

Deploying IPv6

kurtis@netnod.se

What we wanted to do

- Deploy IPv6
 - I.e make it work like IPv4 - just with IPv6
 - Seemed like such a straight-forward idea :-)
- The important thing
 - It is meant to be production, so it had to be production quality
 - All monitoring and statistics had to work too...

Problem split

- We have different categories of services that needed to be ported
 - The IX LANs
 - The services at the IX(es)
 - i.root-servers.net / TDLs / Anycast

IX LANs

- Very straightforward
- Apply for your IPv6 allocation from RIPE
 - Will be a /64 for a single LAN IX
 - Or /48
- We then assigned /64 per VLAN (there are two VLANs per site, MTU1500 and MTU4470)

IX LANs

- We are running IPv6 on the same infrastructure as IPv4
- Each ISP will get a static IPv6 address from each /64
 - For example 2001:7f8:d:ff::73/64
 - The last “octet” matches that of IPv4

Netnod services

- Netnod as LIR “acquired” an initial allocation of a /32 from RIPE under the current policy
- Then we made an address plan
 - Somewhat interesting starting from scratch
 - We split out one /48 per city using binary-chop
 - And left space for (services-) customer allocations

The steps...

- I. Enabled IPv6 on the infrastructure
 - Loopback interfaces
 - P2P Links
 - LANs
 - Established iBGP sessions and set up OSPFv3

The steps

- 2. Enabled IPv6 on office LANs, and some servers
 - Services where given static addresses. I.e 2a01:3f0:1:3::101.
 - Office LAN is given addresses via RA
 - DNS resolving is done over IPv4 (we are looking at DHCPv6 - now we just need a client...)
 - Added AAAA for public names, i.e www, mail, etc

The steps

- 3. Monitoring
 - Nagios is used for monitoring using 2.10
 - Initially ran into problems with perl libraries (Net::DNS) needed to support IPv6 addresses
 - Doubles the number of alarms
- 4. Then the problems started...

Problems

- Routing
 - Native or not - it's a jungle. We are still seeing really wired routing issues. And we have spent significant efforts on debugging routing all around the world
- Customers reporting faults
 - Related to above, but it turns out to be really hard to be (seemingly) alone in trying to offer a production quality service. I.e first of all to find someone at a provider that knows they are running IPv6 followed by convincing them to fix problems is hard

The problems

- Vendors
 - Needs to be able to check the “IPv6 ready” box for US DoD contracts
 - Only assures the box actually forwards packets with IPv6 packet headers
 - Seems like vendor regression testing consists of “it compiles”
 - We still lack those 20 years of usage that made RFCs into software that you can operate
 - A lot of relearning needs to be done
 - I am ignoring all software bugs we found

Lessons learnt

- The IX side is very straight forward
- But some better debugging for example NDP snooping would have been nice
- Discussion on same vs. dedicated VLAN
 - Again, we see this as production and want our customers to do the same
 - 18 members have IPv6 IX addresses

Lessons learnt

- EUI64 addresses are meaningless except perhaps on workstations
 - Impossible to debug, maintain reverse DNS for and changes when you replace hardware
- Router interfaces, loopbacks, and servers all have static interfaces.
- P2P links set to /64. No real reason...

Lessons learnt

- Make sure you don't fall into known traps as you start adding AAAA records
- We found it easier to threat IPv6 just as IPv4. Same processes, same set-up etc.
- Seems to be great interest in the topic among members
- Organising a workshop on Wednesday with 75 attendees - for multicast we had 20

Left to do

- Upgrade standard kernels on machines running Quagga
- Install / Activate IPv6 in production for i.root-servers.net (already running for Unicast DNS zones)
- Pick a set of IPv6 capable anycast hosts and start anycasting

Left to do

- The other cities will be interesting
 - Either announce more specifics, i.e /48s and hope they get propagated (likely)
 - Or build tunnels back to Stockholm where we announce the /32
 - Or most likely - both
- Training, documentation, debugging, and training
 - <http://www.6diss.org/e-learning/>

Summary

- It wasn't hard
- It does take planning
- It will take a lot more time than you think - due to details annoying differences and debugging
- It actually works...
- We will do a more detailed writeup and post to the Wiki at <http://www.civile-tounge.org/6and4>

Questions?