

IPv6 Investigation: Progress to December 2015

IPv6 investigation areas

- Initial questions
- DNS
- Edge switches
- Core switches and routing
- iptables
- Linux routing
- (DHCP is a followup project)

DNS

- Development meeting decided no need to link IPv4 and IPv6 address
- New tool written to make forward/reverse zone creation simpler
- `rfe dns/inf6`
- Simple syntax, described in the file as well as in the final report
- Forward (`inf.ed.ac.uk`) zone is being populated
- Reverse zones are being created but haven't been delegated to our NS yet

Edge switches

- RA parameters set per-VLAN
- RA enabled on switches doing IPv6 forwarding
- Disabled everywhere else (though configured where possible)
- RA-guard enabled on all untrusted ports
- Manager-addrs list can now take IPv6 addresses. Set by hand for now to avoid security holes. Will be done through the tools later.
- MLD-snooping to come

Core switches and routing

- core[012], atc[01], cs[01] are all now doing IPv6 routing.
- All have addresses on some carefully-chosen VLANs
- Speaking OSPFv3 (*not* authenticated)
- Inf-unit instructions written based on the first couple and debugged on the rest

iptables

- Component was mostly already there
- Existing rules audited for IPv6-safety, and new ones created where necessary
- All “generating” files (“g.*”) now test for IPv6 and adjust their output accordingly
- Rules are in <stable>. Not running everywhere yet, as turning them on will cause some IPv6 addresses to appear – it’s slightly too early for that yet.

Linux routing

- Using BIRD, as quagga doesn't appear to do OSPFv3 areas
 - That may not actually be a problem if we speak BGP to the EdLAN routers
- Component written, and running on test routers
- Should support IPv4 too, though not tested yet
- SL6 only for now, with the component starting the daemons
 - But written with init and/or systemd in mind

Still To Do

- EdLAN
 - Routing protocols
 - ABRs
 - Turn it on!
- Extend testing to all managed machines
- Auditing tools
- Another blog article

Follow-on projects

- DHCP for IPv6
 - Too big a job to incorporate into an investigation
 - Not really an investigation any more anyway!
 - Spread the knowledge
- Self-managed machines
 - Depends on DHCP and audit tools

Implications 1

- SL6 machines (mostly) have IPv6 disabled
- SL7 machines have IPv6 enabled
 - So any on a VLAN on which RA has been enabled (32, 33, 202, 216 so far) will acquire global IPv6 addresses
 - They may try to use these to speak to the outside
 - This won't work until we have external routing going, but should be generally OK after that
 - /etc/gai.conf will need some tweaks to suit

Implications 2

- If you want to advertise your machine over IPv6 then you will need to add a DNS entry for it
 - There is *one* DNS namespace with both IPv4 and IPv6 addresses merged together
 - Setting a static IPv6 address on a machine is currently inconvenient (MPU?)
- If there's no IPv6 entry in our DNS then nothing outside will know to speak to an IPv6-enabled machine
 - Including any (SL7) which have autoconfigured as a result of seeing RA multicasts
- Once there is an IPv6 address added for a machine then it's fair game for *any* service on that machine
 - You can't pick and choose
 - Edge filter holes will be automatically created
 - Expect to see about 10% native IPv6

Implications 3

- TEST your services with IPv6!
 - The outside world won't distinguish between IPv4 and IPv6, and will expect both to Just Work
- Any local access controls should be reviewed to ensure they are IPv6-friendly
 - IPv4 and IPv6 may work independently
 - “129.215/16 = EdLAN” will be broken
- Beware of “old methods” which undo your IPv6 settings
- Once we enable IPv6 for self-managed machines, they will start to use it to speak to you too
- Once IPv6 is enabled on your machine, IPv4 and IPv6 have equal status so far as everyone else is concerned.

Bedtime reading

... is linked from the project's index page
<<http://www.dice.inf.ed.ac.uk/units/infrastructure/Projects/352-IPv6/>>