Renumbering still needs work

draft-carpenter-renum-needs-work

http://tinyurl.com/numnum

Brian Carpenter (U of Auckland)
Ran Atkinson (Extreme Networks)
Hannu Flinck (Nokia Siemens Networks)

March 2009

Like it or not, site renumbering is needed (sometimes). The most demanding case would be unplanned automatic renumbering.

The network will be down for cleaning today

Objectives of the draft

Considering both IPv4 and IPv6:

- Summary of existing renumbering mechanisms
- Description of current operational issues with renumbering
- Summary of relevant work in progress
- Gap analysis

➔ May lead to suggestions for future work, and/or operational recommendations.

Existing Host-related Mechanisms

- DHCP and DHCPv6
  - "Strong" IP asset management. Site has a central database with MAC addresses, admin info, plug #, and uses this to generate IP addresses, DHCP, DNS, ACLs...
  - "Weak" IP asset management. No database, FCFS addresses from DHCP, DNS and ACLs maintained manually.
- SLAAC (IPv6 stateless address autoconfig)
  - Hosts inherit subnet prefix from their local router.
  - Designed for unmanaged, unattended automatic configuration.
- PPP
  - IPv4: the server end of PPP assigns subscriber address
  - IPv6: PPP only assigns interface-identifiers. DHCPv6 or SLAAC is used to assign subscriber address.

DNS aspects, SLP

- It's elementary that you shorten DNS TTLs before renumbering
- You want to generate DNS and DHCP from the same source (database)
- Dynamic DNS and DNSSEC are needed if you want real automation
- SLP, or SRV records, should help with server renumbering.

Existing Router-related Mechanisms

- Router renumbering for IPv6 via DHCPv6 Prefix Delegation [RFC3633]
- ICMPv6 extension to allow router renumbering [RFC2894] (not used??)
- IPv6 RAs can carry default router preferences and more-specific routes [RFC4191] (not used??)
- IPv4??

Multi-addressing for IPv6

- IPv6 was designed to allow multiple prefixes per subnet and therefore multiple addresses per host.
- Yes, that has some issues (glitches in RFC3484 address selection rules, and issues for exit router selection, ISP ingress address filtering, and traditional TE).
- But it allows overlap between old and new address plans during renumbering. Avoids a flag day.
- Also allows use of ULAs (unique local addresses) for invariant internal addressing (e.g. for network management, printers)

But there's a basic design flaw

- It's obvious that you should shorten address lifetimes prior to renumbering, but
  - IP addresses do not have a built-in lifetime.
  - Even when an address is leased for a finite time by DHCP or SLAAC, or when it is derived from a DNS record with a finite time to live, this information is lost once the address has been passed to an upper layer by the socket interface.
  - Thus, a renumbering event is almost certain to be an unpredictable surprise from the point of view of any software using the address. Many of the issues below derive from this fact.
  - Don't expect this bug to be fixed any time soon.

Operational issues

- Host-related
- Router-related
- Other
  - NAT state issues
  - Mobility issues
  - Multicast issues
  - Management issues
  - Security issues

Host issues

- Network layer *should* do the right thing when DHCP or SLAAC is updated.
  - With "weak" asset management, some confusion seems inevitable, especially around servers.
  - Note that many DHCP options carry addresses around
  - The M/O bit ambiguity in the interaction between DHCPv6 and SLAAC will cause problems during renumbering
  - Embedded systems may need manual or ROM updates
- TCP and UDP sessions break. SCTP might survive.
- DNS - prone to administrative errors and TTL override
- Applications that remember addresses will break.
  - Notorious example: software licences keyed off the IP address.

Router issues

- RFC2072 (from 1997) discusses issues.
  - Some improvement since then (DHCP was still young)
  - Systematic planning and administrative preparation is needed
  - All forms of configuration file and script must be reviewed
  - Addresses are cached in routers - routers may need to be restarted
  - Addresses used by configured tunnels and VPNs may be overlooked, although secure tunnels configured by FQDN are fully standard [RFC2407, RFC4306].

NAT state issues

- Entries in the state table of any NAT that happens to contain renumbered addresses will become invalid before they time out. (Doesn't matter too much, since TCP and UDP break anyway.)
- A NAT itself may be renumbered and may need a configuration change

Mobility issues

- A Mobile IP node will be affected if either its current care-of address or its home address is renumbered.
- Mobile IPv6 will recover *except* if it is disconnected at the moment of renumbering. In that case, it has to use DNS to find its home agent again.
- Mobile IPv4 will not normally recover until the mobile node is back on its home network again.

Multicast issues

- IPv6 multicast actually helps renumbering due to the SLAAC discovery mechanisms.
- However, there are issues due to use of IPv6 unicast addresses in the Rendezvous Point and Source Specific Multicast mechanisms.
- IPv4 multicast: TBD

Management issues (1)

- Static addresses are routinely embedded in configuration files and network management databases, including MIBs.
  - Ideally, all these would be generated from a "strong" site asset management database.
- Because of routing policies and VPNs, a site may embed addresses from other sites in its own config data. Thus renumbering will cause a ripple effect for a site's neighbours.
- Some config data may be very hard to find, e.g. configs for building routers, printer addresses configured by individual users, and personal firewall configs.

Management issues (2)

- FQDNs rather than IP addresses wherever possible in config files & databases might mitigate the issues.
  - But there's 20 years of history of not doing that.
- Administration issues (i.e., tracking down, recording, and updating all cases where addresses are stored rather than looked up dynamically) are the dominant concern of managers considering renumbering.
  - Only a "strong" IP asset management tool and database can mitigate this.
- There's a risk element stemming from the complex dependencies of renumbering: it is hard to be fully certain that renumbering will not cause unforeseen service disruptions.

Security issues

- IPv6 addresses are intended to be protected against forgery by SEcure Neighbor Discovery (SEND) [RFC3971]. But SEND appears to be very difficult to actually deploy and operate.
- Firewall rules need to be updated, and any other cases where addresses or prefixes are embedded in security components (ACLs, AAA systems, IDS, etc.)
- Problem if an X.509v3 PKI Certificate includes a subjectAltName extension containing an IP Address.
- Spam white lists need to be updated.
- DNSSEC is needed, to make security folk less nervous about using FQDNs.

Mechanisms in the IETF mill

- SHIM6 - intended to help multihoming, but would also simplify address overlap during renumbering
- MANET and AUTOCONF - such networks demand automatic addressing and routing setups. Maybe the mechanisms can be generalised? But this work is going very slowly.
- draft-dec-dhcpv6-route-option
- NETCONF - secure remote config
- NSCP (nameserver control protocol) - based on NETCONF

Gap analysis (preliminary) (1)

- Host related gaps:
  - FQDN based socket API or FQDN based transport layer (to alleviate application layer issues)
  - Multipath survivable transport protocol(s)
  - Single registry per host for all address-based configuration
  - Deploy DHCP FORCERENEW and DHCPv6 RECONFIGURE for bulk renumbering.
  - IPv6 ND M/O flag debate to be resolved
  - IPv6 hosts should be able to learn "liveness" of upstream prefixes

Gap analysis (preliminary) (2)

- Router-related gaps
  - A non-proprietary secure mechanism to allow all address-based configuration to be driven by a central repository for site configuration data. NETCONF might be a suitable basis.
  - A MANET solution
    - solid enough to use on operational small to medium non-mobile sites, for voluntary or involuntary renumbering;
    - possibly also for voluntary renumbering of large sites.
  - Short-term, make [RFC2894] and [RFC3633] router renumbering operable.

Gap analysis (preliminary) (3)

- Operational gaps
  - Deploy DNSSEC and DynDNS
  - Deploy multi-prefix usage of IPv6 (as an aid to renumbering)
  - Document and encourage systematic site databases and secure configuration protocols for network elements and servers (e.g., NETCONF).
  - Document functional requirements for site renumbering tools or toolkits.
  - In general, document renumbering instructions as part of every product manual.
  - Recommend that all IPv6 deployment plans include a strategy for eventual renumbering.

Gap analysis (preliminary) (4)

- Other gaps
  - Secure mechanism for announcing changes of site prefix to peer sites and in public.
  - For Mobile IPv6, better mechanism to handle change of home agent address while mobile is disconnected.

Input requested

- http://tools.ietf.org/id/draft-carpenter-renum-needs-work
- Please read the draft, and email your comments (errors, omissions, suggested text)
  - write to the authors, or the ops-area@ietf.org list