Network Working Group Internet Draft Intended status: Informational Expires: August 9, 2012

IPv6 Enterprise Network Renumbering Scenarios and Guidelines draft-ietf-6renum-enterprise-00.txt

Status of this Memo

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

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 <u>http://datatracker.ietf.org/drafts/current/</u>.

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 August 9, 2012.

Copyright Notice

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

This document is subject to <u>BCP 78</u> and the IETF Trust's Legal Provisions Relating to IETF Documents (<u>http://trustee.ietf.org/license-info</u>) 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 analyzes enterprise renumbering events and describes the best current practice among the existing renumbering mechanisms. According to the different stages of renumbering events, considerations and best current practices are described in three categories: during network design, for preparation of renumbering, and during a renumbering operation. A gap inventory is listed at the end of this document.

Table of Contents

<u>1</u> .	Introduction
<u>2</u> .	Enterprise Network Illustration for Renumbering3
<u>3</u> .	Enterprise Network Renumbering Scenario Categories <u>4</u>
	<u>3.1</u> . Renumbering caused by External Network Factors <u>5</u>
	<u>3.2</u> . Renumbering caused by Internal Network Factors <u>5</u>
4.	Network Renumbering Considerations and Best Current Practise . 5
	4.1. Considerations and Best Current Practice during Network
	Design
	4.2. Considerations and Best Current Practice for the Preparation
	of Renumbering
	4.3. Considerations and Best Current Practice during Renumbering
	Operation <u>10</u>
<u>5</u> .	Gap Inventory <u>12</u>
<u>6</u> .	Security Considerations <u>13</u>
<u>7</u> .	IANA Considerations <u>13</u>
<u>8</u> .	Acknowledgements <u>13</u>
<u>9</u> .	Change Log [RFC Editor please remove] <u>14</u>
<u>10</u>	. References
	<u>10.1</u> . Normative References <u>14</u>
	<u>10.2</u> . Informative References <u>15</u>
Au	thor's Addresses <u>17</u>

1. Introduction

IPv6 site renumbering is considered difficult. Network managers currently prefer to use 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 and PI space is not always available for enterprise according to the RIR (Regional Internet Registry) policies. It is thus desirable to develop tools and practices that may make renumbering a simpler process to reduce demand for IPv6 PI space. In any case, renumbering may be necessary for other reasons.

This document undertakes scenario descriptions, including documentation of current capabilities and existing BCPs, for enterprise networks. It takes [<u>RFC5887</u>] and other relevant documents as the primary input.

The IPv4 and IPv6 are logically separated from the perspective of renumbering, regardless of overlapping of the IPv4/IPv6 networks or devices. This document focuses on IPv6 only, by leaving IPv4 out of scope. Dual-stack network or IPv4/IPv6 transition scenarios are out of scope, too.

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

The concept of enterprise network and a typical network illustration are introduced first. Then, according to the different stages of renumbering events, considerations and best current practices are described in three categories: during network design, for preparation of renumbering, and during renumbering operation. A gap inventory is listed at the end of this document.

2. Enterprise Network Illustration for Renumbering

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

The enterprise network architecture is illustrated in the figure below. Those entities relevant to renumbering are highlighted.

Address reconfiguration is fulfilled either by DHCPv6 or ND protocols. During the renumbering event, the DNS records need to be synchronized while routing tables, ACLs and IP filtering tables in various gateways also need to be updated, too.

Static address issue is described in a dedicated draft [I-D.carpenter-6renum-static-problem]. (Editor's note: some major conclusions would be included in this document if we can get consensus on the discussion of the static address problem.)



Figure 1 Enterprise network illustration

It is assumed that IPv6 enterprise networks are IPv6-only, or dualstack in which a logical IPv6 plane is independent from IPv4. The complicated IPv4/IPv6 co-existence scenarios are out of scope.

This document focuses on the unicast addresses; site-local, linklocal, multicast and anycast addresses are 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 most influential external network factor is the uplink ISP.

o The enterprise network switches to a new ISP. Of course, the prefixes received from different ISPs are different. This is the most common scenario.

Whether there is an overlap time between the old and new ISPs would also influence the possibility whether the enterprise can fulfill renumbering without a flag day [<u>RFC4192</u>].

- o The renumbering event may 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 new uplink(s) for multihoming purposes. This may not a typical renumbering 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 may need to be re-built. This will trigger the internal renumbering.
- o The enterprise network may 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 may reorganize its topology or subnets.

4. Network Renumbering Considerations and Best Current Practices

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

This section tries to give the recommended solutions or strategies for the enterprise renumbering, chosen among existing mechanisms. There are known gaps analyzed by [I-D.liu-6renum-gap-analysis]. If these gaps are filled in the future, the enterprise renumbering may be processed more automatically, with fewer issues.

4.1. Considerations and Best Current Practices during Network Design

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

- Prefix Delegation

In a large or a multi-site enterprise network, the prefix should be carefully managed, particularly during renumbering events. Prefix information needs to be delegated from router to router. The DHCPv6 Prefix Delegation options [RFC3633] [I-D.ietf-dhc-pdexclude] provide a mechanism for automated delegation of IPv6 prefixes. DHCPv6 PD options may also be used between the enterprise routers and their upstream ISPs.

- Usage of FQDN

In general, Fully-Qualified Domain Names (FQDNs) are recommended to be used to configure network connectivity, such as tunnels, whenever possible. The capability to use FQDNs as endpoint names has been standardized in several RFCs, such as [<u>RFC5996</u>], although many system/network administrators do not realize that it is there and works well as a way to avoid manual modification during renumbering.

Service Location Protocol [RFC2608] and multicast DNS with SRV records for service discovery can reduce the number of places that IP addresses need to be configured. But it should be noted that multicast DNS is link-local only.

- Address Types

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

Manual-configured addresses are not scalable in medium to large sites, hence are out of scope. Manual-configured addresses/hosts should be avoided as much as possible.

Unique Local Addresses (ULA, [RFC4193]) may be used for local communications, usually inside of enterprise networks. They can be sufficient for any host that is accessible only inside the enterprise network and has no need for external communication [RFC4864]. Normally, they do not need to be changed during a global prefix renumbering event. However, they may need to be renumbered in some rare scenarios, quite separate from the global prefix renumbering.

- Address configuration models

In IPv6 networks, there are two auto-configuration models for address assignment: Stateless Address Auto-Configuration (SLAAC) by Neighbor Discovery (ND, [RFC4861, <u>RFC4862</u>]) and stateful address configuration by Dynamic Host Configuration Protocol for IPv6 (DHCPv6, [<u>RFC3315</u>]). In the latest work, DHCPv6 can also support host-generated address model by assigning a prefix through DHCPv6 messages [<u>I-D.ietf-dhc-host-gen-id</u>].

ND is considered easier to renumber by broadcasting a Router Advertisement message with a new prefix. DHCPv6 can also trigger the renumbering process by sending unicast RECONFIGURE messages, though it may cause a large number of interactions between hosts and DHCPv6 server.

This document has no preference between ND and DHCPv6 address configuration models. It is network architects' job to decide which configuration model is employed. But it should be noticed that using DHCPv6 and ND together within one network, especially in one subnet, may cause operational issues. For example, some hosts use DHCPv6 as the default configuration model while some use ND. Then the hosts' address configuration model depends on the policies of operating systems and cannot be controlled by the network. <u>Section 5.1</u> of [I-D.liu-6renum-gap-analysis] discusses more details on this topic. So, in general, this document recommends using DHCPv6/SLAAC independently in different subnets.

However, since DHCPv6 is also used to configure many other network parameters, there are ND and DHCPv6 co-existence scenarios. Combinations of address configuration models may coexist within a single enterprise network. [I-D.ietf-savi-mix] provides recommendations to avoid collisions and to review collision handling in such scenarios.

- DNS

It is recommended that the site have an automatic and systematic procedure for updating/synchronising its DNS records, including both forward and reverse mapping [RFC2874]. A manual on-demand updating model does not scale, and increases the chance of errors.

Although the A6 DNS record model [RFC2874] was designed for easier renumbering, it has a lot of unsolved technical issues [RFC3364]. Therefore, it has been moved to experimental status [RFC3363], and will move to historic status by [I-D.jiang-dnsext-a6-to-historic] (It is currently in RFC Editor Queue already). So A6 is not recommended.

In order to simplify the operation procedure, the network architect should combine the forward and reverse DNS updates in a single procedure.

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.

The DNS synchronization may be completed through the Secure DNS Dynamic Update [RFC3007]. Normally, the dynamic DNS update is achieved by DHCPv6 server on behalf of individual hosts. [RFC4704] defined 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 RRs. 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 DHCPv6 server gets this option, it can use the dynamic DNS update on behalf of the client. In this document, we promote to support this FQDN option. But since it's a DHCPv6 option, it implies that only the DHCP-managed networks are suitable for this operation. In a model including SLAAC, host addresses may be registered on an address registration server, which could in fact be a DHCPv6 server; then the server would update corresponding DNS records.

- Security

Any automatic renumbering scheme has a potential exposure to hijacking. Malicious entity in the network can forge prefixes to renumber the hosts. So proper network security mechanisms are needed.

Internet-Draft draft-ietf-6renum-enterprise-00.txt February 2012

For ND, Secure Neighbor Discovery (SEND, [<u>RFC3971</u>]) is a possible solution, but it is complex and there's almost no real deployment so far. Comparing the non-trivial deployment of SEND, RA guard [RFC6105] is a light-weight alternative, however, it also hasn't been widely deployed since it hasn't been published for long.

For DHCPv6, there are built-in secure mechanisms (like Secure DHCPv6 [<u>I-D.ietf-dhc-secure-dhcpv6</u>]), and authentication of DHCPv6 messages [<u>RFC3315</u>] could be utilized. But these security mechanisms also haven't been verified by wide real deployment.

- Miscellaneous

A site or network should also avoid embedding addresses from other sites or networks in its own configuration data. Instead, the Fully-Oualified Domain Names should be used. Thus, these connections can survive after renumbering events at other sites. This also applies to host-based connectivities.

4.2. Considerations and Best Current Practices 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 enterprise renumbering event. By adopting these recommendations, a site could be renumbered more easily. However, these recommendations are not cost free. They might increase the daily 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 may cause issues for renumbering events. Particularly, some offline hosts may 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 TTL on the local DNS server.

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

- Reduce the DNS configuration lifetime on the hosts.

Since the DNS server could be renumbered as well, the DNS configuration lifetime on the hosts should also be reduced if renumbering events are expected. The DNS configuration can be done through either ND [RFC6106] or DHCPv6 [RFC3646].

- Identify long-living sessions

Any applications which 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.

<u>4.3</u>. Considerations and Best Current Practices during Renumbering Operation

Renumbering events are not instantaneous events. Normally, there is a transition period, in which both the old prefix and the new prefix are used in the site. Better network design and management, better pre-preparation and longer transition period are helpful to reduce the issues during 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 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. But 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 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 Router Advertisement 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 enforcement messages, either in Router Advertisement messages or DHCPv6 RECONFIGURE messages.

- Impact to branch/main sites

Renumbering in main/branch site may cause impact on branch/main site communication. The routes, ingress filtering of site's gateways, and DNS may 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 ISP's DNS system, it should report the new host's 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 may 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 may co-exist on the hosts. If the lifetime of DNS configuration is shorter than the transition period, name resolving failure may be reduced to minimum. A notification mechanism may be needed to indicate to the hosts that a renumbering event of local recursive DNS happens or is going to take place.

- 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.

- Tunnel concentrator renumbering

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

- Connectivity session survivability

During the renumbering operations, connectivity sessions in IP layer would break if the old address is deprecated before the session ends. However, the upper layer sessions may survive by using session survivability technologies, such as SHIM6 [<u>RFC5533</u>]. As mentioned above, some long-living applications may need to be handled specially.

<u>5</u>. Gap Inventory

This section lists a few issues that still appear to remain unsolvable (also see [I-D.liu-6renum-gap-analysis]). Some of them may be inherently unsolvable.

- Some environments like embedded systems might not use DHCPv6 or SLAAC and even configuration scripts might not be an option. This creates special problems that no general-purpose solution is likely to address.
- TCP and UDP flows can't survive a renumbering event at either end.
- The embedding of IPv6 unicast addresses into multicast addresses and the embedded-RP (Rendezvous Point) [<u>RFC3956</u>] will cause issues when renumbering.
- Changing the unicast source address of a multicast sender might also be an issue for receivers.
- When a renumbering event takes place, entries in the state table of tunnel concentrator that happen to contain the old addresses will become invalid and will eventually time out. However, this can be considered as harmless though it takes resources on these devices for a while.
- A site that is listed in an IP black list can escape that list by renumbering itself. The site itself of course will not

report its renumbering and the black list may not be able to monitor or discover the renumbering event.

- Multihomed sites, using SLAAC for one address prefix and DHCPv6 for another, would clearly create a risk of inconsistent host behaviour and operational confusion.

6. Security Considerations

As noted, a site that is listed by IP address in a black list can escape that list by renumbering itself.

Any automatic renumbering scheme has a potential exposure to hijacking. Proper network security mechanisms are needed. Although there are existing security mechanisms such as SEND, RA guard, secure DHCPv6 etc., they haven't been widely deployed and haven't been verified whether they are suitable for ensuring security while not bringing too much operational complexity and cost.

Dynamic DNS update may bring risk of DoS attack to the DNS server. So along with the update authentication, session filtering/limitation may also be needed.

The "make-before-break" approach of [RFC4192] requires the routers keep advertising the old prefixes for some time. But if the ISP changes the prefixes very frequently, the co-existence of old and new prefixes may cause potential risk to the enterprise routing system. However, enterprise scenarios may not involve the extreme situation; this issue needs to be identified in the future.

The security configuration updates will need to be made in two stages (immediately before and immediately after the event).

7. IANA Considerations

This draft does not request any IANA action.

8. Acknowledgements

This work is illuminated by RFC 5887, so thank for RFC 5887 authors, Randall Atkinson and Hannu Flinck. Useful ideas were also presented in by documents from Tim Chown and Fred Baker. The authors also want to thank Wesley George, Olivier Bonaventure and other 6renum members for valuable comments.

Internet-Draft <u>draft-ietf-6renum-enterprise-00.txt</u> February 2012

9. Change Log [RFC Editor please remove]

draft-jiang-6renum-enterprise-00, original version, 2011-07-01

<u>draft-jiang-6renum-enterprise-01</u>, Update according to IETF81 and mail list discussions, 2011-10-09

<u>draft-jiang-6renum-enterprise-02</u>, Update according to IETF82 discussions, 2011-12-06

<u>draft-ietf-6renum-enterprise-00</u>, Update according to mail list discussions, 2012-02-06

10. References

<u>10.1</u>. Normative References

- [RFC2608] Guttman, E., Perkins, C., Veizades, J., and M. Day "Service Location Protocol, Version 2", <u>RFC 2608</u>, June 1999.
- [RFC3007] B. Wellington, "Secure Domain Name System (DNS) Dynamic Update", <u>RFC 3007</u>, November 2000.
- [RFC3315] Droms, R., Bound, J., Volz, B., Lemon, T., Perkins, C., and M. Carney, "Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", <u>RFC 3315</u>, July 2003.
- [RFC3633] Troan, O., and R. Droms, "IPv6 Prefix Options for Dynamic Host Configuration Protocol (DHCP) version 6", <u>RFC 3633</u>, December 2003.
- [RFC3646] R. Droms, "DNS Configuration options for Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", <u>RFC 3646</u>, December 2003.
- [RFC3956] Savola, P., and B. Haberman, "Embedding the Rendezvous Point (RP) Address in an IPv6 Multicast Address", <u>RFC 3956</u>, November 2004
- [RFC3971] Arkko, J., Ed., Kempf, J., Zill, B., and P. Nikander "SEcure Neighbor Discovery (SEND)", <u>RFC 3971</u>, March 2005
- [RFC4193] Hinden, R., and B. Haberman, "Unique Local IPv6 Unicast Addresses", <u>RFC 4193</u>, October 2005.

Jiang, et al.Expires August 10, 2012[Page 14]

- [RFC4861] Narten, T., Nordmark, E., Simpson, W., and H. Soliman, "Neighbor Discovery for IP version 6 (IPv6)", <u>RFC 4861</u>, September 2007.
- [RFC4862] Thomson, S., Narten, T., and T. Jinmei, "IPv6 Stateless Address Autoconfiguration", <u>RFC 4862</u>, September 2007.
- [RFC5996] Kaufman, C., Hoffman, P., Nir, Y., and P. Eronen, "Internet Key Exchange Protocol Version 2 (IKEv2)", <u>RFC 5996</u>, September 2010.

<u>10.2</u>. Informative References

- [RFC2874] Crawford, M., and C. Huitema, "DNS Extensions to Support IPv6 Address Aggregation and Renumbering", <u>RFC 2874</u>, July 2000.
- [RFC3364] R. Austein, "Tradeoffs in Domain Name System (DNS) Support for Internet Protocol version 6 (IPv6)", <u>RFC 3364</u>, August 2002.
- [RFC4057] J. Bound, Ed. "IPv6 Enterprise Network Scenarios", <u>RFC</u> 4057, June 2005.
- [RFC4192] Baker, F., Lear, E., and R. Droms, "Procedures for Renumbering an IPv6 Network without a Flag Day", <u>RFC 4192</u>, September 2005.
- [RFC4864] Van de Velde, G., T. Hain, R. Droms, B. Carpenter, E. Klein, Local Network Protection for IPv6", <u>RFC 4864</u>, May 2007.
- [RFC5533] Nordmark, E., and Bagnulo, M., "Shim6: Level 3 Multihoming Shim Protocol for IPv6", <u>RFC 5533</u>, June 2009.

- [RFC5887] Carpenter, B., Atkinson, R., and H. Flinck, "Renumbering Still Needs Work", <u>RFC 5887</u>, May 2010.
- [RFC6105] Levy-Abegnoli, E., Van de Velde, G., Popoviciu, C., and J. Mohacsi, "IPv6 Router Advertisement Guard", RFC 6105, February 2011.
- [I-D.ietf-dhc-secure-dhcpv6] Jiang, S., and S. Shen, "Secure DHCPv6 Using CGAs", working in progress.
- [I-D.ietf-dhc-host-gen-id] S. Jiang, F. Xia, and B. Sarikaya, "Prefix Assignment in DHCPv6", draft-ietf-dhc-host-gen-id (work in progress), April, 2011.
- [I-D.ietf-savi-mix]

Bi, J., Yao, G., Halpern, J., and Levy-Abegnoli, E., "SAVI for Mixed Address Assignment Methods Scenario", working in progress.

- [I-D.ietf-dhc-pd-exclude] J. Korhonen, T. Savolainen, S. Krishnan, O. Troan, "Prefix Exclude Option for DHCPv6-based Prefix Delegation", working in progress.
- [I-D.ietf-6renum-gap-analysis] Liu, B., and Jiang, S., "IPv6 Site Renumbering Gap Analysis", working in progress.
- [I-D.jiang-dnsext-a6-to-historic]

Jiang, S., Conrad, D. and Carpenter, B., "Moving A6 to Historic Status", working in progress.

[I-D.carpenter-6renum-static-problem]

Carpenter, B. and S. Jiang., "Problem Statement for Renumbering IPv6 Hosts with Static Addresses", working in progress.

Author's Addresses

Sheng Jiang Huawei Technologies Co., Ltd Huawei Q14 Building, No.156 Beiging Rd., Zhong-Guan-Cun Environmental Protection Park, Hai-Dian District EMail: jiangsheng@huawei.com

Bing Liu Huawei Technologies Co., Ltd Huawei Q14 Building, No.156 Beiging Rd., Zhong-Guan-Cun Environmental Protection Park, Hai-Dian District 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