

Thorsten Kuefer
JOIN/University of Muenster (DE)
2005-08-01

General problem statement, Prior
art/tools, Backbone renumbering
and BGP issues, SOHO



Project aim

- Why renumber?
 - No provider independent (PI) address space for IPv6
 - New address assignment (RIR -> LIR -> Customer), changed conditions
 - Migration from transition method to native IPv6 connection
 - Thus likely to be more common for IPv6 networks
- Determine relative ease of renumbering IPv6 networks compared to IPv4
 - Has IPv6 design advantages for renumbering?
 - Capabilities of existing tools?
 - Seamless renumbering without network outage possible?
- Examine backbone, SOHO and enterprise environments
 - Test IPv6-specific features supporting renumbering
 - Follow Baker procedure and observe behaviors
(*draft-ietf-v6ops-renumbering-procedure-05*)

Tools for renumbering

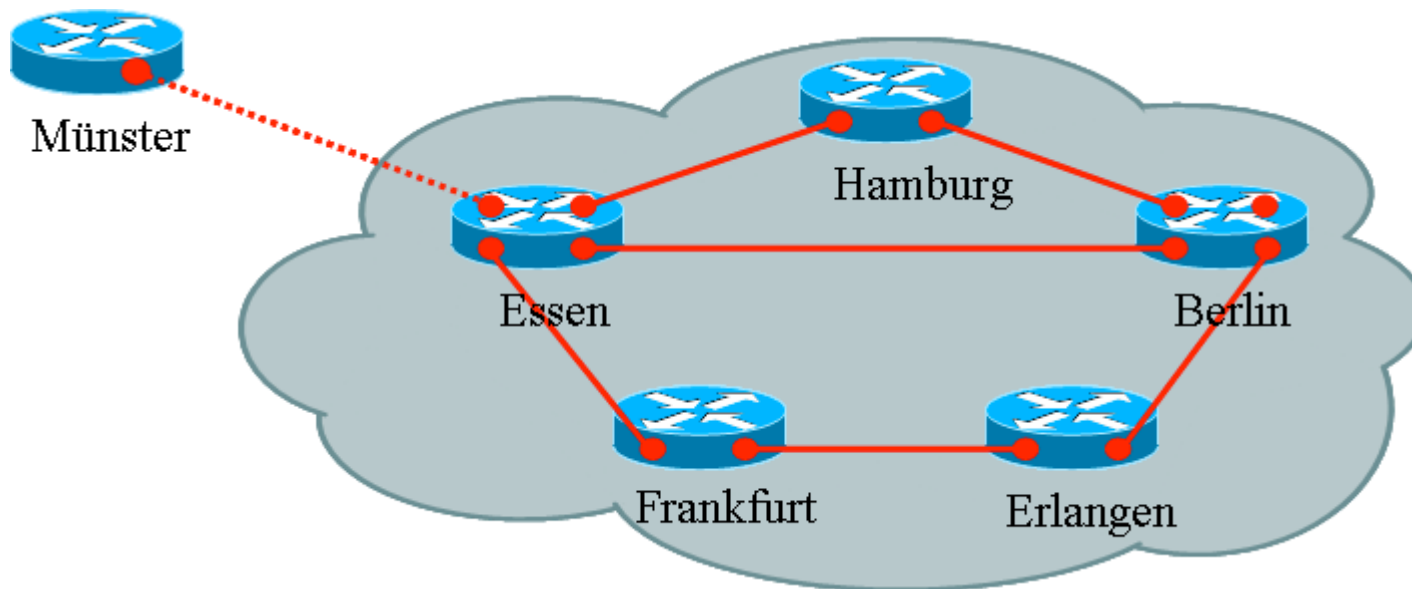
- Searched for prior work in the area
 - IETF PIER WG, c. 1998 => RFC1916, 2071, 2072 (IPv4 renumbering)
 - RFC2072 “Router Renumbering Guide” gives some operational procedures much as they are in Baker’s draft for IPv6.
 - RFC2072: renumbering networks whilst remaining the same hierarchy of subnets is the 'easiest' scenario to renumber; when each "old" (IPv4) prefix can be mapped to a single "new" (IPv6) prefix of same length
 - PIER: transition from IPv4 to IPv6 addressing was considered a renumbering scenario, we strictly consider only the renumbering from IPv4 prefixes to other IPv6 prefixes as renumbering
- IPv6-specific tools
 - Auto-configuration (RFC2462)
 - Multi-addressing (RFC3513)
 - Router advertisement options & lifetimes (RFC2461)
 - Default Address Selection (RFC3484)
 - DHCPv6 Prefix Delegation (RFC3633)
 - IPv6 Router Renumbering (RFC2894)
 - ULAs (draft-ietfipv6-unique-local-addr-09)

Backbone

- Only consists of routers doing data transfer and exchanging routing information
- No services running
- Backbone has /48 prefix, each router has a /56 subnet, routers connect via /112 subnets
- Delegates /48 prefixes to customers connected via IPv6-in-IPv4 tunnels
- Interfaces, routing and DNS must be updated

Testing scenario

- 6WiN: German IPv6 NREN backbone
- Backbone address space 2001:638:f::/48



Results

- Update of DNS (incl. reverse zone and timers) and BGP peerings is a lot of manual work and takes time (third parties)
- Interface IPv6 addresses and routing configuration have to be updated step by step (simply duplicating does not work, creation of new loopback interfaces)
- BGP peerings partly need new update sources to come up
- Static routes have to be updated manually unless using interface names
- Renumbering took four weeks (mainly because of 40 customers), could be possible in 1-2 days

Web resources

- 6NET deliverables

- <http://www.6net.org/publications/deliverables/D3.6.1.pdf>
- <http://www.6net.org/publications/deliverables/D3.6.2.pdf>

- Cisco deliverables

- To be announced

Summary

- Motivations for IPv6 renumbering
- Current tools and experiments
- Backbone & SOHO scenario

- Questions? General Q&A scheduled for later...

- Univ. Southampton will continue with enterprise scenario and recommendations