

IPv6 Transition Strategies



Philip Smith

<philip@apnic.net>

APNIC 44

Taichung, Taiwan

7th – 14th September 2017

Presentation Slides

- Will be available on
 - <http://bgp4all.com/dokuwiki/conferences/>
 - And on the APNIC44 website
- Feel free to ask questions any time

Acknowledgements

- This material originated from the Cisco ISP/IXP Workshop Programme developed by Philip Smith & Barry Greene
 - These slides were originally developed by Cisco's CTO Consulting Engineering Group
- Use of these materials is encouraged as long as the source is fully acknowledged and this notice remains in place
- Bug fixes and improvements are welcomed
 - Please email *workshop (at) bgp4all.com*


Philip Smith

Introduction




Why should we care?

“The times, They are a’ changin’”



Is IPv4 really running out?

- Yes!
 - IANA IPv4 free pool ran out on 3rd February 2011
 - RIR IPv4 free pool is starting to run out now
 - www.potaroo.net/tools/ipv4/
 - (depends on RIR soft-landing policies)
- The runout gadgets and widgets are now watching when the RIR pools will run out:
 - inetcore.com/project/ipv4ec/index_en.html
 - (shows 1 RIR with no IPv4 left, and 3 out of 4 RIRs in run out austerity phase)
 - ipv6.he.net/statistics/



Strategies available for Service Providers

- Do nothing
 - Wait and see what competitors do
 - Business not growing, so don't care what happens
- Extend life of IPv4
 - Force customers to NAT
 - Buy IPv4 address space on the marketplace
- Deploy IPv6
 - Dual-stack infrastructure
 - IPv6 and NATed IPv4 for customers
 - 6rd (Rapid Deploy) with native or NATed IPv4 for customers
 - DS-Lite or 464XLAT with native IPv6 and NATed IPv4 for customers
 - Or other combinations of IPv6, IPv4 and NAT

Definition of Terms



Dual-Stack Networks

- Both IPv4 and IPv6 have been fully deployed across all the infrastructure
 - Routing protocols handle IPv4 and IPv6
 - Content, application, and services available on IPv4 and IPv6
- End-users use dual-stack network transparently:
 - If DNS returns IPv6 address for domain name query, IPv6 transport is used
 - If no IPv6 address returned, DNS is queried for IPv4 address, and IPv4 transport is used instead
 - Recent improvements introduce “happy eye-balls” (RFC6555)
- It is envisaged that the Internet will operate dual-stack for many years to come

IP in IP Tunnels

- A mechanism whereby an IP packet from one address family is encapsulated in an IP packet from another address family
 - Enables the original packet to be transported over network of another address family
- Allows ISP to provide dual-stack service prior to completing infrastructure deployment
- Tunnelling techniques include:
 - IPinIP, GRE, 6to4, Teredo, ISATAP, 6rd, MPLS

Address Family Translation (AFT)

- Refers to translation of an IP address from one address family into another address family
 - e.g. IPv6 to IPv4 translation
 - Usually called NAT64
 - Or IPv4 to IPv6 translation
 - Usually called NAT46, usually using SIIT

Network Address Translation (NAT)

- ❑ NAT is translation of one IP address into another IP address
- ❑ NAT (Network Address & Port Translation) translates multiple IP addresses into one other IP address
 - TCP/UDP port distinguishes different packet flows
- ❑ NAT-PT (NAT – Protocol Translation) is a particular technology which does protocol translation in addition to address translation
 - NAT-PT is has now been made obsolete by the IETF

Carrier Grade NAT (CGN)

- ISP version of subscriber NAT
 - Subscriber NAT can handle only hundreds of translations
 - ISP NAT can handle millions of translations
 - Expensive high performance hardware
- Not limited to just translation within one address family, but does address family translation as well
- Sometimes referred to as Large Scale NAT (LSN)

“Happy Eyeballs” – RFC6555

- The device or application chooses the protocol which will give the user the best experience
- Designed to work around shortcomings in either IPv4 or IPv6 infrastructure, or misconfigured IPv4 or IPv6 destination devices
- Short summary for dual stack device:
 - Application asks for IPv4 and IPv6 address
 - If both are returned, application opens connection using IPv6 and IPv4 simultaneously (or IPv6 first, then IPv4 after a short (few ms) delay)
 - Application uses the transport which responds with a connection first

Aside: NAT Issues (1)

- ❑ How to scale NAT performance for large networks?
 - Limiting tcp/udp ports per user harms user experience
- ❑ CGN deployment usually requires redesign of SP network
 - Deploy in core, or access edge, or border,...?
- ❑ Breaks the end-to-end model of IP
- ❑ Breaks end-to-end network security
- ❑ Breaks non-NAT friendly applications
 - Or NAT has to be upgraded (if possible)

Aside: NAT Issues (2)

- ❑ Makes fast rerouting and multihoming more difficult
 - Moving IPv4 address pools between CGNs for external traffic engineering
- ❑ Address sharing has reputation, reliability and security issues for end-users
- ❑ Layered NAT devices (double or even triple NAT)
- ❑ Mandates that the network keeps the state of the connections
- ❑ Makes the NAT device a target for miscreants due to possible impact on large numbers of users
- ❑ Makes content hosting impossible

Aside: NAT Issues (4)


- ❑ Consumer NAT device:
 - 3000 sessions means only 7 connected devices!
 - “NAT table FULL” error messages
 - “Broken Googlemaps”
 - “Stuck Internet”
- ❑ Carrier Grade NAT device:
 - 20 million sessions (Cisco ASR9001 ISM)
 - Which realistically is 50k users (400 sessions per user)
 - APNIC final /22 only allows 320k users ☹
- ❑ How to support LTE networks?!
 - Number of users? Public IPv4 addresses for CGN?
 - Maintaining LTE performance? Throughput of CGN?

Strategy One



Do Nothing

IPv4 only Network



- The situation for many SPs today:
 - No IPv6 for consumer
 - IPv4 scaling lasts as long as IPv4 addresses are available

IPv4 only: Issues

□ Advantages

- Easiest and most cost effective short term strategy

□ Disadvantages

- Limited to IPv4 address availability (RIRs or marketplace)
- No access to IPv6
- Negative public perception of SP as a laggard
- Strategy will have to be reconsidered once IPv4 address space is no longer available

IPv4 only: Applicability

- For Network Operators who:
 - Have sufficient IPv4 address space for foreseeable future business needs
 - Don't undertake long term planning
 - Are not heeding customer requests regarding IPv6 access
 - Have sufficient funds to purchase IPv4 address space via the marketplace

Strategy Two




Extend life of IPv4 network

Extending life of IPv4 Network

- Two ways of extending IPv4 network
 - Next step along from “Strategy One: Do nothing”
- 1. Force customers to use NAT
 - Customers moved to RFC1918 address space
 - SP infrastructure moved to RFC6598 address space (as public IPv4 needed for CGN pools)
- 2. Acquire IPv4 address space from another organisation
 - IPv4 subnet trading

SP NAT in IPv4-only network



- Next step on from “doing nothing”:
 - SP introduces NAT in core when IPv4 addresses run out
 - No access to IPv6 Internet for IPv6 enabled hosts

SP NAT in IPv4-only network:

Issues

- Advantages
 - ISPs can reclaim global IPv4 addresses from their customers, replacing with non-routable private addresses and NAT
 - Allows continued IPv4 subscriber growth
- Disadvantages
 - SP needs a large NAT device in the aggregation or core layers
 - Has every well known technical drawback of NAT, including prevention of service deployment by customers
 - Double NAT highly likely (customer NAT as well as SP NAT)
 - Sharing IPv4 addresses could have behavioural, security and liability implications
 - Tracking association of port/address and subscriber, not to mention Lawful Intercept issues, are still under study
 - May postpone IPv6 deployment for a couple of years
 - Prevents subscribers from using IPv6 content, services and applications

SP NAT in IPv4-only network:

Applicability

- For Network Operators who:
 - Are happy to purchase and operate CGN devices within their core network
 - Are aware of the operational and performance pitfalls of CGN devices
 - Are aware that their IPv4 network will need to be redesigned to accommodate CGN devices
 - Are aware of suboptimal routing and extra bandwidth requirements
 - Are able to reclaim public addresses from their customers for redeployment in their backbone
 - Are not heeding requests from customers for IPv6 access

IPv4 Subnet Trading

- Today the cost of getting IPv4 address space is low:
 - Service Provider:
 - RIR membership fee
 - Registration service fee (varies according to RIR service region)
 - 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
- In future 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 organisations

IPv4 Subnet Trading: Issues

□ Advantages

- 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

□ Disadvantages

- Market may not materialise, so organisations hoping to benefit may not
- Depending on region, if RIR doesn't register transfer, there may be no routability
- Risk to integrity of routing system, as RIRs no longer authoritative for address records
- Even more rapid growth of routing system
- Financial pressure on ISPs to dispose of IPv4 addresses they still need

IPv4 Subnet Trading: Applicability

- For Network Operators who:
 - Are have sufficient funds to purchase IPv4 address space on the marketplace
 - Are aware of the operational and performance pitfalls of purchased address space
 - Routability (legacy SP filters)
 - Registration (RIR vs not)
 - Reputation (previous user)
 - Are not heeding requests from customers for IPv6 access

Strategy Three



IPv4/v6 Coexistence/Transition
techniques


IPv4/IPv6 coexistence & transition

- Three strategies for IPv6 transition:
 - Dual Stack Network
 - The original strategy
 - Depends on sufficient IPv4 being available
 - 6rd (Rapid Deploy)
 - Special case of & improvement on 6to4 for SP customer deployment
 - Documented in RFC5969
 - 464XLAT or DS-Lite or NAT64 with CGN
 - SP deploys large NAT boxes to do address and/or protocol translation

IPv4/IPv6 coexistence & transition

- Carrier Grade NAT (CGN)
 - Dual-Stack Lite
 - IPv4 to IPv4 over IPv6
 - Documented in RFC6333
 - 464XLAT
 - IPv4 to IPv4 over IPv6
 - Documented in RFC6877
 - NAT64
 - Translation between IPv6 and IPv4
 - Documented in RFC6146

Dual-Stack Network



- The original transition scenario, but dependent on:
 - IPv6 being available all the way to the consumer
 - Sufficient IPv4 address space for the consumer and SP core

Dual-Stack Network: Issues

□ Advantages

- Most cost effective long term model
- Once services are on IPv6, IPv4 can simply be discontinued


□ Disadvantages

- IPv4 growth limited to available IPv4 address space
- Running dual-stack network requires extra staff training
- IPv6 on existing IPv4 infrastructure might cost extra in terms of hardware changes (RIB and FIB memories)
- IPv6-only end-points cannot access IPv4, but given most IPv6 end-points are dual-stack, require IPv4 address too

Dual-Stack Network: Applicability

- For Network Operators who:
 - Have sufficient IPv4 address space for foreseeable future
 - Also may consider purchasing IPv4 address space on the open market
 - Have no legacy equipment or infrastructure which does not support IPv6
 - Do not wish to deploy CGN (NAT44)
 - **Are willing to support dual-stack CPE**
- Note: this is considered the ideal option
- Example:
 - Typical traditional Internet Service Provider deployment

Dual-Stack with SP NAT



- More likely scenario:
 - IPv6 being available all the way to the consumer
 - SP core and customer has to use IPv4 NAT due to v4 depletion

Dual-Stack with SP NAT: Issues

□ Advantages

- ISPs can reclaim global IPv4 addresses from their customers, replacing with non-routable private addresses and NAT
- Allows continued IPv4 subscriber growth
- SP can offer IPv6 connectivity too
- Does not postpone IPv6 deployment
- SP NAT off-load (compared with IPv4-only network)

□ Disadvantages

- SP needs a large NAT device in the aggregation or core layers
- Has every well known technical drawback of NAT, including prevention of service deployment by customers
- Double NAT highly likely (customer NAT as well as SP NAT)
- Sharing IPv4 addresses could have behavioural, security and liability implications
- Tracking association of port/address and subscriber, not to mention Lawful Intercept issues, are still under study
- SP incurs additional investment and operational expenditure by deploying an IPv6 infrastructure

Dual-Stack with SP-NAT:


Applicability

- For Network Operators who:
 - Have do not sufficient IPv4 address space and are content deploying CGN (NAT44) in the core
 - Are able to reclaim public IPv4 address space from customers for redeployment on their backbone infrastructure
 - Have no legacy equipment or infrastructure which does not support IPv6
 - **Are willing to support dual-stack CPE**
- Note: this is considered the realistic best practice
- Example:
 - Typical traditional Internet Service Provider deployment

Aside: SP-NAT Offload

- If 50% of end user traffic is IPv6, then this means 50% less IPv4 traffic which has to be mapped and translated via the SP's CGN installation
 - The greater the proportion of IPv6 traffic (compared with IPv4), the less the load is on the CGN devices, and reduced demand on the public IPv4 address pool
 - CGN is used simply for accessing legacy IPv4 sites
- Operators with high data volumes realise that by deploying IPv6:
 - End users have better Internet experience when traffic is not NAT'ed
 - They have reduced CapEx deploying fewer CGN devices
 - Savings from reduced CGN CapEx are often greater than the additional costs to deploy IPv6 to end-users
- This is called SP-NAT Offload

6rd



- 6rd (Rapid Deploy) used where ISP infrastructure to customer is not IPv6 capable (eg IPv4-only BRAS)
 - Customer has IPv4 Internet access either natively or via NAT
 - Customer IPv6 address space based on ISP IPv4 block

6rd: Issues

□ Advantages

- The service provider has a relatively quick way of providing IPv6 to their customer without deploying IPv6 across their infrastructure
- Subscribers can readily get access to IPv6
- SP NAT off-load (compared with IPv4-only network)
- 6rd relay and CPE are becoming available from vendors
- 6rd operation is completely stateless, does not have the operational drawbacks of 6to4, and does not postpone IPv6 deployment


□ Disadvantages

- 6rd is not a long-term solution for transitioning to IPv6 – one further transition step to remove the tunnels
- CPE needs to be upgraded to support 6rd
- The ISP has to deploy one or several 6rd termination devices
- If customer or SP uses NAT for IPv4, all NAT disadvantages are inherited

6rd: Applicability

- For Network Operators who:
 - Have do not sufficient IPv4 address space and are content deploying CGN (NAT44) in the core
 - Are able to reclaim public IPv4 address space from customers for redeployment on their backbone infrastructure
 - Have legacy equipment or infrastructure which does not support IPv6
 - And realise that it will eventually have to be upgraded
 - Are willing to run a 6rd Border Router
 - Are willing to support dual-stack CPE (with 6rd)
- Example:
 - Broadband operators who have legacy DSLAMs or lease a third party's L2 network

Dual-Stack Lite



- Service Provider deploys IPv6-only infrastructure:
 - IPv6 being available all the way to the consumer
 - IPv4 is tunnelled through IPv6 core to Internet via SP NAT device

Dual-Stack Lite: Issues

□ Advantages

- The SP is using IPv6 across their entire infrastructure, avoiding the IPv4 address pool depletion issue totally
- The SP can scale their infrastructure without any IPv4 dependencies
- Consumers can transition from IPv4 to IPv6 without being aware of any differences in the protocols
- IPv6 packets routed natively
- SP NAT off-load (compared with IPv4-only network)


□ Disadvantages

- SP requires NAT device in core supporting DS-Lite
- Subscriber router needs to be IPv6 capable
- Model has all drawbacks of SP NAT model for IPv4 traffic

Dual-Stack Lite: Applicability


- For Network Operators who:
 - Are considering “green-field” deployments
 - Are content running an IPv6-only backbone
 - Are willing to deploy CGN (DS-Lite) in the core
 - **Are willing to support dual-stack CPE (with DS-Lite)**
- Example:
 - Mobile operators rolling out a brand new network, with handsets which have dual-stack radios

Stateful AFT (NAT64)



- Service Provider deploys IPv6-only infrastructure:
 - Only IPv6 is available to the consumer
 - IPv4 Internet available via Address Family Translation on SP NAT device

Stateful AFT (NAT64) Details



Stateful AFT: Issues

□ Advantages

- Allows IPv6 only consumers access to IPv4 based content without giving them IPv4 address resources
- IPv6 services and applications offered natively to consumers
- SP network runs IPv6 only, avoiding IPv4 dependencies


□ Disadvantages

- SP requires NAT device in core
- SP's DNS infrastructure needs to be modified to support NAT64
- Subscriber router needs to be IPv6 capable
- Subscriber devices need to be IPv6 capable (no legacy support)
- Model has all drawbacks of SP NAT model for IPv4 traffic

Stateful AFT: Applicability

- For Network Operators who:
 - Are considering “green-field” deployments
 - Are content running an IPv6-only backbone
 - Are willing to deploy CGN (NAT64) in the core
 - **Are willing to support IPv6-only CPE**
- Example:
 - Mobile operators rolling out a brand new network, with handsets which have single-stack (IPv6-only) radios

464XLAT



- Service Provider deploys IPv6-only infrastructure:
 - IPv6 being available all the way to the consumer
 - IPv4 is transported through IPv6 core to Internet via SIIT on customer router, and NAT64 on SP NAT device

464XLAT: Issues

□ Advantages

- The SP is using IPv6 across their entire infrastructure, avoiding the IPv4 address pool depletion issue totally
- The SP can scale their infrastructure without any IPv4 dependencies
- Consumers can transition from IPv4 to IPv6 without being aware of any differences in the protocols
- Devices not supporting IPv6 can access IPv6-only networks
- IPv6 packets routed natively
- SP NAT off-load (compared with IPv4-only network)

□ Disadvantages

- SP requires NAT device in core (PLAT & NAT64)
- Subscriber router needs to be IPv6 capable and support IPv4/IPv6 header translation (CLAT – SIIT)
- Model has all drawbacks of SP NAT model for IPv4 traffic

464XLAT: Applicability

- For Network Operators who:
 - Are considering “green-field” deployments
 - Are content running an IPv6-only backbone
 - Are willing to deploy CGN (PLAT) in the core
 - **Are willing to support dual-stack CPE (CLAT)**
- Example:
 - Mobile operators rolling out a brand new network, with handsets which have dual-stack radios

Conclusions & Recommendations



Functionalities and Operational Issues

	IPv4-only network	IPv4-only network with IPv4 NAT	Dual-Stack, no IPv4 NAT	Dual-Stack with IPv4 NAT	6rd, no IPv4 NAT	6rd with IPv4 NAT	DS-Lite	464XLAT	Stateful AFT
Prolongs IPv4	No	Yes	No	Yes	No	Yes	Yes	Yes	Yes
Allows Business Growth	No	Yes (scaling issues if content is mostly IPv6)	Limited to IPv4 address availability	Yes (traffic to IPv4-only servers)	Limited to IPv4 address availability	Yes	Yes	Yes	Yes
Requires IPv6 Deployment	No	No	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Coexists with IPv6 Deployment	No	No	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Complexity of Operation	Low	Low	Low	Moderate	Moderate	Moderate	Moderate	Moderate	Moderate
Complexity of Troubleshooting	Low	Moderate	Low	High	Moderate	High	High	High	Moderate
Breaks End-to-End IPv4	No	Yes	No	Yes	No	Yes	Yes	Yes	N/A
NAT Scalability issues to IPv4 services	No	Yes	No	Yes	No	Yes	Yes	Yes	Yes
NAT Scalability issues to IPv6 services	N/A	Yes	No	No	No	No	No	No	No
DNSSEC issues	No	Yes	No	Yes for IPv4 No for IPv6	No	Yes for IPv6 No for IPv4	Yes for IPv4 No for IPv6	Yes for IPv4 No for IPv6	Yes for IPv4 No for IPv6
Lawful Intercept issues	No	Yes	No	Yes for IPv4	No	Yes for IPv4	Yes for IPv4	Yes for IPv4	Yes for IPv4

Functionalities and Operational Issues

- Complexity of operation:
 - Moderate in the case of a single network with two address families
- Complexity of troubleshooting:
 - Running two address families and/or tunnels is assumed to be more complex
- Breaks end-to-end connectivity in IPv4:
 - Subscribers sharing a CGN will have little to no hurdles in their communication
 - Subscribers separated by one or several CGN will experience some application issues

Comparing where changes will occur

	IPv4-only network	IPv4-only network with IPv4 NAT	Dual-Stack, no IPv4 NAT	Dual-Stack with IPv4 NAT	6rd, no IPv4 NAT	6rd with IPv4 NAT	DS-Lite	464XLAT	Stateful AFT
Change CPE	No	No	Only if customer wants IPv6	Only if customer wants IPv6	Yes	Yes	Yes	Yes	Yes
CPE to do AFT to access IPv6	No	No	No	No	No	No	No	No	No
IPv4 NAT in core/edge	No	Yes	No	Yes	No	Yes	Yes	Yes	No
AFT in core/edge to access IPv6	Yes	Yes	No	No	No	No	No	No	Yes

Conclusions

Potential Scenarios

1. Most of the content and applications move to IPv6 only;
2. Most of the content and applications are offered for IPv4 and IPv6;
3. Most of the users move to IPv6 only
 - Especially mobile operators offering LTE handsets in emerging countries
4. No change (the contents/applications stay IPv4 and absence of pro-IPv6 regulation), SP customer expectations devolve to double-NAT;
5. No change (the contents/applications stay IPv4) but SP customer expectations do not devolve to double-NAT (or they are ready to pay for peer-to-peer connectivity).
 - Perhaps well established broadband markets like US or Europe

Conclusions

Potential Techniques

Scenario	Potential Techniques
Content and Applications move to IPv6	IPv6 only network; Dual-Stack, 6rd, 464XLAT or DS-lite as migration techniques
Content and Applications on IPv4 and IPv6	Dual-Stack (if enough IPv4) or 6rd; SP IPv4-NAT; DS-lite or 464XLAT (for greenfield) *
Users are IPv6 only	Stateful AFT to get to IPv4 content *
No change (double NAT)	SP IPv4-NAT *
No change (no double NAT)	Do nothing *

* Transfer Market applicable

Recommendations

1. Start deploying IPv6 as long term strategy
2. Evaluate current addressing usage to understand if IPv4 to IPv4 NAT is sufficient for transition period
3. Prepare a translation mechanism from the IPv4 Internet to the IPv6 Internet
4. Educate your user base on IPv6 introduction, the use cases and troubleshooting

Recommendations

- Mobile operator:
 - 464XLAT (support IPv4 and IPv6)
 - NAT64 (IPv6 only)
- Access provider:
 - Dual stack core and access (if supported to end-user)
 - 6rd (if legacy IPv4 infrastructure or legacy 3rd party L2)
- Enterprise service provider:
 - Dual stack access and core
- Content provider:
 - Dual stack access and core

IPv6 Transition Planning



A Review of IPv4-IPv6 Co-
Existence Techniques