INTERNET DRAFT   C. Huitema
<draft-huitema-ipv6-renumber-00.txt>   Microsoft
Expires January 12, 2002   July 12, 2001

IPv6 Site Renumbering

Status of this memo

Abstract

There has been recently a lot of the discussion in the IPNG, NGTRANS and DNSEXT working group about the level at which IPv6 shall support renumbering. A specific question is whether we need special support in the DNS to enable renumbering, as specified in [RFC2874], or if the simpler mechanisms specified in [RFC1886] are sufficient. In order to organize the discussion, this memo presents a set of realistic renumbering scenarios, discusses the possible frequency at which such scenarios can be repeated, presents some tools that can be used to organize the renumbering, and summarizes the operational requirements that have to be met by any renumbering solution.

1 Introduction

There has been recently a lot of the discussion in the IPNG, NGTRANS and DNSEXT working group about the level at which IPv6 shall support renumbering. A specific question is whether we need special support in the DNS to enable renumbering, as specified in [RFC2874], or if the simpler mechanisms specified in [RFC1886] are sufficient. In order to facilitate the discussion, this memo presents a set of realistic renumbering scenarios, and then analyzes requirements and potential solutions.

The purpose of the exercise is to evaluate how the current "IPv6 and

DNS toolbox" can be used to facilitate renumbering. In addition to the two possible DNS formats mentioned above, the toolbox includes

"IPv6 Stateless Address Autoconfiguration" [RFC2462] and "Router Renumbering for IPv6" [RFC2894]. We do not detail here the operation of the DNS in each of the scenarios - this is left as an exercise for the DNSEXT working group.

2 Description of the scenarios

In this version of the memo, the renumbering scenarios are broadly sketched. For example, we say nothing of "static address filters" used for QoS and security purposes; we may guess that these could be updated as a side effect of router renumbering, but we would be better off with a real specification. The purpose of the exercise, however, is to provide five realistic renumbering scenarios.

2.1 Scenario 1, first connection

A site is currently isolated. The internal subnets have been numbered using "site local" addresses. The site joins the IPv6 Internet. The site managers use "Router Renumbering for IPv6" [RFC2894] to automatically inform the internal routers that they should start advertising the new prefix. The hosts receive a router advertisement and automatically create a global address as specified in [RFC2462].

Most sites will do this once. In many case, the first connection of a site to the IPv6 Internet will be through a tunneling solution, such as "6to4". For the purpose of the exercise, we will consider a tunneled connection as just another connection, that happens to use a virtual link instead of a dedicated interface.

2.2 Scenario 2, disconnection

A site is currently connected to the Internet. The site managers plan to disconnect. This occurs in two phases, first deprecating the old prefix, then removing it. Both phases are implemented using "Router Renumbering for IPv6" [RFC2894] and "IPv6 Stateless Address Autoconfiguration" [RFC2462].

2.3 Scenario 3, multi-homing

A site is connected to the Internet through a single provider. The site managers set a contract with another provider, and obtain a new prefix. The site managers use "Router Renumbering for IPv6" [RFC2894] to automatically inform the internal routers that they should start advertising the new prefix. The hosts receive a router advertisement and automatically create a second global address as specified in "IPv6 Stateless Address Autoconfiguration" [RFC2462].

2.4 Scenario 4, removing a provider

A site is connected to the Internet through two providers. The site managers want to terminate the contract with one of these providers.

This occurs in two phases, first deprecating the old prefix, then removing it. Both phases are implemented using "Router Renumbering for IPv6" [RFC2894] and "IPv6 Stateless Address Autoconfiguration" [RFC 2462].

2.5 Scenario 5, time-of-day preference

A site is connected to the Internet through two providers. These providers use different tariffs. The site managers desire that one of the providers be preferred during working hours, say from 9:00 am to 5:00 pm, and another be preferred during the rest of the day. They use "Router Renumbering for IPv6" [RFC2894] at critical times (9:00 am, 5:00 pm) to deprecate one of the global prefixes and promote the other. The hosts receive router advertisements and heed them as specified in "IPv6 Stateless Address Autoconfiguration" [RFC2462].

There are few sub cases here:

- time of day preference along with actual renumbering of hosts
- time of day preference only reflected in DNS (no renumbering)
- time of day preference for subset of services
- load balancing by advertising subset of links

The second case is the one that probably is going to be used most as that has the lowest impact and eliminates the DNS TTL issue that a host can remove address before the last cached DNS entry has expired.

3 Renumbering requirements

The discussion of the renumbering scenarios in the IPNG and NGTRANS working group unearthed a number of operational requirements that must be met by any renumbering solution. These requirements include continuous addressability, DNS security, network stability, and also a minimum frequency. In addition, we list here a non-requirement, the automatic support of the fusion between several sites.

3.1 Continuous addressability

Since DNS records may linger in various caches for the duration of their TTL, IPv6 addresses should remain valid for at least as long as the TTL of the DNS record. In fact, we observe that it is also desirable to maintain usability of an "old" prefix for some time after it has ceased being advertised in the DNS, in order to allow existing connections to terminate.

This is addressed in the scenarios 2 and 4 through a two phase approach: first deprecate a prefix, then at the end of the TTL remove it. In IPv6, if an address prefix is deprecated, the host can continue using it for existing connections or existing associations, but should use only the preferred prefixes when initiating new connections. In order to meet the TTL requirements, the hosts not refuse new connections on deprecated prefixes.

3.2 DNS server load upon renumbering

When it comes to creating new addresses, or deprecating them, we really have two choices. One possibility is to let the hosts use dynamic DNS updates to create or update AAAA or A6 records on the fly; another possibility is to have the site managers update the AAAA or A6 records in a reference file. We have to analyse the benefit/cost of AAAA/A6 in this context.

A particularly nasty consequence can occur if many hosts create new addresses and attempt almost simultaneous DNS updates. This phenomenon is discussed in [DISPRE], and a possible solution is presented in conjunction with the use of A6 records.

3.3 The DNS security requirement

One additional concern is the interaction between renumbering, AAAA, and DNSSEC -- specifically, the cost of re-signing a zone with new addresses. The careful system administrator would do this after the new prefix was known, but before the new prefix started to be used. The effort required to re-sign scales linearly with the number of RR's changed. Fortunately, this task is parallelizable; however, the processors doing the work must be trusted with the zone's private key. Folks with appropriate levels of paranoia likely won't want to do much else with this hardware besides maintain the zone(s).

As an unscientific test, during the DNSEXT meeting in Minneapolis Bill Sommerfeld took the mit.edu zone (with about 82000 hosts), synthesized AAAA records for all hosts, and signed it using the tools included with a recent bind 9 release. His recollection was that signing the synthesized zone took roughly 90 minutes on his laptop -- a 333mhz Celeron, which averages to about 1000 signatures per minute on this system, or maybe 3000 signatures per minute per GHz of CPU. In the absence of DNAME, a roughly similar re-signing effort is required for PTR zones: we would thus need two signatures per address, one for AAAA, one on PTR. In these conditions, we are down to 1500 addresses per minute per GHZ of CPU.

Renumbering a million-address network would take a bit over 11 GHz-hours of cpu time just for the dnssec signatures alone; whether anyone would actually want to renumber a million-nodes network is indeed debatable.

Note that resigning needs to be complete before the RR's can be replaced -- i.e., the time for renumbering to be complete is the resigning time plus the TTL...

3.4 Name servers

To avoid cyclical references or "can opener-in-can" situations for records pointed to by NS records, name servers basically must have address records that provide their entire address, i.e. either AAAA record or A6 records with a null length prefix. This means that, in any renumbering scenario, individual records will have to be published for the site's name servers, even if A6 is used.

3.5 Non requirement: fusion of sites

During the mailing list discussions, it was decided to not consider a very specific type of renumbering: the merging of independent sites. This is a noticeably different case, where the internal numbering of a site may need to be radically altered, and a new addressing plan needs to be created.

Fortunately though, this one generally is a rare event, is usually known and can be planned for well in advance, and the disruptions that occur are usually occurring in all kinds of other fields as well, not just the network, so people tend to be a little more forgiving (eg: people are more likely to curse when the merged payroll division doesn't manage to get anyone's salary paid on time, than when the net is flaky for half a day due to the numbering changes not having propagated properly).

4 Security Considerations

This memo presents renumbering scenarios. Renumbering has implications on security, since it forces the use of new addresses and may invalidate previous bindings between names and addresses. Secure bindings may require the use of DNS security; the effects of renumbering on DNS security is discussed in section 3.3.

5 IANA Considerations

This document does not call for an IANA action.

6 Copyright

The following copyright notice is copied from RFC 2026 [Bradner, 1996], Section 10.4, and describes the applicable copyright for this document.

7 Intellectual Property

The following notice is copied from RFC 2026 [Bradner, 1996], Section 10.4, and describes the position of the IETF concerning intellectual property claims made against this document.

The IETF takes no position regarding the validity or scope of any intellectual property or other rights that might be claimed to pertain to the implementation or use other technology described in this document or the extent to which any license under such rights might or might not be available; neither does it represent that it has made any effort to identify any such rights. Information on the IETF's procedures with respect to rights in standards-track and standards-related documentation can be found in BCP-11. Copies of claims of rights made available for publication and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the IETF Secretariat.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights which may cover technology that may be required to practice this standard. Please address the information to the IETF Executive Director.

8 Acknowledgements

This memo incorporates text submitted to the working group lists by Bill Sommerfeld, Robert Elz, and Jun-ichiro itojun Hagino. Olafur Gudmundsson was instrumental in prodding the author to submit it

before the deadline.

9 References

[RFC2874] M. Crawford, C. Huitema, "DNS Extensions to Support IPv6 Address Aggregation and Renumbering", RFC 2874, July 2000.

[RFC1886] S. Thomson, C. Huitema, "DNS Extensions to support IP version 6", RFC 1886, December 1995.

[RFC2462] S. Thomson, T. Narten, "IPv6 Stateless Address Autoconfiguration", RFC 2462, December 1998.

[RFC2894] M. Crawford, "Router Renumbering for IPv6", RFC 2894, August 2000.

[DISPRE] M. Crawford, "Discovery of Resource Records Designating IPv6 Address prefixes", Work in progress, November 2000.

10 Author's Addresses

Christian Huitema
Microsoft Corporation
One Microsoft Way
Redmond, WA 98052-6399

Email: huitema@microsoft.com