

WELCOME



IPv6 TRANSITION TECHNOLOGIES

Alastair (AJ) JOHNSON
November 2012

alastair.johnson@alcatel-lucent.com

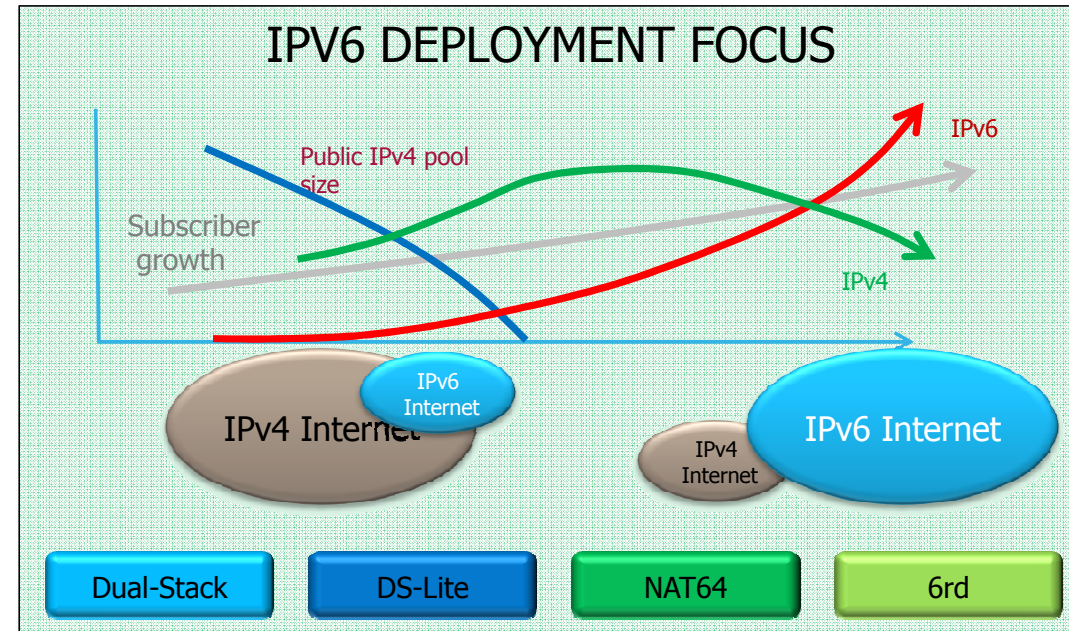
..... Alcatel-Lucent 

COPYRIGHT © 2011 ALCATEL-LUCENT. ALL RIGHTS RESERVED.

INTRODUCTION

WHAT ARE TRANSITION TECHNOLOGIES

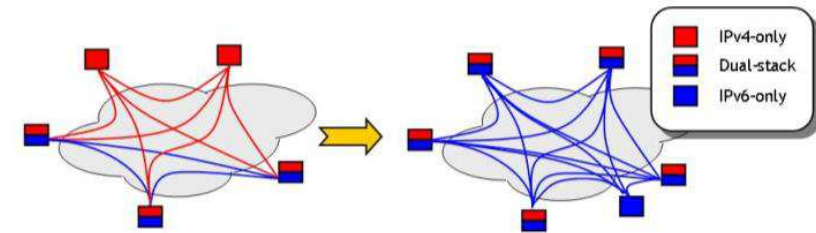
- Access Transition technologies are mechanisms that allow operators to deploy and migrate their subscriber-base to IPv6
- Transition technologies have been developed by the IPv6 community and vendors to help accelerate IPv6 deployment, and reduce barriers to IPv6 uptake
- All transition technologies should be evaluated carefully to identify which technology or technologies are the best fit for any given operator to deploy
- Some transition technologies have a 'long term life', others are seen as interim solutions to deploy IPv6 quickly while investment or technology catches up
- CPE is one of the most important domains for IPv6 deployment – to support any transition technology, long term strategy, and managing cost.
- **Avoiding multiple CPE swaps and migrations should be a key goal for any operator**



INTRODUCTION

WHY USE TRANSITION TECHNOLOGIES

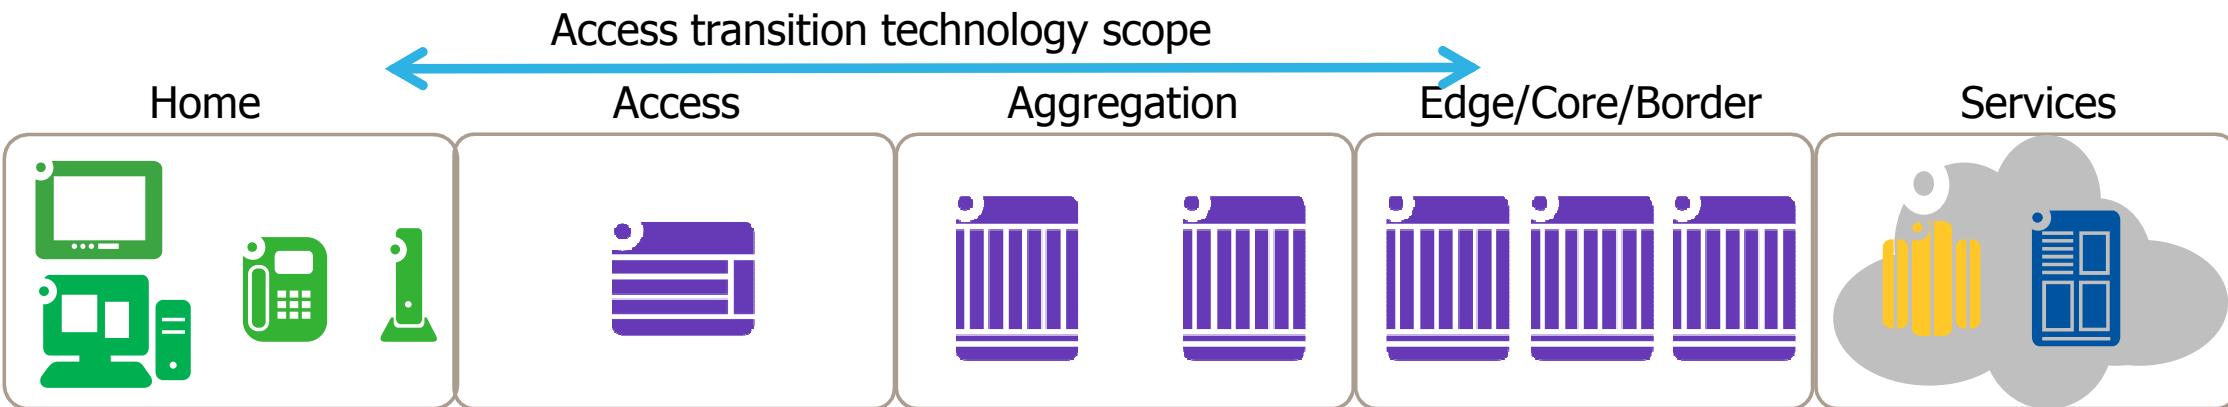
- If IPv6 services are desired, why not deploy native IPv6?
 - This is still the best approach if possible!
- Not always possible:
 - Technology constraints in the network may make native IPv6 deployment difficult or impossible without equipment replacement
 - Wholesale environments might not support IPv6 services currently
 - Desire to roll out IPv6 services as quickly as possible (trial, or overlay services)
 - Previous network architecture decisions may make native IPv6 deployment difficult without network changes (design, test, etc)
 - CPE support and replacement concerns
- **Transition technologies may allow operators to deploy IPv6 services in environments where native deployment is not possible; or to deploy IPv6 services quicker**



INTRODUCTION

TRANSITION TECHNOLOGIES IN OTHER DOMAINS

- The focus of this presentation is about IPv6 transition technologies and the role they may play in giving subscribers access to IPv6 services
- Other transition technologies may be present in the domains of a typical service provider network, but are not in the scope of this presentation
 - E.g. 6PE/6VPE, service/datacenter ALGs, ABGW-F, etc.



INTRODUCTION

LARGE SCALE NAT

- Large Scale NAT (LSNAT), Carrier Grade NAT (CGNAT), or any other type of service provider IPv4-to-IPv4 based NAT platforms and technologies are **not a transition mechanism to IPv6**
- These technologies are **IPv4 continuity solutions**
- LSNAT is one of several mechanisms that an operator may use to manage IPv4 exhaustion in their network while deploying IPv6 services
- This presentation will not discuss LSNAT beyond this slide



IPv6 TRANSITION TECHNOLOGIES

IPv6 TRANSITION TECHNOLOGIES

What will be discussed:

1. Native Dual Stack
2. Dual-Stack Lite (DS-Lite)
3. NAT64
4. 6rd

What won't be discussed:

1. (d)IVI
2. 464XLAT
3. 6to4
4. 6in4
5. 6over4
6. 4rd
7. ISATAP
8. NAT-PT/NAPT-PT
9. ...and many others

THE IMPORTANCE OF TRANSITION TECHNOLOGY SELECTION

- The Internet community (vendors, operators, IETF, subscribers) have done an excellent job in making sure we have a large number of transition technologies to pick from
- It is important to make sure the technology picked aligns with your engineering, operational and business needs and optimizes investment
- Understand the industry direction and what this means for support of transition technologies
- With so many to choose from, analyzing vendor support in both the CPE and translator domains is critical
 - Avoiding or minimizing vendor lock-in
 - Ensuring long-term support for features in your software deployed in the network

NATIVE DUAL-STACK

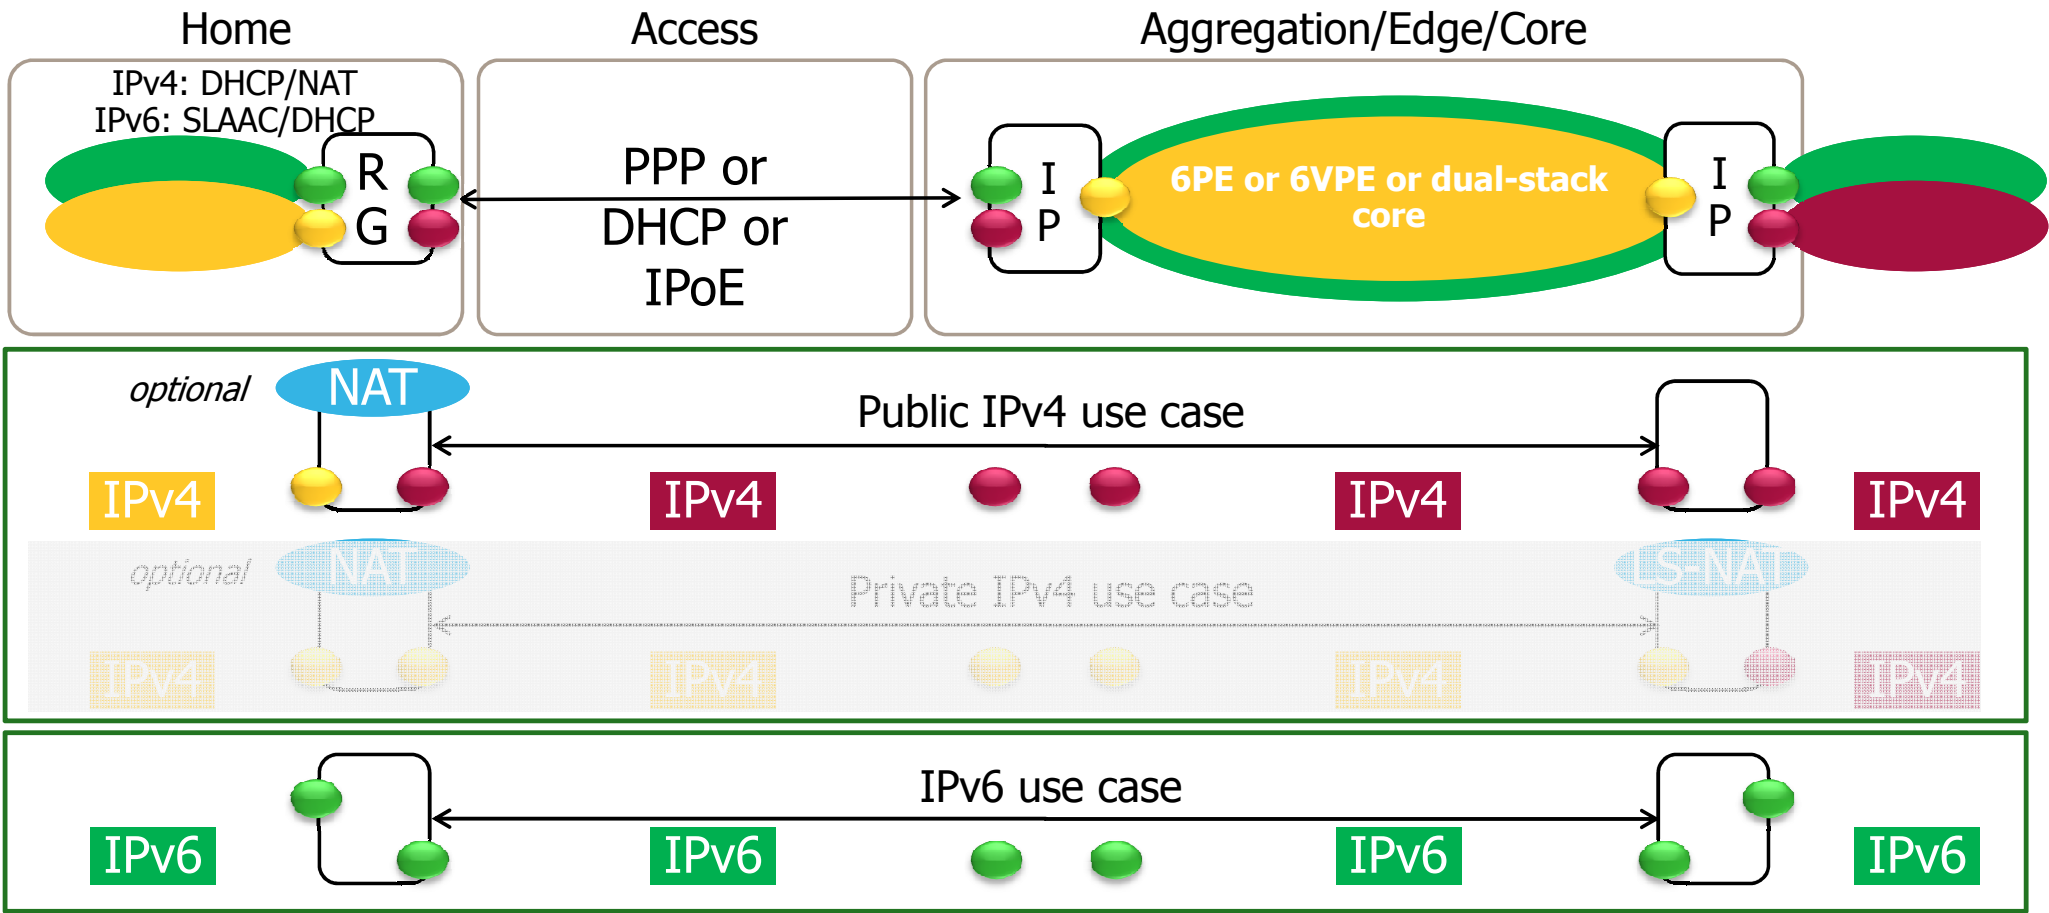
INTRODUCTION



- Deploying IPv6 services as native dual-stack is the best case approach for most operators and subscribers
 - However, it is often the most difficult
- No special encapsulation or tunneling is required
- Native IPv4 and IPv6 services are offered in parallel in the same subscriber session
 - Consistent service edge behavior between IPv4 and IPv6
 - IPv4 addressing is still provided to the subscriber with a potential for very long term sunsetting
- Deployment complexity levels vary in different environments
 - Some networks with minimal or no legacy equipment may find deploying native dual stack services very easy
 - Other networks with older or legacy equipment may find dual stack is not possible due to equipment constraints
 - CPE support is increasing significantly for dual-stack services on PPPoE and IPoE interfaces, including DHCPv6 (with prefix delegation) and SLAAC WAN support
- Ongoing operational considerations
 - What's the impact of running two parallel stacks on the network? Twice the monitoring, reporting, etc...

DUAL-STACK THROUGHOUT THE WIRELINE NETWORK

- IPv6 address
- Public IPv4 address
- Private IPv4 address



NATIVE DUAL-STACK

DOMAIN IMPACT

ACCESS

- Zero impact in PPPoE environments
 - PPPoE encapsulates traffic; RGs will be enable IPv6 when supported
- Medium to high impact in IPoE environments
 - N:1 VLANs may require network rearchitecture and rely on features in the access network
- Access node support (DSLAM, OLT, CMTS) becomes very important

SUBSCRIBER EDGE

- High impact – need to support IPv6 services:
 - Subscriber management, queuing, accounting, DHCP-PD, SLAAC(*), etc
- Scaling may be impacted when enabling IPv6 in BRAS/BNG
- Equivalency of features in the subscriber edge node is required – IPv4 & IPv6 should feel the same
- * SLAAC for subscriber management is an interesting issue, general industry trend is DHCPv6 based

HOME NETWORK

- Still the most complex domain to manage
- Customer Gateway most likely needs to be replaced
- BBF TR-124i2 specifies the requirements for IPv6 residential gateways
- Vendor support for IPv6 WAN/LAN is increasing significantly, e.g. Technicolor and D-Link
- Home network components need to support IPv6
- Internal addressing structure for the home network needs to be considered too



NATIVE DUAL-STACK

REFERENCE SLIDE
NOT PRESENTED

- Dual-stack is the 'best-case' transition design for IPv6 deployment which allows full coexistence of IPv6 and IPv4 services on an incremental deployment basis (e.g. subscribers can take up IPv6 services as their CPE is replaced, after network-wide deployment)
- Subscriber experience is identical regardless of IPv6 or IPv4 service, which are terminated on the same equipment (CPE, BNG) and share queues, SLA, and authorization and accounting policies
- Impact to the customer side of the network is high due to the CPE swap requirement – however significant number of CPE today are now IPv6 capable (including many transition technologies – refer to CPE link in references)
- Broadband Forum TR-177 and TR-187 along with TR-124i2 give excellent references for operators looking to deploy dual-stack services into existing TR-101 and PPP based environments, and provide requirements for RG behavior
- Depending on topology (IPoE v. PPPoE) the impact in the access/aggregation is variable:
 - PPPoE is very straightforward to deploy IPv6 on and allows easy customer uptake.
 - IPoE does require some changes in the access network, particularly if Lightweight DHCPv6 Relay Agent (LDRA) support is required and what the access architecture looks like. N:1 VLAN architecture does place some requirements on CPE behavior and potential requirements around handling Duplicate Address Detection. 1:1 VLAN architecture is preferred for IPoE broadband deployment
- Debate over SLAAC vs. DHCPv6 in the access attachment continues, however general recommendation and approach is DHCPv6 based to align with DHCPv4 model in existing networks
- Impact in the subscriber edge (BNG) is variable: impact to some legacy BNGs may be substantial when dual stack service is enabled impacting scalability, or lack of features for full equivalent IPv4 deployment. Operators need to investigate this carefully, however modern BNGs should have no issues when deploying dual-stack services at high subscriber scale
- Dual-stack does have drawbacks in that it does require potential capital investment if equipment forklift upgrades are required, as well as the impact of monitoring two address families in the network (twice the link monitoring, etc).
- Dual-stack does provide an interesting and easy approach to an IPv6-only network by simply turning IPv4 off in the future (and potentially using NAT64, etc).
- Allows status-quo to remain for non-Internet services (e.g. VoIP ATA, CPE/RG management, IPTV services etc) as existing IPv4 path is retained

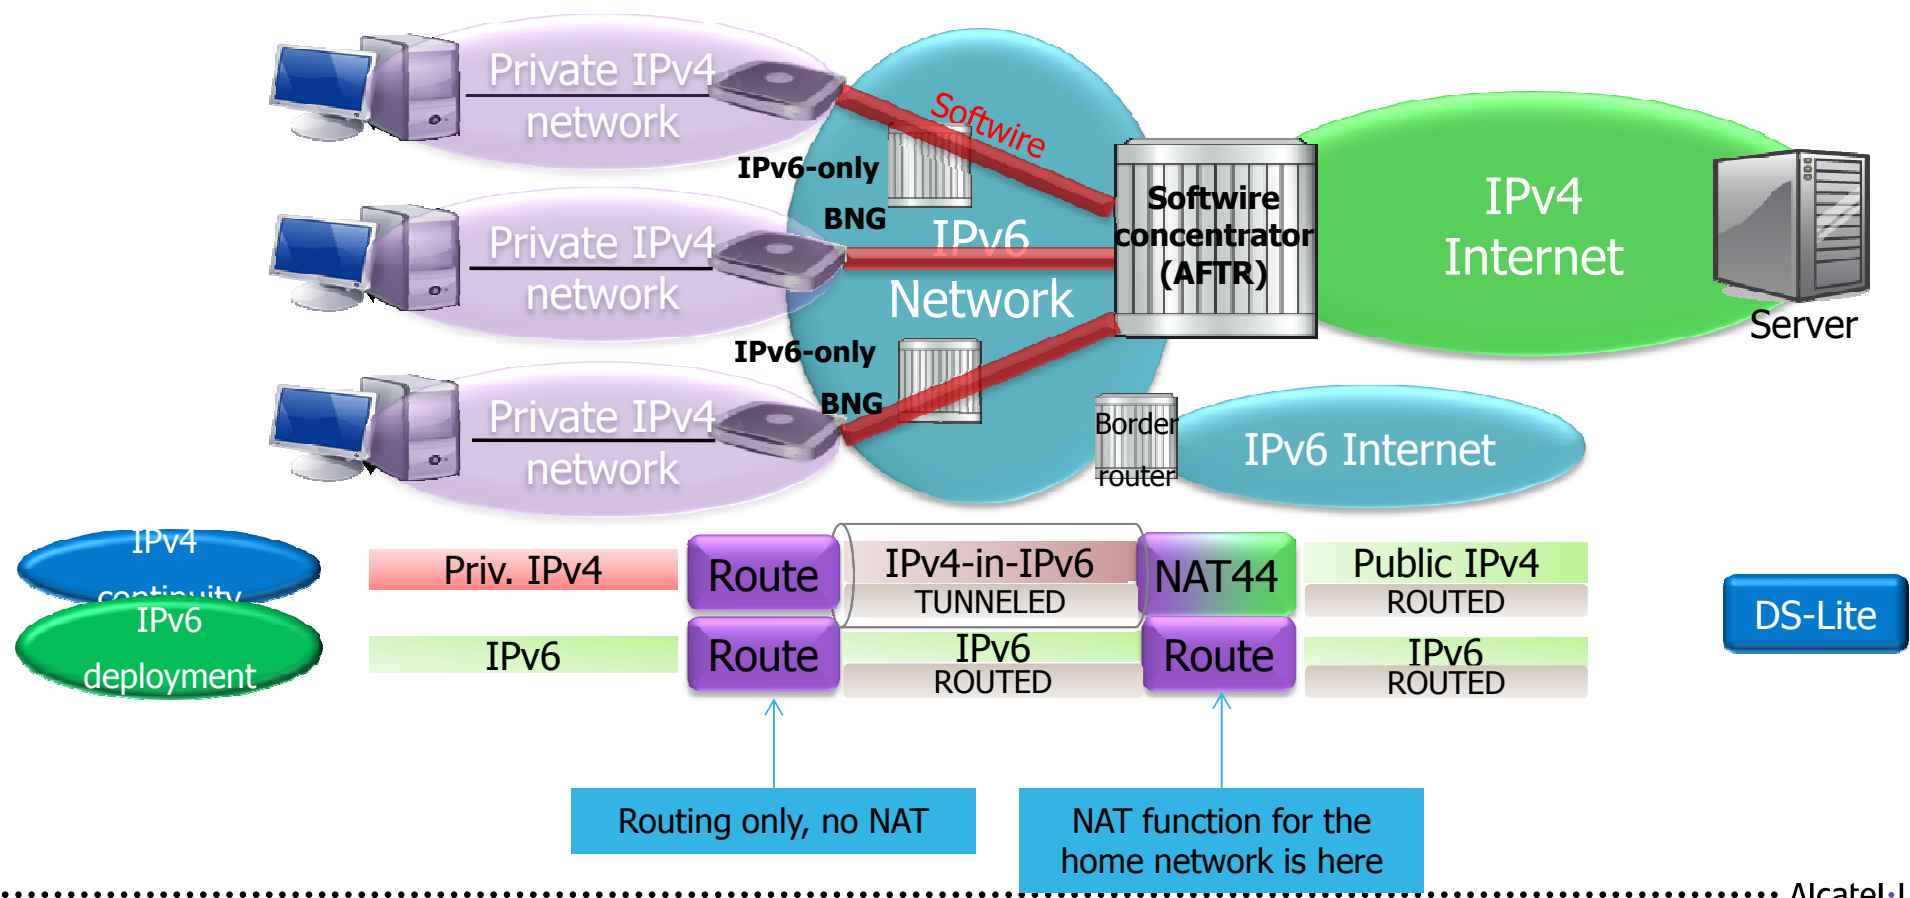
DS-LITE

INTRODUCTION

DS-LITE

- Defined by RFC6333 *Dual-Stack Lite Broadband Deployments Following IPv4 Exhaustion*
- Addresses operators that want IPv6-only access networks while providing support for IPv4-only nodes
 - Supporting the view that removing IPv4 from the access network is more efficient than supporting two stacks
- CPE encapsulates IPv4 traffic into IPv4-over-IPv6 tunnel using RFC2473
- Softwire concentrator (AFTR) decapsulates IPv4 packet and performs NAT44 using unique IPv6 transport address for NAT mapping (LSNAT)
- IPv4 traffic is routed by CPE (Basic Bridging Broad Band element [B4]) to IPv4-over-IPv6 tunnel and is subject to a single NAT operation at the softwire concentrator (or Address Family Transition Router [AFTR])
- IPv6 traffic is routed natively by CPE and BNG
- There is no protocol translation between IPv4 and IPv6

DUAL-STACK LITE



DS-LITE

DOMAIN IMPACT

ACCESS

- Medium to high impact depending on topology and technology
- Access network becomes **single stack IPv6 only**
 - All upgrades that a native dual-stack scenario requires are also required for DS-Lite
- Removing IPv4 from the access network becomes interesting
 - All CPE attaching to the network must support DS-Lite and IPv6 attachment

SUBSCRIBER EDGE

- AFTR node(s) are needed in the network
 - May be colocated in the BNG or a dedicated element
 - LSNAT and support infrastructure is required
- BNG must support all requisites for implementing IPv6 subscriber management
- Older equipment that does not support IPv6 will need to be replaced
- Considerations for Lawful Intercept and DPI are needed

HOME NETWORK

- Still the most complex domain to manage
- Customer Gateway (DSL modem/router, cable modem, etc) most likely needs to be replaced – must support IPv6-only WAN and DS-Lite – vendor support increasing
- Home network components need to support IPv6
- Internal addressing structure for the home network needs to be considered too
- IPv4 NAT at the customer gateway is removed

DUAL-STACK LITE

REFERENCE SLIDE
NOT PRESENTED

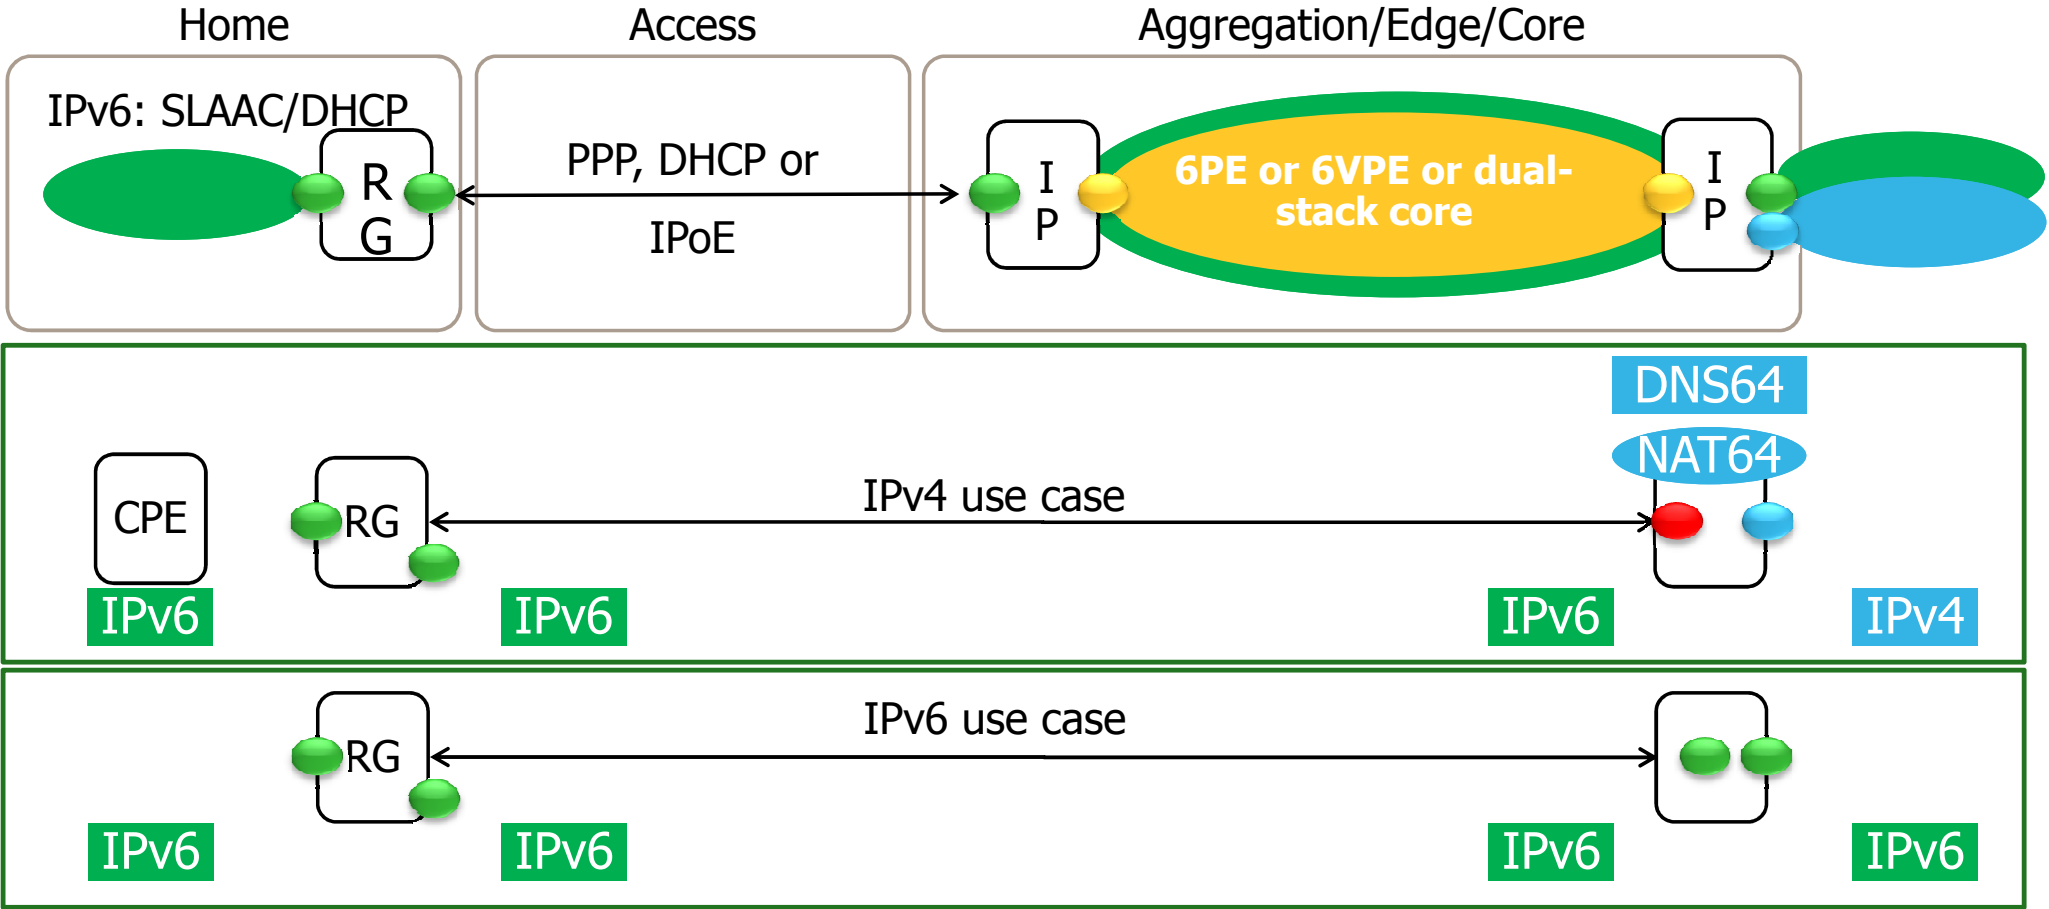
- DS-Lite specifically targets the case where operators wish to immediately remove IPv4 from the access –aggregation and subscriber management edge and run single-stack IPv6 while continuing to support IPv4 connectivity through a classic NAT44 capability rather than address family translation. This view was developed ~2007 and has started to gain deployment traction in 2012
- Significant impact in the CPE domain as the CPE must be upgraded to support IPv6 WAN and all associated connectivity (management, VoIP, IPTV etc), however NAT function is removed from CPE which potentially reduces cost (CPU/memory) in maintaining NAT state in the CPE. CPE are commercially available today that support DS-Lite and vendor support is continuing to increase
- Access network and subscriber management edge must support IPv6 in the same manner that dual-stack deployment. DS-Lite typically assumes an IPoE deployment but could be used in the PPP case as well
- Debate over SLAAC vs. DHCPv6 in the access attachment continues, however general recommendation and approach is DHCPv6 based to align with DHCPv4 model in existing networks
- As the operator must now deploy an AFTR, this node needs to be located near subscriber traffic (e.g. in or adjacent to the BNG) to avoid hauling traffic to centralized locations in the network which may impact TE or interface scaling in the network core. A potential drawback to non-BNG located AFTR is that any DPI or other IPv4 classification may be forced to occur at AFTR or elsewhere in the network, potentially stranding existing investment
- As DS-Lite moves the NAT44 function out of the RG and into the service provider environment, the service provider must support transaction logging for the LS-NAT as the subscribers share a common LAN IPv4 prefix (192.0.0.0/29) for the inside prefix
- DS-Lite does force re-architecture of existing service offerings such as VoIP and IPTV which may need to be moved to native IPv6 services to avoid transiting AFTR nodes in the network which may present a significant bandwidth bottleneck (in particular with multicast traffic!)
- Deployment of DS-Lite generally implies a significant migration stage where entire Access Nodes (or regions) are migrated at once, rather than an incremental migration on a per-subscriber basis – however this is up to individual service provider deployment approach
- DS-Lite provides an interesting and easy approach to an IPv6-only network by simply turning IPv4 off in the future when it is no longer required

INTRODUCTION

- Addresses operators who want IPv6-only access networks, but providing support for IPv4-only servers or content
 - Implies a well behaved, well understood CPE/UE – and ideally a minimal set of applications
 - Does **not** support IPv4-only hosts attaching to the network
- CPE/UE connects to IPv4 hosts through a synthesized IPv6 address, provided by a DNS64 engine
 - Well known prefix 64:ff9b::/96 is used to map IPv4 server addresses
 - Any client that cannot use a DNS64 server or provide local DNS64 resolution will not be able to connect to the IPv4 server, e.g. no more connecting by IP address
- IPv6 traffic is routed natively by CPE and BNG
- NAT64 can be used as a PLAT in the 464XLAT architecture

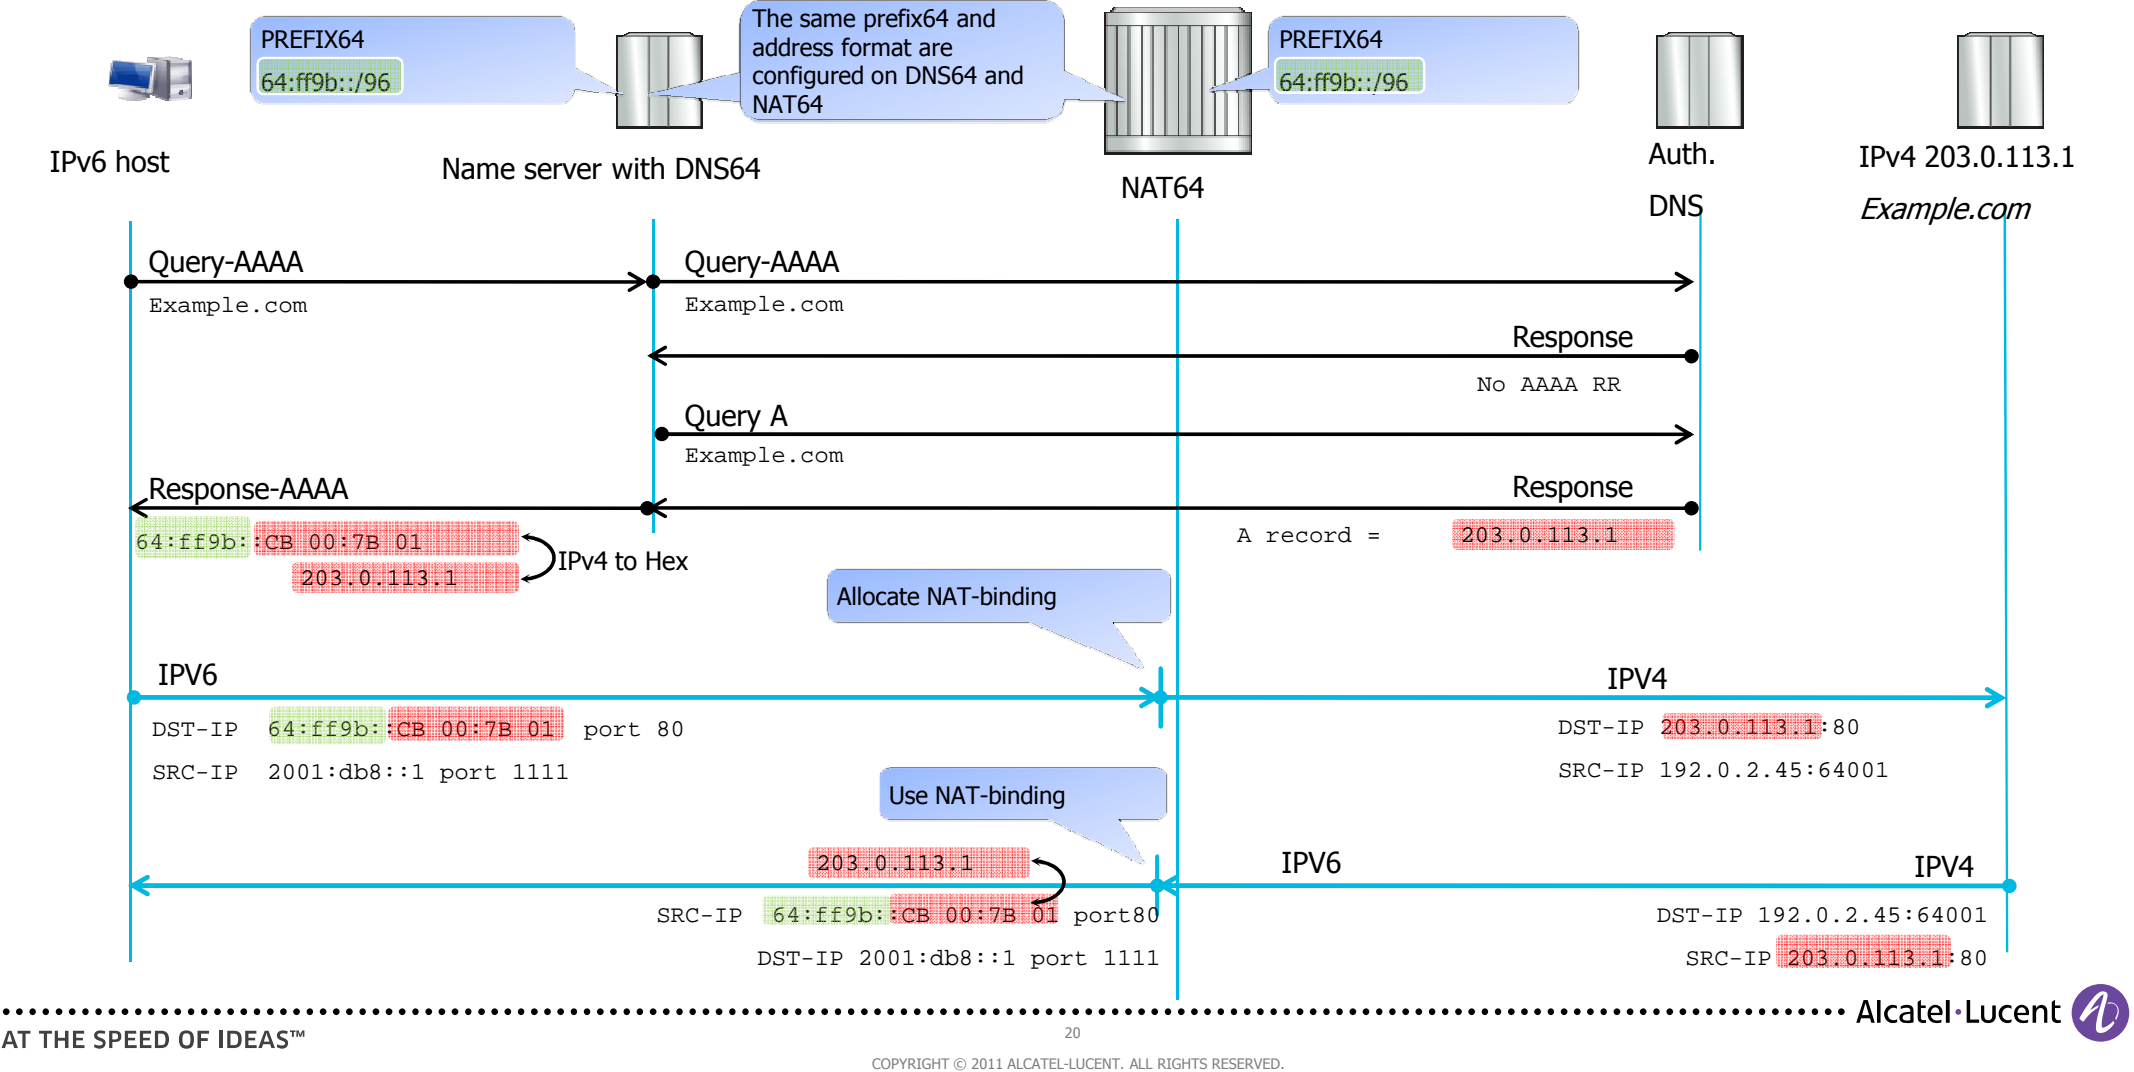
NAT64 IN THE WIRELINE NETWORK

- IPv6 address
- Public IPv4 address
- Pref64 IPv6 address



NAT64

FLOWCHART



NAT64

DOMAIN IMPACT

ACCESS

- Medium to high impact depending on topology and technology
- Access network becomes **single stack IPv6 only** so all upgrades that a native dual-stack scenario requires are also required for NAT64
 - All devices attaching to the network must support IPv6, including in-home

SUBSCRIBER EDGE

- NAT64 is needed in the network
 - May be colocated in the BNG or a dedicated element
 - If a separate element, load balancing of NAT64 traffic must be considered
- DNS64 node must also be deployed
- BNG must support all requisites for implementing IPv6 subscriber management
- Lawful Intercept and DPI must be considered

HOME NETWORK

- Customer Gateway (DSL modem/router, cellphone, cable modem, etc) most likely needs to be replaced – must support IPv6-only WAN
- **Home network components must support IPv6**
- Internal addressing structure for the home network needs to be considered too
- IPv4 NAT at the customer gateway is removed – and direct IPv4 support may be removed
- DNSSEC support will break with a DNS64 in the middle of the DNS chain
- Typically only useful or talked about for wireless environments at the moment

NAT64

REFERENCE SLIDE
NOT PRESENTED

- NAT64 specifically targets the case where operators wish to immediately remove IPv4 from the access–aggregation and subscriber management edge and run single-stack IPv6 while continuing to support IPv4 connectivity. **NAT64 most closely aligns with wireless deployment models rather than wireline, given the drawbacks in NAT64 for application translation** and the wider number of applications found in wireline environments vs. wireless
- Significant impact in the CPE domain as the CPE must be upgraded to support IPv6 WAN and all associated connectivity (management, VoIP, IPTV etc), however NAT function is removed from CPE which potentially reduces cost (CPU/memory) in maintaining NAT state in the CPE
- Access network and subscriber management edge must support IPv6 in the same manner that dual-stack deployment. NAT64 typically assumes an IPoE deployment but could be used in the PPP case as well
- Debate over SLAAC vs. DHCPv6 in the access attachment continues, however general recommendation and approach is DHCPv6 based to align with DHCPv4 model in existing networks
- As the operator must now deploy a NAT64, this node needs to be located near subscriber traffic (e.g. in or adjacent to the BNG) to avoid hauling traffic to centralized locations in the network which may impact TE or interface scaling in the network core. All DPI and classification on the IPv4 side of the NAT64 should be translated into the IPv6 side as well to preserve end-to-end behavior in the service provider network
- The operator must also deploy a DNS64 node that can provide the DNS synthesis by translating DNS responses with only A-records into AAAA-records with the well-known Pref64 prefix. Major DNS vendors support DNS64 translation today
- NAT64 will break a number of applications that rely on IPv4-literals (e.g. attempt to establish a socket directly to 192.0.2.1) and applications that will not traverse NAT environments happily. Some experiments have been conducted with IPv6-only networks and NAT64 environments and document the broken applications – refer to reference slide
- NAT64 does force re-architecture of existing service offerings such as VoIP and IPTV which may need to be moved to native IPv6 services to avoid transiting NAT64 nodes in the network which may present a significant bandwidth bottleneck (in particular with multicast traffic!)
- Deployment of NAT64 generally implies a significant migration stage where entire Access Nodes (or regions) are migrated at once, rather than an incremental migration on a per-subscriber basis – however this is up to individual service provider deployment approach
- NAT64 provides an interesting and easy approach to an IPv6-only network by simply turning IPv4 off in the future when it is no longer required

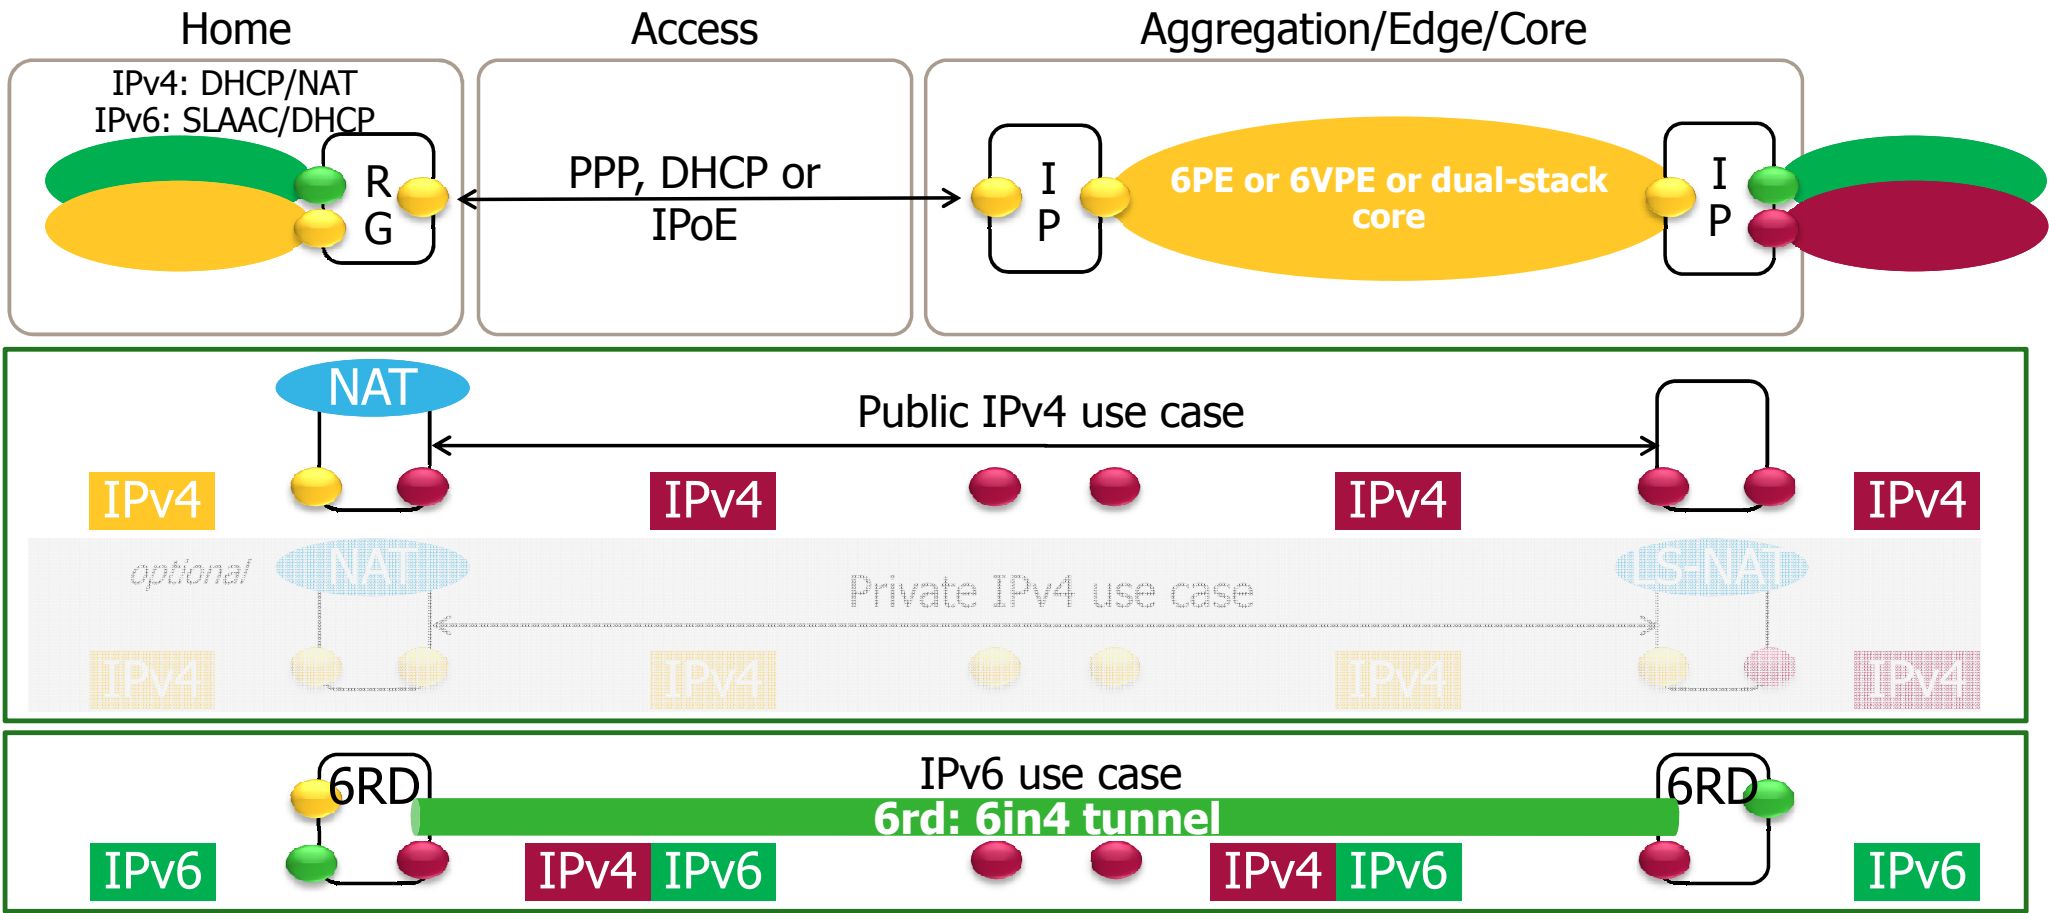
6rd

INTRODUCTION

- 6 **R**apid **D**eployment – RFC5969 *IPv6 Rapid Deployment on IPv4 Infrastructures*
- A tunneling technology based loosely on 6to4
- 6rd allows IPv6 to be deployed over existing IPv4-only access networks, without any forklift upgrades to the access, aggregation, or subscriber management networks
- All addresses are automatically discovered by the CPE, while the BR address may be statically configured or discovered via a variety of mechanisms (e.g. dhcp option)
- Fits well for wireline network environments where a CPE swap or upgrade is easy, but access networks are complex or expensive to modify (or are third party)
- Device-to-device traffic may be routed directly, and not through the BR when staying within a 6rd domain
- 6rd has plans under discussion for eventual sunseting in favor of native IPv6 (dual or single stack)

6rd

- IPv6 address
- Public IPv4 address
- Private IPv4 address



6rd

DOMAIN IMPACT

ACCESS

- No impact for 6rd – access network remains **exactly the same** for the initial deployment
- Subsequent migrations may still be required to get to ultimate end-state
 - E.g. native dual-stack, DS-Lite, or similar

SUBSCRIBER EDGE

- Border Relay is needed in the network
 - May be colocated in the BNG or a dedicated element
 - Load balancing of elements should be considered, as well as traffic engineering
- No change to subscriber management at the BNG
- Potential loss of visibility of tunneled traffic at BNG
- Lawful Intercept and DPI need to be considered

HOME NETWORK

- Customer Gateway (DSL modem/router, cable modem, etc) most likely needs to be replaced or upgraded – must support 6rd.
- Many RGs are shipping 6rd support today
- Home network components need to support IPv6 for native services
- IPv4 NAT at the customer gateway is still present
- Potential MTU impact for tunnels – potentially higher WAN MTU or frag-support required
- Useful for environments where the access network can't be touched (wholesale..)

6rd

REFERENCE SLIDE
NOT PRESENTED

- 6rd specifically targets the case where operators wish to immediately deploy IPv6 to their subscriber base, but cannot enable it in the native access. As 6rd encapsulates IPv6 in IPv4, it can be deployed across any existing IPv4 network
 - Some constraints faced by operators that drive 6rd as the technology include legacy Access Nodes (e.g. DSLAMs) that cannot support forwarding IPv6 packets, or older access technologies (e.g. DOCSIS 1.1) that cannot support IPv6
 - L3 wholesale access environments that cannot support IPv6 are another common barrier to deployment
- Significant impact in the CPE domain as the CPE must be upgraded to support 6rd. CPE are commercially available today that support 6rd and vendor support is continuing to increase
- Access network and subscriber management edge face no changes
- As the operator must now deploy a 6rd BR, this node needs to be located near subscriber traffic (e.g. in or adjacent to the BNG) to avoid hauling traffic to centralized locations in the network which may impact TE or interface scaling in the network core. A potential drawback to non-BNG located 6rd BR is that any DPI or other IPv4 classification may be forced to occur at 6rd BR or elsewhere in the network, potentially stranding existing investment or impacting service provider operations
- As 6rd may automatically derive the subscriber prefix with variable length subnetting (e.g. 48-56-64) based on the IPv4 address, the operator must consider exactly how many IPv4 bits they wish to stuff into the IPv6 prefix, and how this impacts any RIR allocated IPv6 prefixes. There are multiple approaches for managing the IPv6 addressing in 6rd environments
- 6rd does not force re-architecture of existing service offerings such as VoIP and IPTV which may remain on the existing IPv4 service
- 6rd can be deployed incrementally with no impact to the subscriber base as and when CPE are upgraded to support 6rd
- 6rd does not solve the long term problem of removing IPv4 from the access network or moving to native IPv6 services, however discussion for this is being undertaken in the IETF currently (refer reference slide)
- Potential MTU issues may occur with the tunnel, but may be mitigated by increasing WAN MTU or implementing fragmentation in the 6rd BR and CPE

METHODS OF TRANSITION

Home device	Access network	Destination	Solutions
IPv4	IPv4	IPv4 Internet	Dual-Stack
IPv6	IPv6	IPv6 Internet	
IPv4	IPv6	IPv4 Internet	DS-Lite
IPv6	IPv6	IPv4 Internet	NAT64 Stateful
IPv6	IPv4	IPv6 Internet	6RD

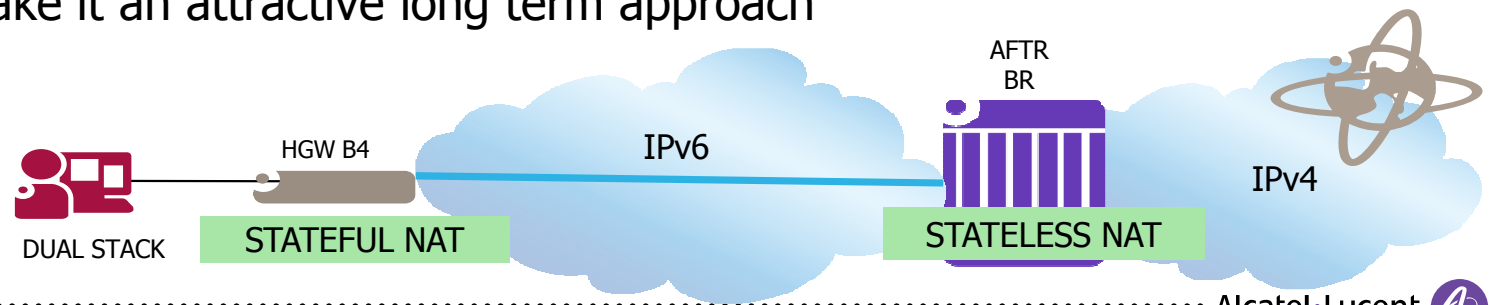
SUMMARIES AND COMPARISON

	Native Dual Stack	DS-Lite	NAT64	6rd
CPE	Almost always CPE change	CPE change and support for DS-Lite	CPE change (IPv6 only)	CPE change
End user impact	OK – not much changes	OK – not much changes	NOK – any IPv4-only devices (or partial-IPv6) are impacted. No non-DNS64 support	OK – not much changes
Pro	'Simple' technology with no transition or tunneling involved	Single address family in the access network	Single address family in the access network	Single address family in the access network Quick to deploy
Con	Cost of supporting dual-stack networks Device support Deployment time	All the effort of deploying dual-stack + extra Extra DS-Lite AFTR needed Traffic obfuscation in the network Device support	Application brokenness with IPv4-literals NAT logging required Will only work for IPv6-supporting hosts	Traffic obfuscation in the network Device support Not necessarily a 'long term' solution
Most suitable for...	Deployment everywhere! Best long term option that gives the widest support for both address families → Wireline, Wireless	New build environments where both removing IPv4 from and deploying IPv6-only access is feasible. → Wireline	New build environments where IPv6-only access is acceptable and the majority of content will work through NAT64/DNS64 → Wireless environments	Legacy environments that cannot support native IPv6 access, and are willing to trade-off multi-stage migrations over the long term → Wireline environments

Every transition technology employs translation – applications **will** be affected

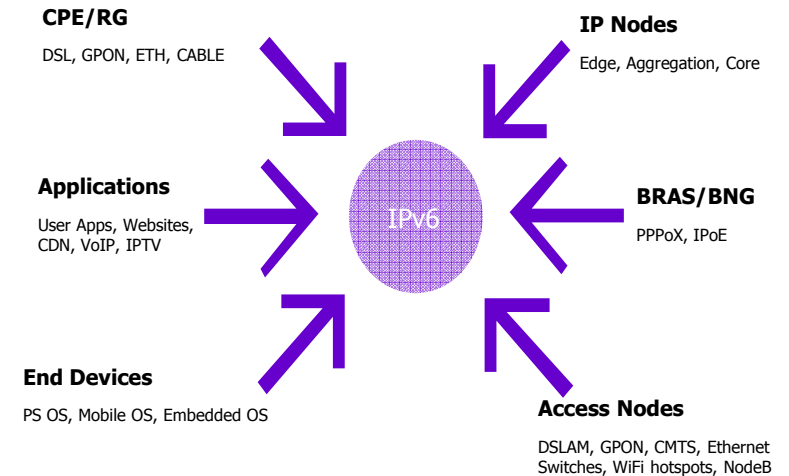
TRANSITION TECHNOLOGY EVOLUTION

- Mapping Address and Port (MAP) Encapsulation (MAP-E) proposes a stateless architecture for translation technologies
 - Standards track IETF draft
 - Has a relative, MAP-T
- Based on mapping an IPv4 address and port range into the IPv6 address of the subscriber
- Forwarding performs a stateless mapping lookup of the (IP+Port) and identifies the subscriber to forward to over IPv6
- Encapsulates IPv4 traffic in IPv6
- Scaling properties of MAP make it an attractive long term approach



CONCLUSION

- IPv6 deployment and transition technologies are a multi-dimensional problem
- There are a lot of transition technologies available with varying levels of support
- Operators should carefully evaluate which technology is most appropriate to meet their needs
 - The field is still changing!
- **The transition technology should align with the long term vision of the operator – generally this should look towards native IPv6 support**
- It might take multiple iterations to get to a long term view of native IPv6 (with transitional support for IPv4) – but it is important to try minimizing this from an investment and complication perspective



REFERENCES

Document	Location
Broadband Forum TR-177 IPv6 in Context of TR-101	http://www.broadband-forum.org/technical/download/TR-177.pdf
Broadband Forum TR-187 IPv6 for PPP Broadband Access	http://www.broadband-forum.org/technical/download/TR-187.pdf
Broadband Forum TR-124i2 Functional Requirements for Residential Gateway Devices	http://www.broadband-forum.org/technical/download/TR-124_Issue-2.pdf
RFC6333 Dual-Stack Lite Broadband Deployments	http://tools.ietf.org/html/rfc6333
RFC6052 IPv6 Addressing of IPv4/IPv6 Translators	http://tools.ietf.org/html/rfc6052
RFC6146 Stateful NAT64	http://tools.ietf.org/html/rfc6146
RFC6147 DNS64 DNS Extensions	http://tools.ietf.org/html/rfc6147
RFC5969 IPv6 Rapid Deployment on IPv4 Infrastructures	http://tools.ietf.org/html/rfc5969
draft-despres-softwire-4rd-u IPv4 Residual Deployment	http://tools.ietf.org/html/draft-despres-softwire-4rd-u-06
RFC6204 Basic Requirements for IPv6 Customer Edge Routers	http://tools.ietf.org/html/rfc6204 https://datatracker.ietf.org/doc/draft-ietf-v6ops-6204bis/
RFC6586 Experiences from an IPv6 Only Network	http://tools.ietf.org/html/rfc6586
464XLAT Experiences from T-Mobile USA	https://sites.google.com/site/tmoipv6/464xlat
RFC6219 CERNET IVI Translation Design and Deployment	http://tools.ietf.org/html/rfc6219
RFC6145 IP/ICMP Translation Algorithm	http://tools.ietf.org/html/rfc6145
RFC6144 Framework for IPv4/IPv6 Translation	http://tools.ietf.org/html/rfc6144
draft-townsley-v6ops-6rd-sunsetting Sunsetting for 6rd	http://tools.ietf.org/html/draft-townsley-v6ops-6rd-sunsetting-00
IPv6 CPE at the ARIN GetIPv6 Wiki	http://getipv6.info/index.php/Broadband_CPE
RIPE IPv6 CPE Survey	https://www.ripe.net/data-tools/ripe-labs/ipv6-cpe-survey

AT
THE
SPEED
OF
IDEAS™

AT
THE
SPEED
OF
IDEAS™

www.alcatel-lucent.com