introduces operational complexity and expense to your network, but is a short-mid term solution to IPv4 network growth

Continuing IPv6 deployment and overcoming the challenges is the only viable, and cost-effective, long term plan

LARGE SCALE NAT STANDARDS

- IETF
  - BEHAVE Working Group
    - Large Scale NAT
    - NAT64
    - NAT recommendations
    - STUN/TURN/etc
  - PCP Working Group
    - PCP-Base
    - Extensions to PCP (multiple ports, interworking functions, DHCP/RADIUS config, etc)
  - draft-weil-shared-transition-space-request
- Broadband Forum
  - WT-242
- CableLabs CGN testing

AT THE SPEED OF IDEAS

Alcatel·Lucent

www.alcatel-lucent.com