Network Working Group   F. Templin, Ed.
Internet-Draft   Boeing Research & Technology
Intended status: Informational   September 5, 2017
Expires: March 9, 2018

IPv6 Prefix Delegation for Hosts
draft-templin-v6ops-pdhost-06.txt

Abstract

IPv6 prefixes are typically delegated to requesting routers which then use them to number their downstream-attached links and networks. This document considers both this traditional case, and the case when the "requesting router" is actually a simple host which receives a delegated prefix that it can use for its own internal multi-addressing purposes. This latter method can be employed in a wide variety of use cases to allow ample address availability without impacting link performance.

Table of Contents

1. Introduction 2
2. Terminology 5
3. Multi-Addressing Considerations 5
4. Multi-Addressing Alternatives for Delegated Prefixes 6
5. MLD/DAD Implications 6
6. IPv6 Neighbor Discovery Implications 7
7. ICMPv6 Implications 7
8. IANA Considerations 8
9. Security Considerations 8
10. Acknowledgements 8
11. References 8
11.1. Normative References 8
11.2. Informative References 9
Author's Address 10

1. Introduction

IPv6 Prefix Delegation (PD) entails 1) the communication of a prefix from a delegating authority to a requesting node, 2) a representation of the prefix in the routing system, and 3) a control messaging service to maintain delegated prefix lifetimes. Following delegation, the prefix is available for the requesting node's exclusive use and is not shared with any other nodes. An example IPv6 PD service is DHCPv6 PD [RFC3315][RFC3633].

Using any available prefix delegation service, a Delegating Router 'D' delegates a prefix 'P' to a Requesting Node 'R'' as shown in Figure 1:

```
                        +---------------------+
                        |Delegating Router 'D'|
                        |   (Delegate 'P')    |
                        +----------+----------+
                                   |
                             .-(::::::::)
                           .-(:::: IP ::::)-.
                          (:: Internetwork ::)
                           `-(::::::::::::)-'
                              `-(::::::)-'
                                   | WAN Interface
                        +----------+----------+
                        |    (Receive 'P')    |
                        |  Requesting Node 'R'|
                        +----------+----------+
                                   | LAN Interface
       X----+-------------+--------+----+---------------+---X
            |             |      LAN    |               |
       +---++-+--+   +---++-+--+   +---++-+--+     +---++-+--+
       |   |A1|  |   |   |A2|  |   |   |A3|  |     |   |An|  |
       |   +--+  |   |   +--+  |   |   +--+  |     |   +--+  |
       | Host H1 |   | Host H2 |   | Host H3 | ... | Host Hn |
       +---------+   +---------+   +---------+     +---------+
```

Figure 1: Prefix Delegation Model

In this figure, when Delegating Router 'D' delegates prefix 'P', the prefix is injected into the IP Internetwork routing system in some fashion to ensure that IPv6 packets with destination addresses covered by 'P' are unconditionally forwarded to Requesting Node 'R'. Meanwhile, 'R' receives 'P' via its "WAN" interface and sub-delegates 'P' to its downstream-attached links via one or more "LAN" interfaces. Hosts 'H(i)' on a LAN subsequently receive addresses 'A(i)' taken from 'P' via an address autoconfiguration service such as IPv6 Stateless Address Autoconfiguration (SLAAC) [RFC4862]. 'R' then acts as a router between hosts 'H(i)' and correspondents reachable via the WAN interface.

This document also considers the case when 'R' is actually a simple host, and receives a prefix delegation 'P' as if it were a router. The host need not have any LAN interfaces, and can use the prefix solely for its own internal addressing purposes. 'R' can act as a host under the weak end system model [RFC1122] if it can assign addresses taken from 'P' to its own internal virtual interfaces (e.g., a loopback) as shown in Figure 2:

```
                     +--------------------+
                     |Delegating Router 'D'|
                     |   (Delegate 'P')    |
                     +----------+----------+
                                |
                           .-(::::::::)
                         .-(:::: IP ::::)-.
                        (:: Internetwork ::)
                         `-(::::::::::::)-'
                            `-(::::::)-'
                                | WAN Interface
                     +----------+----------+
                     |    (Receive 'P')    |
                     | Requesting Node 'R' |
                     +---------------------+
                     | Loopback Interface  |
                     +--+-+--+-++-+-----+--+
                     |A1| |A2| |A3| ... |An|
                     +--+-+--+-+--+-----+--+
```

Figure 2: Weak End System Model

'R' could instead function as a host under the strong end system model [RFC1122] by assigning IPv6 addresses taken from prefix 'P' to the WAN interface as shown in Figure 3:

```
                     +--------------------+
                     |Delegating Router 'D'|
                     |   (Delegate 'P')    |
                     +----------+----------+
                                |
                           .-(::::::::)
                         .-(:::: IP ::::)-.
                        (:: Internetwork ::)
                         `-(::::::::::::)-'
                            `-(::::::)-'
                                | WAN Interface
                     +--+-+--+-++-+-----+--+
                     |A1| |A2| |A3| ... |An|
                     +--+ +--+ +--+     +--+
                     |    (Receive 'P')    |
                     | Requesting Node 'R' |
                     +---------------------+
```

Figure 3: Strong End System Model

The major benefit for a host managing a delegated prefix in either the weak or strong end system models is multi-addressing. With

multi-addressing, the host can configure an unlimited supply of addresses to make them available for local applications without requiring coordination with any other nodes.

The following sections present multi-addressing considerations for hosts that employ prefix delegation mechanisms.

2. Terminology

The terminology of the normative references apply. The following terms are defined for the purposes of this document:

shared prefix
: an IPv6 prefix that may be advertised to more than one node on the same link, e.g., in a Prefix Information Option (PIO) included in a Router Advertisement (RA) message [RFC4861]. The shared prefix property applies not only on multiple access links (e.g., multicast-capable, NBMA, shared media, etc.), but also on point-to-point links where the shared prefix is visible to both ends of the link.

delegated prefix
: a prefix that is delegated to a requesting node solely for its own use, and is not delegated to any other nodes on the link.

3. Multi-Addressing Considerations

IPv6 allows nodes to assign multiple addresses to a single interface. [RFC7934] discusses options for multi-addressing as well as use cases where multi-addressing may be desirable. Address configuration options for multi-addressing include SLAAC [RFC4862], stateful DHCPv6 address configuration [RFC3315] and any other address formation methods (e.g., manual configuration).

Nodes that use SLAAC and/or DHCPv6 address configuration configure addresses from a shared prefix and assign them to the interface over which the prefix was received (e.g., in an RA). When this happens, the node is obliged to use Multicast Listener Discovery (MLD) to join the appropriate solicited-node multicast group(s) and to use the Duplicate Address Detection (DAD) algorithm [RFC4862] to ensure that no other node that receives the shared prefix configures a duplicate address.

In contrast, a node that uses address configuration from a delegated prefix can assign addresses without invoking MLD/DAD on the WAN interface, since the prefix has been delegated to the node for its own exclusive use and is not shared with any other nodes.

4. Multi-Addressing Alternatives for Delegated Prefixes

When a node receives a prefix delegation, it has many alternatives for the way in which it can provision the prefix. [RFC7278] discusses alternatives for provisioning a prefix obtained by a User Equipment (UE) device under the 3rd Generation Partnership Program (3GPP) service model. This document considers the more general case when the node receives a prefix delegation in which the prefix is delegated for its own exclusive use.

When the node receives the prefix, it can distribute the prefix to downstream-attached networks via its LAN interfaces and configure one or more addresses for itself on a LAN interface. The node then acts as a router on behalf of its downstream-attached networks and configures a default route that points to a router on the WAN link. This approach is often known as the "tethered" configuration.

The node could instead use the delegated prefix for its own multi-addressing purposes. In a first alternative, the node can receive the prefix acting as a requesting node over the WAN interface but then assign the prefix to an internal virtual interface (e.g., a loopback interface) and assign one or more addresses taken from the prefix to the virtual interface. In that case, applications on the node can use the assigned addresses according to the weak end system model.

In a second alternative, the node can receive the prefix as a requesting node over the WAN interface but then assign one or more addresses taken from the prefix to the WAN interface. In that case, applications on the node can use the assigned addresses according to the strong end system model.

In both of these latter two cases, the node acts as a host internally even though it behaves as a router from the standpoint of prefix delegation and neighbor discovery over the WAN interface. The host can configure as many addresses for itself as it wants.

5. MLD/DAD Implications

When a node configures addresses for itself using either SLAAC or DHCPv6 from a shared prefix, the node performs MLD/DAD by sending multicast messages to test whether there is another node on the link that configures a duplicate address from the shared prefix. When there are many such addresses and/or many such nodes, this could result in substantial multicast traffic that affects all nodes on the link.

When a node configures addresses for itself using a delegated prefix, the node can configure as many addresses as it wants but does not perform MLD/DAD for any of the addresses over the WAN interface. This means that arbitrarily many addresses can be assigned without causing any multicast messaging over the WAN link that could disturb other nodes.

6. IPv6 Neighbor Discovery Implications

The node acts as a simple host to send Router Solicitation (RS) messages over the WAN interface the same as described in Section 4.2 of [RFC7084].

In order to maintain the appearance of a router (i.e., even though it is acting as a simple host), the node sets the "Router" flag to TRUE in any Neighbor Advertisement messages it sends. This ensures that the "isRouter" flag in the neighbor cache entries of any neighbors remains TRUE.

The node initially has only a default route pointing to a router on the WAN link. This means that packets sent over the node's WAN interface will initially go through a default router even if there is a better first-hop node on the link. In that case,a Redirect message can update the node's neighbor cache, and future packets can take the more direct route without disturbing the default router. The Redirect can apply either to a singleton destination address, or to an entire destination prefix as described in [I-D.templin-6man-rio-redirect].

7. ICMPv6 Implications

The Internet Control Message Protocol for IPv6 (ICMPv6) includes a set of control message types [RFC4443] including Destination Unreachable (DU). Routers return DU messages with code 0 ("No route to destination") when a packet arrives for which there is no matching entry in the routing table, and with code 3 ("Address unreachable") when the IPv6 destination address cannot be resolved into a link-layer address. Hosts return DU messages with code 3 to internal applications when an address cannot be resolved, and with code 4 ("Port unreachable") to the sender if the transport protocol has no listener.

A node that obtains a prefix delegation for "tethering" purposes acts like a router in all respects and returns DU messages the same as for any router.

A node that obtains a prefix delegation for its own multi-addressing purposes (whether weak or strong end system) should act like a host

and refrain from sending DU messages with code 0 or 3 when it receives a packet from a sender with an unknown IPv6 destination address. That is to say that the node should silently drop any IPv6 packets with a destination address that matches the delegated prefix but does not match any of its configured addresses.

8. IANA Considerations

This document introduces no IANA considerations.

9. Security Considerations

Security considerations are the same as specified for DHCPv6 Prefix Delegation in [RFC3633] and for IPv6 Neighbor Discovery in[RFC4861].

10. Acknowledgements

This work was motivated by recent discussions on the v6ops list. Mark Smith pointed out the need to consider MLD as well as DAD for the assignment of addresses to interfaces. Ricardo Pelaez-Negro, Edwin Cordeiro, Fred Baker and Naveen Lakshman provided useful comments that have greatly improved the draft.

11. References

11.1. Normative References

[RFC0791] Postel, J., "Internet Protocol", STD 5, RFC 791, DOI 10.17487/RFC0791, September 1981, <https://www.rfc-editor.org/info/rfc791>.

[RFC1122] Braden, R., Ed., "Requirements for Internet Hosts - Communication Layers", STD 3, RFC 1122, DOI 10.17487/RFC1122, October 1989, <https://www.rfc-editor.org/info/rfc1122>.

[RFC2460] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", RFC 2460, DOI 10.17487/RFC2460, December 1998, <https://www.rfc-editor.org/info/rfc2460>.

[RFC3315] Droms, R., Ed., Bound, J., Volz, B., Lemon, T., Perkins, C., and M. Carney, "Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", RFC 3315, DOI 10.17487/RFC3315, July 2003, <https://www.rfc-editor.org/info/rfc3315>.

[RFC3633] Troan, O. and R. Droms, "IPv6 Prefix Options for Dynamic Host Configuration Protocol (DHCP) version 6", RFC 3633, DOI 10.17487/RFC3633, December 2003, <https://www.rfc-editor.org/info/rfc3633>.

[RFC4443] Conta, A., Deering, S., and M. Gupta, Ed., "Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification", STD 89, RFC 4443, DOI 10.17487/RFC4443, March 2006, <https://www.rfc-editor.org/info/rfc4443>.

[RFC4861] Narten, T., Nordmark, E., Simpson, W., and H. Soliman, "Neighbor Discovery for IP version 6 (IPv6)", RFC 4861, DOI 10.17487/RFC4861, September 2007, <https://www.rfc-editor.org/info/rfc4861>.

[RFC4862] Thomson, S., Narten, T., and T. Jinmei, "IPv6 Stateless Address Autoconfiguration", RFC 4862, DOI 10.17487/RFC4862, September 2007, <https://www.rfc-editor.org/info/rfc4862>.

[RFC7084] Singh, H., Beebee, W., Donley, C., and B. Stark, "Basic Requirements for IPv6 Customer Edge Routers", RFC 7084, DOI 10.17487/RFC7084, November 2013, <https://www.rfc-editor.org/info/rfc7084>.

[RFC7278] Byrne, C., Drown, D., and A. Vizdal, "Extending an IPv6 /64 Prefix from a Third Generation Partnership Project (3GPP) Mobile Interface to a LAN Link", RFC 7278, DOI 10.17487/RFC7278, June 2014, <https://www.rfc-editor.org/info/rfc7278>.

11.2. Informative References

[I-D.templin-6man-rio-redirect] Templin, F. and j. woodyatt, "Route Information Options in IPv6 Neighbor Discovery", draft-templin-6man-rio-redirect-04 (work in progress), August 2017.

[RFC7934] Colitti, L., Cerf, V., Cheshire, S., and D. Schinazi, "Host Address Availability Recommendations", BCP 204, RFC 7934, DOI 10.17487/RFC7934, July 2016, <https://www.rfc-editor.org/info/rfc7934>.

Author's Address

Fred L. Templin (editor)
Boeing Research & Technology
P.O. Box 3707
Seattle, WA 98124
USA

Email: fltemplin@acm.org