

Network Working Group
Internet Draft
Intended status: Informational
Expires: July 18, 2013

S. Jiang
B. Liu
Huawei Technologies Co., Ltd
B. Carpenter
University of Auckland
January 15, 2013

IPv6 Enterprise Network Renumbering Scenarios,
Considerations and Methods
draft-ietf-6renum-enterprise-06.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 July 18, 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 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.

Table of Contents

1. Introduction	3
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	6
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	12
5. Security Considerations	14
6. IANA Considerations	14
7. Acknowledgements	14
8. References	15
8.1. Normative References	15
8.2. Informative References	16
Author's Addresses	18

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 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 network or IPv4/IPv6 transition scenarios are out of scope, too.

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, 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.)

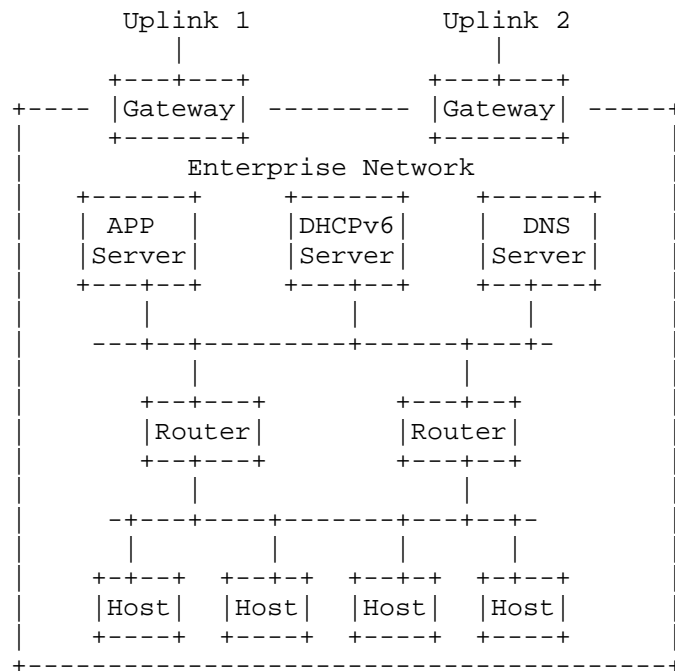


Figure 1 Enterprise network illustration

Address reconfiguration is fulfilled either by the Dynamic Host configuration Protocol for IPv6 (DHCPv6) or Neighbor Discovery for IPv6 (ND) 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 draft [I-D.ietf-6renum-static-problem].

The emerging cloud-based enterprise network architecture might be different with Figure 1. But it is out of the scope of this document since the 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 co-existence 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 stop numbering its resources from the prefix allocated by the old ISP and rennumbers 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 for 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 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 re-built. 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 solutions or strategies for enterprise renumbering, chosen among existing mechanisms. There are known gaps analyzed by [I-D.ietf-6renum-gap-analysis] and [I-D.ietf-6renum-static-problem]. If these gaps are filled in the future, enterprise renumbering can be processed more automatically, with fewer issues.

4.1. Considerations and Current Methods 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] and [RFC6603] provide a mechanism for automated delegation of IPv6 prefixes. Normally, DHCPv6 Prefix Delegation (PD) options are used between the internal enterprise routers, for example, a router receives 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 FQDN

In general, Fully-Qualified Domain Names (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, for example for IPsec [RFC5996], although many system/network administrators do not realize that it is there and works well as a way to avoid manual modification during renumbering.

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

Service discovery protocols such as Service Location Protocol [RFC2608], multicast DNS with SRV records and DNS Service Discovery [I-D.cheshire-dnsext-dns-sd] use names and can reduce the number of places that IP addresses need to be configured. But 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 built-in dependency on IP addresses instead of FQDNs [I-D.ietf-6renum-static-problem].

- 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 in CLI (command line interface) or other local configurations. 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 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 ULA

Unique Local Addresses (ULAs) are defined in [RFC4193] as provider-independent prefixes. Since there is a 40 bits pseudo random 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 Provider Aggregated (PA) addresses can provide a logically local routing plane 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.

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.

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, ULA is an effective tool for the internal-only nodes.

ULA 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 ULA 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 should be avoided for both network elements and application servers [I-D.ietf-6renum-static-problem].

- Address configuration models

In IPv6 networks, there are two auto-configuration models for address assignment after each host obtains a link-local address: Stateless Address Auto-Configuration (SLAAC, [RFC4862]) by Neighbor Discovery (ND, [RFC4861]) and stateful address

configuration by Dynamic Host Configuration Protocol for IPv6 (DHCPv6, [RFC3315]). In the latest work, DHCPv6 may also support the host-generated address model by assigning a prefix through DHCPv6 messages [I-D.ietf-dhc-host-gen-id].

SLAAC is considered to support easy renumbering by broadcasting a Router Advertisement 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 architects' job to decide which configuration model is employed. But 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 hosts' address configuration model depends on the policies of operating systems and cannot be controlled by the network. Section 5.1 of [I-D.ietf-6renum-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 co-existence scenarios. Combinations of address configuration models might 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

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, or Secure Dynamic DNS Update [RFC3007], or both, could be used.

Dynamic DNS update can be provided by the DHCPv6 client or by the 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 (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 DHCPv6 server gets this option, it can use 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 Secure Dynamic DNS Update directly, or to register addresses on a registration server. This could in fact be a DHCPv6 server (as described in [I-D.ietf-dhc-addr-registration]); 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 below.

- Miscellaneous

A site or network should also avoid embedding addresses from other sites or networks in its own configuration data. Instead, the Fully-Qualified Domain Names 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 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 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.

Recent research [BA2011] [JSBM2002] indicates that it is both practical and reasonable for A, AAAA, and PTR records 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 on 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 RDNSS 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 similar effect.

- 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.

4.3. Considerations and Current Methods 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 DHCPv6 RECONFIGURE messages. For some operating systems such as Windows 7, if the hosts receive RA messages with ManagedFlag=0, they'll release the DHCPv6 addresses and do SLAAC according to the prefix information in the RA messages, so this could be another enforcement method for some specific scenarios.

- Impact to branch/main sites

Renumbering in main/branch site might cause impact on branch/main site communication. The routes, ingress filtering of 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 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 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 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.

- 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.

For IPsec, 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 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 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 non-trivial deployment of SEND, RA Guard [RFC6105] is a lightweight alternative, which focuses on preventing rogue router advertisements in a network. However, it was also not widely deployed at the time when this memo was published.

For DHCPv6, there are built-in secure mechanisms (like Secure DHCPv6 [I-D.ietf-dhc-secure-dhcpv6]), and authentication of DHCPv6 messages [RFC3315] could be utilized. But 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 black list can escape that list by renumbering itself. However, the new prefix might be back on a black list rather soon, if the root cause for being added to such a list is not corrected. In practice, the cost of renumbering will be typically much larger than the cost of getting off the black list.

Dynamic DNS update might bring risk of 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 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 might cause potential risk to the enterprise routing system, since the old address relevant route path might already invalid and the routing system just doesn't know it. However, normally enterprise scenarios don't involve the extreme situation.

6. IANA Considerations

This draft does not request any IANA action.

7. Acknowledgements

This work is inspired by RFC5887, 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, Lee Howard, Ronald Bonica, other 6renum members, and several reviewers for valuable comments.

8. References

8.1. Normative References

- [RFC2608] Guttman, E., Perkins, C., Veizades, J., and M. Day
"Service Location Protocol, Version 2", RFC 2608, June
1999.
- [RFC3007] B. Wellington, "Secure Domain Name System (DNS) Dynamic
Update", RFC 3007, November 2000.
- [RFC3315] Droms, R., 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] R. Droms, "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] J. Bound, 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] B. Volz, "The Dynamic Host Configuration Protocol for IPv6
(DHCPv6) Client Fully Qualified Domain Name (FQDN) Option",
RFC 4706, 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., Ed., Park, S., Beloeil, L., and S. Madanapalli
"IPv6 Router Advertisement Option for DNS Configuration",
RFC 6106, November 2011.

8.2. Informative References

- [RFC2874] Crawford, M., and C. Huitema, "DNS Extensions to Support
IPv6 Address Aggregation and Renumbering", RFC 2874, July
2000.
- [RFC3364] R. Austein, "Tradeoffs in Domain Name System (DNS) Support
for Internet Protocol version 6 (IPv6)", RFC 3364, August
2002.
- [RFC4116] J. Abley, K. Lindqvist, E. Davies, B. Black, 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 Bagnulo, M., "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 Carpenter, B., "Moving A6 to
Historic Status", RFC 6563, May 2012.
- [RFC6603] J. Korhonen, T. Savolainen, S. Krishnan, O. Troan, "Prefix
Exclude Option for DHCPv6-based Prefix Delegation", RFC
6603, May 2012.
- [I-D.ietf-dhc-secure-dhcpv6]
Jiang, S., and S. Shen, "Secure DHCPv6 Using CGAs",
working in progress, March 2012.

- [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), August, 2012.
- [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, April 2012.
- [I-D.ietf-dhc-addr-registration]
Jiang, S., Chen, G., "A Generic IPv6 Addresses Registration Solution Using DHCPv6", working in progress, May 2012.
- [I-D.ietf-6renum-gap-analysis]
Liu, B., and Jiang, S., "IPv6 Site Renumbering Gap Analysis", working in progress, August 2012.
- [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.cheshire-dnsext-dns-sd]
Cheshire, S. and M. Krochmal, "DNS-Based Service Discovery", draft-cheshire-dnsext-dns-sd-11 (work in progress), December 2011.
- [I-D.cheshire-dnsext-multicastdns]
Cheshire, S. and M. Krochmal, "Multicast DNS", draft-cheshire-dnsext-multicastdns-15 (work in progress), December 2011.
- [BA2011] Bhatti, S. and R. Atkinson, "Reducing DNS Caching", Proc. 14th IEEE Global Internet Symposium (GI2011), Shanghai, China. 15 April 2011.
- [JSBM2002] J. Jung, E. Sit, H. Balakrishnan, & R. Morris, "DNS Performance and the Effectiveness of Caching", IEEE/ACM Transactions on Networking, 10(5):589-603, 2002.

Author's 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

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

B. Liu
S. Jiang
Huawei Technologies Co., Ltd
B. Carpenter
University of Auckland
S. Venaas
Cisco Systems
W. George
Time Warner Cable
June 9, 2013

IPv6 Site Renumbering Gap Analysis
draft-ietf-6renum-gap-analysis-08.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 December 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. Its main content is a gap analysis that provides a basis for future works 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.

Table of Contents

1. Introduction	4
2. Overall Requirements for Renumbering	4
3. Existing Components for IPv6 Renumbering	5
3.1. Relevant Protocols and Mechanisms	5
3.2. Management Tools	6
3.3. Procedures/Policies	7
4. Managing Prefixes	7
4.1. Prefix Delegation	7
4.2. Prefix Assignment	8
5. Address Configuration	8
5.1. Host Address Configuration	8
5.2. Router Address Configuration	9
6. Updating Address-relevant Entries	10
6.1. DNS Records Update	10
6.2. In-host Server Address Update	11
6.3. Address update in scattered configurations	11
7. Renumbering Event Management	13
7.1. Renumbering Notification	13
7.2. Synchronization Management	14
7.3. Renumbering Monitoring	14
8. Miscellaneous	14
8.1. Multicast	14
8.2. Mobility	16
9. Gap Summary	17
9.1. Managing Prefixes	17
9.2. Address configuration	17
9.3. Address relevant entries update	17
9.4. Renumbering event management	18
9.5. Miscellaneous	19
10. Gaps considered unsolvable	19

10.1. Address Configuration	19
10.2. Address-relevant Entries Update	19
10.3. Miscellaneous	20
11. Security Considerations	20
12. IANA Considerations.....	21
13. Acknowledgments	21
14. References	22
14.1. Normative References	22
14.2. Informative References	23

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 [RFC4984]. It is thus desirable to develop tools and practices that may make renumbering a simpler process to reduce demand for IPv6 PI space.

Building upon the IPv6 enterprise renumbering scenarios described in [RFC6879], 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 according to 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. Renumbering event management is presented independently from the steps of a renumbering process, in order to identify some operational and administrative gaps in renumbering.

This document starts from existing work in [RFC5887] and [RFC4192]. It does further analysis and identifies the valuable and solvable issues, digs out of some undiscovered gaps, and gives some solution suggestions. This document considers make-before-break approach as a premise for the gap analysis, so readers should be familiar with [RFC4192].

Renumbering nodes with static addresses has a particular set of problems, thus discussion of that space has been covered in a related document [RFC6866].

This document does not cover the un-planned emergency renumbering cases.

2. Overall Requirements for Renumbering

This section introduces the overall 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.

The automation can be divided into four aspects as follows. (Detailed analysis of the four aspects is presented respectively in section 4 to section 7.)

- o Prefix delegation and delivery should be automatic and accurate in aggregation and coordination.
- o Address reconfiguration should be automatically achieved through standard protocols with minimum human intervention.
- o Address-relevant entry updates should be performed together and without error.
- o Renumbering event management is needed to provide the functions of renumbering notification, synchronization, and monitoring.

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 be solved within renumbering context only. However, with the [RFC4192] make-before-break approach, and the address lifetime mechanisms in IPv6 Stateless Address Autoconfiguration (SLAAC) and Dynamic Host Configuration Protocol for IPv6 (DHCPv6), a smooth transition mechanism from old to new prefixes is applicable. In most of the cases, since we can set the transition period long enough to cover the on-going sessions, we consider this mechanism sufficient for avoiding session brokenness issue in IPv6 site renumbering. (Please note that if multiple addresses are running simultaneously on hosts, the address selection [RFC6724] needs to be carefully handled.)

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

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

- o 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.
- o Hosts that are configured through DHCPv6 [RFC3315] obtain new addresses through the renewal process or when they receive the reconfiguration messages initiated by the DHCPv6 servers.
- o DHCPv6-PD (Prefix Delegation) [RFC3633] enables automated delegation of IPv6 prefixes using the DHCPv6.
- o [RFC2894] defined standard ICMPv6 messages for router renumbering. This is a dedicated protocol for renumbering, but we are not aware of it being used in real network deployment.

3.2. Management Tools

Some renumbering operations 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.

- o IP address management (IPAM) tools. There are both commercial and open-source solutions. IPAM tools are used to manage IP address plans, and usually integrate 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.
- o Some organizations use third-party tools to push configuration to devices. This is sometimes used as a supplement to vendor specific solutions. A representative of such third-party tool is [cfengine].
- o [LEROY] proposed a mechanism of macros to automatically update the address-relevant entries/configurations inside the DNS, firewall, etc. The macros can be delivered through SOAP protocol from a network management server to the managed devices.

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

3.3. Procedures/Policies

- o [RFC4192] proposed a procedure for renumbering an IPv6 network without a flag day. The document includes a set of operational suggestions which can be followed step by step by network administrators. It should be noted that the administrators need to carefully deal with the address selection issue while the old and new prefixes are both available during the overlapping period in [RFC4192] procedure. And the address selection policies might need to be updated after renumbering. So administrator could leverage the address selection policy distribution mechanism in [I-D.ietf-6man-addr-select-opt].
- o [RFC6879] 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 IPv6 enterprise site, the key procedural issue is switching the old prefix (es) to the new one(s). A new short prefix may be divided into longer ones for subnets. So we need to carefully manage the prefixes to ensure they are synchronized and coordinated in the whole network.

4.1. Prefix Delegation

For big enterprises, the new short prefix(es) usually comes down through off-line human communication. But for the SOHO style SMEs (Small & Medium Enterprises), the prefixes might be dynamically 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.

The delegation routers might need to renumber themselves with the new delegated prefixes. So there should be a mechanism informing the router to renumber themselves by delegated prefixes; and there also

should be a mechanism for the routers to derive addresses automatically based on the delegated prefixes.

4.2. Prefix Assignment

When subnet routers receive the longer prefixes, they can advertise a prefix on a link to which hosts are connected. Host address configuration, rather than routers, is the primary concern for prefix assignment which is described in the following section 5.1.

5. Address Configuration

5.1. Host Address Configuration

- o SLAAC/DHCPv6 interaction problems

Both of the DHCPv6 and Neighbor Discovery (ND) protocols have IP address configuration function. They are suitable for different scenarios. 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 [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 additional complexity for both the hosts and the network management.

With the flags defined in RA (ManagedFlag indicating the DHCPv6 service available in the network; OtherConfigFlag indicating other configurations such as DNS/routing), the two separated address configuration modes are 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 flags. Different operating systems have followed different approaches. (For more details, please refer to [I-D.liu-bonica-dhcpv6-slaac-problem] and [I-D.liu-6renum-dhcpv6-slaac-switching].)

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

- DHCPv6-configured hosts might not be able to be renumbered by RA

It is unclear whether a DHCPv6 configured host will accept address configuration through RA messages, especially when M flag

transitioning from 1 to 0; this depends on the implementation of the operating system. It might not be possible for administrators to only use RA messages for renumbering, since renumbering might fail on some already DHCPv6-configured hosts. It means administrators have to use DHCPv6 reconfiguration for some DHCPv6-configured hosts. It is not convenient and DHCPv6 reconfiguration is not suitable for bulk usage as analyzed in below.

- DHCPv6-configured hosts might not be able to learn new RA prefixes

[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. Relevant standards ([RFC4862],[RFC3315]) are ambiguous about what approach should be taken by a DHCPv6-configured host when it receives RA messages containing a new prefix. Current behavior depends on the operating system of the host and cannot be predicted or controlled by the network.

- SLAAC-configured hosts might not be able to add DHCPv6 address(es)

The behavior when the host receives RA messages with M flag set is unspecified.

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.

5.2. Router Address Configuration

o Learning new prefixes

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

o Restart after renumbering

As [RFC2072] mentioned, some routers cache IP addresses in some situations, so routers might need to be restarted as a result of site renumbering. While most modern systems support a cache-clear function that eliminates the need for restarts, there are always exceptions that must be taken into account.

- o Router naming

[RFC4192] suggests 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 currently it is present in most IPSec VPN implementations. However, many administrators may need to be alerted to the fact that it could be utilized to avoid manual modification during renumbering.

6. Updating Address-relevant Entries

In conjunction with renumbering the nodes, any configuration or data store containing previous addresses must be updated as well. Some examples include DNS records and filters in various entities such as ACLs in firewalls/gateways.

6.1. DNS Records Update

- o Secure Dynamic DNS update

In real network operations, DNS update is normally achieved 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 implementations, but it hasn't been widely deployed. Normal hosts are not suitable to do the update mainly because of the complexity of key management issues inherited from secure DNS mechanisms, so current practices usually assign the DHCP servers to act as DNS clients to request the DNS server updating relevant records [RFC4704]. This server-oriented approach is applicable for large numbers of hosts' using secure DDNS. (In some commercial solutions, DNS service could be integrated with DHCP service provided by the same vendor so that the secure DDNS might be silently enabled as default.) However, there is still a gap here, since the DHCP servers have to learn the relevant hosts have changed their addresses and thus trigger the DDNS

update. If the hosts were numbered and also renumbered by DHCP, then it is easy for the DHCP servers to learn the address changes; but if the hosts were numbered by SLAAC, then there could be trouble. [I-D.ietf-dhc-addr-registration] proposed a address registration mechanism which could be used to address the latter issue; however, it has not been deployed yet.

6.2. In-host Server Address Update

While DNS stores addresses of hosts in servers, hosts are also configured with addresses of servers such as DNS server, radius server. While renumbering, the hosts must update these addresses if the server addresses changed.

In principle, the addresses of DHCPv6 servers do not need to be updated, since they could be dynamically discovered through DHCPv6 relevant multicast messages. But in practice, most relay agents have the alternative of being configured with DHCPv6 server address rather than sending to a multicast address. So the DHCP server addresses update might be an issue in practice.

6.3. Address update in scattered configurations

Besides the DNS records and the in-host server address entries, there are also many places in which IP addresses are configured, for example, filters such as ACL and routing policies. 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.

In renumbering, it is annoying and error-prone to update the IP addresses in all the above mentioned places. We lack a "one-stop" mechanism to trigger the updates for all the subsystems on a host/server, and all the external databases that refer to the IP address update. We decompose the general "one-stop" gap into the following two aspects.

o 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 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.

o Unified Configuration Management among Devices

This refers to 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, it is hard to get all the relevant configurations through one place.

- Configuration Update Automation

As mentioned in section 3.2, [LEROY] proposed a mechanism which can automatically update the configurations. The mechanism utilizes macros suitable for various devices such as routers, firewalls. to update the configurations 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. Although there are some vendor-independent tools as mentioned in section 3.2, a standard and comprehensive way of uniformly updating configurations in multi-vendor devices is still 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.

- o A notification mechanism may be needed to indicate to hosts that a renumbering event has changed some DNS records in DNS servers (normally, in an enterprise it/they is/are (a) local recursive DNS server(s).), and then the hosts might want to refresh the DNS cache. That mechanism may also need to indicate that such a change will happen at a specific time in the future.
- o As suggested in [RFC4192], if the DNS service can be given prior notice about a renumbering event, then people could reduce the delay in the transition to new IPv6 addresses, thus improve the efficiency of renumbering.
- o Router awareness: in a site with multiple domains which are connected by border routers, all border routers should be aware of renumbering in one domain or multiple domains, and update the internal forwarding tables and the address/prefix based filters accordingly to correctly handle inbound packets.
- o Ingress filtering: ISPs normally enable ingress filter to drop packets with source addresses from other ISPs at the end site routers to prevent IP spoofing [RFC2827]. In a multihomed site with multiple PA prefixes, the ingress router of ISP A should be notified if the ISP B initiates a renumbering event in order to properly update its filters to permit the new legitimate prefix(es). For large enterprises, it might be applicable to indicate this new legitimate prefix information through human communication, however, for the millions of small enterprises, an automated notification mechanism is needed.

- o In the NMS (network management system), logs collected through syslog, SNMP notification, IPFIX, etc. usually treat the UDP message source IP addresses as the host or router IDs. When one source IP address is changed, the log collectors will consider that a new device appeared in the network. So a mechanism is needed for the NMS applications to learn the renumbering event, so that they could correlate the old and new addresses in the logs.

7.2. Synchronization Management

- o DNS update synchronization

The DNS changes must be coordinated with the changes of node address configuration. DNS has a latency issue of propagating information from the server to the resolver. The latency is mainly caused by TTL assigned to individual DNS records and the timing of updates from primary to secondary servers [RFC4192].

Ideally, during a renumbering operation, the DNS TTLs should always be shorter than any other lifetime associated with an address. If the TTLs were set correctly, then the DNS latency could be well controlled. However, there might be some exceptional situations in which the DNS TTLs were already set too long for the time available to plan and execute a renumbering event. In these situations, there currently are no mechanisms to deal with the already configured long DNS TTLs.

7.3. Renumbering Monitoring

While treating renumbering as a network event, mechanisms to monitor the renumbering process might be needed to inform the administrators whether the renumbering has been successfully done. Considering the address configuration operation might be stateless (if ND is used for renumbering), it is difficult for monitoring.

8. Miscellaneous

Since multicast and mobility are special use cases which might not be included in routine/common renumbering operations, they are separately discussed as miscellaneous in this section.

8.1. Multicast

In the perspective of interface renumbering operations, renumbering a multicast address is the same with renumbering a unicast address. So this section mainly discusses the issues from the perspective of the impact to the multicast application systems caused by renumbering.

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.

- o 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.

- o 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 is 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.

In summary, currently the multicast renumbering issues are basically handled by application-specific methods. There is no standard way to guarantee multicast service could live across a renumbering event.

8.2. Mobility

As described in [RFC5887], if a mobile node's home address changes unexpectedly, the node can be informed of the new global routing prefix used at the home site through the Mobile Prefix Solicitation and Mobile Prefix Advertisement ICMPv6 messages [RFC6275]. But if the mobile node is unfortunately disconnected at the time of home address renumbering, it will no longer know a valid subnet anycast address for its home agent, leaving it to deduce a valid address on the basis of DNS information.

So, for Mobile IP, we need a better mechanism to handle change of home agent address while mobile is disconnected.

9. Gap Summary

9.1. Managing Prefixes

- o A mechanism informing the router to renumber themselves by delegated prefixes
- o A mechanism for the routers to derive addresses automatically based on the delegated prefixes.

9.2. Address configuration

- o Host address configuration
 - DHCPv6-configured hosts might not able to be renumbered by RA on some of current implementations
 - DHCPv6-configured hosts might not able to learn new RA prefixes on some of current implementations
 - SLAAC-configured hosts might not able to add DHCPv6 address(es) on some of current implementations
- o Router address configuration
 - A mechanism for interior routers in multihomed site to learn which upstream providers and prefixes were currently reachable
 - Cache-clear might need restart (rarely in modern routers)
 - Using router domain names is not widely learned/deployed by administrators

9.3. Address relevant entries update

- o DNS records update
 - For key management scalable issue, secure dynamic DNS update is usually done by DHCP servers on behalf of the hosts, so it might not be applicable for SLAAC-configured hosts to do secure DDNS.
- o In-host server address update

- DHCP relays might be configured with DHCP server addresses rather than sending multicast messages to discover the DHCP server dynamically, so the DHCP server addresses update might be an issue in practice.
- o Address update in scattered configurations
 - Devices don't support parameterized configuration, administrators need to touch every places where IP addresses were configured
 - It is hard to get all the address-relevant configurations spread in various devices through one place
 - Uniformly update configurations in multi-vendor devices is a big gap currently

9.4. Renumbering event management

- o Renumbering notification
 - A mechanism to indicate hosts local recursive DNS is going to be renumbered
 - A prior notice about a renumbering event for DNS
 - A mechanism for border routers to know internal partial renumbering
 - For multihomed sites, a mechanism to notify the egress router of ISPA that egress router connecting to ISPB initiates renumbering is needed.
 - A mechanism is needed for the NMS applications to learn the renumbering event, so that they could correlate the old and new addresses in the logs.
- o Synchronization management
 - DNS information propagating latency issue
- o Renumbering monitoring
 - Mechanisms to monitor the process and feedback of renumbering might be needed.

9.5. Miscellaneous

- o Multicast
 - Mechanism for SSM receivers to learn the source addresses when multicast sources are renumbered.
- o Mobility
 - A better mechanism to handle change of home agent address while mobile is disconnected.

10. Gaps considered unsolvable

This section lists gaps have been identified by other documents but are considered unsolvable.

10.1. Address Configuration

- o RA prefix lifetime limitation

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

10.2. Address-relevant Entries Update

- o DNS authority

In an enterprise that hosts servers on behalf of collaborators and customers, it is often the case that DNS zones outside the administrative control of the hosting enterprise maintain resource records concerning addresses for hosted nodes within its address space. When the hosting enterprise renumbers, it does not have sufficient authority to change those records.

This is an operational and policy issue. It is out of scope for this document to consider a technical solution or to propose an additional protocol or mechanism to standardize the interaction between DNS systems for authority negotiations.

- o DNS entries commonly have matching Reverse DNS entries which will also need to be updated during renumbering. It might not be possible to combine forward and reverse DNS entries update in one procedure.

- o DNS data structure optimization

[RFC2874] proposed an A6 record type for DNS recording of IPv6 address and prefix. Several extensions to DNS query and processing were also proposed. A6 was designed to be a replacement for AAAA record. The changes were designed to facilitate network renumbering and multihoming. 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.

However, the A6 record has not been widely used, and has been shown to have various problems and disadvantages (see section 2 in [RFC6563]). It has been deprecated and moved to historic status by [RFC6563]. The idea of a structured record to separate prefix and suffix is still potentially valuable for renumbering, but avoiding the problems of the A6 record would require a major development effort.

10.3. Miscellaneous

- o For transport layer, [RFC5887] said that TCP connections and UDP flows are rigidly bound to a given pair of IP addresses.
- o For application layer, in general, we can assert that any implementation is at risk from renumbering if it does not check whether an address is valid each time it starts session resumption (e.g. a laptop wakes from sleep state). It is also more or less risky when it opens a new communications session by using cached addresses.

We considered the above two points (ID/Locator overloading in transport layer & address caching in app layer) are fundamental issues that might not be proper to deal with them just in terms of renumbering.

11. Security Considerations

- o 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 be needed to validate prefixes.

o Influence on Security Controls

During renumbering, security controls (e.g. ACLs) protecting legitimate resources should not be interrupted. For example, if some addresses are in a blacklist, they should not escape from the blacklist due to renumbering.

If there are DHCPv6 authentication keys associated with an IP address then the keys need to be changed for continually working when the addresses are renumbered.

Addresses in SEND certificates are going to need to get updated when renumbering. During the overlap between old and new addresses, both certificates must remain valid.

o Security Protection for Renumbering Notification

Section 7.1 mentions possible notification mechanisms to signal a change in the DNS system or in the border routers related to a renumbering event. Since DNS system and border routers are key elements in any network, and they might take action according to the notification, a security authentication for the renumbering notification is needed.

o Security Protection for Configuration Update

Automated configuration update approaches like [LEROY] would increase the risk since a bad actor with the right permission could cause havoc to the networks.

12. IANA Considerations

This draft does not request any IANA actions.

13. Acknowledgments

This work adopts significant amounts of content from [RFC5887] and particularly the "DNS Authority" topic in section 10.2 is from

[draft-chown-v6ops-renumber-thinkabout]. Both of the two documents are important input for this work, that some principles/considerations applied in this work are implicitly inherited from them. 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 Ted Lemon, Lee Howard, Robert Sparks, S. Moonesamy, Fred Baker, Sean Turner, Benoit Claise, Stephen Farrell, Brian Haberman, Joel Jaeggli, Eric Vyncke, Phillips Matthew, Benedikt Stockebrand, Gustav Reinsberger, Teco Boot and other members of 6renum WG.

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

14. References

14.1. Normative References

- [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.

14.2. Informative References

- [RFC2072] H. Berkowitz, "Router Renumbering Guide", RFC2072, January 1997.
- [RFC2827] Ferguson, P. and D. Senie, "Network Ingress Filtering: Defeating Denial of Service Attacks which employ IP Source Address Spoofing", RFC 2267, January 1998.
- [RFC2874] Crawford, M., and C. Huitema, "DNS Extensions to Support IPv6 Address Aggregation and Renumbering", RFC 2874, July 2000.
- [RFC2894] M. Crawford, "Router Renumbering for IPv6", RFC 2894, August 2000.
- [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.
- [RFC4704] Volz, B., "The Dynamic Host Configuration Protocol for IPv6 (DHCPv6) Client Fully Qualified Domain Name (FQDN) Option", RFC 4704, October 2006.
- [RFC4714] Enns, R., "NETCONF Configuration Protocol", RFC 4714, December 2006.
- [RFC4984] Meyer, D., Ed., Zhang, L., Ed., and K. Fall, Ed., "Report from the IAB Workshop on Routing and Addressing", RFC 4984, September 2007.
- [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.
- [RFC6724] Thaler, D., Ed., Draves, R., Matsumoto, A., and T. Chown, "Default Address Selection for Internet Protocol Version 6 (IPv6)", RFC 6724, September 2012.

- [RFC6275] Johnson, D., Perkins, C., and J. Arkko, "Mobility Support in IPv6", RFC 3775, June 2004.
- [RFC6866] Carpenter, B. and S. Jiang, "Problem Statement for Renumbering IPv6 Hosts with Static Addresses in Enterprise Networks", RFC 6866, February 2013.
- [RFC6879] Jiang, S., Liu, B., and B. Carpenter, "IPv6 Enterprise Network Renumbering Scenarios, Considerations, and Methods", RFC 6879, February 2013.
- [I-D.ietf-6man-addr-select-opt]
Matsumoto, A.M., Fujisaki T.F., and T. Chown, "Distributing Address Selection Policy using DHCPv6", Working in Progress, April 2013.
- [I-D.ietf-dhc-secure-dhcpv6]
Jiang, S., and Shen S., "Secure DHCPv6 Using CGAs", working in progress, March 2012.
- [I-D.ietf-dhc-addr-registration]
Jiang, S., Chen G., and S. Krishnan, "A Generic IPv6 Addresses Registration Solution Using DHCPv6", working in progress, February 2013.
- [I-D.liu-6renum-dhcpv6-slaac-switching]
Liu, B., "DHCPv6/SLAAC Address Configuration Switching for Host Renumbering", Working in Progress, July 2012.
- [I-D.liu-bonica-dhcpv6-slaac-problem]
Liu, B., and R. Bonica, "DHCPv6/SLAAC Address Configuration Interaction Problem Statement", Working in Progress, February 2013.
- [cfengine]<http://cfengine.com/what-is-cfengine>
- [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