ENIGMA CONSULTING

IPv6 from 50 thousand feet

Niall Richard Murphy
Enigma Consulting Ltd
HEAnet Conference 2002

Bargain

- No diagrams of header formats
- No slides reading "Drivers for IPv6"
- No marketing
  - Well, maybe just a little bit
  - But only of IPv6
- In return, your attention and suspended disbelief before dinner

What I'm going to talk about

- Why?
  - Not the usual reasons
    - Running out of address space
    - "It's, like, cool"
  - Savings: Money & Time
  - Useful features

What I'm going to talk about

- How?
  - Deployment or experimentation
  - Existing deployments
    - With a nod to our hosts
- When?
  - Recommendations

Why?

- The usual reasons
  - IPv4 has problems
    - Some of these will get worse over time
      - Address space exhaustion
      - Management scalability
  - Do it **now** because it's the latest and greatest thing

Why?

- My reasons
  - IPv6 saves money and time
    - Key features increase
      - Simplicity
      - Manageability
    - IPv6 has other features worth having
      - Security
      - QoS

Problems with IPv4

- Address space exhaustion
  - Prior to CIDR (~93) extremely inefficent address allocation
  - Better now, but the writing is on the wall
    - 36% of address space left (RIPE 43)
    - We have till 2004-2018 depending on whose curves you believe

Problems with IPv4

- Address space exhaustion
  - This is an ultimate deadline
  - Eventually you will need more and there won't be any there
    - More a problem for commercial ISPs than university end-sites
    - But it's still a problem for the people you are trying to reach, if not for you

Problems with IPv4

- Address space exhaustion
  - Key concept:
    - You will need to look at this for your network eventually
    - It might be within your tenure or might not, but why not be nice to your successor?

Problems with IPv4

- Management
  - "Man in a van" problem
    - Address management,
- Security
  - Not designed in
- Complexity
  - End-to-end philosophy is broken
  - Subtle capabilities have been lost

Problems with IPv4

- Management
  - Address management needs human oversight
  - Renumbering handled badly
  - Manual configuration generally required

Address allocation

- Address management is necessary
  - Automatic address assignment is broken
    - DHCP is deficient
    - DHCP servers require management
  - Prevent address collision
  - Security implications
  - Many deployments are statically addressed
  - Need staff to do this

Renumbering

- Renumbering doesn't happen often...
- But when it does
  - It's a network down-time event
  - It's very costly in terms of time and staff

Hands on configuration – why?

- Key concepts
  - Address management is broken
  - Renumbering requires manual attention
  - Nodes generally require some element of administrator attention

Problems with IPv4

- Security
  - Spoofing
    - Address spoofing
    - Packet synthesis
    - Plaintext authentication
    - Etc...
  - No built-in authentication/encryption

Problems with v4

- Complexity
  - IPv4 has been around a long long time
    - Feature creep
  - Philosophy of simplicity has been undermined
    - NATs
    - Firewalls, etc
  - We need to start again, taking into account lessons learned

IPv6 solves IPv4 problems

- Address space exhaustion
  - Solves it twice
    - By being bigger (128 bits)
    - By being cleverer about using address space(s)
      - Site- and link- locals
- Management
  - If it can be done automatically, it is
    - Stateless (and stateful) autoconfiguration
    - Router advertisement and solicitation

IPv6 solves IPv4 problems

- Security
  - IPSEC **mandatory**
  - Secures everything...
    - DNS, router solicitations, HTTP, everything
    - As long as the end-point supports it
- Complexity
  - Simplification of standards, header formats
  - No requirement for workarounds like NAT

Another way to put this

- IPv6 saves you money and time
  - Less hands-on management required
  - Management moved to routers and infrastructure servers rather than end hosts
- This directly translates to smaller staff costs
  - Less time is less money

Key IPv6 features to achieve this

- Stateless autoconfiguration
- Router solicitation and discovery
- Prefix deprecation
- Automatic key negotiation
- 6-to-4 transition mechanism

Stateless autoconfiguration

- Nodes generate addresses automatically
  - Combine network prefixes with MAC addresses
  - Network prefixes include "well-known" prefixes that mean
    - local link only
    - local network only
    - and any others you choose to provide
  - Routers provide prefix information

Router solicitation/discovery

- Nodes can automatically discover routing tables
  - Send "router solicitation" messages
  - Receive prefixes, MTU, TTL etc
  - Routers can deprecate prefixes for renumbering purposes

Key negotiation

- Inbuilt (IPsec) security has a notion of automatic keying
  - Nodes can transmit periodically (re)generated cryptographic keys automatically
  - Communication can be automatically secured without administrator intervention

6-to-4 transition mechanism

- An address space has been reserved for this interoperability method
- Embeds the IPv4 address in the IPv6 address
- Communicate via encapsulation
- All works automatically
- Downside: latency

How do I start understanding this?

- Decide whether to run trial production services or experiment
- You need a plan
- Recommendations
  - Start small
  - Very easy to incrementally add and change
  - Very easy to run in parallel

Do I do this quickly?

- There won't be an flag-day
  - Don't worry about having to synchronise deployment
  - Don't worry about having to solve all the problems in one go
  - Go for phased, gradual ushering-in

Production or experimentation?

- I'm not talking about moving your network over to IPv6 tomorrow
- I'm talking about the ease of running the two in parallel
  - It is surprisingly easy
- Question is whether you introduce in production environment or not

Production services

- Running production services
  - Most if not all IP protocol services can be run over IPv6
    - HTTP, DNS, SMTP...
    - All the core ones, most of the rest (including IM and other frivolities)
  - IPv6 can in many cases be introduced without affecting the stability of the service
    - Particularly using 6-to-4

Production services

- Gotchas
  - Apache 2 should be used for IPv6
    - Does not currently support mod_php or other important modules
    - Can run the two in parallel as necessary
  - BIND (for DNS) listens on all available IPv6 interfaces
    - Can't restrict it to (say) internal interfaces
    - There are other DNS servers

Production services

- Gotchas
  - Hosts require dual stacks
  - Older hosts will require patching/OS updates
  - Commercial products in many cases fall behind the free ones in implementing IPv6

Experimentation

- Minimal deployment of IPv6
  - One host and a tunnel
- Minimal useful deployment of IPv6
  - Several hosts, one router, and a tunnel
  - Allows you to experiment with the network management features
- Recommendation
  - Get addresses from upstream now
  - But don't let it delay you

Start Small

- Tunnel Brokers
- 6-to-4 transition mechanism
- Get addresses from upstream(s)

Tunnel Brokers

- A way to get address space to an end-node quickly and easily
  - HEAnet will be offering this shortly

6-to-4 transition mechanism

- Makes it very easy to start offering a few services
  - You don't even need an allocation from your upstream
  - You can embed the AAAA records in the DNS and see how this affects production networks

Grow when necessary

- New shipments of hosts will have v6 stacks pre-supplied
  - XP has it installed but you have to enable it
  - Next version of Windows will have it pre-enabled (fear!)
- You don't need to do anything for new equipment

Grow when necessary

- IPv6 stacks will learn most of what they have to know from their environment
  - Once you have a router announcing network policy, the clients will keep in line
  - Services can be offered locally and remotely as you wish

Grow when necessary

- Commercial support available
  - Consulting
    - Usual vendors
  - Products
    - Usual vendors
  - Your own staff may know more than you think
    - In-house support only may be viable

When to start?

- What are your production services?
  - If they are easy to switch, why not do that now?
  - You'll have to do it eventually, and the more experience you have at the time of switch, the better
  - It can be done at very low cost

When do I benefit?

- The more of a server farm which has IPv6 the better
  - The sooner you will feel the management benefits
  - Ultimately it will begin to be more cost-effective to offer services over IPv6 when the clients want it
  - Overhead you pay in software/training etc amortized over lifetime of protocol

When to finish?

- The IPv4 Internet will be with us for a long(++) time
  - This is a process with no flag-day
- A salutary thought:
  - 80% of everything is HTTP
  - 80% of HTTP is top 20 websites
  - Only the top 20 websites have to be reachable by IPv6 for over 50% of your traffic to be V6

Who else is doing this?

- Difficult to measure
  - Internal deployment only
  - Worldwide organisations with allocations may not be offering service here
- Methods
  - 6Bone
  - DNS
  - HEAnet MRTG

Irish deployment

- Our hosts...
  - http://mrtg.hea.net/mrtg/ip6/
    - ~250kbps at TCD, for example
  - Services offered over both protocols
    - ftp.heanet.ie for example
  - http://www.sixxs.net/ will be using some HEAnet space
- IPv6 assignments to most Irish ISPs

Irish deployment

- Plans afoot for
  - Eircom 6to4 router
  - INEX IPv6 traffic exchange
  - IEDR recently began testing AAAA DNS records
- Important to note
  - You aren't alone
  - There is community and upstream support

Worldwide deployment

- Commercial services in
  - America
  - Japan

References

- http://aso.icann.org/meetings/ga-2/presentations/mylotte/sld019.html

Thank you!

- Contact details:
  - niallm@enigma.ie
  - +353 87 220 8587
- I am available for childrens parties