

6RENUM
Internet-Draft
Intended status: Informational
Expires: April 20, 2012

B. Carpenter
Univ. of Auckland
S. Jiang
Huawei Technologies Co., Ltd
October 18, 2011

Problem Statement for Renumbering IPv6 Hosts with Static Addresses
draft-carpenter-6renum-static-problem-00

Abstract

This document analyses the problems of updating the IPv6 addresses of hosts in enterprise networks that for operational reasons require static addresses.

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 April 20, 2012.

Copyright Notice

Copyright (c) 2011 IETF Trust and the persons identified as the document authors. All rights reserved.

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

Table of Contents

1. Introduction	3
2. Analysis	3
2.1. Static Addresses Imply Static Prefixes	3
2.2. Other Hosts Need Literal Address	4
2.3. Static Server Addresses	4
2.4. Static Virtual Machine Addresses	5
2.5. Asset Management and Security Tracing	5
2.6. Primitive Software Licensing	6
2.7. Network Elements	6
3. Summary of Problem Statement	6
4. Security Considerations	7
5. IANA Considerations	7
6. Acknowledgements	7
7. Change log [RFC Editor: Please remove]	8
8. Informative References	8
Authors' Addresses	9

1. Introduction

A problem that is frequently mentioned in discussions of renumbering enterprise networks [RFC5887] [I-D.jiang-6renum-enterprise] [I-D.liu-6renum-gap-analysis] is that of statically assigned addresses. Static addressing often implies manual address assignment, including manual preparation of configuration scripts. An implication of hosts having static addresses is that subnets must have static prefixes, which also requires analysis.

Although static addressing is in general problematic for renumbering, hosts inside an enterprise may have static addresses for a number of operational reasons:

- o For some reason, other hosts need to be configured with a literal numeric address for the host in question, so its address must be static.
- o Even if a site has local DNS support and this is normally used to locate servers, some operators wish their servers to have static addresses so that issues of address lifetime and DNS TTL cannot affect connectivity.
- o Some approaches to virtual server farms require static addressing.
- o On some sites, the network operations staff require hosts to have static addresses for asset management purposes and for address-based backtracking of security incidents.
- o Certain software licensing mechanisms have existed which are based on IP addresses. [Question: Is this still relevant for IPv6 addresses?]
- o Network elements such as routers are usually assigned static addresses, which are also configured into network monitoring and management systems.

This document analyses these issues in more detail and presents a problem statement.

2. Analysis

2.1. Static Addresses Imply Static Prefixes

Host addresses can only be static if subnet prefixes are also static. Static prefixes are such a long-established practice in enterprise networks that it is hard to discern the reason for them. Originally, before DHCP became available, there was simply no alternative. Thus it became accepted practice to assign subnet prefixes manually and build them into static router configurations. Today, the static nature of subnet prefixes has become a diagnostic tool in itself, at least in the case of IPv4 where prefixes can easily be memorised. If

several users sharing a subnet prefix report problems, the fault can readily be localised.

This model is being challenged for the case of unmanaged home IPv6 networks, in which it is possible to assign subnet prefixes automatically, at least in a cold start scenario [I-D.baker-homenet-prefix-assignment]. For an enterprise network, the question arises whether automatic subnet prefix assignment can be made using the "without a flag day" approach to renumbering. [RFC4192] specifies that "the new prefix is added to the network infrastructure in parallel with (and without interfering with) the old prefix." Any method for automatic prefix assignment needs to support this.

2.2. Other Hosts Need Literal Address

This issue commonly arises in small networks without local DNS support, for devices such as printers that all other hosts need to reach. In this case, not only does the host in question have a static address, but that address is also configured in the other hosts. It is long established practice in small IPv4 networks that printers in particular are manually assigned a fixed address (typically an [RFC1918] address) and that users are told to manually configure printer access using that fixed address. For a small network the work involved in doing this is much less than the work involved in doing it "properly" by setting up DNS service, whether local or hosted by an ISP, to give the printer a name. It is also unusual to enable the Service Location Protocol [RFC2608] for the same purpose. In consequence, if the printer is renumbered for any reason, the manual configuration of all users' hosts must be updated.

In the case of IPv6, exactly the same situation would be created by numbering the printer statically under the site's ULA prefix [RFC4193]. The disadvantage compared to IPv4 is that an IPv6 address is harder to communicate reliably, compared to something as simple as "10.1.1.10". The process will be significantly more error-prone for IPv6.

If such a host is numbered out of a prefix that is potentially subject to renumbering, then a renumbering event will require a configuration change in all hosts using the device in question, and the configuration data are by no means stored in the network layer.

2.3. Static Server Addresses

On larger sites, it is safe to assume that servers of all kinds, including printers, are identified in user configurations and applications by DNS names. However, it is very widespread

operational practice that servers have static IP addresses. If they did not, whenever an address assigned by stateless address auto-configuration [RFC4862] or DHCPv6 [RFC3315] expired, and if the address actually changed for some extraneous reason, sessions in progress might fail (depending on whether the address deprecation period was long enough). Also, since a dynamic DNS update [RFC3007] would be required, remote users would attempt to use the wrong address until the DNS time-to-live expired.

Such server addresses can be managed centrally even if they are static, by using DHCPv6 in stateful mode, and by generating both DHCPv6 data and DNS data from a common configuration database. This does normally carry the implication that the database also contains the hardware (MAC) addresses of the relevant LAN interfaces on the servers, so that the correct IPv6 address can be delivered whenever a server requests an address. Not every operator wishes to maintain such a costly database, however, and some sites are very likely today to fall back on manual configuration of server addresses as a result.

In the event of renumbering of the prefix covering such servers, the situation should be manageable if there is a common configuration database; the "without a flag day" procedure [RFC4192] could be followed. However, if there is no such database, a manual procedure would have to be adopted.

2.4. Static Virtual Machine Addresses

According to [I-D.narten-nvo3-overlay-problem-statement], "Placement and movement of VMs in a network effectively requires that VM IP addresses be fixed and static." Otherwise, when a VM is migrated to a different physical server, its IP address would change and transport sessions in progress would be lost. In effect this is a special case of the previous one.

If VMs are numbered out of a prefix that is subject to renumbering, there is a direct conflict with transport session continuity, unless a procedure similar to [RFC4192] is followed.

2.5. Asset Management and Security Tracing

There are some large (campus-sized) sites that not only capture the MAC addresses of servers in a configuration system, but also do so for desktop client machines with wired connections, that are then given static IP addresses. Such hosts are not normally servers, so the two preceding cases do not apply. One motivation for this approach is straightforward asset management (who has which computer, connected to which cable?). Another, more compelling, reason is security incident handling. If, as occurs with reasonable frequency

on any large network, a particular host is found to be generating some form of unwanted traffic, it is urgent to be able to track back from its IP address to its physical location, so that an appropriate intervention can be made.

Such users will not in most circumstances be significantly inconvenienced by prefix renumbering, as long as it follows the [RFC4192] procedure. The address deprecation mechanism would allow for clean termination of current sessions, including those in which their machine was actually operating as a server, e.g., for a peer-to-peer application. The only users who would be seriously affected would be those running extremely long transport sessions that might outlive the address deprecation period.

Note that such large campus sites generally allocate addresses dynamically to wireless hosts, since (in an IPv4 world) addresses are scarce and allocating static addresses to intermittent users is not acceptable. Also, a wireless user may appear on different subnets at different times, so cannot be given a single static address. These users will in most circumstances only be slightly inconvenienced, if at all, by prefix renumbering.

2.6. Primitive Software Licensing

TBD if relevant.

2.7. Network Elements

Each interface of a router needs an IP address, and so do other network elements such as firewalls, proxies, and load balancers. Since these are critical infrastructure, they must be monitored and in some cases controlled by a network management system. A conventional approach to this is to assign the necessary IP addresses statically, and also to configure those addresses in the monitoring and management systems. It is quite common practice that some such addresses will have no corresponding DNS entry. If these addresses need to be changed, there will be considerable ramifications. A restart of the network element might be needed, interrupting all user sessions in progress. Simultaneously, the monitoring and management system configurations must be updated, and in the case of a default router, its clients must be informed. To avoid such disruption, network elements must be renumbered according to an [RFC4192] procedure, like any other host.

3. Summary of Problem Statement

If subnet prefixes are statically assigned, various network elements

and the network management system must be informed when they are renumbered. Alternatively, can automatic subnet prefix assignment be performed without interruption to user sessions?

If a printer or similar local server is statically addressed out of a non-ULA prefix, and has no DNS name, prefix renumbering will require configuration changes in all hosts using that server. Most likely, these changes will be manual. Even if the server is under a ULA prefix, any subnet rearrangement that causes it to be renumbered will have the same effect.

If a server with a DNS name is statically addressed via a common configuration database that supports both DHCPv6 and DNS, then it can be renumbered "without a flag day" by following RFC 4192. However, if there is no common configuration database, then present technology requires manual intervention. Similar considerations apply to virtual servers with static addresses.

If client computers such as desktops are statically addressed via a common configuration database and stateful DHCPv6, they can also be renumbered "without a flag day." But other statically addressed clients will need manual intervention.

If network elements have static addresses, the network management system and affected client hosts must be informed when they are renumbered. Alternatively, can automatic network element renumbering be performed without interruption to user sessions?

4. Security Considerations

This document defines no protocol, so does not introduce any new security exposures.

5. IANA Considerations

This document requests no action by IANA.

6. Acknowledgements

Valuable comments and contributions were made by ...

This document was produced using the xml2rfc tool [RFC2629].

7. Change log [RFC Editor: Please remove]

draft-carpenter-6renum-static-problem-00: original version,
2011-10-18.

8. Informative References

- [I-D.baker-homenet-prefix-assignment]
Baker, F. and R. Droms, "IPv6 Prefix Assignment in Small Networks", draft-baker-homenet-prefix-assignment-00 (work in progress), October 2011.
- [I-D.jiang-6renum-enterprise]
Jiang, S., Liu, B., and B. Carpenter, "IPv6 Enterprise Network Renumbering Scenarios and Guidelines", draft-jiang-6renum-enterprise-01 (work in progress), September 2011.
- [I-D.liu-6renum-gap-analysis]
Liu, B. and S. Jiang, "IPv6 Site Renumbering Gap Analysis", draft-liu-6renum-gap-analysis-01 (work in progress), July 2011.
- [I-D.narten-nvo3-overlay-problem-statement]
Narten, T. and M. Sridharan, "Problem Statement: Using L3 Overlays for Network Virtualization", draft-narten-nvo3-overlay-problem-statement-00 (work in progress), September 2011.
- [RFC1918] Rekhter, Y., Moskowitz, R., Karrenberg, D., Groot, G., and E. Lear, "Address Allocation for Private Internets", BCP 5, RFC 1918, February 1996.
- [RFC2608] Guttman, E., Perkins, C., Veizades, J., and M. Day, "Service Location Protocol, Version 2", RFC 2608, June 1999.
- [RFC2629] Rose, M., "Writing I-Ds and RFCs using XML", RFC 2629, June 1999.
- [RFC3007] Wellington, B., "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.

- [RFC4192] Baker, F., Lear, E., and R. Droms, "Procedures for Renumbering an IPv6 Network without a Flag Day", RFC 4192, September 2005.
- [RFC4193] Hinden, R. and B. Haberman, "Unique Local IPv6 Unicast Addresses", RFC 4193, October 2005.
- [RFC4862] Thomson, S., Narten, T., and T. Jinmei, "IPv6 Stateless Address Autoconfiguration", RFC 4862, September 2007.
- [RFC5887] Carpenter, B., Atkinson, R., and H. Flinck, "Renumbering Still Needs Work", RFC 5887, May 2010.

Authors' Addresses

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

Email: brian.e.carpenter@gmail.com

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

Email: jiangsheng@huawei.com

Network Working Group
Internet Draft
Intended status: Best Current Practice
Expires: March 18, 2012

S. Jiang
B. Liu
Huawei Technologies Co., Ltd
B. Carpenter
University of Auckland
September 29, 2011

IPv6 Enterprise Network Renumbering Scenarios and Guidelines
draft-jiang-6renum-enterprise-01.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 March 18, 2012.

Copyright Notice

Copyright (c) 2011 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 enterprise renumbering events and describes the best current practice among the existing renumbering mechanisms. According to the different stages of renumbering events, considerations and best current practices are described in three categories: during network design, for preparation of renumbering, and during a renumbering operation. A gap inventory is listed at the end of this document.

Table of Contents

1. Introduction	3
2. Enterprise Network Illustration for Renumbering	3
3. Enterprise Network Renumbering Scenario Categories	4
3.1. Renumbering caused by External Network Factors.....	4
3.2. Renumbering caused by Internal Network Factors.....	5
4. Network Renumbering Considerations and Best Current Practise..	5
4.1. Considerations and Best Current Practice during Network Design	6
4.2. Considerations and Best Current Practice for the Preparation of Renumbering	9
4.3. Considerations and Best Current Practice during Renumbering Operation	10
5. Gap Inventory	12
6. Security Considerations	12
7. IANA Considerations	13
8. Acknowledgements	13
9. Change Log [RFC Editor please remove]	13
10. References	13
10.1. Normative References	13
10.2. Informative References	14
Author's Addresses	16

1. Introduction

IPv6 site renumbering is considered difficult. Network managers currently prefer to Provider Independent (PI) addressing for IPv6 to attempt to minimize the need for future renumbering. However, widespread use of PI may create very serious BGP4 scaling problems. It is thus desirable to develop tools and practices that may make renumbering a simpler process to reduce demand for IPv6 PI space. In any case, renumbering may be necessary for other reasons.

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

This document focuses on IPv6 only, by leaving IPv4 out of scope. Dual-stack network or IPv4/IPv6 transition scenarios are out of scope, too.

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

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

2. Enterprise Network Illustration for Renumbering

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

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

Address reconfiguration is fulfilled either by DHCPv6 or ND protocols. Static address assignment is not considered in this version. During the renumbering event, the DNS records need to be synchronized while routing tables, ACLs and IP filtering tables in various gateways also need to be updated, too.

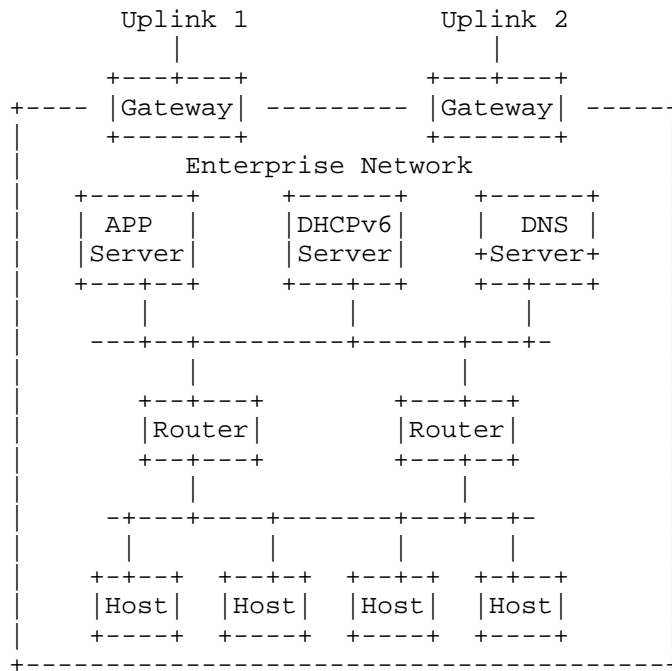


Figure 1 Enterprise network illustration

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

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

3. Enterprise Network Renumbering Scenario Categories

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

3.1. Renumbering caused by External Network Factors

The most influential external network factor is the uplink ISP.

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

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

- o The renumbering event may be initiated by receiving new prefixes from the same uplink. The typical scenario is that the DHCPv6 server in the ISP delegates a new prefix to the enterprise network. This might happen if the enterprise network is switched to a different location within the network topology of the same ISP due to various considerations, such as commercial, performance or services reasons, etc. Alternatively, the ISP itself might be renumbered due to topology changes or migration to a different or additional prefix. These ISP renumbering events would initiate enterprise network renumbering events, of course.
- o The enterprise network adds new uplink(s) for multihoming purposes. This may not a typical renumbering because the original addresses will not be changed. However, initial numbering may be considered as a special renumbering event. If the administrators only want part of the network to have multiple prefixes, the renumbering process should be carefully managed.
- o The enterprise network removes uplink(s) or old prefixes.

3.2. Renumbering caused by Internal Network Factors

- o As companies split, merge, grow, relocate or reorganize, the enterprise network architectures may need to be re-built. This will trigger the internal renumbering.
- o The enterprise network may proactively adopt a new address scheme, for example by switching to a new transition mechanism or stage of a transition plan.
- o The enterprise network may reorganize its topology or subnets.

4. Network Renumbering Considerations and Best Current Practices

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

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

4.1. Considerations and Best Current Practices during Network Design

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

- Prefix Delegation

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

- Usage of FQDN

It is recommended that Fully-Qualified Domain Names (FQDNs) should be used to configure network connectivity, such as tunnels, whenever possible. The capability to use FQDNs as endpoint names has been standardized in several RFCs, such as [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.

Service Location Protocol [RFC2608] and multicast DNS with SRV records for service discovery can reduce the number of places that IP addresses need to be configured.

- Address Types

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

Manual-configured addresses are not scalable in medium to large sites, hence are out of scope. However, some hosts such as servers may need static addresses. Manual-configured addresses/hosts should be avoided as much as possible.

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

- Address configuration models

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

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

In principle, an enterprise network should choose only one address configuration model and employ either ND or DHCPv6. This document has no preference between ND and DHCPv6 address configuration models. It is network architects' job to decide which configuration model is employed. Even in a large network that contains several subnets, it is better not to mix the two address configuration models, though using them independently in different subnets may partly reduce the risk.

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

- DNS

It is recommended that the site have an automatic and systematic procedure for updating/synchronising its DNS records, including both forward and reverse mapping [RFC2874]. A manual on-demand

updating model is considered as a harmful source of problems in a renumbering event.

Although the A6 DNS record model [RFC2874] was designed for easier renumbering, it has a lot of unsolved technical issues [RFC3364, I-D.jiang-dnsexst-a6-to-historic]. Therefore, it has been moved to experimental status [RFC3363]. It is not recommended.

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

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

The DNS synchronization may be completed through the Secure DNS Dynamic Update [RFC3007]. Alternatively, a DHCPv6 server could update host DNS records following the operations defined by [RFC4704]. In a model including SLAAC, host addresses may be registered on an address registration server, which could in fact be a DHCPv6 server; then the server would update corresponding DNS records.

- Security

Any automatic renumbering scheme has a potential exposure to hijacking at the moment that a new address is announced. Proper network security mechanisms should be employed. Secure Neighbor Discovery (SEND, [RFC3971]), which is not widely deployed, is recommended to replace ND if this is considered to be a serious threat. DHCPv6 built-in secure mechanisms, like Secure DHCPv6 [I-D.ietf-dhc-secure-dhcpv6] or authentication of DHCPv6 messages [RFC3315] are recommended.

- 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, these connectivities can survive after renumbering events at other sites. This also applies to host-based connectivities.

4.2. Considerations and Best Current Practices for the Preparation of Renumbering

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

This section describes several recommendations for the preparation of enterprise renumbering event. By adopting these recommendations, a site could be renumbered more easily. However, these recommendations are not cost free. They might increase the daily burden of network operation. Therefore, only those networks that are expected to be renumbered soon or very frequently should adopt these recommendations, with balanced consideration between daily cost and renumbering cost.

- Reduce the address preferred time or valid time or both.

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

- Reduce the DNS record TTL on the local DNS server.

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

- Reduce the DNS configuration lifetime on the hosts.

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

- Identify long-living sessions

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

4.3. Considerations and Best Current Practices during Renumbering Operation

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

- Within/without a flag day

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

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

- Transition period

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

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

- Network initiative enforced renumbering

If the network has to enforce renumbering before address leases expire, the network should initiate enforcement messages, either in Router Advertisement messages or DHCPv6 RECONFIGURE messages.

- Impact to branch/main sites

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

gateways, and DNS may need to be updated. This needs careful planning and organizing.

- DNS record update and DNS configuration on hosts

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

DNS configuration on hosts should be updated if local recursive DNS servers are renumbered. During the transition period, both old and new DNS server addresses may co-exist on the hosts. If the lifetime of DNS configuration is shorter than the transition period, name resolving failure may be reduced to minimum. A notification mechanism may be needed to indicate to the hosts that a renumbering event of local recursive DNS happens or is going to take place.

- Router awareness

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

- Border filtering

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

- Tunnel concentrator renumbering

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

- Connectivity session survivability

During the renumbering operations, connectivity sessions in IP layer would break if the old address is deprecated before the session ends. However, the upper layer sessions may survive by using session survivability technologies, such as SHIM6 [RFC5533].

As mentioned above, some long-living applications may need to be handled specially.

5. Gap Inventory

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

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

6. Security Considerations

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

Any automatic renumbering scheme has a potential exposure to hijacking at the moment that a new address is announced. Proper network security mechanisms should be employed. SEND is recommended

to replace ND. Alternatively, certain lightweight renumbering specific security mechanism may be developed in the future. DHCPv6 build-in secure mechanisms, like Secure DHCPv6 [I-D.ietf-dhc-secure-dhcpv6] or authentication of DHCPv6 messages [RFC3315] are recommended.

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

7. IANA Considerations

This draft does not request any IANA action.

8. Acknowledgements

This work is illuminated by 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 and other 6renum members for valuable comments.

9. Change Log [RFC Editor please remove]

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

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

10. References

10.1. 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.
- [RFC3956] Savola, P., 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
- [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.

10.2. Informative References

- [RFC2874] Crawford, M., and C. Huitema, "DNS Extensions to Support IPv6 Address Aggregation and Renumbering", RFC 2874, July 2000.
- [RFC3363] R. Bush, A. Durand, B. Fink, O. Gudmundsson, T. Hain, "Representing Internet Protocol version 6 (IPv6) Addresses in the Domain Name System (DNS)", RFC 3363, August 2002.
- [RFC3364] R. Austein, "Tradeoffs in Domain Name System (DNS) Support for Internet Protocol version 6 (IPv6)", RFC 3364, August 2002.

- [RFC4057] J. Bound, Ed. "IPv6 Enterprise Network Scenarios", RFC 4057, June 2005.
- [RFC4192] Baker, F., Lear, E., and R. Droms, "Procedures for Renumbering an IPv6 Network without a Flag Day", RFC 4192, September 2005.
- [RFC4864] Van de Velde, G., T. Hain, R. Droms, B. Carpenter, E. Klein, Local Network Protection for IPv6", RFC 4864, May 2007.
- [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.
- [I-D.ietf-dhc-secure-dhcpv6]
Jiang, S., and S. Shen, "Secure DHCPv6 Using CGAs", working in progress.
- [I-D.ietf-dhc-host-gen-id]
S. Jiang, F. Xia, and B. Sarikaya, "Prefix Assignment in DHCPv6", draft-ietf-dhc-host-gen-id (work in progress), April, 2011.
- [I-D.ietf-savi-mix]
Bi, J., Yao, G., Halpern, J., and Levy-Abegnoli, E., "SAVI for Mixed Address Assignment Methods Scenario", working in progress.
- [I-D.ietf-dhc-pd-exclude]
J. Korhonen, T. Savolainen, S. Krishnan, O. Troan, "Prefix Exclude Option for DHCPv6-based Prefix Delegation", working in progress.
- [I-D.liu-6renum-gap-analysis]
Liu, B., and Jiang, S., "IPv6 Site Renumbering Gap Analysis", working in progress.
- [I-D.jiang-dnsext-a6-to-historic]
Jiang, S., Conrad, D. and Carpenter, B., "Moving A6 to Historic Status", working in progress.

Author's Addresses

Sheng Jiang
Huawei Technologies Co., Ltd
Huawei Building, No.3 Xinxu Rd.,
Shang-Di Information Industry Base, Hai-Dian District, Beijing
P.R. China
EMail: jiangsheng@huawei.com

Bing Liu
Huawei Technologies Co., Ltd
Huawei Building, No.3 Xinxu Rd.,
Shang-Di Information Industry Base, Hai-Dian District, Beijing
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: May 1, 2012

B. Liu
S. Jiang
Huawei Technologies Co., Ltd
B . Carpenter
University of Auckland
October 31, 2011

IPv6 Site Renumbering Gap Analysis
draft-liu-6renum-gap-analysis-02.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 May 1, 2012.

Copyright Notice

Copyright (c) 2011 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 could be utilized by IPv6 site renumbering and envisions the effort could be done. This document tries to cover the most explicit issues and

requirements of IPv6 renumbering. Through the gap analysis, the document provides a basis for future work to identify and develop solutions or to stimulate such development as appropriate. The gap analysis is presented following a renumbering event procedure clue.

Table of Contents

1. Introduction	3
2. Overall Requirements for Renumbering	3
3. Existing Components for IPv6 Renumbering	4
3.1. Relevant Protocols and Mechanisms	4
3.2. Management Tools	5
3.3. Procedures/Policies	5
4. Managing Prefixes	6
4.1. Prefix Delegation	6
4.2. Prefix Assignment	6
5. Address Configuration	6
5.1. Host Address Configuration	6
5.2. Router Address Configuration	8
5.3. Static Address Configuration	9
6. Address Relevant Entries Update	9
6.1. DNS Records Update	9
6.2. In-host Server Address Update	10
6.3. Filters	10
7. Renumbering Event Management	11
7.1. Renumbering Notification	11
7.2. Synchronization Management	11
7.3. Renumbering Monitoring	11
8. Miscellaneous	12
9. Gap Summary	12
9.1. Managing Prefixes	12
9.2. Address configuration	12
9.3. Address relevant entries update	13
9.4. Renumbering event management	14
10. Security Considerations	14
11. IANA Considerations.....	14
12. References	15
12.1. Normative References	15
12.2. Informative References	15
13. Acknowledgments	16
Appendix A. Other Documented Gaps	17

1. Introduction

As introduced in [RFC5887], renumbering, especially for medium to large sites and networks, is currently viewed as an expensive, painful, and error-prone process, avoided by network managers as much as possible. If IPv6 site renumbering continues to be considered difficult, network managers will turn to Provider Independent (PI) addressing for IPv6 to attempt to minimize the need for future renumbering. However, widespread use of PI may create very serious BGP4 scaling problems. It is thus desirable to develop tools and practices that may make renumbering a simpler process to reduce demand for IPv6 PI space.

This document performs a gap analysis to provide a basis for future work to identify and develop solutions or to stimulate such development as appropriate. The gap analysis is organized by the main steps of renumbering process, which include prefix management, node address (re)configuration, and updating relevant address entries in various gateways, routers and servers, etc. Besides these steps, the aspect of renumbering event management is presented independently, which intends to help the operational/administrative process. It is expected that these steps and management could cover all the aspects of an renumbering event.

This document draws on existing work in (at least) [RFC5887], [I-D.chown-v6ops-renumber-thinkabout] and [RFC4192]. Contributions from [I-D.jiang-6renum-enterprise] are incorporated into the more detailed analysis. Lots of issues were analyzed in RFC5887 & [I-D.chown-v6ops-renumber-thinkabout], but many of them are out of 6renum scope or unsolvable. This document intends to identify the valuable and solvable issues, dig out of some undiscovered gaps, and tries to give solution suggestions.

2. Overall Requirements for Renumbering

This section introduces the overall ultimate goals we want to achieve in a renumbering event. In general, we need to leverage renumbering automation to avoid human intervention as much as possible at reasonable cost. Some existing mechanisms have already provided useful help. Further efforts may be achieved in the future.

The automation can be divided into four aspects as follows, which are also the gap analysis topics.

- 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 Updating relevant address entries should be performed integrally and without error. [Open Question]Is it necessary to develop automatic entries update mechanisms? If necessary, do we need standard protocols/interface for it?
- o Renumbering event management is needed to provide the functions of renumbering notification, synchronization, and monitoring .etc.

Besides automation, session survivability is another important issue during renumbering since application outage is one of the most obvious impacts that make renumbering painful and expensive. We have an enormous advantage in IPv6 which is the ability to overlap the old and new prefixes and to use the address lifetime mechanisms in SLAAC and DHCPv6. That is fully described in [RFC4192]. We consider this mechanism is sufficient for session survivability issue in most of the cases.

[Open Question]Should we consider the case of very long-lived application sessions (days or weeks) which cannot be resolved by [RFC4192]?

3. Existing Components for IPv6 Renumbering

Since renumbering is not a whole new issue, some protocols/mechanisms have been already utilized or even be developed dedicated for renumbering. However, generally current renumbering is achieved by existing protocols rather than dedicated renumbering protocols. This section briefly reviews these existing protocols/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 basic mechanisms for host configuration.
- o When a host is renumbered, it may use SLAAC [RFC4862] for address configuration with the new prefix. 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 configured through DHCPv6 [RFC3315] can reconfigure addresses by initialing RENEW sessions when the current addresses' lease time are expired or 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 has not been widely used.

3.2. Management Tools

Some operations of renumbering could be automatically processed by management tools in order to make the renumbering process more efficient and accurate. The tools may be designed dedicated for renumbering or just common tools could be utilized for some operations in renumbering.

Following are samples of the tools.

- o Address management tools. There are both commercial and open-source, and even home-made solutions.
[Further work is needed to identify what an address management tool should be able to do for improving the ability of managing a network through a renumbering event.]
- 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 though 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.

- o [I-D.jiang-6renum-enterprise] analyzes the enterprise renumbering events and gives the recommendations among the existing renumbering mechanisms. According to the different stages, renumbering considerations are described in three categories: considerations and recommendations during network design, for preparation of enterprise network renumbering, and during renumbering operation

4. Managing Prefixes

When renumbering an enterprise site, a short prefix may be divided into longer prefixes for subnets. So we need to carefully manage the prefixes for prefix delivery, delegation, aggregation, synchronization, coordination, etc.

4.1. Prefix Delegation

Usually, the short prefix comes down from the operator and received by DHCPv6 server or router inside the enterprise network. The short prefix could be automatically delegated through DHCPv6-PD. Then the downlink DHCP servers or routers can begin advertising the longer prefixes to the subnets.

For the delegation routers, they may need to renumber themselves with the delegated prefixes. We need to consider the router renumbering issue which cannot be covered by DHCP-PD only.

4.2. Prefix Assignment

When subnet routers receive the longer prefixes, they can directly assign them to the hosts. The prefix assignment overlaps with the host address configuration, which is described in the following section 5.1.

5. Address Configuration

5.1. Host Address Configuration

Both of the DHCPv6 and ND protocols have IP address configuration function. They are suitable for different scenarios respectively. During renumbering, the SLAAC-configured hosts can reconfigure IP addresses by receiving ND Router Advertisement (RA) messages containing new prefix information (It should be noted that, the prefix delivery could be achieved through DHCP according to the new IETF DHC WG document [I.D ietf-dhc-host-gen-id]). The DHCPv6-configured hosts can reconfigure addresses by initiating RENEW sessions when the current addresses' lease time are expired or

receiving the reconfiguration messages initiated by the DHCPv6 servers.

- o SLAAC and DHCPv6 address configuration co-existence

While an IPv6 site is being renumbered, both DHCPv6 and ND may be used to reconfigure the host addresses. The co-existence issue mainly includes following aspects:

- Dynamic upstream learning

[RFC5887] mentioned that, DHCP-configured hosts may want to learn about the upstream availability of new prefixes or loss of prior prefixes dynamically by deducing from periodic RA messages. But there is no standard specifying what approach should be taken by a DHCPv6-configured host when it receives RA messages containing new prefix. It depends on the operation system of the host and cannot be predicted or controlled by the network.

- DHCP-managed hosts receiving RA messages

It is unclear that whether a DHCP-managed host would accept configuration through RA messages, it depends on the policies in the host's operating system. If it ignores the RA messages and there are no DHCPv6 reconfiguration messages received either, the renumbering would fail.

- SLAAC-configured hosts finding DHCPv6 in use

[RFC5887] mentioned RA message of ND protocol contains a "Managed Configuration" flag to indicate DHCPv6 is in use. But it is unspecified what behavior should be taken when the host receives RA messages with "M" set to 1. The gap of standard will cause ambiguous host behavior because it depends on the operation system of the host.

The host may start a DHCPv6 session and receives the DHCPv6 address configuration. It is also possible that the host finds the DHCPv6 assigned prefix is different from the prefix in the RA messages, which means multiple uplinks are available or there is a serious network configuration error.

Another possibility is that the host may receive no response from any DHCPv6 servers, which means the DHCPv6 service is not available and the "Managed Configuration" flag was mis-configured.

- o DHCPv6 reconfigure bulk usage

[RFC5887] mentioned that ''DHCPv6 reconfiguration doesn't appear to be widely used for bulk renumbering purposes''.

The reconfiguration defined in [RFC3315] needs to establish a session between DHCP server and client. This could be considered as a stateful approach which needs much resource on the server to maintain the renumbering sessions. This is probably one of the reasons that DHCP reconfiguration is not suitable for bulk usage.

Another limitation of reconfiguration is that it only allows the messages to be delivered to unicast addresses. So if we want to use it for bulk renumbering, stateless DHCPv6 reconfiguration with multicast may be needed. However, this may involve protocol modification.

5.2. Router Address Configuration

- o Learning new prefixes

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

- o Restart after renumbering

"Some routers cache IP addresses in some situations. So routers might need to be restarted as a result of site renumbering" [RFC2072].

- o Router naming

In [RFC4192], it is suggested that "To better support renumbering, switches and routers should use domain names for configuration wherever appropriate, and they should resolve those names using the DNS when the lifetime on the name

expires."

As [RFC5887] described, this capability is not new, and at least it is present in most IPsec VPN implementations. But many administrators do not realize that it could be utilized to avoid manual modification during renumbering.

In enterprise scenario, the requirement of router naming is not as strong as that in ISP. So for the administrators, the motivation of using router naming for easing renumbering may be not enough.

[Open Question]Whether it is not easy to use or just suitable in few situations needs further investigation.

5.3. Static Address Configuration

There is another draft dedicated to the static address issue. Please refer to [I-D.carpenter-6renum-static-problem].

6. Updating Relevant Address Entries

When nodes in a site have been renumbered, then all the entries in the site which contain the nodes' addresses must be updated. The entries mainly include DNS records and filters in various entities such as ACLs in firewalls/gateways.

6.1. DNS Records Update

o Dynamic DNS update

For DNS records update, most sites will achieve it by maintaining a DNS zone file and loading this file into the site's DNS server(s). Synchronization between host renumbering and the updating of its A6 or AAAA record is hard. [RFC5887] mentioned that an alternative is to use Secure Dynamic DNS Update [RFC3007], in which a host informs its own DNS server when it receives a new address.

Secure Dynamic DNS Update has been widely supported by the major DNS systems, but it hasn't been widely deployed, especially in the host. Current practices mainly involve the DHCP servers which act as clients to request the DNS server to update relevant records. Normal hosts are not suitable to do this mainly because of the complexity of key management issues inherited from secure DNS mechanisms.

6.2. In-host Server Address Update

While DNS records addresses of hosts in servers, hosts also record addresses of servers such as DNS server, radius server, etc. While renumbering, the hosts must update the records if the server addresses changed. Addresses of DHCPv6 servers do not need to be updated. They are dynamically discovered using DHCPv6 relevant multicast addresses.

- o The DNS server addresses for hosts are configured by DHCPv6. But current DHCPv6 messages do not indicate to hosts the lifetimes of DNS. If the DNS lifetime expired and has been renumbered, the hosts may still use the old addresses. DHCPv6 should be extended to indicate to hosts the associated DNS lifetimes when making DNS configuration. How the DHCP server could know about the DNS lifetime is another issue.

6.3. Filters

- o Filters Management

Filters 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 filter management.

In [LEROY], a server is used for managing filter update in various devices. But identifying where and which of the filters need to be updated during renumbering is still a gap.

- o Filter Update Automation Operation

As mentioned in section 3.2, [LEROY] proposed a mechanism which can automatically update the filters. The mechanism utilizes macros suitable for various devices such as routers, firewalls etc. to update the filter entries based on the new prefix. Such automation tool is valuable for renumbering because it can reduce manual operation which is error-prone and inefficiency.

Besides the macros, [LEROY] also proposed to use SOAP to deliver the macros to the devices. As well as SOAP we may consider

whether it is possible and suitable to use other standardized protocols such as NETCONF.

Update of filters based on prefixes and filters based on addresses may have different requirements and methods. Address-based filters may be mainly with regard to domain names while prefix-based filters be relevant to more abstract entity (mask e.g.). Thus, we may consider different ways to update the two kinds of filters, for example, the prefix-based filters may consider to be updated through DHCPv6 server, which may provide better efficient.

7. Renumbering Event Management

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

7.1. Renumbering Notification

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

- o A notification mechanism may be needed to indicate the hosts that a renumbering event of local recursive DNS happen or is going to take place.
- o [RFC4192] suggests that "reducing the delay in the transition to new IPv6 addresses applies when the DNS service can be given prior notice about a renumbering event." Reducing delay could improve the efficiency of renumbering.

7.2. Synchronization Management

- o DNS update synchronization

DNS update synchronization focuses on the coordinating between DNS and other entities/mechanisms, for example, as described in [RFC5887], synchronizing the SLAAC and DNS updates, and of reducing the SLAAC lease time and DNS TTL.

7.3. Renumbering Monitoring

While treating renumbering a network event, mechanisms to monitor the renumbering process may be needed. Considering the address

configuration operation may be stateless (if ND is used for renumbering), it is difficult for monitoring. But for the DNS and filter update, it is quite possible to monitor the whole process.

8. Miscellaneous

[TBD]

9. Gap Summary

9.1. Managing Prefixes

[TBD]

9.2. Address configuration

o Host address configuration

- SLAAC and DHCPv6 address configuration co-existence

- Dynamic upstream learning:

- DHCP-configured host may want to learn about the upstream availability of new prefixes or loss of prior prefixes dynamically by deducing from periodic RA messages.

- DHCP-managed hosts receiving RA messages:

- It is unclear that whether a DHCP-managed host would accept configuration through RA messages, it depends on the policies in the host's operating system.

- SLAAC-configured hosts find DHCPv6 is in use:

- RA messages contain a "Managed Configuration" flag to indicate DHCPv6 is in use. But it is unspecified what behavior should be taken when the host receives RA messages with "M" set to 1. The gap of standard will cause ambiguous host behavior.

- DHCPv6 reconfigure bulk usage

- If we want to use it for bulk renumbering, stateless DHCPv6 reconfiguration with multicast may be needed. However, this may involve protocol modification.

o Router address configuration

- Learning new prefixes

If the site is multihoming, the interior routers would need a mechanism to learn which upstream providers and prefixes were currently reachable (and valid).

- Restart after renumbering

Some routers cache IP addresses in some situations. So routers might need to be restarted as a result of site renumbering.

- Router naming

Using domain names for routers is suitable in some scenarios but has not been widely deployed. The motivation of using router naming for easing renumbering in enterprise networks may be not enough.

- o Static address configuration

Please refer to [I-D.carpenter-6renum-static-problem].

9.3. Address relevant entries update

- o DNS records update

- DNS update automation

Synchronization between host renumbering and the updating of its DNS records is hard. Forward and reverse DNS entries update may need to be combined together.

- o In-host server address update

Hosts also record addresses of servers such as DNS server addresses, radius server address, etc. While renumbering, the host must update the records if these server addresses changed.

- o Filters

- Filters management

There is a gap of filter management to identify where and which of the filters need to be updated during renumbering. Manageable filter update may be also needed.

- Filter update automation operation

Automation update tool is valuable for renumbering, and there may be requirement of protocol standardization to deliver of facility the tools.

9.4. Renumbering event management

- o Renumbering notification

If hosts/servers are aware of a renumbering event happening, it may help the relevant process. A basic way is to extend current protocol messages to carry the renumbering notification.

- o Synchronization management

- DNS update synchronization

An example is synchronizing the SLAAC and DNS updates, and of reducing the SLAAC lease time and DNS TTL. [TBD]

- o Renumbering monitoring

Mechanisms to monitor the process and feedback of renumbering may be needed. [TBD]

10. 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 SeND ([RFC3971], Secure Neighbor Discovery) deployment may need to validate prefixes.

- o Influence to Security Controls

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

11. IANA Considerations

None.

12. References

12.1. Normative References

- [RFC2894] M. Crawford, "Router Renumbering for IPv6", RFC 2894, August 2000.
- [RFC2874] Crawford, M., and C. Huitema, "DNS Extensions to Support IPv6 Address Aggregation and Renumbering", RFC 2874, July 2000.
- [RFC3007] B. Wellington, "Secure Domain Name System (DNS) Dynamic Update", RFC 3007, November 2000.
- [RFC3315] R. Droms, Bound, J., Volz, B., Lemon, T., Perkins, C., and M. Carney, "Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", RFC 3315, July 2003.
- [RFC3633] Troan, O. and R. Droms, "IPv6 Prefix Options for Dynamic Host Configuration Protocol (DHCP) version 6", RFC 3633, December 2003.
- [RFC3956] P. Savola, and B. Haberman. "Embedding the Rendezvous Point (RP) Address in an IPv6 Multicast Address.", RFC 3956, November 2004.
- [RFC3971] Arkko, J., Ed., Kempf, J., Zill, B., and P. Nikander "SECure Neighbor Discovery (SEND)", RFC 3971, March 2005.
- [RFC4861] Narten, T., Nordmark, E., Simpson, W., and H. Soliman, "Neighbor Discovery for IP version 6 (IPv6)", RFC 4861, September 2007.
- [RFC4862] Thomson, S., Narten, T., and T. Jinmei, "IPv6 Stateless Address Autoconfiguration", RFC 4862, September 2007.

12.2. Informative References

- [RFC2072] H. Berkowitz, "Router Renumbering Guide", RFC2072, January 1997.
- [RFC4192] Baker, F., Lear, E., and R. Droms, "Procedures for Renumbering an IPv6 Network without a Flag Day", RFC 4192, September 2005.
- [RFC4714] Enns, R., "NETCONF Configuration Protocol", RFC 4714, December 2006.

- [RFC5887] Carpenter, B., Atkinson, R., and H. Flinck, "Renumbering Still Needs Work", RFC 5887, May 2010.
- [I-D.ietf-dhc-secure-dhcpv6]
Jiang, S., and Shen S., "Secure DHCPv6 Using CGAs", working in progress.
- [I-D.chown-v6ops-renumber-thinkabout]
Chown, T., "Things to think about when Renumbering an IPv6 network", Work in Progress, September 2006.
- [I-D.jiang-6renum-enterprise]
Jiang, S., and B. Liu, "IPv6 Enterprise Network Renumbering Scenarios and Guidelines ", Working in Progress, July 2011.
- [I-D.carpenter-6renum-static-problem]
Carpenter, B., and S. Jiang, "Problem Statement for Renumbering IPv6 Hosts with Static Addresses", Working in Progress, October 2011.
- [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>>

13. Acknowledgments

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

Useful comments and contributions were made by Wesley George, and others.

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

Appendix A. Other Documented Gaps

This appendix lists gaps have been documented but are considered not in the scope of this gap analysis.

A.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. [Open Question]This limitation requires renumbering to be planned in advance so that an immediate renumbering event is impossible. Should it be considered as a standard gap for renumbering?

A.2 Address Relevant Entries Update

- o DNS entries commonly have matching Reverse DNS entries which will also need to be updated during renumbering.
- o [Open Question]So synchronizing the procedures of forward and reverse DNS or even combining forward and reverse DNS updates in a single procedure also need to be considered.
- o DNS data structure optimization

[RFC2874] proposed a new A6 record type for DNS recording IPv6 address/prefix. And several extensions on query and processing were also proposed. With the A6 record and the extensions, an IPv6 address can be defined by using multiple DNS records. This feature increases the complexity of resolver but reduce the cost of zone file maintenance. So renumbering could be easier than AAAA record. But the [RFC2874] has not been widely used.

[Open Question]Is the DNS data structure optimization such as [RFC2874] necessary for easing renumbering? If necessary, is the optimization in [RFC2874] enough?

- o DNS authority

As described in [I-D.chown-v6ops-renumber-thinkabout], "it is often the case in enterprises that host web servers and application servers on behalf of collaborators and customers that DNS zones out of the administrative control of the host maintain resource records concerning addresses for nodes out of their control. When the service host renumbers, they do not have sufficient authority to change the records."

[Open Question]Whether it is only an operational issue or additional protocol/mechanism is needed to standardize the interaction between DNS systems needs to be considered.

A.3 Miscellaneous

A.3.1 Multicast

- o The embedding of IPv6 unicast addresses into multicast addresses and the embedded-RP (Rendezvous Point)[RFC3956] will cause issues when renumbering.

As [I-D.chown-v6ops-renumber-thinkabout] described, "If the RP address changes, then the group addresses must also be changed. This may happen not only when a site is renumbered, but also if a site is restructured or the RP is moved within the site. The embedded address is used by routers to determine the RP address. Applications must use new group addresses once the RP is not available on the old address."

- o Changing the unicast source address of a multicast sender might also be an issue for receivers.

As [I-D.chown-v6ops-renumber-thinkabout] described, "If a site's unicast prefix changes, then one will also need to change the multicast addresses. By way of example, a site renumbering away from prefix 2001:DB8:BEEF::/48" might have globally-scoped multicast addresses in use under the prefix "FF3E:30:2001:DB8:BEEF::/96". One may continue using the old addresses for a while, but this should be avoided since another site may inherit the prefix and they may end up using the same multicast addresses."

A.3.2 Mobility

- o [RFC5887] suggested that, for Mobile IP, define a better mechanism to handle change of home agent address while mobile is disconnected.

A.3.3 Non-network issues

- o For transport layer, [5887] said that TCP connections and UDP flows are rigidly bound to a given pair of IP addresses.
- o For application layer, as [5887] said, in general, we can assert that any implementation is at risk from renumbering if it does not check that an address is valid each time it opens a new communications session.

Authors' Addresses

Bing Liu
Q14-4-A Building
Huawei Technologies Co., Ltd
Zhong-Guan-Cun Environment Protection Park, No.156 Beiqing Rd.
Hai-Dian District, Beijing
P.R. China

Email: leo.liubing@huawei.com

Sheng Jiang
Q14-4-A Building
Huawei Technologies Co., Ltd
Zhong-Guan-Cun Environment Protection Park, No.156 Beiqing Rd.
Hai-Dian District, Beijing
P.R. China

Email: shengjiang@huawei.com

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

Email: brian.e.carpenter@gmail.com

