Internet Engineering Task Force (IETF)
Request for Comments: 7010
Category: Informational
ISSN: 2070-1721

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

September 2013

IPv6 Site Renumbering Gap Analysis

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 associated with IPv6 renumbering. The content is mainly 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.

Status of This Memo

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

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

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at http://www.rfc-editor.org/info/rfc7010.

Table of Contents

1. Introduction

As introduced in [RFC5887], renumbering, especially for medium to large sites and networks, is currently viewed as expensive and painful. This error-prone process is 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 as an attempt to minimize the need for future renumbering. However, widespread use of PI addressing may create very serious BGP4 scaling problems [RFC4984]. It is thus desirable to develop tools and practices that make renumbering a simpler process and reduces 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 includes prefix management, node address (re)configuration, and updates to 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 further analyzes and identifies the valuable and solvable issues, digs out of some undiscovered gaps, and gives some solution suggestions. This document considers the 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 unplanned emergency renumbering cases.

2. Overall Requirements for Renumbering

This section introduces the overall goals of a renumbering event. In general, we need to leverage renumbering automation to avoid human intervention as much as possible at a reasonable cost. Some existing mechanisms already provide useful capabilities.

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

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

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

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

- 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 a renumbering context only. However, the [RFC4192] make-before-break approach, which uses the address lifetime mechanisms in IPv6 Stateless Address Autoconfiguration (SLAAC) and Dynamic Host Configuration Protocol for IPv6 (DHCPv6), allows for a smooth transition mechanism from old to new prefixes. In most cases, since we can set the transition period to be long enough to cover the ongoing sessions, we consider this mechanism sufficient to avoid broken sessions in IPv6 site renumbering. (Please note that if multiple addresses are running on hosts simultaneously, 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 this purpose. There are also some dedicated protocols and mechanisms that have been 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

- Router Advertisement (RA) messages, defined in [RFC4861], are used to deprecate prefixes that are old or announce prefixes that are new, and to advertise the availability of an upstream router. In renumbering, RA is one of the basic mechanisms for host configuration.

- When renumbering a host, SLAAC [RFC4862] may be used for address configuration with the new prefix(es). Hosts receive RA messages that contain a routable prefix(es) and the address(es) of the default router(s); then hosts can generate an IPv6 address(es) by themselves.

- 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.

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

- [RFC2894] defines standard ICMPv6 messages for router renumbering. This is a dedicated protocol for renumbering, but we are not aware of 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 such tools:

- 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.

- Third-party tools. 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 a third-party tool is [CFENGINE].

- Macros. [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 the SOAP protocol from a network management server to the managed devices.

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

3.3. Procedures and Policies

- [RFC4192] proposed a procedure for renumbering an IPv6 network without a flag day. The document includes a set of operational suggestions that 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 as described in the procedures in [RFC4192]. The address selection policies might need to be updated after renumbering, so the administrator could leverage the address-selection-policy distribution mechanism as described in [6MAN-ADDR-OPT].

- [RFC6879] analyzes the enterprise renumbering events and makes recommendations based on the existing renumbering mechanisms. According to the different stages, renumbering considerations are described in three categories: considerations and recommendations during network design, for the preparation of enterprise network renumbering, and during the renumbering operation.

4. Managing Prefixes

When renumbering an IPv6 enterprise site, the key procedural issue is switching the old prefix(es) to a 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 within the whole network.

4.1. Prefix Delegation

For big enterprises, the new short prefix(es) usually comes down through offline human communication. But, for the SOHO-style (Small Office, Home Office) 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 could 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 to inform the routers to renumber themselves by delegated prefixes; there should also 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 Section 5.1.

5. Address Configuration

5.1. Host Address Configuration

- o SLAAC and DHCPv6 Interaction Problems

  Both DHCPv6 and Neighbor Discovery (ND) protocols have an IP address configuration function, which 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 [PREFIX-DHCPv6]). The DHCPv6-configured hosts can reconfigure addresses by initiating RENEW sessions [RFC3315] 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 be available in the same network. This would add additional complexity for both the hosts and network management.

  With the flags defined in RA (ManagedFlag (M) indicating the DHCPv6 service available in the network; OtherConfigFlag (O) indicating other configurations such as DNS/routing), the two separated address configuration modes are correlated. However, the ND protocol does 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 [DHCPv6-SLAAC] and [6RENUM-SLAAC].)

  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 though RA messages, especially when the M flag transitions 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. This 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 below.

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

  [RFC5887] mentions 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 a DHCPv6 address(es)

  The behavior when the host receives RA messages with the 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. Whether the hosts start DHCPv6 configuration is outside the control of the network side.

5.2. Router Address Configuration

- 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."

- Restarting After Renumbering

  As [RFC2072] mentions, 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.

- Router Naming

  [RFC4192] states 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 can 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 Access Control Lists (ACLs) in firewalls/gateways.

6.1. DNS Records Update

- Secure Dynamic DNS (DDNS) Update

  In real network operations, a 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 AAAA record is hard. [RFC5887] discusses an alternative that uses the Secure Dynamic DNS Update [RFC3007], in which a host informs its own DNS server when it receives a new address.

  The 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 complex key-management issues inherited from secure DNS mechanisms, so current practices usually assign DHCP servers to act as DNS clients to request that the DNS server update relevant records [RFC4704]. The key-management problem is tractable in the case of updates for a limited number of servers, so Dynamic DNS

updates could serve as a suitable solution for keeping server DNS records up to date on a typical enterprise network. However, this solution is not easily applicable to hosts in general.

To address the larger use case of arbitrary non-server hosts being renumbered, DHCP servers have to learn that the relevant hosts have changed their addresses and thus trigger the DDNS update. If the hosts were numbered and also renumbered by DHCP, it would be easy for the DHCP servers to learn the address changes; however, if the hosts were numbered by SLAAC, then there could be trouble.

6.2. In-Host Server Address Update

While DNS stores the addresses of hosts in servers, hosts are also configured with the addresses of servers, such as DNS and RADIUS servers [RFC2865]. While renumbering, the hosts must update these addresses if the server addresses change.

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 option of being configured with a DHCPv6 server address rather than sending to a multicast address. Therefore, 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, updating the IP addresses in all the above mentioned places is burdensome and error-prone. 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 break the general "one-stop" gap into the following two aspects.

- Self-Contained Configuration in Individual Devices

  Ideally, 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 individually updated, 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 current devices, some routers support defining multiple loopback interfaces that can be called in other configurations. For example, when defining a tunnel, it can call the defined loopback interface to use its address as the local address of the tunnel. This can be considered as a kind of parameterized self-contained configuration. However, this only applies to 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 other configurations. Parameterized self-contained configuration is still a gap that needs to be filled.

- 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 changes are pushed to the individual devices. This issue contains two aspects, as follows:

  - Configuration Aggregation

    Configuration data based on addresses or prefixes are usually spread out in various devices. As [RFC5887] describes, some address configuration data might be widely dispersed and much harder to find. Some will inevitably be found only after the renumbering event. Because there’s a big gap in configuration aggregation, it is hard to get all the relevant configuration data together in one place.

  - Configuration Update Automation

    As mentioned in Section 3.2, [LEROY] proposes a mechanism that can automatically update the configurations. The mechanism utilizes macros suitable for various devices such as routers and firewalls to update configurations based on the new prefix. Such an automation tool is valuable for renumbering because it can reduce manual operation, which is error-prone and inefficient.

    Besides the macros, [LEROY] also proposes the use of SOAP to deliver the macros to the devices. Along with SOAP, we may consider whether it is possible and suitable to use other standardized protocols, such as the Network Configuration Protocol (NETCONF) [RFC6241].

In current real networks, most 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 to uniformly update configurations in multi-vendor devices is still missing.

7. Renumbering Event Management

From the perspective of network management, renumbering is an event that may need additional processes to make it easier and more manageable.

7.1. Renumbering Notification

The process of renumbering could benefit from hosts or servers being made aware of an occurrence of a renumbering event. Following are several examples of additional processes that may ease renumbering.

- 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 is 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.

- As suggested in [RFC4192], if the DNS service can be given prior notice about a renumbering event, then delay in the transition to new IPv6 addresses could be reduced and thus improve the efficiency of renumbering.

- Router awareness: In a site with multiple domains that 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.

- Ingress filtering: ISPs normally enable an 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 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 practical to manage this new legitimate prefix information through human communication. However, for the millions of small enterprises, an automated notification mechanism is needed.

- Log collectors: In the NMS (network management system), log collectors that collect logs 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. Therefore, a mechanism is needed for the NMS applications to learn the renumbering event including the mappings of old and new IP addresses for each host/router, so that they could update the address-relevant mappings as described in Section 7.2.

7.2. Synchronization Management

- DNS Update Synchronization

  The DNS changes must be coordinated with node address configuration changes. DNS has a latency issue of propagating information from the server to the resolver. The latency is mainly caused by the Time to Live (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 lifetimes 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 are currently no mechanisms to deal with the already configured long DNS TTLs.

- NMS Address-Relevant Mapping Synchronization

  As described in Section 7.1, the NMS needs to learn the renumbering event and thus correlate the old and new address in the logs. If the NMS applies unique IDs for the hosts or routers, then the mappings between the unique IDs and IP addresses also need to be updated after renumbering.

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 successful. Considering that the address configuration operation might be stateless (if ND is used for renumbering), it is difficult to monitor.

8. Miscellaneous

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

8.1. Multicast

From the perspective of interface renumbering operations, renumbering a multicast address is the same as 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.

- 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 not to 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, for example, need to get new SDP (Session Description Protocol) information with the new source address. This is similar to the process for learning a new group address; see the "Renumbering of Multicast Addresses" topic below.

- Renumbering of Rendezvous-Point

  If the unicast addresses of routers in a network are renumbered, then the RP (Rendezvous-Point) address is also likely to change for at least some groups. An RP address is needed by PIM-SM (Protocol Independent Multicast - Sparse Mode) to provide 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 to

distribute group-to-RP mappings, the change can be fairly quick and any impact time should be very brief. However, if routers are statically configured, the time impacted 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 the change of IP address. A network operator should be careful not to 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 described in "Unicast-Prefix-based IPv6 Multicast Addresses" [RFC3306] and "Embedded-RP IPv6 Multicast Addresses" [RFC3956]. In those cases, 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 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, multicast renumbering issues are basically handled by application-specific methods. There is no standard way to guarantee that 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]. However, if the mobile node is 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 the change of home agent address while the mobile address is disconnected.

9. Gap Summary

The following is a summary of the gaps identified previously in this document that are considered solvable, but may require process or protocol changes to resolve.

9.1. Managing Prefixes

- A mechanism informing the routers to renumber themselves by delegated prefixes.

- A mechanism for the routers to derive addresses automatically based on the delegated prefixes.

9.2. Address Configuration

- Host Address Configuration

  - DHCPv6-configured hosts might not be able to be renumbered by RA on some current implementations.

  - DHCPv6-configured hosts might not be able to learn new RA prefixes on some current implementations.

  - SLAAC-configured hosts might not be able to add DHCPv6 address(es) on some current implementations.

- o Router Address Configuration
  - A mechanism for interior routers in a multihomed site to learn which upstream providers and prefixes are currently reachable.
  - Cache-clear might need to restart (rarely in modern routers).
  - Use of router domain names is not widely learned or deployed by administrators.

9.3. Address-Relevant Entries Update

- o DNS Records Update
  - For key-management scalability issues, secure dynamic DNS update is usually done by DHCP servers on behalf of the hosts, so it might not be practical 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 by sending multicast messages to discover the DHCP server dynamically, so updating the DHCP server addresses might be an issue in practice.
- o Address Update in Scattered Configurations
  - For devices that don’t support parameterized configuration, administrators need to individually update all devices in which IP addresses were previously configured.
  - It is hard to get all the address-relevant configurations spread in various devices through one place.
  - Uniformly updating configurations in multi-vendor devices is currently a big gap that needs to be filled.

9.4. Renumbering Event Management

- o Renumbering Notification
  - A mechanism to indicate a host’s 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 is needed to notify the egress router connecting to ISP A that the egress router connecting to ISP B has initiated renumbering.
  - 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, and update the unique ID of the device and address mappings.
- o Synchronization Management
  - DNS information propagation latency issue.
  - Mechanisms are needed for the NMS applications to correlate the old and new addresses in logs and to update the unique ID of the device and address mappings.
- o Renumbering Monitoring
  - Mechanisms to monitor the process and feedback of renumbering might be needed.

9.5. Miscellaneous

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

10. Gaps Considered Unsolvable

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

10.1. Address Configuration

- RA Prefix Lifetime Limitation

  Section 5.5.3 of [RFC4862] states "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 to less than two hours. This limitation is to prevent denial-of-service attacks.

10.2. Address-Relevant Entries Update

- 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.

- DNS Entries

  DNS entries commonly have matching reverse DNS entries that will also need to be updated during renumbering. It might not be possible to combine forward and reverse DNS entry updates in one procedure where addresses are not being managed using DHCP.

- DNS Data Structure Optimization

  [RFC2874] proposed an A6 record type for DNS recording of the IPv6 address and prefix. Several extensions to DNS query and processing were also proposed. A6 was designed to be a replacement for the 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.

[RFC2874] has been deprecated and moved to Historic status by [RFC6563]. The A6 record has not been widely used and has been shown to have various problems and disadvantages (see Section 2 in [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

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

- For the 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 and address caching in application layer) fundamental issues that might not be proper to deal with just in terms of renumbering.

11. Security Considerations

- Prefix Validation

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

  In the LAN, Secure DHCPv6 [SECURE-DHCPv6] or Secure Neighbor Discovery (SEND) [RFC3971] deployment may be needed to validate prefixes.

- 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.

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

- 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 the 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.

- 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. Acknowledgments

This work adopts significant amounts of content from [RFC5887]. In addition, it draws largely from the "DNS Authority" topic in Section 10.2 from [IPv6-RENUM-THINK]. Both documents offer such important input for this work that some principles and 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 the 6renum WG.

13. References

13.1. Normative References

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

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

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

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

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

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

13.2. Informative References

[RFC2072] Berkowitz, H., "Router Renumbering Guide", RFC 2072, January 1997.

[RFC2827] Ferguson, P. and D. Senie, "Network Ingress Filtering: Defeating Denial of Service Attacks which employ IP Source Address Spoofing", BCP 38, RFC 2827, May 2000.

[RFC2865] Rigney, C., Willens, S., Rubens, A., and W. Simpson, "Remote Authentication Dial In User Service (RADIUS)", RFC 2865, June 2000.

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

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

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

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

[RFC6241] Enns, R., Ed., Bjorklund, M., Ed., Schoenwaelder, J., Ed., and A. Bierman, Ed., "Network Configuration Protocol (NETCONF)", RFC 6241, June 2011.

[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.

[RFC6275] Perkins, C., Ed., Johnson, D., and J. Arkko, "Mobility Support in IPv6", RFC 6275, July 2011.

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

[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.

[6MAN-ADDR-OPT] Matsumoto, A., Fujisaki T., and T. Chown, "Distributing Address Selection Policy using DHCPv6", Work in Progress, August 2013.

[6RENUM-SLAAC]
Liu, B., "DHCPv6/SLAAC Address Configuration Switching for Host Renumbering", Work in Progress, January 2013.

[CFENGINE] CFEngine, <http://cfengine.com/what-is-cfengine>.

[DHCPv6-SLAAC]
Liu, B. and R. Bonica, "DHCPv6/SLAAC Address Configuration Interaction Problem Statement", Work in Progress, February 2013.

[IPv6-RENUM-THINK]
Chown, T., Thompson, M., Ford, A., and S. Venaas, "Things to think about when Renumbering an IPv6 network", Work in Progress, September 2006.

[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>

[PREFIX-DHCPv6]
Jiang, S., Xia, F., and B. Sarikaya, "Prefix Assignment in DHCPv6", Work in Progress, February 2013.

[SECURE-DHCPv6]
Jiang, S. and Shen S., "Secure DHCPv6 Using CGAs", Work in Progress, September 2012.

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
United States
EMail: stig@cisco.com

Wesley George
Time Warner Cable
13820 Sunrise Valley Drive
Herndon, VA 20171
United States
Phone: +1 703-561-2540
EMail: wesley.george@twcable.com