Study of alternatives to renumber IPv6 networks

Juan Francisco Rodríguez Hervella
Departamento de Ingeniería Telemática. Universidad Carlos III de Madrid
Avenida de la Universidad, 30. Edificio Torres Quevedo.
E-28911 Leganés (Madrid)
Teléfono: 91 624 87 56
E-mail: jrh@it.uc3m.es

Abstract *Renumbering can be defined as the transition from the use of an existing prefix to a new prefix in a network. This paper analyzes different scenarios where planned IPv6 network renumbering events might happen and it points out different mechanisms that are available to the system administrator, in order to minimize the burden of managing the renumbering process. This job is focused on both the renumbering of site routing infraestructure and the renumbering of the related DNS information in IPv6 managed networks. Pros and cons are identified for every presented solution. Before describing the mechanisms currently available for IPv6 renumbering, a brief summary of the studies that have been developed for the IPv4 world are introduced.*

1 Introduction

In [1] renumbering is defined as the transition from the use of an existing prefix to a new prefix in a network. [2] also gives a simple definition of renumbering for IPv4 networks, saying that the exercise of renumbering a network consists of changing the IP host addresses, and perhaps the network mask, of each device within the network that has an address associated with it.

The Internet Architecture Board (IAB) recommends in [13] that to make renumbering more feasible, "all designs and implementations should minimize the cases in which the IP addresses are stored in non-volatile storage maintained by humans, such as configuration files". Note that for some applications, addresses might be used during a short lifetime as if they were hardcoded, though it is also recommended to rely on the DNS infrastructure whenever it is possible, at least when the communication is being stablished.

The paper is organized as follows. In the first section a qualitative evaluation of the studies that have been made about this topic for IPv4 networks, mostly developed inside the Internet Engineering Task Force (IETF)[1], are discussed. The next section is devoted to introducing the problem in the context of the new IPv6 protocol. Two mechanisms are briefly examined, the router renumbering protocol [12] and the DNS extensions to aid the renumbering process, defined in [12]. As the last section regarding to the IPv6 renumbering "state of the art", a paper which has been recently published as a personal draft inside the IETF is also commented. This paper is intended to be an operational guide to develop a smooth renumbering transition in a step-by-step way. Finally, conclusions and future works are addressed from the point of view of a system administrator, which is the person in charge of facing this problem.

2 Renumbering in IPv4

Renumbering in IPv4 was identified as an operational issue when CIDR [3] [4] [5] came up, and the "Procedures for Internet/Enterprise Renumbering" working group was created inside the IETF (PIER WG) to help administrators to better manage their networks. This working group focused in "identifying processes and procedures, tools and techniques for renumbering in both the IPv4 and IPv6 environments". It was conclude in June, 1998 with three "request for comments" (RFCs):

- "Enterprise Renumbering: Experience and Information Solicitation" [6]
- "Network Renumbering Overview: Why would I want it and what is it anyway?" [2]
- "Router Renumbering Guide" [7]

These documents are all informational, and they only talk about IPv6 on the corner. Besides, The PIER working group states that it is not concerned about the development of protocols to aid the renumbering process, so they do not define any automatic mechanism to ease this process.

[6] requested help from the internet community to know how renumbering was undertaken in IPv4, focusing on the specific case of "single-homed networks that are not transit providers", which was

[1] This research paper is based on the public information which is available at the IETF, and documentation outside this scope is not analized.

present its own mechanisms to face the renumbering problem. [6] also pointed out areas where some tools could be used to make renumbering easier, asking for input from the internet community as if it was a survey to disclose what tools either were being used or could be used.

[2] states the reasons by which a site should be renumbered. Some of these reasons are still valid for IPv6 as well as for the actual IPv4.

Finally, [7] talks about router renumbering but it is limited to identify what things should be upgraded on routers.

To sum it up, the IETF is nowadays not aware of any automatic mechanism/tools to aid the renumbering process, though the requisites, objetives and desired outcomes are fully defined.

3 Renumbering in IPv6

Currently there is almost no documentation about how to make renumbering a successful story in the IPv6 world. The one good thing about repeating your mistakes is that you know when to cringe. It is believed that the only mistake that has been done with the treatment of renumbering in IPv4 is that there is no automatic mechanism to help administrators who are in the situation of facing the renumbering process.

On the other hand, IPv6 has some automatic mechanisms such as address and router autoconfiguration, which try to ease administrator's lives by the means of fighting against the burden of the typical IP set up process. With IPv6, there is an opportunity to both define and specify tools and protocols to try to cope with a lot of hand-made tasks, including the processes and steps involved in a renumbering event.

3.1 Why renumbering in IPv6 might be necessary

This section describes the reasons that push the renumbering process in IPv6 networks.

[2] states that

> "end users cannot assume they 'own' address allocations (...). Rather, end users will "borrow" part of the address space of an upstream provider's allocation."

which introduces the concept of "address lending". This was not the situation in the early days of IPv4, but routing scalability issues plus shortage of addresses made the policy of address asignement stiffer. Althought [2] talks about the IPv4 world, in IPv6 we currently face the same music. [2]

capabilities of CIDR are based on the assumption that address allocation reflects network topology as much as possible, to enable both sites and providers to act as aggregators [8]. While CIDR does not require every site that changes its providers to renumber, currently there are not any provider independant address schemes for IPv6, implying that IPv6 addresses are delegated secuencially from one provider to the next in a top down manner. This address space is called "provider aggregatable", and it is defined in [9]. To sum it up, this means that portability of addresses from one provider to another is not allowed, so sites that change of ISP will have to renumber its whole network.

Besides the change of internet service provider, other minor reasons to re-address are defined in [2]. In short:

- Change in organizational structure or network topology.
- Change from private addressing schemes to public/global addressing ones.

3.2 Premises

[2] refers specifically to IPv6 pointing out that

> "at the very least, DNS hosts will need to be reconfigured to resolve new host names and addresses, and routers will need to be reconfigured to advertise new prefixes."

The author believes that the list of network devices that should be renumbered to get a minimun network functionality as well as a smooth transition to the new addressing scheme are the following items:

1. Routers.
2. Hosts.
3. DNS servers and information related to it.

In the next section, it is supposed that the hosts get their addresses using the stateless autoconfiguration mechanism which is provided by the transmition of router advertisements [10], meaning that with the mechanism defined in [11], both routers and hosts can be easily re-addressed.

[2] It should be noted that this policy is an operational issue that could be changed if better solutions to the routing scalability problem came up, because address exhaustion does not apply to IPv6.

The router renumbering specification, defined in [11] and commonly abreviated as "RR", defines a protocol to automatically change the prefix of routers in a IPv6 managed network. An IPv6 managed network is a kind of network where all the routers and hosts belong to the same control authority, which has full knowledge of the number and topological distribution of the devices. Common policies can be applied and there is full control over the resources. This kind of enviroment is what [11] addresses.

Although this protocol defines a mechanism for informing a set of routers of renumbering operations they are going to perform, this mechanism can also be applied when the number of routers is unknown. Nevertheless, reliably informing all the routers when the actual number of them is unknown is described as "a difficult problem". It is recognized that both implementation and operational experience will be needed to fully understand the applicability and scalability aspects of the mechanisms defined in the specification.

Note that the combination of router renumbering plus stateless address autoconfiguration of hosts allows both hosts and routers to be smoothly renumbered. Also it is important to understand that this protocol is not focused on the problem of connection survival across a renumbering event. It is expected that the renumbering process of a set of related networks will be driven by some external planning which should be taken over by a well-defined policy.

This specification defines two modes of operation, both of them are set up by a managed host, which has all the information regarding the number of routers which have to be renumbered. Also is supposed that the actual network is working correctly before starting the process. The modes of operation are the following:

- A multicast mode: in this case, the managed host only needs to issue a "starting" packet, which is addressed to the all-routers site multicast address, ff05::1.
- A unicast mode: in this case, the managed host is in charge of seding at least one renumbering packet per router.

The protocol runs over ICMPv6, so the specification has a complete analisys about the likehood of a packet lost and the number of packet retransmitions that should be sent based on the number of routers and other network parameters.

The structure of the route renumbering packets are the following:

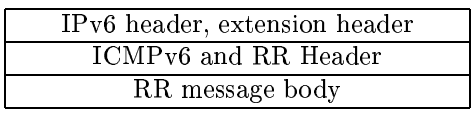

Fig. 1, Router renumbering packet format.

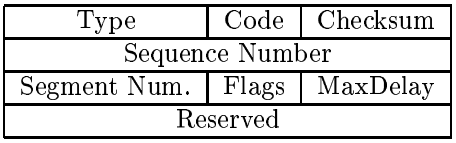

Fig. 2, RR header format.

The sequence number field is used to avoid replays attacks, while the segment number field is related to the fragmentation issues that might ocurr in large packets. The rest of the fields in this figure and the rest of the figures that constitute this section are not explained, to get a detailed information the reader is encouraged to examined [11].

The specification also defines three types of messages, which are identified by means of the ICMPv6 code field:

- Commands: for changing router's state. Data authentication and message integrity must be provided by IPsec [15] means.
- Results: for notifying the sender about the outcome of a command message. The processing of these kind of messages is "implementation-defined".
- Sequence number reset: for using in exceptional situations when it might be necessary to reboot the router's sequence numbers.

The RR command message is made of a serie of "prefix control operations" (from 1 to N), also called "PCOs" in short, which have the following internal structure:

- An operation: there are three actions defined, ADD, CHANGE and SET-GLOBAL. the last one specifies that every assigned prefix should be changed.
- A "match prefix": it is used for comparing the currently assigned prefix with the prefix we want to modify, if this match is successful, the operation will be carried out.
- Zero or more "use prefixes": these are the new prefixes that are going to be installed, if the match with the "match prefix" happens.

This messsage is flexible enough to not only change or add new prefixes, but also to combine new and old prefixes to make up a mixed prefix as a result of it. Besides, other fields are defined for the command messages, as can be seen in the following picture:

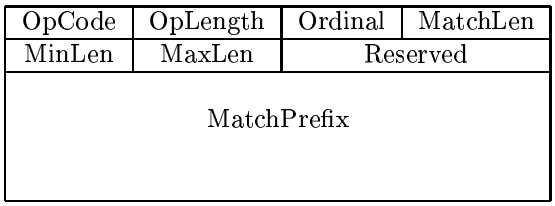

Fig. 3, RR command message format.

it matches on each interface. Its structure is the following:

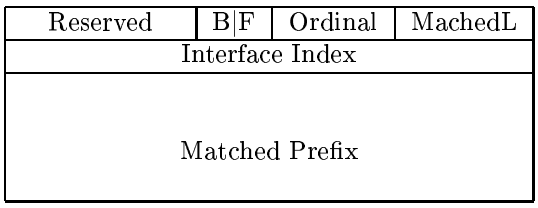

Fig. 4, RR result message format.

The behaviour that [11] specifies for the processing of RR command messages can be described like the following algorithm:

1. Header check.
2. Bounds check.
3. Execution:
   (a) For each address, for each PCO, do:
      i. Check the address against "MatchPrefix".
      ii. If there is a "match":
         - Execute the action associated with the RR message.

Regarding to the security of the RR command messages, [11] is aware of the powerful of the mechanism and that it is important to avoid spoofing of the command messages. Replay of old messages must be prevented. As it is said in the protocol specification document:

> "What constitutes a sufficiently strong authentication algorithm may change from time to time, but algorithms should be chosen which are strong against current key-recovery and forgery attacks."

The use of IPsec [15] is achieved based on the following rule:

> The security policy database of a router implementing this specification MUST cause incoming router renumbering command packets to either be discarded or have IPsec applied.

At the time of this writing, there exist no multicast key management protocol for IPsec, so it is expected that this security associations will be manually configured. Also, [11] says nothing about how the renumbering process might affect the IPsec policy that is being applied. A smooth transition from an old prefix to a new one means that only when there is no transmited packets on

In short, router renumbering mechanism is a powerful automatic tool to renumber routers and by extension hosts in multiple administrative-controlled sites. More experience and feedback from the system administrators is needed to improve and fix some elements of the protocol, provided that the community counts on this mechanism. If this tool is not used, we can conclude that something might not be good enough. From the point of view of the message exchange and protocol definition, router renumbering is considered a robust and flexible mechanism, based on IPsec, ICMPv6 and optionally making use of multicast infrastructure. Currently there are implementations for linux and *BSD systems.

3.4 DNS extensions

Regarding to the DNS infrastructure, the renumbering process should deal with the following issues:

1. How to upgrade the DNS resource records of our hosts ("forward zone").
2. How to upgrade the parent zone to point to our new location.
3. How to upgrade the client's DNS configurations.
4. How to upgrade the "reverse zone".

[12] specifies some modifications to the existing DNS protocol to workaround the points numbered 1, 2 and 4. These extensions introduce new problems on the over-loaded DNS hierarchy and consecuently it has been dropped to experimental for IPv6. Regarding issue number 3, various ways of upgrading the information could be developed. For example, if the client's resolvers are automatically configured using DHCP, we could update the deprecated information using DHCP as well.

In the following sections a short description of [12] is explained. Although the implementation and management costs are too high to implement this solution, it worths looking at it as an example of what need to be done to ease the DNS renumbering process. Pros and cons of these extensions are pointed out inline.

3.4.1 Extensions to address the renumbering of the forward zone

Maintenance of address information in the DNS is one of several obstacles which have prevented site renumbering from being feasible in IPv4. To support the storage of IPv6 addresses without impeding renumbering, [12] defines a new resource record [3] called "A6" which allows to make referrals to parts of the looked-up address. This implies

[3] [12] defines other addons to the DNS infrastructure which are not commented in this paper. For more information, the interested reader is encouraged to read the full specification.

latency. Moreover, it is quite easy to misconfigure the zone to point to blackholes in other zones, and even it is likely that this misconfiguration can end up in recursive loops which would hurt the resolvers in such a way that some restriction on the number of recursive calls should be implemented, to avoid memory exhaustion. All these problems have moved out [12] to the experimental RFC state, and although most of the implementations currently support these extensions, it is not recommended its used.

The new resource record:

> "expedites network renumbering and updated definitions of existing query types that return internet addresses as part of additional section processing."

The meaning of the "A6" resource records is better explained using an example:

1. Suppose the DNS authoritative server of the zone "ipv6.it.uc3m.es" has the following configuration:

```
host1 A6 64 ::02d0:09ff:fef7:6d2c \
     ispA.ipv6.net.
```

2. Suppose that a client host's resolver wants to resolve the name "host1.ipv6.it.uc3m.es". The resolver starts asking the root name servers for the authoritative server of ".es", following by "uc3m.es", "it.uc3m.es" and finally it is reached the DNS of "ipv6.it.uc3m.es". Then the A6 resource record is found, and it is sent back to the resolver.

3. The resolver saves the address included in the answer. The A6 resource record informs about the number of bits that belong to the suffix of the address, and it also gives those bits of information, in this case 64 bits, as well as pointing out to the name of the name server which is in charge of holding the rests of the bits.

4. So the resolver still has to make a new resolution of the name "ispA.ipv6.net" to get the address that it is looking for.

5. Eventually the resolver gets an answer of "ispA.ipv6.net" which includes the bits that represent the prefix part of the address (the remaining part).

6. Finally, the resolver merges both parts to get a full sized address, which returns to the application.

At large sites, AAAA renumbering changes a huge number of records, while A6 renumbering would be needed is a new entry in the "ispB" authoritative zone to specify (using "A6" semantics) the new prefix part of the address. it is said that this new resource record allows servers to construct a "provider-independent" forward zone.

3.4.2 Extensions to address the renumbering of the reverse zone

In a similar way as it is just been explained in the previous section, the reverse zone can be made "provider-independent" using the "DNAME" resource record, which is defined in [14]. This resource record:

> "allows a zone to be used without modification for parallel copies of an address space, as for a multihomed provider or site, and across network renumbering events."

The inverse tree is rooted in IP6.ARPA. if we made use of DNAME records we could redirect queries on the reverse tree to our name server, and the pain of changing prefixes could be aleviated using a DNAME resource record in the new part of the inverse tree, which would redirect the queries to the existing old infrastructure. The operation and structure of the DNAME is quite similar to the A6 resource records, and interesting people is encouraged to read [12].

3.4.3 Problems

To sum it up, if we were able to use this mechanism, we could get the following benefits:

- Ease to renumber the forward zone.
- Permits the own management of the reverse tree.
- Permits the address aggregation to appear at bits boundaries.

Besides what it is been described in the previous sections, this specification also defines the following modifications to the DNS system:

- Existing queries that perform additional section processing to locate IPv4 addresses are redefined to do that processing for both IPv4 and IPv6 addresses.
- A new domain, IP6.ARPA, is defined to support lookups based on IPv6 address.

Nevertheless, this attempt to ease the renumbering process of the information held in the DNS system is flaky because of the following aspects:

- Memory problems in the resolvers, due to the unlimited state that in the worst case should be kept to resolve the forward zone.

- Delays, due to the recursive name resolutions.
- Complexity of distributed administration, which means that the DNS administration should not be complicated, because misconfigurations could be very tough of finding out, and misconfigurations has been demostrated to be quite a common issue regarding the DNS infrastructure.

All these factors make clear that this solution is not recommended for wide use on the current internet, and the IETF has dropped [12] to experimental. More work is needed to achieve the benefits described in this solution without these drawbacks.

3.5 Operational procedures

[1] describes the steps in a procedure that can be used to transition from the use of an existing prefix to a new prefix in a network. It uses IPv6's intrinsic ability to assign multiple addresses to a network interface to provide continuity of network services, as well as addressing naming and configuration management issues. It also uses other IPv6 features to minimize the effort and time required to complete the transition from the old prefix to the new prefix.

3.5.1 Summary of what must be changed

The draft identifies some elements that should be renumbered, specifically:

- Link prefixes and IPv6 addresses.
- Propagated routing information.
- Ingress and egress filters.
- Active control lists (ACLs) and other embedded addresses.
- DNS entries.
- Configuration information provided by DHCP.
- IPv6 addresses embedded in configuration files, applications and "elsewhere".

3.5.2 Steps to accomplish the renumbering procedure

The remaining sections of [1] are devoted to describe a procedural and well defined group of steps that the system administrator should follow to make the renumbering process as clean as possible. The template is the following:

1. Initial condition: the prefix that is going to be changed is currently being used without problems, and everything in the network is working as it is expected.
2. Extending the network with the new prefix: this step involves the following tasks:
   - Getting the prefix and inverse zone delegation.
   - Assigning a subprefix of the new prefix to every link, which should be guided by an enterprise addressing plan.
   - Reducing the TTL of DNS resource records, to achive a better response when they have to be changed.
   - Lowering the leasing time of DCHP configured information.
3. Configuring network elements for the new prefix:
   - The network still gives service to the old prefix.
   - Routing, ACLs and the rest of the services must be working for both prefixes.
   - Finally, announcing the new prefix to the outer world.
4. Adding new host addresses: this implies to upgrade the DNS with the new information.
5. Stable use of either prefixes: this has to be done in all the elements that play an active role in the network, like hosts, routers and any other fundamental service.
6. Transition from use of the old prefix to the new prefix: there are two things that must be thought of:
   - Favouring the packets that use the new prefix over the old ones.
   - Only when packets using the old prefix are not seen on the network, continue.
7. Removing the old prefix: when this task is performed following the previous steps, it is assured that everything should work exactly as if it was before. Note that this task, in general, is the most risky of all the steps involved in a renumbering procedure.
8. Final condition: stable using the new prefix.

Finally, [1] ends up giving some good advices of how to avoid the most common mistakes that can make up when we are on the renumbering process. Also, the document recognizes that there is a lack of automatic tools:

> "Sadly, there are several mechanisms that either have not been automated, or have not been automated consistently across platforms."

The problem of renumbering of networks can be considered as a very difficult process. There have been some attempts to address this problem in the IPv4 world, but the success has been limited in its scope and applicability.

The IPv6 world has not got into the public internet yet, so there is still a chance that some modifications to the protocol could get to ease the renumbering process. Otherwise, the level of automation that should be achieved is not defined, and it seems clear that a solution that fits everybody is not a goal. Nevertheless, the autoconfiguration and related advantages of IPv6 over the current version of IP might lead the research over the next years to focuse in this subject.

Regarding to the DNS information that has to be changed when we face a renumbering process, the author believes that some mechanism should be developed in a compatible way with the existing managed infrastructure to make renumbering easier. The automatic mechanism defined by DHCP and the security extensions for updating DNS information are the starting points on which a renumbering solution should be based.

Inside the IETF, it is expected that a WG will be created, though at the moment of this writing there is no official information that confirms this gossip, but if people wants a standarized solution to the renumbering problem that would be freely adopted by the internet community as a whole, this and only this way is the path that must be followed.

Aknowlegments

This is the easy part, but I am always surprised how many people forget it. There is no need to give concrete names in this section, they already know that I am talking about them. To those who wander here, thank you for backing me up, specially when I didn't come up with my duties at time. I have been extraordinarily grateful for the community here, during all these years.

I wish there were a better way to show gratefulness, but until I've got some real whuffie to pass around, thanks is the best I have. I will try to do my best, as usually.

[1] Baker, F. "Procedures for Renumbering an IPv6 Network without a Flag Day", draft-baker-ipv6-renumber-procedure-01.txt, Octuber 2003.

[2] P. Ferguson, H. Berkowitz, "Network Renumbering Overview: Why would I want it and what is it anyway?", RFC 2071, January 1997.

[3] P. Gross, P. Almquist, "ESG Deliberations on Routing and Addressing", RFC 1380, November 1992.

[4] Y. Rekhter, T. Li, "An Architecture for IP Address Allocation with CIDR", RFC 1518, September 1993.

[5] V. Fuller, T. Li, J. Yu, "Classless Inter-Domain Routing (CIDR): an Address Assignment and Aggregation Strategy", RFC 1519, September 1993.

[6] H. Berkowitz, P. Ferguson, W. Leland, "Enterprise Renumbering: Experience and Information Solicitation", RFC 1916, February 1996.

[7] H. Berkowitz, "Router Renumbering Guide", RFC 2072, January 1997.

[8] B. Carpenter, H. Crowcroft, "IPv4 Address Behaviour Today", RFC 2101, February 1997.

[9] R. Hinden, S. Deering, "IP Version 6 Addressing Architecture", RFC 2373, July 1998.

[10] T. Narten, E. Nordmark, W. Simpson, "Neighbor Discovery for IP Version 6 (IPv6)", RFC 2461, December 1998.

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

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

[13] B. Carpenter, Y. Rekhter, "Renumbering Needs Work", RFC 1900, February 1996

[14] M. Crawford, "Non-Terminal DNS Name Redirection", RFC 2672, August 1999.

[15] S. Kent, R. Atkinson, "Security Architecture for the Internet Protocol", RFC 2401, November 1998.