Internet Engineering Task Force
Internet-Draft
Updates: 4861, 5942 (if approved)
Intended status: Standards Track
Expires: February 17, 2016

M. Smith
IMOT

August 16, 2015

Indicating Link-Local Unicast Destinations are Off-Link
draft-smith-6man-link-locals-off-link-00

Abstract

Certain link-layers limit reachability for one set of nodes, while permitting full reachability for a different set of nodes, for unicast, multicast and broadcast traffic. If IPv6 hosts are members of the first set of nodes, and IPv6 routers are members of the second, Link-Local traffic between IPv6 hosts will fail, due to the default on-link assumption for Link-Local destinations. This memo describes the use of a Link-Local Prefix Information Option to indicate to these hosts that Link-Local destinations are "off-link", and are reachable via their default router(s).

Table of Contents

1. Introduction 2
1.1. Requirements Language 2
2. Constrained Broadcast Multi-Access (CBMA) Links 3
3. IPv6 over CBMA Links 3
4. Link-Local traffic over CBMA Links 4
5. Indicating Link-Local Destinations are Off-Link 5
5.1. Link-Local Router Advertisement Prefix Information Option 5
5.2. Host Link-Local Prefix Information Option Processing . . 5
5.2.1. Upon Receipt of a Link-Local PIO 5
5.2.1.1. Validation 5
5.2.1.2. Processing 6
5.2.2. Upon Expiry of Link-Local Off-Link Information . . . 6
6. Updates to RFC4861 6
7. Updates to RFC5942 7
8. Security Considerations 8
9. Acknowledgements 8
10. Change Log [RFC Editor please remove] 8
11. References 8
11.1. Normative References 8
11.2. Informative References 8
Author's Address 9

1. Introduction

Certain link-layers limit reachability for one set of nodes, while permitting full reachability for a different set of nodes, for unicast, multicast and broadcast traffic. If IPv6 hosts are members of the first set of nodes, and IPv6 routers are members of the second, Link-Local traffic between IPv6 hosts will fail, due to the default on-link assumption for Link-Local destinations. This memo describes the use of a Link-Local Prefix Information Option to indicate to these hosts that Link-Local destinations are "off-link", and are reachable via their default router(s).

1.1. Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119 [RFC2119].

2. Constrained Broadcast Multi-Access (CBMA) Links

A variety of multi-access link-layers operate or can be configured to constrain layer 2 reachability between attached nodes. Typically, one set of nodes can reach all other attached nodes, while a second disjoint set of nodes are only able to reach members of the first set. Members of the second set are isolated from each other. These constraints are applied to link-layer unicast, multicast and broadcast traffic.

Examples of these types of links are Broadband Forum TR-101 VLANs following the N:1 forwarding model, IEEE 802.11 Wifi Networks with station isolation switched on, and "Private VLANs" [RFC5517].

This memo uses the more general term "Constrained Broadcast Multi-Access" (CBMA) to describe these types of links. These types of links are distinct from Non-Broadcast Multi-Access (NBMA) links.

3. IPv6 over CBMA Links

When IPv6 is operated over a CBMA link, one or more IPv6 routers would be selected as the link-layer nodes that can reach all other nodes, while IPv6 hosts would be selected as the link-layer nodes limited to only being able to reach the IPv6 routers.

Unless informed otherwise via Router Advertisement Prefix Information Options (PIOs) with the L or on-link flag switched on [RFC4861], IPv6 hosts are to consider all non-Link-local destinations off-link, including destination addresses that fall within the prefix their own addresses are assigned from [RFC5942]. Unicast destination addresses attached to the same link will be reached by hosts sending their packets towards one of their default routers, which will then forward the packets back over the same link towards the final destinations.

This "hair-pin" or "trombone" forwarding between IPv6 hosts attached to the same link, via one or more default routers, allows the router(s) to be used to perform functions in addition to standard IPv6 forwarding, such as traffic inspection for security purposes or per-host traffic accounting. Security examples might be to prevent unauthorised nodes emitting Router Advertisements or acting as unauthorised DHCPv6 servers.

Note that multicast IPv6 traffic is normally sent to all link-layer reachable nodes, possibly limited to interested hosts using MLD Snooping [RFC4541]. On a CBMA link, IPv6 hosts' multicasts will be further limited to only reaching the IPv6 routers. These IPv6 routers may choose to drop this multicast traffic if they're not

interested in it or perform proxy functions for other hosts attached to the link (e.g., DAD Proxy [RFC6957]).

The author is not aware of a more general method of multicast forwarding that could be used by the routers to allow all hosts attached to the CBMA link to receive other CBMA link hosts' multicasts should they pass any multicast security policies applied by the CBMA router(s).

(When hosts receive multicasts over an interface, do they check if the multicast source address is one of their own, and ignore the multicast if so (as distinct from multicasts looped back locally within the host, enabled by a socket API call)? If so, the CBMA routers could send multicasts back onto the CBMA link (after passing other security checks), and the originating host would ignore them. Perhaps hosts could perform this sort of source address checking if they receive the Link-Local Prefix Information Option, described below, while it remains valid. There isn't a forwarding loop in the link-layer, it is just that originating hosts would receive their own multicasts that need to be ignored. Only one of the routers on the CBMA link would send multicasts back onto the link; the router with the lowest value Link-Local address could be the one, using the set of routers' RA and RS Link-Local source addresses to choose (similar to a multicast router querier election)). The non-elected router would not forward multicast traffic it receives back onto the CBMA link.

4. Link-Local traffic over CBMA Links

Due to the on-link assumption for Link-Local unicast destinations [RFC5942], attempts to send Link-Local traffic to other hosts attached to the CBMA link will fail, as the inter-host reachability has been constrained by the CBMA link. The only Link-Local destinations reachable by the hosts are the Link-Local addresses of the default routers.

This limited Link-Local reachability can be detrimental to the operation of IPv6 applications, as IPv6 applications are permitted to use Link-Local addresses for their connectivity [RFC4007], and if multiple scopes of addresses are available for the application to use, Link-Local addresses will be preferred over all others with exception to the loopback address, due to the general rule of preferring addresses with the smaller scopes [RFC6724].

If hosts could be informed that Link-Local destinations are to also be considered "off-link", reachability to all Link-Local destinations on the CBMA link would be restored. Hosts would send traffic to all Link-Local destinations via their default router(s), with the chosen

default router then forwarding the traffic back onto the CBMA link [RFC4007] towards the final Link-Local destination.

5. Indicating Link-Local Destinations are Off-Link

5.1. Link-Local Router Advertisement Prefix Information Option

To signal to hosts that they should consider Link-Local destinations "off-link", a router sends a Link-Local Prefix Information Option in its Router Advertisements [RFC4861], with the following PIO field values:

- Prefix Length: 10 [RFC4291]
- L or On-Link Flag: 0 (Off)
- A or Autonomous Address-Configuration Flag: 0 (Off)
- Valid Lifetime: Length of time Link-Local destinations are to be considered off-link
- Preferred Lifetime: 0xffffffff (representing infinity) [RFC4862]
- Prefix: fe80:: [RFC4291]

5.2. Host Link-Local Prefix Information Option Processing

5.2.1. Upon Receipt of a Link-Local PIO

5.2.1.1. Validation

When a host receives a Link-Local Prefix Information Option, it MUST perform the following validation steps:

1. Verifies the Prefix field value is fe80::. If not, this is not a LL PIO, and should be processed as a conventional PIO.
2. Verifies the Prefix Length field value is 10. If not, ignores the LL PIO.
3. Verifies the L or On-Link Flag value is 0. If not, ignores the LL PIO.
4. Ignores the A or Autonomous Address-Configuration Flag value, as Link-Local addresses always use Autonomous Address-Configuration, and are formed when an interface becomes enabled [RFC4862].

5. Verifies the Preferred Lifetime field value is Infinity (0xffffffff). If not, ignores the LL PIO.

If any of the above validation steps fail, in addition to ignoring the LL PIO, an implementation MAY choose to log an informational or debugging severity level system message about the malformed LL PIO, appropriately rate limited.

5.2.1.2. Processing

Once the LL PIO has been successfully validated, the Link-Local prefix is removed from the host's Prefix List [RFC5942]. A count down to zero timer is started with the LL PIO's Valid Lifetime value.

While the timer is still running, the host sends all Link-Local destined traffic for the interface it received the LL PIO on to either the router it received the LL PIO from, or to any of the default routers on the link, achieving an amount of load-sharing [RFC4311].

As Link-Local destinations are now being reached via the host's default router(s), Neighbor Cache entries for Link-Local destinations, excepting Link-Local entries with the IsRouter flag set [RFC4861], should be removed immediately, regardless of their resolution state. Any active related Neighbor Unreachability Detection procedures should also be terminated.

5.2.2. Upon Expiry of Link-Local Off-Link Information

If Link-Local "Off-Link" information expires, as it has not been refreshed by receiving a LL PIO from any of the link's routers, the Link-Local prefix is returned to the host's Prefix List for the corresponding interface, meaning that Link-Local destinations return to being considered on-link. Subsequent transmissions to Link-Local destinations should trigger Neighbor Discovery [RFC4861], despite the link possibly continuing to be a CBMA-type link.

6. Updates to RFC4861

The following statement in Section 6.3.4 of Neighbor Discovery in IPv6 [RFC4861]:

"Note, however, that a Prefix Information option with the on-link flag set to zero conveys no information concerning on-link determination and MUST NOT be interpreted to mean that addresses covered by the prefix are off-link."

is replaced by

"Note, however, that a Prefix Information option with the on-link flag set to zero conveys no information concerning on-link determination and MUST NOT be interpreted to mean that addresses covered by the prefix are off-link, with exception to a Prefix Information Option for the Link-Local prefix. The Link-Local prefix is considered on-link by default [RFC5942]."

"The reception of a Prefix Information option for the Link-Local prefix, with the L-bit set to 0, MUST be interpreted by a host as meaning that Link-Local destinations are to considered off-link, and are to be reached by one of the host's available default routers, while the Prefix Information option information for the Link-Local prefix remains valid [draft-smith-6man-link-locals-off-link]."

The following statement in Section 6.3.4 of Neighbor Discovery in IPv6 [RFC4861]:

"If the prefix is the link-local prefix, silently ignore the Prefix Information option."

is removed.

7. Updates to RFC5942

The following statement in the Introduction of IPv6 Subnet Model: The Relationship between Links and Subnet Prefixes [RFC5942]:

"In IPv6, by default, a host treats only the link-local prefix as on-link."

is replaced by

"In IPv6, by default, a host treats only the link-local prefix as on-link, unless updated by a Prefix Information option for the link-local prefix, indicating the link-local prefix is to be considered off-link [draft-smith-6man-link-locals-off-link]."

The following statement in the Section 3 of IPv6 Subnet Model: The Relationship between Links and Subnet Prefixes [RFC5942]:

"(The link-local prefix is effectively considered a permanent entry on the Prefix List.)"

is deleted.

8. Security Considerations

The security benefit of operating IPv6 over a CBMA link-layer is the insertion of an IPv6 traffic forwarding device between each host and all other possible destinations, including those attached to the same CBMA link. This allows the forwarding device to be used to perform security functions on all CBMA attached host originated traffic, in addition to performing normal IPv6 forwarding.

Allowing Link-Local source and destination addresses to be used in an IPv6 over CBMA network does not reduce this security benefit.

9. Acknowledgements

Thanks to Chris Chaundy for asking the author about the on-link status of the Link-Local prefix when the author was describing the purpose of the L-bit in Prefix Information Options. Chris's question prompted the thinking behind and writing of this memo.

Review and comments were provided by YOUR NAME HERE!

This memo was prepared using the xml2rfc tool.

10. Change Log [RFC Editor please remove]

draft-smith-6man-link-locals-off-link-00, initial version, 2015-08-16

11. References

11.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <http://www.rfc-editor.org/info/rfc2119>.

11.2. Informative References

[RFC4007] Deering, S., Haberman, B., Jinmei, T., Nordmark, E., and B. Zill, "IPv6 Scoped Address Architecture", RFC 4007, DOI 10.17487/RFC4007, March 2005, <http://www.rfc-editor.org/info/rfc4007>.

[RFC4291] Hinden, R. and S. Deering, "IP Version 6 Addressing Architecture", RFC 4291, DOI 10.17487/RFC4291, February 2006, <http://www.rfc-editor.org/info/rfc4291>.

[RFC4311] Hinden, R. and D. Thaler, "IPv6 Host-to-Router Load Sharing", RFC 4311, DOI 10.17487/RFC4311, November 2005, <http://www.rfc-editor.org/info/rfc4311>.

[RFC4541] Christensen, M., Kimball, K., and F. Solensky, "Considerations for Internet Group Management Protocol (IGMP) and Multicast Listener Discovery (MLD) Snooping Switches", RFC 4541, DOI 10.17487/RFC4541, May 2006, <http://www.rfc-editor.org/info/rfc4541>.

[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>.

[RFC5517] HomChaudhuri, S. and M. Foschiano, "Cisco Systems' Private VLANs: Scalable Security in a Multi-Client Environment", RFC 5517, DOI 10.17487/RFC5517, February 2010, <http://www.rfc-editor.org/info/rfc5517>.

[RFC5942] Singh, H., Beebee, W., and E. Nordmark, "IPv6 Subnet Model: The Relationship between Links and Subnet Prefixes", RFC 5942, DOI 10.17487/RFC5942, July 2010, <http://www.rfc-editor.org/info/rfc5942>.

[RFC6724] Thaler, D., Ed., Draves, R., Matsumoto, A., and T. Chown, "Default Address Selection for Internet Protocol Version 6 (IPv6)", RFC 6724, DOI 10.17487/RFC6724, September 2012, <http://www.rfc-editor.org/info/rfc6724>.

[RFC6957] Costa, F., Combes, J-M., Ed., Pougnard, X., and H. Li, "Duplicate Address Detection Proxy", RFC 6957, DOI 10.17487/RFC6957, June 2013, <http://www.rfc-editor.org/info/rfc6957>.

Author's Address

Mark Smith
In My Own Time
PO BOX 521
HEIDELBERG, VIC 3084
AU

Email: markzzzsmith@gmail.com