IPv6 update from a broadband ISP

Rev Adrian Kennard
Director

Andrews & Arnold Ltd
www.aa.net.uk

FireBrick Ltd
www.firebrick.ltd.uk

Where have we got to?

- What is IPv6 anyway?
- IPv6 has been around 15 years – what's the hold up?
- What does IPv6 mean for an ISP and its customers?
- What have AAISP been doing?
- What are we doing now exactly?
- What is happens next?

What is IPv6 anyway?

- The current version of Internet Protocol (IP)
  - Previous version was IPv4
- Main feature is the bigger addresses
  - 128 bits rather than 32
  - 340 trillion trillion trillion addresses
  - Some other protocol changes as well
- In principle it just slots in, in place of IPv4

IPv6 addresses

- 128 bits split in to eight 16 bit values shown in hex with :
  - 2001:08b0:0000:0000:0203:97FF:FE12:3456
- Dropping any leading zeros in any part
  - 2001:8b0:0:0:203:97FF:FE12:3456
- Replace one string of :0:'s with ::
  - 2001:8b0::203:97FF:FE12:3456
- ::1 is an address (loopback) all zeros and a 1
- :: is an address (the *undefined* address) all zeros
- Can also show last 4 bytes in IPv4 format

In practice, IPv6

- Not quite "slotting in", but close...
- Applications have to understand there is a difference
- Older function calls work, but were typically written for IPv4 (e.g. gethostbyname)
- Newer functions are more version agnostic, but still need coding to consider them
- Newer applications work in more generic way
  - Trying a list of IPs for a name in order
- Easier if ever changes in future?

IPv6 allocations

- ISP gets a /32, well in fact first /32 of a /29 reserved for ISP
- End user normally allocated /48
  - 65536 customers in a /32
- Each subnet normally /64
  - 64 bit network part
  - 64 bit host part
  - 65536 subnets of /64 size in a /48
- That is 18 million trillion addresses per subnet!

64 bit host part

- Can be random, or allocated automatically or manually
- Typically based on MAC
- Take first 3 bytes
  - Invert bit 1, i.e. XOR first byte with 0x02
- The FFFE (meaning we based on MAC)
- Then last 3 bytes
  - e.g. MAC of 000397123456
  - Makes host part of 0203:97FF:FE12:3456
- e.g. 2001:8b0::203:97FF:FE12:3456 as a full IPv6 address

64 bit network part

- Router announcements (SLAAC)
  - Tells devices the network part
  - Tells them they can pick their host part
  - Can tell them DNS servers as well
  - Provides a gateway MAC for default route
  - Can tell devices to use DHCPv6 for some or all config
  - Is the default on most devices and stupidly easy to set up
- DHCPv6
  - Can be used for stateful address allocation
  - Can allocate DNS, and many other settings like DHCP

Link local addressing

- Slightly controversial
- All devices have a link local address
- Within FE80::/10
- Uses MAC to create a link local address
- Only for communications on same link
- E.g. used for DHCPv6, neighbour discovery, etc.
  - i.e. DHCPv6 uses link local and multicast addresses rather than 0.0.0.0 and 255.255.255.255 used by IPv4 DHCP which is marginally more sane

Neighbour discovery

- AKA neighbor discovery!
- Used just like ARP
- Converts IPv6 address to MAC on Ethernet LAN
- Uses ICMPv6 rather than separate protocol
- Uses Link Local (FE80::/10) addresses and multicast addresses
- Technically multicast not broadcast
- Additional parameters, defined timeout, better than ARP

Site local addressing

- Idea of a site wide addressing plan
- FEC0::/10
- Deprecated
- Still seen on some older windows installs as generic site local DNS server addresses

Applications using IPv6

- Applications can work in a mode where sockets are treated as IPv6 even if the underlying traffic is IPv4!
- Simplifies application coding in some ways
- IPv4 appear as ::FFFF:x.x.x.x address
- Can be confusing in some logs showing like this rather than as a proper IPv4 address
- Some applications have to know difference anyway if they ever textually quote the IP (e.g. SIP) or have different messages for IPv6 (e.g. ftp)

IPv6 literals

- In IPv4 you can have http://81.187.81.187/
- You can even have http://123456789/ !!!
- In IPv6 you use [and]
- e.g. http://[2001:8b0::102]/

Fragments

- IP packets are bigger than Ethernet frames
  - IP up to 65535, Ethernet typically 1500
- IPv4 allowed fragmentation on the way
  - DF flag to stop this and return an error, path discovery
- IPv6 insists endpoint does the fragmentation
  - Effectively always DF
  - Much more efficient for TCP
  - Needs ICMPv6 "TOO BIG" error
- Some numpties blocking ICMPv6 even now

Why use IPv6?

- IPv4 ran out in February 2011
- More and more erosion of basic principles with NAT
- CGNAT and proxies used more and more
- IPv6 allows IP to work as intended
  - Every endpoint has a globally unique address
- Does not mean no firewalls
  - Obviously
- Allows innovation and development stifled by IPv4 NAT
- Ultimately necessary as IPv4 will not cope forever

Why IPv6 now?

- Now is as good a time to start as any
  - IPv4 finally ran out – so that focuses the mind somewhat
- Easier to start using IPv6 while IPv4 still works
  - Gives you time to deploy and test in parallel
- Think of IPv6 in any coding and development now
- Think dual stack, IPv4 will be with us a long time
- The more we delay the more NAT will take a grip
- ISPs finally able to offer IPv6 sensibly this year

What's the hold up?

- AAISP have been doing IPv6 for nearly 9 years
- Techie customers able to make it work their end
- Only now is it "standard" and enabled by default
- We still don't have cheap domestic DSL routers
  - That's the hold up!
- Very close – several manufacturers almost ready
  - Still too buggy to be usable, sadly
  - Give it a couple of months
- AAISP will ship routers with IPv6 enabled as standard

What does IPv6 mean for an ISP?

- For the ISP, IPv6 is not rocket science...
  - Core routers need IPv6 configuring – routers can do that!
  - Transit providers need IPv6 – that is available now
  - ISPs need management systems and staff training
  - Some technical aspects, DNS, reverse DNS, RADIUS, etc
  - Concerns over whether customers are ready?
  - Will it be seamless for the end user if IPv6 is turned on?

A few gotchas for ISPs

- The IPv6 routing is not the hard or tedious bit!
- Log files, searching, etc.
- Management systems, IP allocation systems
- Config files generated from databases
- Support staff training
- End user confusion
- Access control lists (including 2002::/16 addresses)
- DNS and Reverse DNS
- Old routers not understanding IP6CP

Tunnels

- Simple wrapper to send IPv6 over IPv4 (proto 41)
- Mapping of IPv4 addresses to 2002::/16
  - Auto tunnelled IPv6 to proto 41 to IPv4
  - Reply tunnels via 192.88.99.1
  - Allows IPv6 without ISP support
  - End user equipment, PCs, etc, using this
- Has led to customers using IPv6 without realising it!
  - Hence needing 2002:: blocks on access lists

What does IPv6 mean to customers?

- Ideally IP is "behind the scenes" anyway
- IPv4 + IPv6 dual stack should be transparent to users
- Some bugs in some systems have led to some concerns
- Big issue for customers is what happens if we do nothing!
  - More and more NAT breaking things
  - More and more problems with networks with CGNAT
  - Just using IPv4 means changing the way IP works
- IPv6 brings back the original design principles
- IPv6 allows new and innovative products and services

What have AAISP been doing?

- Providing IPv6 connectivity for years
- Badgering router manufacturers to get ready
- Telling people about IPv6
- With FireBrick, making solutions for ISPs
- Testing with customers
  - Large and small
  - Techie and non techie
- Making it work!

Technical – IPv6 routers

- Simplest way is a stand alone DSL modem/router
- Chipset manufacturers are coming out with IPv6 now
  - Small solutions with DSL, WiFi, and even 3G
- Router manufacturers starting to use these
  - Very difficult to get then to fix anything
- Standards still being defined even now
  - RFC proposal on PPP IP6CP changes
- Routers just too flaky and badly implemented
- We hope this will take only months to get better

Technical – PPPoE bridge

- Most DSL routers can work as a PPPoE bridge
- Most ISPs can do PPPoEoAoDSL bridged
  - BT can, Be could but not new lines (only PPPoA)
- Allows PPP termination on device able to do IPv6
  - Linux box, pppd (some changes by A&A customers)
  - Stand alone routers (routerboard, etc)
  - FireBrick FB2700 (can also work with VLAN switch)
- Problem if you try and bridge multiple lines
  - Bridging does not work to same LAN segment

Technical – PPPoE modem

- Some modems will work as PPPoE endpoints
  - i.e. PPPoA or PPPoE on DSL to PPPoE on Ethernet
- Vigor V.120 is simple example, only PPPoA on DSL
- Same advantages as PPPoE bridging
- Allows multiple lines on same LAN as not actually bridged
- Allows use on PPPoA DSL (e.g. Be lines)

Problems with PPPoE

- Not a *one box* solution, no so good for consumers
- Multiple lines and PPPoE bridging is a problem
- PPPoE on DSL side is less efficient than PPPoA
- PPPoE on DSL is not always an option (new Be lines)
- PPPoE normally restricts to 1492 byte MTU
  - Reduced MTU has all sorts of issues
  - Can MSS fix-up TCP SYN packets
  - Even needed on IPv6 (arrrg!)
  - Can be avoided in some cases with baby jumbo frames
  - Lack of kit to allow larger frames though

Advantages of PPPoE

- No longer tying purchase of modem chipset to purchase of good router
- Modem/routers often poor IP stack and poor features
  - e.g. lack of IPv6 for a start!
- Router can often be free / open-source
- PPPoE is also standard for BT FTTC and FTTP
  - Support 1500 byte MTU, which is nice
- Single place to configure – ideal for multiple lines
  - i.e. easy to have *no config* spares of the actual modem

What are we doing now exactly?

- We want cheap domestic DSL routers that handle IPv6
- All new orders get an IPv6 allocation automatically
  - We have made changes to allow some older broken routers to still work (IPv4) when offered IPv6
- For businesses we have a PPPoE based solution
  - Using FireBrick FB2700
- For consumers we want the "free with the service" router to have IPv6 and enable the customers (firewalled) IPv6 LAN by default – still looking!

What is happens next?

- Crystal balls are not that reliable
- We hope there will be a lot of interest and take up of IPv6
  - In a few months all DSL routers do it anyway
  - ISPs starting to offer it, as new routers make it easy
  - End users using IPv6 without even realising
- What of major players, BT? Virgin?
- What of mobile networks?
  - 5 billion mobile phones!
- What if we do not get enough momentum?

IPv4 forever?

- The problem is that you can change the way the Internet works, abandoning key principles, and make it just web pages and email
- If you do that, you can make it all work with IPv4 and proxies and mapping and NAT and all sorts of mess
- Such things will work, so people will not change
- But they stop innovation and development of new services. Anything not compatible with this "dumbed down" Internet cannot take off

IPv6: The way forward

- The DSL routers are working on proper IPv6 blocks
  - They do have auto-delegation from the ISP
  - They do have firewalls by default
- As people replace routers anyway they will get IPv6
  - DSL routers do not last forever :-)
- Stuff being broken on IPv4 will help people asking ISPs
- NAT was the lazy solution for ISPs and customers
- For IPv6, proper IPs is the lazy solution, this is good!

Questions

Adrian Kennard:-

25 years experience in communications protocols

10 years experience with FireBrick (co founder)

- Coding Ethernet drivers, IP stack, TCP, HTTP, BGP, etc.

14 years experience with Andrews & Arnold (owner)

- Design and coding of systems and network
- Using linux and quagga routers initially
- Now all FireBrick based