Network Working Group
Internet Draft
Intended status: Informational
Expires: June 7, 2012

B. Liu
S. Jiang
Huawei Technologies Co., Ltd
B. Carpenter
University of Auckland
December 5, 2011

IPv6 Site Renumbering Gap Analysis
draft-liu-6renum-gap-analysis-03.txt

Status of this Memo

Copyright Notice

Abstract

This document briefly introduces the existing mechanisms could be utilized by IPv6 site renumbering and envisions the effort could be done. This document tries to cover the most explicit issues and requirements of IPv6 renumbering. Through the gap analysis, the document provides a basis for future work to identify and develop solutions or to stimulate such development as appropriate. The gap analysis is presented following a renumbering event procedure clue.

Table of Contents

1. Introduction

As introduced in [RFC5887], renumbering, especially for medium to large sites and networks, is currently viewed as an expensive, painful, and error-prone process, avoided by network managers as much as possible. If IPv6 site renumbering continues to be considered difficult, network managers will turn to Provider Independent (PI) addressing for IPv6 to attempt to minimize the need for future renumbering. However, widespread use of PI may create very serious BGP4 scaling problems. It is thus desirable to develop tools and practices that may make renumbering a simpler process to reduce demand for IPv6 PI space.

This document performs a gap analysis to provide a basis for future work to identify and develop solutions or to stimulate such development as appropriate. The gap analysis is organized by the main steps of renumbering process, which include prefix management, node address (re)configuration, and updating relevant address entries in various gateways, routers and servers, etc. Besides these steps, the aspect of renumbering event management is presented independently, which intends to help the operational/administrative process. It is expected that these steps and management could cover all the aspects of an renumbering event.

This document draws on existing work in (at least) [RFC5887], [I-D.chown-v6ops-renumber-thinkabout] and [RFC4192]. Contributions from [I-D.jiang-6renum-enterprise] are incorporated into the more detailed analysis. Lots of issues were analyzed in RFC5887 & [I-D.chown-v6ops-renumber-thinkabout], but many of them are out of 6renum scope or unsolvable. This document intends to identify the valuable and solvable issues, dig out of some undiscovered gaps, and tries to give solution suggestions.

2. Overall Requirements for Renumbering

This section introduces the overall ultimate goals we want to achieve in a renumbering event. In general, we need to leverage renumbering automation to avoid human intervention as much as possible at reasonable cost. Some existing mechanisms have already provided useful help. Further efforts may be achieved in the future.

The automation can be divided into four aspects as follows, which are also the gap analysis topics.

- o Prefix delegation and delivery should be automatic and accurate in aggregation and coordination.

- o Address reconfiguration should be automatically achieved through standard protocols with minimum human intervention.

- o Updating relevant address entries should be performed integrally and without error. [Open Question]Is it necessary to develop automatic entries update mechanisms? If necessary, do we need standard protocols/interface for it?

- o Renumbering event management is needed to provide the functions of renumbering notification, synchronization, and monitoring .etc.

Besides automation, session survivability is another important issue during renumbering since application outage is one of the most obvious impacts that make renumbering painful and expensive. We have an enormous advantage in IPv6 which is the ability to overlap the old and new prefixes and to use the address lifetime mechanisms in SLAAC and DHCPv6. That is fully described in [RFC4192]. We consider this mechanism is sufficient for session survivability issue in most of the cases.

[Open Question]Should we consider the case of very long-lived application sessions (days or weeks) which cannot be resolved by [RFC4192]?

3. Existing Components for IPv6 Renumbering

Since renumbering is not a whole new issue, some protocols/mechanisms have been already utilized or even be developed dedicated for renumbering. However, generally current renumbering is achieved by existing protocols rather than dedicated renumbering protocols. This section briefly reviews these existing protocols/mechanisms to provide a basis for the gap analysis.

3.1. Relevant Protocols and Mechanisms

- o RA messages, defined in [RFC4861], are used to deprecate/announce old/new prefixes and to advertise the availability of an upstream router. In renumbering, it is one of basic mechanisms for host configuration.

- o When a host is renumbered, it may use SLAAC [RFC4862] for address configuration with the new prefix. Hosts receive RA messages which contain routable prefix(es) and the address(es) of the default router(s), then hosts can generate IPv6 address(es) by themselves.

- o Hosts configured through DHCPv6 [RFC3315] can reconfigure addresses by initialing RENEW sessions when the current addresses' lease time are expired or they receive the reconfiguration messages initiated by the DHCPv6 servers.

- o DHCPv6-PD (Prefix Delegation) [RFC3633] enables automated delegation of IPv6 prefixes using the DHCPv6.

- o [RFC2894] defined standard ICMPv6 messages for router renumbering. This is a dedicated protocol for renumbering, but has not been widely used.

3.2. Management Tools

Some operations of renumbering could be automatically processed by management tools in order to make the renumbering process more efficient and accurate. The tools may be designed dedicated for renumbering or just common tools could be utilized for some operations in renumbering.

Following are samples of the tools.

- o Address management tools. There are both commercial and open-source, and even home-made solutions.
  [Further work is needed to identify what an address management tool should be able to do for improving the ability of managing a network through a renumbering event.]

- o [LEROY] proposed a mechanism of macros to automatically update the address relevant entries/configurations inside the DNS, firewall, etc. The macros can be delivered though SOAP protocol from a network management server to the managed devices.

- o Asset management tools/systems. These tools may provide the ability of managing configuration files in nodes so that it is convenient to update the address relevant configuration in these nodes.

3.3. Procedures/Policies

- o [RFC4192] proposed a procedure for renumbering an IPv6 network without a flag day. The document includes a set of operational suggestions which can be followed step by step by network administrators.

o [I-D.jiang-6renum-enterprise] analyzes the enterprise renumbering events and gives the recommendations among the existing renumbering mechanisms. According to the different stages, renumbering considerations are described in three categories: considerations and recommendations during network design, for preparation of enterprise network renumbering, and during renumbering operation

4. Managing Prefixes

When renumbering an enterprise site, a short prefix may be divided into longer prefixes for subnets. So we need to carefully manage the prefixes for prefix delivery, delegation, aggregation, synchronization, coordination, etc.

4.1. Prefix Delegation

Usually, the short prefix comes down from the operator and received by DHCPv6 server or router inside the enterprise network. The short prefix could be automatically delegated through DHCPv6-PD. Then the downlink DHCP servers or routers can begin advertising the longer prefixes to the subnets.

For the delegation routers, they may need to renumber themselves with the delegated prefixes. We need to consider the router renumbering issue which cannot be covered by DHCP-PD only.

4.2. Prefix Assignment

When subnet routers receive the longer prefixes, they can directly assign them to the hosts. The prefix assignment overlaps with the host address configuration, which is described in the following section 5.1.

5. Address Configuration

5.1. Host Address Configuration

Both of the DHCPv6 and ND protocols have IP address configuration function. They are suitable for different scenarios respectively. During renumbering, the SLAAC-configured hosts can reconfigure IP addresses by receiving ND Router Advertisement (RA) messages containing new prefix information (It should be noted that, the prefix delivery could be achieved through DHCP according to the new IETF DHC WG document [I.D ietf-dhc-host-gen-id]). The DHCPv6-configured hosts can reconfigure addresses by initiating RENEW sessions when the current addresses' lease time are expired or

receiving the reconfiguration messages initiated by the DHCPv6 servers.

o SLAAC and DHCPv6 address configuration co-existence

While an IPv6 site is being renumbered, both DHCPv6 and ND may be used to reconfigure the host addresses. The co-existence issue mainly includes following aspects:

- Dynamic upstream learning

  [RFC5887] mentioned that, DHCP-configured hosts may want to learn about the upstream availability of new prefixes or loss of prior prefixes dynamically by deducing from periodic RA messages. But there is no standard specifying what approach should be taken by a DHCPv6-configured host when it receives RA messages containing new prefix. It depends on the operation system of the host and cannot be predicted or controlled by the network.

- DHCP-managed hosts receiving RA messages

  It is unclear that whether a DHCP-managed host would accept configuration though RA messages, it depends on the policies in the host's operating system. If it ignores the RA messages and there are no DHCPv6 reconfiguration messages received either, the renumbering would fail.

- SLAAC-configured hosts finding DHCPv6 in use

  [RFC5887] mentioned RA message of ND protocol contains a "Managed Configuration" flag to indicate DHCPv6 is in use. But it is unspecified what behavior should be taken when the host receives RA messages with "M" set to 1. The gap of standard will cause ambiguous host behavior because it depends on the operation system of the host.

  The host may start a DHCPv6 session and receives the DHCPv6 address configuration. It is also possible that the host finds the DHCPv6 assigned prefix is different from the prefix in the RA messages, which means multiple uplinks are available or there is a serious network configuration error.

  Another possibility is that the host may receive no response from any DHCPv6 servers, which means the DHCPv6 service is not available and the "Managed Configuration" flag was mis-configured.

o DHCPv6 reconfigure bulk usage

[RFC5887] mentioned that "DHCPv6 reconfiguration doesn't appear to be widely used for bulk renumbering purposes".

The reconfiguration defined in [RFC3315] needs to establish a session between DHCP server and client. This could be considered as a stateful approach which needs much resource on the server to maintain the renumbering sessions. This is probably one of the reasons that DHCP reconfiguration is not suitable for bulk usage.

Another limitation of reconfiguration is that it only allows th e messages to be delivered to unicast addresses. So if we want to use it for bulk renumbering, stateless DHCPv6 reconfiguration with multicast may be needed. However, this may involve protocol modification.

5.2. Router Address Configuration

o Learning new prefixes

As described in [RFC5887], "if a site wanted to be multihomed using multiple provider-aggregated (PA) routing prefixes with one prefix per upstream provider, then the interior routers would need a mechanism to learn which upstream providers and prefixes were currently reachable (and valid). In this case, their Router Advertisement messages could be updated dynamically to only advertise currently valid routing prefixes to hosts. This would be significantly more complicated if the various provider prefixes were of different lengths or if the site had non-uniform subnet prefix lengths."

o Restart after renumbering

"Some routers cache IP addresses in some situations. So routers might need to be restarted as a result of site renumbering" [RFC2072].

o Router naming

In [RFC4192], it is suggested that "To better support renumbering, switches and routers should use domain names for configuration wherever appropriate, and they should resolve those names using the DNS when the lifetime on the name

expires."
As [RFC5887] described, this capability is not new, and at least it is present in most IPsec VPN implementations. But many administrators do not realize that it could be utilized to avoid manual modification during renumbering.

In enterprise scenario, the requirement of router naming is not as strong as that in ISP. So for the administrators, the motivation of using router naming for easing renumbering may be not strong.

5.3. Static Address Configuration

There is another draft dedicated to the static address issue. Please refer to [I-D.carpenter-6renum-static-problem].

6. Updating Relevant Address Entries

When nodes in a site have been renumbered, then all the entries in the site which contain the nodes' addresses must be updated. The entries mainly include DNS records and filters in various entities such as ACLs in firewalls/gateways.

6.1. DNS Records Update

o Dynamic DNS update

For DNS records update, most sites will achieve it by maintaining a DNS zone file and loading this file into the site's DNS server(s). Synchronization between host renumbering and the updating of its A6 or AAAA record is hard. [RFC5887] mentioned that an alternative is to use Secure Dynamic DNS Update [RFC3007], in which a host informs its own DNS server when it receives a new address.

Secure Dynamic DNS Update has been widely supported by the major DNS systems, but it hasn't been widely deployed, especially in the host. Current practices mainly involve the DHCP servers which act as clients to request the DNS server to update relevant records. Normal hosts are not suitable to do this mainly because of the complexity of key management issues inherited from secure DNS mechanisms.

6.2. In-host Server Address Update

While DNS records addresses of hosts in servers, hosts also record addresses of servers such as DNS server, radius server, etc. While renumbering, the hosts must update the records if the server addresses changed. Addresses of DHCPv6 servers do not need to be updated. They are dynamically discovered using DHCPv6 relevant multicast addresses.

- The DNS server addresses for hosts are configured by DHCPv6. But current DHCPv6 messages do not indicate to hosts the lifetimes of DNS. If the DNS lifetime expired and has been renumbered, the hosts may still use the old addresses. DHCPv6 should be extended to indicate to hosts the associated DNS lifetimes when making DNS configuration. How the DHCP server could know about the DNS lifetime is another issue.

6.3. Filters

- Filters Management

  Filters based on addresses or prefixes are usually spread in various devices. As [RFC5887] described, some address configuration data might be widely dispersed and much harder to find, even will inevitably be found only after the renumbering event. So there's a big gap for filter management.

  In [LEROY], a server is used for managing filter update in various devices. But identifying where and which of the filters need to be updated during renumbering is still a gap.

- Filter Update Automation Operation

  As mentioned in section 3.2, [LEROY] proposed a mechanism which can automatically update the filters. The mechanism utilizes macros suitable for various devices such as routers, firewalls etc. to update the filter entries based on the new prefix. Such automation tool is valuable for renumbering because it can reduce manual operation which is error-prone and inefficiency.

  Besides the macros, [LEROY] also proposed to use SOAP to deliver the macros to the devices. As well as SOAP we may consider whether it is possible and suitable to use other standardized protocols such as NETCONF.

  Update of filters based on prefixes and filters based on addresses may have different requirements and methods. Address-

based filters may be mainly with regard to domain names while prefix-based filters be relevant to more abstract entity (mask e.g.). Thus, we may consider different ways to update the two kinds of filters, for example, the prefix-based filters may consider to be updated though DHCPv6 server, which may provide better efficient.

7. Renumbering Event Management

From the perspective of network management, renumbering is a kind of event which may need additional process to make the process more easy and manageable.

7.1. Renumbering Notification

If hosts or servers are aware of a renumbering event happening, it may help the relevant process. Following are several examples of such additional process may ease the renumbering. Further contributions are expected.

- A notification mechanism may be needed to indicate the hosts that a renumbering event of local recursive DNS happen or is going to take place.

- [RFC4192] suggests that "reducing the delay in the transition to new IPv6 addresses applies when the DNS service can be given prior notice about a renumbering event." Reducing delay could improve the efficiency of renumbering.

7.2. Synchronization Management

- DNS update synchronization

  DNS update synchronization focuses on the coordinating between DNS and other entities/mechanisms, for example, as described in [RFC5887], synchronizing the SLAAC and DNS updates, and of reducing the SLAAC lease time and DNS TTL.

7.3. Renumbering Monitoring

While treating renumbering a network event, mechanisms to monitor the renumbering process may be needed. Considering the address configuration operation may be stateless (if ND is used for renumbering), it is difficult for monitoring. But for the DNS and filter update, it is quite possible to monitor the whole process.

8. Miscellaneous

8.1. Mobility

As [RFC5887] suggested, for Mobile IP, we need a better mechanism to handle change of home agent address while mobile is disconnected.

9. Gaps considered unsolvable

This section lists gaps have been documented but are considered unsolvable or out of the scope of 6renum working group.

9.1. Address Configuration

o RA prefix lifetime limitation

In section 5.5.3 of [RFC4862], it is defined that "If the received Valid Lifetime is greater than 2 hours or greater than RemainingLifetime, set the valid lifetime of the corresponding address to the advertised Valid Lifetime." So when renumbering, if the previous RemainingLifetime is longer than two hours, it is impossible to reduce a prefix's lifetime less than two hours. This limitation is to prevent denial-of-service attack.

9.2. Address Relevant Entries Update

o DNS entries commonly have matching Reverse DNS entries which will also need to be updated during renumbering.

o DNS data structure optimization

[RFC2874] proposed a new A6 record type for DNS recording IPv6 address/prefix. And several extensions on query and processing were also proposed. With the A6 record and the extensions, an IPv6 address can be defined by using multiple DNS records. This feature increases the complexity of resolver but reduce the cost of zone file maintenance. So renumbering could be easier than AAAA record. But the [RFC2874] has not been widely used.

o DNS authority

As described in [I-D.chown-v6ops-renumber-thinkabout], "it is often the case in enterprises that host web servers and application servers on behalf of collaborators and customers that DNS zones out of the administrative control of the host maintain resource records concerning addresses for nodes out of

their control. When the service host renumbers, they do not have sufficient authority to change the records."

It is an operational issue and this document considers it not suitable to be solved through bring additional protocol/mechanism to standardize the interaction between DNS systems needs to be considered.

9.3. Miscellaneous

- For transport layer, [5887] said that TCP connections and UDP flows are rigidly bound to a given pair of IP addresses.

- For application layer, as [5887] said, in general, we can assert that any implementation is at risk from renumbering if it does not check that an address is valid each time it opens a new communications session.

10. Security Considerations

- Prefix Validation

Prefixes from the ISP may need authentication to prevent prefix fraud. Announcing changes of site prefix to other sites (for example, those that configure routers or VPNs to point to the site in question) also need validation.

In the LAN, Secure DHCPv6 ([I-D.ietf-dhc-secure-dhcpv6]) or SeND ([RFC3971], Secure Neighbor Discovery) deployment may need to validate prefixes.

- Influence to Security Controls

During renumbering, security controls (e.g. ACLs) blocking access to legitimate resources should not be interrupted.

11. IANA Considerations

None.

12. References

12.1. Normative References

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

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

[RFC3007] B. Wellington, "Secure Domain Name System (DNS) Dynamic Update", RFC 3007, November 2000.

[RFC3315] R. Droms, Bound, J., Volz, B., Lemon, T., Perkins, C., and M. Carney, "Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", RFC 3315, July 2003.

[RFC3633] Troan, O. and R. Droms, "IPv6 Prefix Options for Dynamic Host Configuration Protocol (DHCP) version 6", RFC 3633, December 2003.

[RFC3956] P. Savola, and B. Haberman. "Embedding the Rendezvous Point (RP) Address in an IPv6 Multicast Address.", RFC 3956, November 2004.

[RFC3971] Arkko, J., Ed., Kempf, J., Zill, B., and P. Nikander "SEcure Neighbor Discovery (SEND)", RFC 3971, March 2005.

[RFC4861] Narten, T., Nordmark, E., Simpson, W., and H. Soliman, "Neighbor Discovery for IP version 6 (IPv6)", RFC 4861,September 2007.

[RFC4862] Thomson, S., Narten, T., and T. Jinmei, "IPv6 Stateless Address Autoconfiguration", RFC 4862, September 2007.

12.2. Informative References

[RFC2072] H. Berkowitz, "Router Renumbering Guide", RFC2072, January 1997.

[RFC4192] Baker, F., Lear, E., and R. Droms, "Procedures for Renumbering an IPv6 Network without a Flag Day", RFC 4192, September 2005.

[RFC4714] Enns, R., "NETCONF Configuration Protocol", RFC 4714, December 2006.

[RFC5887] Carpenter, B., Atkinson, R., and H. Flinck, "Renumbering Still Needs Work", RFC 5887, May 2010.

[I-D.ietf-dhc-secure-dhcpv6] Jiang, S., and Shen S., "Secure DHCPv6 Using CGAs", working in progress.

[I-D.chown-v6ops-renumber-thinkabout]
Chown, T., "Things to think about when Renumbering an IPv6 network", Work in Progress, September 2006.

[I-D.jiang-6renum-enterprise]
Jiang, S., and B. Liu, "IPv6 Enterprise Network Renumbering Scenarios and Guidelines ", Working in Progress, July 2011.

[I-D.carpenter-6renum-static-problem]
Carpenter, B., and S. Jiang, "Problem Statement for Renumbering IPv6 Hosts with Static Addresses", Working in Progress, October 2011.

[LEROY] Leroy, D. and O. Bonaventure, "Preparing network configurations for IPv6 renumbering", International of Network Management, 2009, <http://inl.info.ucl.ac.be/system/files/dleroy-nem-2009.pdf>

13. Acknowledgments

This work adopts significant amounts of content from [RFC5887] and [I-D.chown-v6ops-renumber-thinkabout], so thank for Brian Carpenter, Randall Atkinson, Hannu Flinck, Tim Chown, Mark Thompson, Alan Ford, and Stig Venaas. Some useful materials were provided by Oliver Bonaventure and his student Damien Leroy, thanks for them, too.

Useful comments and contributions were made by Wesley George, and others.

This document was prepared using 2-Word-v2.0.template.dot.

Authors' Addresses

Bing Liu
Q14-4-A Building
Huawei Technologies Co., Ltd
Zhong-Guan-Cun Environment Protection Park, No.156 Beiqing Rd.
Hai-Dian District, Beijing
P.R. China
Email: leo.liubing@huawei.com

Sheng Jiang
Q14-4-A Building
Huawei Technologies Co., Ltd
Zhong-Guan-Cun Environment Protection Park, No.156 Beiqing Rd.
Hai-Dian District, Beijing
P.R. China
Email: shengjiang@huawei.com

Brian Carpenter
Department of Computer Science
University of Auckland
PB 92019
Auckland, 1142
New Zealand
EMail: brian.e.carpenter@gmail.com