

Internet Engineering Task Force (IETF)
Request for Comments: 6879
Category: Informational
ISSN: 2070-1721

S. Jiang
B. Liu
Huawei Technologies Co., Ltd.
B. Carpenter
University of Auckland
February 2013

IPv6 Enterprise Network Renumbering Scenarios,
Considerations, and Methods

Abstract

This document analyzes events that cause renumbering and describes the current renumbering methods. These are described in three categories: those applicable during network design, those applicable during preparation for renumbering, and those applicable during the renumbering operation.

Status of This Memo

This document is not an Internet Standards Track specification; it is published for informational purposes.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Not all documents approved by the IESG are a candidate for any level of Internet Standard; see Section 2 of RFC 5741.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <http://www.rfc-editor.org/info/rfc6879>.

Copyright Notice

Copyright (c) 2013 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.

Table of Contents

1. Introduction	2
2. Enterprise Network Illustration for Renumbering	3
3. Enterprise Network Renumbering Scenario Categories	5
3.1. Renumbering Caused by External Network Factors	5
3.2. Renumbering Caused by Internal Network Factors	5
4. Network Renumbering Considerations and Current Methods	6
4.1. Considerations and Current Methods during Network Design ...	6
4.2. Considerations and Current Methods for the Preparation of Renumbering	10
4.3. Considerations and Current Methods during Renumbering Operation	11
5. Security Considerations	13
6. Acknowledgements	14
7. References	14
7.1. Normative References	14
7.2. Informative References	15

1. Introduction

Site renumbering is difficult. Network managers frequently attempt to avoid future renumbering by numbering their network resources from Provider-Independent (PI) address space. However, widespread use of PI address space would aggravate BGP4 scaling problems [RFC4116] and, depending on Regional Internet Registry (RIR) policies, PI space is not always available for enterprises of all sizes. Therefore, it is desirable to develop mechanisms that simplify IPv6 renumbering for enterprises.

This document is an analysis of IPv6 site renumbering for enterprise networks. It undertakes scenario descriptions, including

documentation of current capabilities and existing practices. The reader is assumed to be familiar with [RFC4192] and [RFC5887]. Proposals for new technology and methods are out of scope.

Since IPv4 and IPv6 are logically separate from the perspective of renumbering, regardless of overlapping of the IPv4/IPv6 networks or devices, this document focuses on IPv6 only, leaving IPv4 out of scope. Dual-stack networks or IPv4/IPv6 transition scenarios are out of scope as well.

This document focuses on enterprise network renumbering; however, most of the analysis is also applicable to ISP network renumbering. Renumbering in home networks is out of scope, but it can also benefit from the analysis in this document.

The concept of an enterprise network and a typical network illustration are introduced first. Then, current renumbering methods are introduced according to the following categories: those applicable during network design, those applicable during preparation for renumbering, and those applicable during the renumbering operation.

2. Enterprise Network Illustration for Renumbering

An Enterprise Network, as defined in [RFC4057], is a network that has multiple internal links, has one or more router connections to one or more Providers, and is actively managed by a network operations entity.

Figure 1 provides a sample enterprise network architecture for a simple case. Those entities mainly affected by renumbering are illustrated:

- * Gateway (Border router, firewall, web cache, etc.)
- * Application server (for internal or external users)
- * DNS and DHCP servers
- * Routers
- * Hosts (desktops, etc.)

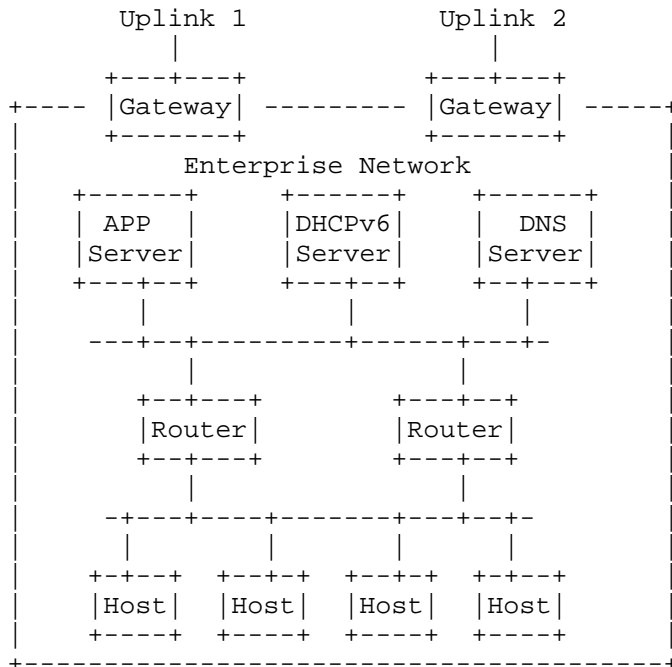


Figure 1. Enterprise Network Illustration

Address reconfiguration is fulfilled either by the Dynamic Host Configuration Protocol for IPv6 (DHCPv6) or by Neighbor Discovery (ND) for IPv6 protocols. During a renumbering event, the Domain Name Service (DNS) records need to be synchronized while routing tables, Access Control Lists (ACLs), and IP filtering tables in various devices also need to be updated. It is taken for granted that applications will work entirely on the basis of DNS names, but any direct dependencies on IP addresses in application-layer entities must also be updated.

The issue of static addresses is described in a dedicated document [RFC6866].

The emerging cloud-based enterprise network architecture might be different than Figure 1. However, it is out of the scope of this document since it is far from mature and has not been widely deployed yet.

It is assumed that IPv6 enterprise networks are IPv6-only or dual-stack in which a logical IPv6 plane is independent from IPv4. As mentioned above, IPv4/IPv6 coexistence scenarios are out of scope.

This document focuses on routable unicast addresses; link-local, multicast, and anycast addresses are also out of scope.

3. Enterprise Network Renumbering Scenario Categories

In this section, we divide enterprise network renumbering scenarios into two categories defined by external and internal network factors, which require renumbering for different reasons.

3.1. Renumbering Caused by External Network Factors

The following ISP uplink-related events can cause renumbering:

- o The enterprise network switches to a new ISP. When this occurs, the enterprise stops numbering its resources from the prefix allocated by the old ISP and renumbers its resources from the prefix allocated by the new ISP.

When the enterprise switches ISPs, a "flag day" renumbering event [RFC4192] may be averted if, during a transitional period, the enterprise network may number its resources from either prefix. One way to facilitate such a transitional period is for the enterprise to contract service from both ISPs during the transition.

- o The renumbering event can be initiated by receiving new prefixes from the same uplink. This might happen if the enterprise network is switched to a different location within the network topology of the same ISP due to various considerations, such as commercial, performance or services reasons, etc. Alternatively, the ISP itself might be renumbered due to topology changes or migration to a different or additional prefix. These ISP renumbering events would initiate enterprise network renumbering events, of course.
- o The enterprise network adds a new uplink(s) for multihoming purposes. This might not be a typical renumbering case because the original addresses will not be changed. However, initial numbering may be considered as a special renumbering event. The enterprise network removes uplink(s) or old prefixes.

3.2. Renumbering Caused by Internal Network Factors

- o As companies split, merge, grow, relocate, or reorganize, the enterprise network architectures might need to be rebuilt. This will trigger partial or total internal renumbering.

- o The enterprise network might proactively adopt a new address scheme, for example, by switching to a new transition mechanism or stage of a transition plan.
- o The enterprise network might reorganize its topology or subnets.

4. Network Renumbering Considerations and Current Methods

In order to carry out renumbering in an enterprise network, systematic planning and administrative preparation are needed. Careful planning and preparation could make the renumbering process smoother.

This section describes current considerations and methods for enterprise renumbering, chosen among existing mechanisms. There are known gaps analyzed by [GAP-ANALYSIS] and [RFC6866]. If these gaps are filled in the future, enterprise renumbering could be processed more automatically, with fewer issues.

4.1. Considerations and Current Methods during Network Design

This section describes the considerations or issues relevant to renumbering that a network architect should carefully plan when building or designing a new network.

- Prefix Delegation (PD)

In a large or a multisite enterprise network, the prefix should be carefully managed, particularly for renumbering events. Prefix information needs to be delegated from router to router. The DHCPv6 PD options ([RFC3633] and [RFC6603]) provide a mechanism for automated delegation of IPv6 prefixes. Normally, DHCPv6 PD options are used between the internal enterprise routers; for example, a router receives a prefix(es) from its upstream router (a border gateway or edge router, etc.) through DHCPv6 PD options and then advertises it (them) to the local hosts through Router Advertisement (RA) messages.

- Usage of Fully Qualified Domain Names (FQDNs)

In general, FQDNs are recommended to be used to configure network connectivity, such as tunnels, servers, etc. The capability to use FQDNs as endpoint names has been standardized in several RFCs (e.g., for IPsec [RFC5996]) although many system/network administrators do not realize that it is there and it works well as a way to avoid manual modification during renumbering.

Note that using FQDNs would rely on DNS systems. For a link-local network that does not have a DNS system, multicast DNS [RFC6762] could be utilized. For some specific circumstances, using FQDNs might not be chosen if adding DNS service in the node/network would cause undesired complexity or issues.

Service discovery protocols such as the Service Location Protocol [RFC2608], multicast DNS with Service Records (SRVs), and DNS Service Discovery [RFC6763] use names and can reduce the number of places that IP addresses need to be configured. However, it should be noted that these protocols are normally used link-local only.

Network designers generally have little control over the design of application software. However, it is important to avoid any software that has a built-in dependency on IP addresses instead of FQDNs [RFC6866].

- Usage of Parameterized Address Configuration

Besides DNS records, IP addresses might also be configured in many other places such as ACLs, various IP filters, various kinds of text-based configuration files, etc.

In some cases, one IP address 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 CLI (command line interface) or other local configuration. Among the real current devices, some routers support defining multiple loopback interfaces that 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 kind of parameterized address configuration is recommended, since it makes managing a renumbering event easier by reducing the number of places where a device's configuration must be updated.

- Usage of Unique Local Addresses (ULAs)

ULAs are defined in [RFC4193] as PI prefixes. Since there is a 40-bit pseudorandom field in the ULA prefix, there is no practical risk of collision (please refer to Section 3.2.3 in [RFC4193] for more detail). For enterprise networks, using ULA simultaneously with PA addresses can provide a local routing plane logically separated from the global routing plane. The benefit is to ensure stable and specific local communication regardless of any ISP uplink failure. This benefit is especially meaningful for renumbering. It mainly includes three use cases described below.

- o During the transition period, it is desirable to isolate local communication changes in the global routing plane. If we use ULA for the local communication, this isolation is achieved.
- o Enterprise administrators might want to avoid the need to renumber their internal-only, private nodes when they have to renumber the PA addresses of the whole network because of changing ISPs, ISPs restructuring their address allocation, or any other reasons. In these situations, a ULA is an effective tool for the internal-only nodes.
- o ULAs can be a way of avoiding renumbering from having an impact on multicast. In most deployments, multicast is only used internally (intra-domain), and the addresses used for multicast sources and Rendezvous Points need not be reachable nor routable externally. Hence, one may, at least internally, make use of ULAs for multicast-specific infrastructure.

- Address Types

This document focuses on the dynamically configured global unicast addresses in enterprise networks. They are the targets of renumbering events.

Manually configured addresses are not scalable in medium to large sites; hence, they should be avoided for both network elements and application servers [RFC6866].

- Address configuration models

In IPv6 networks, there are two autoconfiguration models for address assignment after each host obtains a link-local address: Stateless Address Autoconfiguration (SLAAC) [RFC4862] by ND [RFC4861] and stateful address configuration by DHCPv6 [RFC3315]. In the latest work, DHCPv6 may also support the host-generated address model by assigning a prefix through DHCPv6 messages [PREFIX-DHCPV6].

SLAAC is considered to support easy renumbering by broadcasting an RA message with a new prefix. DHCPv6 can also trigger the renumbering process by sending unicast RECONFIGURE messages, though it might cause a large number of interactions between hosts and the DHCPv6 server.

This document has no preference between the SLAAC and DHCPv6 address configuration models. It is the network architect's job to decide which configuration model is employed. However, it should be noticed that using DHCPv6 and SLAAC together within one network, especially in one subnet, might cause operational issues. For example, some

hosts use DHCPv6 as the default configuration model while some use ND. Then, the host's address configuration model depends on the policies of operating systems and cannot be controlled by the network. Section 5.1 of [GAP-ANALYSIS] discusses more details on this topic. So, in general, this document recommends using DHCPv6 or SLAAC independently in different subnets.

However, since DHCPv6 is also used to configure many other network parameters, there are ND and DHCPv6 coexistence scenarios. Combinations of address configuration models might coexist within a single enterprise network. [SAVI] provides recommendations to avoid collisions and to review collision handling in such scenarios.

- DNS

Although the A6 DNS record model [RFC2874] was designed for easier renumbering, it left many unsolved technical issues [RFC3364]. Therefore, it has been moved to Historic status [RFC6563] and should not be used.

Often, a small site depends on its ISP's DNS system rather than maintaining its own. When renumbering, this requires administrative coordination between the site and its ISP.

It is recommended that the site have an automatic and systematic procedure for updating/synchronizing its DNS records, including both forward and reverse mapping. In order to simplify the operational procedure, the network architect should combine the forward and reverse DNS updates in a single procedure. A manual on-demand updating model does not scale and increases the chance of errors. Either a database-driven mechanism, a secure dynamic DNS update [RFC3007], or both could be used.

A dynamic DNS update can be provided by the DHCPv6 client or by the server on behalf of individual hosts. [RFC4704] defines a DHCPv6 option to be used by DHCPv6 clients and servers to exchange information about the client's FQDN and about who has the responsibility for updating the DNS with the associated AAAA and PTR (Pointer Record) RRs (Resource Records). For example, if a client wants the server to update the FQDN-address mapping in the DNS server, it can include the Client FQDN option with proper settings in the SOLICIT with Rapid Commit, REQUEST, RENEW, and REBIND message originated by the client. When the DHCPv6 server gets this option, it can use a secure dynamic DNS update on behalf of the client. This document suggests use of this FQDN option. However, since it is a DHCPv6 option, only the DHCP-managed hosts can make use of it. In SLAAC mode, hosts need either to use a secure dynamic DNS update

directly, or to register addresses on a registration server. This could in fact be a DHCPv6 server (as described in [ADDR-REG]); then the server would update corresponding DNS records.

- Security

Any automatic renumbering scheme has a potential exposure to hijacking. A malicious entity in the network could forge prefixes to renumber the hosts, so proper network security mechanisms are needed. Further details are in the Security Considerations section below.

- Miscellaneous

A site or network should also avoid embedding addresses from other sites or networks in its own configuration data. Instead, the FQDNs should be used. Thus, connections can be restored after renumbering events at other sites. This also applies to host-based connectivity.

4.2. Considerations and Current Methods for the Preparation of Renumbering

In ND, it is not possible to reduce a prefix's lifetime to below two hours. So, renumbering should not be an unplanned sudden event. This issue could only be avoided by early planning and preparation.

This section describes several recommendations for the preparation of an enterprise renumbering event. By adopting these recommendations, a site could be renumbered more easily. However, these recommendations might increase the daily traffic, server load, or burden of network operation. Therefore, only those networks that are expected to be renumbered soon, or very frequently, should adopt these recommendations, with balanced consideration between daily cost and renumbering cost.

- Reduce the address preferred time or valid time or both

Long-lifetime addresses might cause issues for renumbering events. Particularly, some offline hosts might reconnect using these addresses after renumbering events. Shorter, preferred lifetimes with relatively long valid lifetimes may allow short transition periods for renumbering events and avoid frequent address renewals.

- Reduce the DNS record Time to Live (TTL) on the local DNS server

The DNS AAAA RR TTL on the local DNS server should be manipulated to ensure that stale addresses are not cached.

Recent research [BA2011] [JSBM2002] indicates that it is both practical and reasonable for A, AAAA, and PTRs that belong to leaf nodes of the DNS (i.e., not including the DNS root or DNS top-level domains) to be configured with very short DNS TTL values, not only during renumbering events but also for longer-term operation.

- Reduce the DNS configuration lifetime on the hosts

Since the DNS server could be renumbered as well, the DNS configuration lifetime of the hosts should also be reduced if renumbering events are expected. In ND, the DNS configuration can be done through reducing the lifetime value in the Recursive DNS Server (RDNS) option [RFC6106]. In DHCPv6, the DNS configuration option specified in [RFC3646] doesn't provide a lifetime attribute, but we can reduce the DHCPv6 client lease time to achieve a similar effect.

- Identify long-living sessions

Any applications that maintain very long transport connections (hours or days) should be identified in advance, if possible. Such applications will need special handling during renumbering, so it is important to know that they exist.

4.3. Considerations and Current Methods during Renumbering Operation

Renumbering events are not instantaneous events. Normally, there is transition period in which both the old prefix and the new prefix are used in the site. Better network design and management, better preparation, and a longer transition period are helpful to reduce the issues during a renumbering operation.

- Within/Without a flag day

As is described in [RFC4192] "a 'flag day' is a procedure in which the network, or a part of it, is changed during a planned outage, or suddenly, causing an outage while the network recovers".

If a renumbering event is processed within a flag day, the network service/connectivity will be unavailable for a period until the renumbering event is completed. It is efficient and provides convenience for network operation and management. However, a network outage is usually unacceptable for end users and enterprises. A renumbering procedure without a flag day provides smooth address switching, but much more operational complexity and difficulty is introduced.

- Transition period

If a renumbering transition period is longer than all address lifetimes, after which the address leases expire, each host will automatically pick up its new IP address. In this case, it would be the DHCPv6 server or RA itself that automatically accomplishes client renumbering.

Address deprecation should be associated with the deprecation of associated DNS records. The DNS records should be deprecated as early as possible, before the addresses themselves.

- Network initiative enforced renumbering

If the network has to enforce renumbering before address leases expire, the network should initiate DHCPv6 RECONFIGURE messages. For some operating systems such as Windows 7, if the hosts receive RA messages with ManagedFlag=0, they will release the DHCPv6 addresses and utilize SLAAC according to the prefix information in the RA messages, so this could be another enforcement method for some specific scenarios.

- Impact on main and branch sites

Renumbering in the main site might cause impact on branch site communications, and vice versa. The routes, ingress filtering of the site's gateways, and DNS might need to be updated. This needs careful planning and organizing.

- DNS record update and DNS configuration on hosts

DNS records on the local DNS server should be updated if hosts are renumbered. If the site depends on an ISP's DNS system, it should report the new hosts' DNS records to its ISP. During the transition period, both old and new DNS records are valid. If the TTLs of DNS records are shorter than the transition period, an administrative operation might not be necessary.

DNS configuration on hosts should be updated if local recursive DNS servers are renumbered. During the transition period, both old and new DNS server addresses might coexist on the hosts. If the lifetime of DNS configuration is shorter than the transition period, name resolving failure may be reduced to a minimum.

- Tunnel concentrator renumbering

A tunnel concentrator itself might be renumbered. This change should be reconfigured in relevant hosts or routers, unless the configuration of the tunnel concentrator was based on FQDN.

For IPsec, Internet Key Exchange Protocol version 2 (IKEv2) [RFC5996] defines the ID_FQDN Identification type, which could be used to identify an IPsec VPN concentrator associated with a site's domain name. For current practice, the community needs to change its bad habit of using IPsec in an address-oriented way, and renumbering is one of the main reasons for that.

- Connectivity session survivability

During the renumbering operations, connectivity sessions in the IP layer would break if the old address is deprecated before the session ends. However, the upper-layer sessions can survive by using session survivability technologies, such as Stanza Headers and Internet Metadata 6 (SHIM6) [RFC5533]. As mentioned above, some long-living applications may need to be handled specially.

- Verification of success

The renumbering operation should end with a thorough check that all network elements and hosts are using only the new prefixes and that network management and monitoring systems themselves are still operating correctly. A database clean up may also be needed.

5. Security Considerations

Any automatic renumbering scheme has a potential exposure to hijacking by an insider attack. For attacks on ND, SEcure Neighbor Discovery (SEND) [RFC3971] is a possible solution, but it is complex and there is almost no real deployment at the time of writing. Compared to the nontrivial deployment of SEND, RA-Guard [RFC6105] is a lightweight alternative that focuses on preventing rogue router advertisements in a network. However, it is also not widely deployed at the time when this memo was published.

For DHCPv6, there are built-in secure mechanisms (like Secure DHCPv6 [SECURE-DHCPV6]), and authentication of DHCPv6 messages [RFC3315] could be utilized. However, these security mechanisms also have not been verified by widespread deployment at the time of writing.

A site that is listed by IP address in a blacklist can escape that list by renumbering itself. However, the new prefix might be back on a blacklist rather soon if the root cause for being added to such a list is not corrected. In practice, the cost of renumbering will typically be much larger than the cost of getting off the blacklist.

A Dynamic DNS update might bring risk of a Denial-of-Service (DoS) attack to the DNS server. So, along with the update authentication, session filtering/limitation might also be needed.

The "make-before-break" approach of [RFC4192] requires the routers to keep advertising the old prefixes for some time. However, if the ISP changes the prefixes very frequently, the coexistence of old and new prefixes might cause potential risk to the enterprise routing system, since the old address relevant route path might already be invalid and the routing system just doesn't know it. However, normally, enterprise scenarios don't involve this extreme situation.

6. Acknowledgements

This work is inspired by RFC 5887; thank you to the authors (Randall Atkinson and Hannu Flinck). Useful ideas were also presented in documents by Tim Chown and Fred Baker. The authors also want to thank Wesley George, Olivier Bonaventure, Lee Howard, Ronald Bonica, other 6renum members, and several reviewers for their valuable comments.

7. References

7.1. Normative References

- [RFC2608] Guttman, E., Perkins, C., Veizades, J., and M. Day, "Service Location Protocol, Version 2", RFC 2608, June 1999.
- [RFC3007] Wellington, B., "Secure Domain Name System (DNS) Dynamic Update", RFC 3007, November 2000.
- [RFC3315] Droms, R., Ed., 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.

- [RFC3646] Droms, R., Ed., "DNS Configuration options for Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", RFC 3646, December 2003.
- [RFC3971] Arkko, J., Ed., Kempf, J., Zill, B., and P. Nikander, "SEcure Neighbor Discovery (SEND)", RFC 3971, March 2005.
- [RFC4057] Bound, J., Ed., "IPv6 Enterprise Network Scenarios", RFC 4057, June 2005.
- [RFC4193] Hinden, R. and B. Haberman, "Unique Local IPv6 Unicast Addresses", RFC 4193, October 2005.
- [RFC4704] Volz, B., "The Dynamic Host Configuration Protocol for IPv6 (DHCPv6) Client Fully Qualified Domain Name (FQDN) Option", RFC 4704, October 2006.
- [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.
- [RFC5996] Kaufman, C., Hoffman, P., Nir, Y., and P. Eronen, "Internet Key Exchange Protocol Version 2 (IKEv2)", RFC 5996, September 2010.
- [RFC6106] Jeong, J., Park, S., Beloeil, L., and S. Madanapalli, "IPv6 Router Advertisement Options for DNS Configuration", RFC 6106, November 2010.

7.2. Informative References

- [RFC2874] Crawford, M. and C. Huitema, "DNS Extensions to Support IPv6 Address Aggregation and Renumbering", RFC 2874, July 2000.
- [RFC3364] Austein, R., "Tradeoffs in Domain Name System (DNS) Support for Internet Protocol version 6 (IPv6)", RFC 3364, August 2002.
- [RFC4116] Abley, J., Lindqvist, K., Davies, E., Black, B., and V. Gill, "IPv4 Multihoming Practices and Limitations", RFC 4116, July 2005.

- [RFC4192] Baker, F., Lear, E., and R. Droms, "Procedures for Renumbering an IPv6 Network without a Flag Day", RFC 4192, September 2005.
- [RFC5533] Nordmark, E. and M. Bagnulo, "Shim6: Level 3 Multihoming Shim Protocol for IPv6", RFC 5533, June 2009.
- [RFC5887] Carpenter, B., Atkinson, R., and H. Flinck, "Renumbering Still Needs Work", RFC 5887, May 2010.
- [RFC6105] Levy-Abegnoli, E., Van de Velde, G., Popoviciu, C., and J. Mohacsi, "IPv6 Router Advertisement Guard", RFC 6105, February 2011.
- [RFC6563] Jiang, S., Conrad, D., and B. Carpenter, "Moving A6 to Historic Status", RFC 6563, March 2012.
- [RFC6603] Korhonen, J., Ed., Savolainen, T., Krishnan, S., and O. Troan, "Prefix Exclude Option for DHCPv6-based Prefix Delegation", RFC 6603, May 2012.
- [RFC6762] Cheshire, S. and M. Krochmal, "Multicast DNS", RFC 6762, February 2013.
- [RFC6763] Cheshire, S. and M. Krochmal, "DNS-Based Service Discovery", RFC 6763, February 2013.
- [RFC6866] Carpenter, B., and S. Jiang, "Problem Statement for Renumbering IPv6 Hosts with Static Addresses in Enterprise Networks", RFC 6866, February 2013.
- [ADDR-REG] Jiang, S., Chen, G., and S. Krishnan "A Generic IPv6 Addresses Registration Solution Using DHCPv6", Work in Progress, February 2013.
- [BA2011] S. Bhatti, and R. Atkinson, "Reducing DNS Caching", Proc. 14th IEEE Global Internet Symposium (GI2011), Shanghai, China, April 15 2011.
- [GAP-ANALYSIS] Liu, B., Jiang, S., Carpenter, B. Venaas, S., and W. George, "IPv6 Site Renumbering Gap Analysis", Work in Progress, December 2012.

- [JSBM2002] J. Jung, E. Sit, H. Balakrishnan, and R. Morris, "DNS Performance and the Effectiveness of Caching", *IEEE/ACM Transactions on Networking*, 10(5):589-603, 2002.
- [PREFIX-DHCPV6]
Jiang, S., Xia, F., and B. Sarikaya, "Prefix Assignment in DHCPv6", Work in Progress, February 2013.
- [SAVI] Bi, J., Yao, G., Halpern, J., and E. Levy-Abegnoli, "SAVI for Mixed Address Assignment Methods Scenario", Work in Progress, November 2012.
- [SECURE-DHCPV6]
Jiang, S., and S. Shen, "Secure DHCPv6 Using CGAs", Work in Progress, March 2012.

Authors' Addresses

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

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

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

EMail: brian.e.carpenter@gmail.com