Network Working Group
Internet Draft
Intended status: Informational
Expires: June 11, 2013

B. Liu
S. Jiang
Huawei Technologies Co., Ltd
B. Carpenter
University of Auckland
S. Venaas
Cisco Systems
W. George
Time Warner Cable
December 12, 2012

IPv6 Site Renumbering Gap Analysis
draft-ietf-6renum-gap-analysis-05.txt

Status of this Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at http://datatracker.ietf.org/drafts/current/.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on June 11, 2013.

Copyright Notice

Copyright (c) 2012 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (http://trustee.ietf.org/license-info) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Abstract

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

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 a renumbering process, which include prefix management, node address (re)configuration, and updating address-relevant entries in various devices such as firewalls, routers and servers, etc. Besides these steps, the aspect of renumbering event management is also presented independently, which intends to identify some operational/administrative gaps in renumbering.

This document starts from existing work in [RFC5887], [I-D.chown-v6ops-renumber-thinkabout] and [RFC4192]. This document does further analysis and identifies the valuable and solvable issues, digs out of some undiscovered gaps, and gives some 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 ability. Further efforts may be achieved in the future.

The automation can be divided into four aspects as follows.

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

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

- Address-relevant entry updates should be performed integrally and without error.

- 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. Session survivability is a fundamental issue that cannot solved within renumbering context only; however, 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 IPv6 Stateless Address Autoconfiguration (SLAAC) and Dynamic Host Configuration Protocol for IPv6 (DHCPv6). That is fully described in [RFC4192], which provides a smooth transition mechanism from old to new prefixes. In most of the cases, since we can set the transition period long enough to cover the on-going sessions, we consider this mechanism is sufficient for avoiding session brokenness issue in IPv6 site renumbering.

3. Existing Components for IPv6 Renumbering

Since renumbering is not a new issue, some protocols and mechanisms have already been utilized for renumbering. There were also some dedicated protocols and mechanisms developed for renumbering. This section briefly reviews these existing protocols and mechanisms to provide a basis for the gap analysis.

3.1. Relevant Protocols and Mechanisms

- 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 the basic mechanisms for host configuration.

- When renumbering a host, SLAAC [RFC4862] may be used for address configuration with the new prefix(es). 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.

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

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

- [RFC2894] defined standard ICMPv6 messages for router renumbering. This is a dedicated protocol for renumbering, but it 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 specifically for renumbering, or common tools could be utilized for some of the renumbering operations.

Following are examples of these tools.

- IP address management (IPAM) tools. There are both commercial and open-source solutions. An IPAM is basically used to manage an IP address plan, and it integrates the DHCPv6 and DNS services together as a whole solution. Many mature commercial tools can support management operations, but normally they do not have dedicated renumbering functions. However, the integrated DNS/DHCPv6 services and address management function can obviously facilitate the renumbering process.

- [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.

- 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

- [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.

- [I-D.ietf-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(es) come down from the operator(s) and are received by DHCPv6 servers or routers inside the enterprise networks. The short prefix(es) could be automatically delegated through DHCPv6-PD. Then the downlink DHCPv6 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

- SLAAC/DHCPv6 co-existence

  Both of the DHCPv6 and Neighbor Discovery (ND) protocols have an 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 DHCPv6 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 times are expired or when they receive reconfiguration messages initiated by the DHCPv6 servers.

Sometimes the two address configuration modes may both be available in one network. This would add more or less additional complexity for both the hosts and the network management. In ND protocol, there is an M (ManagedFlag) flag defined in RA message, which indicates the hosts the DHCPv6 service status in the network. And there is another O (OtherConfigFlag) flag indicating the host to configure information such as DNS/routing other than addresses.

Using these flags, the two separated address configuration modes are somehow correlated. However, the ND protocol did not define the flags as prescriptive but only as advisory. This has led to variation in the behavior of hosts when interpreting the M flag. Different operating systems follow different approaches [I-D.liu-6renum-dhcpv6-slaac-switching].

The impact of ambiguous M/O flags mainly includes the following aspects:

- SLAAC/DHCPv6 preference

  Ideally, it should be possible to choose either SLAAC or DHCPv6 as the default address configuration mechanism when a host goes online, but this is not always available in reality. On some current operating systems, DHCPv6 procedure will not start until they receive RA messages with M=1. This has an implication that RA messages are required for DHCPv6 management. This may cause operational issues since it needs additional complexity in ND/DHCPv6 co-existence to achieve unified management in a network.

- DHCPv6-configured hosts receiving RA messages

  It is unclear whether a DHCPv6 configured host will accept configuration though RA messages; this 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, renumbering would fail.

  [RFC5887] mentioned that DHCPv6-configured hosts may want to learn about the upstream availability of new prefixes or loss of prior prefixes dynamically by deducing this 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 a new prefix. It depends on the operating system of the host and cannot be predicted or controlled by the network.

- SLAAC-configured hosts finding DHCPv6 in use

It is unspecified what behavior should be taken when the host receives RA messages with "M" set.

The host may start a DHCPv6 session and receive the DHCPv6 address configuration, or it may just ignore the messages. If the network side wants the hosts to start DHCPv6 configuration, it is just out of control of the network side.

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 DHCPv6 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 DHCPv6 reconfiguration is not suitable for bulk usage.

Another limitation of DHCPv6 reconfiguration is that it only allows the 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."

- Restart after renumbering

  As [RFC2072] mentioned, some routers cache IP addresses in some situations, so routers might need to be restarted as a result of site renumbering.

- 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 document dedicated to the static address issue. Please refer to [I-D.ietf-6renum-static-problem].

6. Updating Address-relevant 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

- Dynamic DNS update

  For DNS records update, the most popular DNS system, BIND, 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. But for some commercial DNS systems, the Secure Dynamic DNS Update issue may be much easier, since it could be integrated with services like DHCP provided by the same vendor so that the dynamic DNS update could be silently enabled.

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.

o 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 a host has been renumbered, other 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. Parameterized IP-specific Configuration

Besides the DNS records and the in-host server address entries, there are also many places in which the IP addresses are configured, for example, filters such as ACL and routing policies, etc. There are even more sophisticated cases where the IP addresses are used for deriving values, for example, using the unique portion of the loopback address to generate an ISIS net ID.

Ideally, a layer of abstraction for IP-specific configuration within various devices (routers, switches, hosts, servers, etc.) is needed. However, this cannot be achieved in one step. One possible improvement is to make the IP-specific configuration parameterized. Two aspects of parameterized configuration could be considered to achieve this.

- Self-contained Configuration in Individual device

  In an ideal way, the IP addresses can be defined as a value once, and then the administrators can use either keywords or variables to call the value in other places such as a sort of internal inheritance in CLI (command line interface) or other local configurations. This makes it easier to manage a renumbering event by reducing the number of places where a device's configuration must be updated. However, it still means that every device needs to be touched, but only once instead of having to inspect the whole configuration to ensure that none of the separate places that a given IP address occurs is missed.

  Among the real current devices, some routers support defining multiple loopback interfaces which can be called in other configurations. For example, when defining a tunnel, it can call the defined loopback interface to use its address as the local address of the tunnel. This can be considered as a kind of parameterized self-contained configuration. But this only applies certain use cases; it is impossible to use the loopback interfaces to represent external devices and it is not always possible to call loopback interfaces in many other configurations. Parameterized self-contained configuration is still a gap for current devices.

- Unified Configuration Management among Devices

  This is a more formalized central configuration management system, where all changes are made in one place and the system manages how to push the changes to the individual devices. This issue contains two aspects as the following.

  - Configuration Aggregation

    Configuration 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 configuration aggregation.

  - Configuration Update Automation

    As mentioned in section 3.2, [LEROY] proposed a mechanism which can automatically update the entries. The mechanism utilizes macros suitable for various devices such as routers, firewalls etc. to update the 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 [RFC4714].

But in current real networks, most of the devices use vendor-private protocols to update configurations, so it is not necessarily valid to assume that there is going to be a formalized configuration management system to leverage. It is a big gap currently.

7. Renumbering Event Management

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

7.1. Renumbering Notification

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

- A notification mechanism may be needed to indicate the hosts that a renumbering event of local recursive DNS happens 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.

- Router awareness: in a site with multiple border routers, all border routers should be aware of partial renumbering in order to correctly handle inbound packets. Internal forwarding tables need to be updated.

- Border filtering: in a multihomed site, an egress router to ISP A could normally filter packets with source addresses from other ISPs. The egress router connecting to ISP A should be notified if the egress router connecting to ISP B initiates a renumbering event in order to properly update its filter function.

7.2. Synchronization Management

o 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 as 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. Multicast

Renumbering also has an impact on multicast. Renumbering of unicast addresses affects multicast even if the multicast addresses are not changed. There may also be cases where the multicast addresses need to be renumbered.

o Renumbering of multicast sources

If a host that is a multicast source is renumbered, the application on the host may need to be restarted for it to successfully send packets with the new source address.

For ASM (Any-Source Multicast) the impact on a receiver is that a new source appears to start sending, and it no longer receives from the previous source. Whether this is an issue depends on the application, but we believe it is likely to not be a major issue.

For SSM (Source-Specific Multicast) however, there is one significant problem. The receiver needs to learn which source addresses it must join. Some applications may provide their own method for learning sources, where the source application may somehow signal the receiver.

Otherwise, the receiver may e.g. need to get new SDP information with the new source address. This is similar to how to learn a

new group address, see the "Renumbering of multicast addresses" topic below.

- Renumbering of Rendezvous-Point

  If the unicast addresses of routers in a network are renumbered, then the RP (Rendezvous-Point) address is also likely to change for at least some groups. An RP address is needed by PIM-SM for providing ASM, and for Bidir PIM. Changing the RP address is not a major issue, except that the multicast service may be impacted while the new RP addresses are configured. If dynamic protocols are used for distributing group-to-RP mappings, the change can be fairly quick, and any impact should be only for a very brief time. However, if routers are statically configured, this depends on how long it takes to update all the configurations.

  For PIM-SM one typically switches to SPT (Shortest-Path-Tree) when the first packet is received by the last-hop routers. Forwarding on the SPT should not be impacted by change of IP address. Network operator should be careful not deprecate the previous mapping before configuring a new one, because implementations may revert to Dense Mode if no RP is configured.

  The impact of this is minimal. The main concern is that while the peering is unavailable, one may not receive updates about new sources.

  It may help to configure a new peering before taking down the existing one.

- Renumbering of multicast addresses

  In general multicast addresses can be chosen independently of the unicast addresses, and the multicast addresses can remain fixed even if the unicast addresses are renumbered. However, for IPv6 there are useful ways of deriving multicast addresses from unicast addresses, such as unicast-prefix-based IPv6 Multicast Addresses [RFC3306] and Embedded-RP IPv6 Multicast Addresses [RFC3956]. In that case the multicast addresses used may have to be renumbered.

  Renumbering group addresses may be complicated. For multicast, it is common to use literal addresses, and not DNS. When multicast addresses are changed, source applications need to be reconfigured and restarted.

Multicast receivers need to learn the new group addresses to join.

Note that for SSM, receivers need to learn which multicast channels to join. A channel is a source and group pair. This means that for an SSM application, a change of source address is likely to have the same effect as a change of group address.

Some applications may have dynamic methods of learning which groups (and possibly sources) to join. If not, the application may have to be reconfigured and restarted.

One common way for receivers to learn the necessary parameters are by use of SDP. SDP information may be distributed via various application protocols, or it may be from a file. An SDP file may be distributed via HTTP, email etc. If a user is using a web browser and clicking on a link to launch the application with the necessary data, it may be a matter of closing the current application, and re-clicking the link.

8.2. 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 this document.

9.1. Address Configuration

- 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

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

- DNS data structure optimization

  [RFC2874] proposed an experimental A6 record type for DNS recording of IPv6 address and prefix. Several extensions to DNS query and processing were also proposed. With the A6 record and the extensions, an IPv6 address could be defined by using multiple DNS records. This feature would increase the complexity of resolvers but reduce the cost of zone file maintenance, so renumbering could be easier than with the AAAA record. But the A6 record has not been widely used, and it has been moved to historic status [RFC6563].

- 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 and policy issue and this document considers it not suitable to be solved through technical approach or bring additional protocol/mechanism to standardize the interaction between DNS systems for authority negotiations.

9.3. Static address issues

In [I-D.ietf-6renum-static-problem], there are several open issues listed at the end, which are as below:

- Is minor residual loss of ongoing transport sessions during renumbering operationally acceptable?

- Can automatic network element renumbering can be performed without interrupting user sessions?

- Do any software licensing systems require manual intervention?

The former two questions are about session survivability which is not only static address specific. As described in section 2, this document generally considers [RFC4192] is sufficient for avoiding session brokenness issue in IPv6 site renumbering in most of the cases.

9.4. Miscellaneous

- For transport layer, [RFC5887] said that TCP connections and UDP flows are rigidly bound to a given pair of IP addresses.

- For application layer, as [RFC5887] 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 Secure Neighbor Discovery (SEND, [RFC3971]) deployment may need to validate prefixes.

- Influence on Security Controls

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

11. IANA Considerations

This draft does not request any IANA action.

12. Acknowledgments

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

Many useful comments and contributions were made by Lee Howard, and members of 6renum WG.

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

13. References

13.1. 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.

13.2. Informative References

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

[RFC3306] B. Haberman, D. Thaler, "Unicast-Prefix-based IPv6 Multicast Addresses", RFC 3306, August 2002.

[RFC3956] P. Savola, B. Haberman, "Embedding the Rendezvous Point (RP) Address in an IPv6 Multicast Address", RFC 3965, November 2004.

[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.

[RFC6563] Jiang, S., Conrad, D., and B. Carpenter, "Moving A6 to Historic Status", RFC 6563, May 2012.

[I-D.ietf-dhc-secure-dhcpv6]
Jiang, S., and Shen S., "Secure DHCPv6 Using CGAs", working in progress, March 2012.

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

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

[I-D.ietf-6renum-static-problem]
Carpenter, B., and S. Jiang, "Problem Statement for Renumbering IPv6 Hosts with Static Addresses", Working in Progress, August 2012.

[I-D.liu-6renum-dhcpv6-slaac-switching]
Liu, B., "DHCPv6/SLAAC Address Configuration Switching for Host Renumbering", Working in Progress, July 2012

[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>

Authors' Addresses

Bing Liu
Huawei Technologies Co., Ltd
Q14, Huawei Campus
No.156 Beiqing Rd.
Hai-Dian District, Beijing 100095
P.R. China

Email: leo.liubing@huawei.com

Sheng Jiang
Huawei Technologies Co., Ltd
Q14, Huawei Campus
No.156 Beiqing Rd.
Hai-Dian District, Beijing 100095
P.R. China

Email: jiangsheng@huawei.com

Brian Carpenter
Department of Computer Science
University of Auckland
PB 92019
Auckland, 1142
New Zealand

EMail: brian.e.carpenter@gmail.com

Stig Venaas
Cisco Systems
Tasman Drive
San Jose, CA 95134
USA

Email: stig@cisco.com

Wesley George
Time Warner Cable
13820 Sunrise Valley Drive
Herndon, VA 20171
USA

Phone: +1 703-561-2540
Email: wesley.george@twcable.com