Network Working Group F. Templin, Ed.
Internet-Draft Boeing Research & Technology
Intended status: Informational June 27, 2016
Expires: December 29, 2016

IPv6 Prefix Delegation for Hosts
draft-templin-v6ops-pdhost-02.txt

Abstract

IPv6 prefixes are typically delegated to requesting routers which then use them to number their downstream-attached links and networks. The requesting router then acts as a router between the downstream-attached hosts and the upstream provider network, and can also act as a host under the weak end system model. This document considers the case when the "requesting router" is actually a simple host, and receives a delegated prefix that it can use for multi-addressing purposes. The host does not connect any downstream-attached networks, and uses the prefix solely for its own multi-addressing purposes.

Table of Contents

1. Introduction

IPv6 provides a prefix delegation service using the Dynamic Host Configuration Protocol for IPv6 (DHCPv6) [RFC3315]. Using DHCPv6 Prefix Delegation (PD) [RFC3633], a requesting router asks for a prefix from a delegating router. When the prefix is delegated, the requesting router sub-delegates the prefix to its downstream-attached links via one or more "LAN" interfaces. The requesting router then acts as a router between hosts on the LAN interfaces and the upstream provider network (i.e., the "WAN" interface), and can also act as a host under the weak end system model [RFC1122]. This document considers the case when the "requesting router" is actually a simple host, and receives a prefix delegation as if it were a router. The host need not have any LAN interfaces, and can use the prefix solely for its own multi-addressing purpose.

2. Multi-Addressing

IPv6 allows for assignment of multiple addresses to a single interface. [I-D.ietf-v6ops-host-addr-availability] discusses options for multi-addressing as well as use cases where multi-addressing may be desirable. Multi-addressing options include Stateless Address Autoconfiguration (SLAAC) [RFC4862] or stateful DHCPv6 address delegation [RFC3315], as well as assignment of multiple addresses from a delegated prefix.

SLAAC and DHCPv6 address delegation typically obtain addresses from an on-link prefix configured on the link over which the addresses are obtained. When this happens, the address recipient is obliged to use Multicast Listener Discovery (MLD) to join the appropriate solicited-node multicast group(s) and the Duplicate Address Detection (DAD) algorithm [RFC4862] to ensure that no other node on the link configures a duplicate address. Alternatively, address delegation from a delegated prefix can be used by a node under either the weak or strong end system models [RFC1122]. In that case, the MLD/DAD procedure is not necessary, since the prefix has been delegated to the node for its own exclusive use and the prefix is not assigned to the link over which the prefix was obtained.

3. Multi-Addressing Alternatives

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 the exclusive use of the prefix recipient.

When the node receives the prefix (e.g., a /64), it can sub-delegate the prefix to its LAN interfaces and configure multiple addresses for itself on a LAN interface. The node uses link-local-only addressing on the WAN interface, and configures a default route that points to a router on the WAN link. The node can then act as both a host for its own applications and a router for any downstream-attached hosts. This approach is often known as the "tethered" configuration.

When the node does not have any LAN interfaces, it may still wish to obtain a prefix solely for multi-addressing purposes. In a first alternative, the node can receive the prefix acting as a requesting router 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 router over the WAN interface but then assign the prefix to a loopback interface and 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.

4. MLD/DAD Implications

When a node configures addresses for itself using either SLAAC or DHCPv6 address delegation from a prefix that is on-link on the WAN interface, the node performs MLD/DAD by sending multicast messages to test whether another node that configures a duplicate address is on the link. 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 having any multicast messaging over the WAN link that could disturb other nodes. Note however that nodes that assign the addresses directly to the WAN interface must be capable of disabling MLD/DAD on the WAN interface, i.e., they must set DupAddrDetectTransmits to zero [RFC4862].

5. IPv6 Neighbor Discovery Implications

The node acts as a simple host to send Router Solicitation 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 AERO [I-D.templin-aerolink].

6. IANA Considerations

This document introduces no IANA considerations.

7. Security Considerations

TBD.

8. 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.

9. References

9.1. Normative References

[RFC0791] Postel, J., "Internet Protocol", STD 5, RFC 791, DOI 10.17487/RFC0791, September 1981, <http://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, <http://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, <http://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, <http://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, <http://www.rfc-editor.org/info/rfc3633>.

[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, <http://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, <http://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, <http://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, <http://www.rfc-editor.org/info/rfc7278>.

9.2. Informative References

[I-D.ietf-v6ops-host-addr-availability] Colitti, L., Cerf, V., Cheshire, S., and D. Schinazi, "Host address availability recommendations", draft-ietf-v6ops-host-addr-availability-07 (work in progress), May 2016.

[I-D.templin-aerolink] Templin, F., "Asymmetric Extended Route Optimization (AERO)", draft-templin-aerolink-67 (work in progress), June 2016.

Author's Address

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

Email: fltemplin@acm.org