Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: October 16, 2017

T. Lemon
Nominum, Inc.
April 14, 2017

The Smear64 Mechanism
draft-lemon-smear64-00

Abstract

This document describes a way of sharing a /64 prefix among multiple links in a leaf network scenario, such as a home network or small office network. This provides a way to provide global addressing to all nodes on such a network, allowing for end-to-end communication without address translation.

Status of This Memo

Copyright Notice

Table of Contents

1. Introduction

The smear64 protocol provides a means for addressing the use case where an ISP provides only a single /64 to a CE router. In this case, if the network behind the router has more than one link, there are only two ways to provide addressing on the local network:

- o Advertise the /64 on every link and provide a means for preventing duplicate addresses across links, and a means for routing between hosts on the same prefix but on different links.

- o Number the local network with a ULA and use NAT66 to provide addressing.

The NAT66 solution feels attractive in this case because it seems simple, but in practice the two solutions are not really very different. The requirements for the NAT66 solution are as follows:

Generate and maintain a ULA for the CP network.

Provide a means for allocating prefixes for each link on the CP network.

Discover devices that wish to communicate externally by noticing packets from those devices routed to global addresses.

Maintain a table of such devices, mapping between ULAs and addresses allocated from the ISP-provided /64.

Maintain an in-kernel translator on the router to translate between ULAs and GUAs.

To solve the problem of sharing a /64 across multiple links, the following need to be done:

Announce the shared /64 on all links, with the L bit (See section 4.6.2 of [RFC4861]) clear.

Each link has a single router responsible for advertising the shared /64 on that link, so that if two routers are connected to the same link, only one announces the prefix. That router will also be responsible for detecting nodes using addresses on the /64

on that link, and maintaining state in the global neighbor table for that link.

Every router must listen for router solicitation messages. When a router solicitation message comes from an address on the advertised /64 that has not been seen before, that address, and the link on which it was received, is recorded in a global neighbor table, which is quickly propagated amongs all routers on the CP network.

Every router must listen for neighbor solicitation messages. When a neighbor solicitation message is received for an address which appears in the global neighbor table, the router checks to see if the two addresses are on the same link. If they are not, the router responds to the neighbor solicitation message.

In addition, when a neighbor solicitation comes from a source address on the shared /64 that is not in the global neighbor table, it is added to the global neighbor table.

Addresses in the global neighbor table are maintained as described in section 7.3 of [RFC4861]. Unreachability detection is performed by the router responsible for doing unreachability detection on the link to which the node had been communicating.

When a router receives a message with a destination address on the shared /64, it consults the global neighbor table to deliver it. If the source and destination addresses are on the same link, the router sends a redirect to the source of the message, as well as [instead of?] forwarding it.

The difference between these two solutions is that one requires a protocol for maintaining the global neighbor table. This can be done using HNCP [RFC7788]. HNCP can also be used to elect the router responsible for doing node discovery on that link.

2. Normative References

[RFC7788] Stenberg, M., Barth, S., and P. Pfister, "Home Networking Control Protocol", RFC 7788, DOI 10.17487/RFC7788, April 2016, <http://www.rfc-editor.org/info/rfc7788>.

[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>.

Author's Address

Ted Lemon
Nominum, Inc.
800 Bridge Parkway
Redwood City, California 94065
United States of America

Phone: +1 650 381 6000
Email: ted.lemon@nominum.com