MANET H. Rogge
Internet-Draft Fraunhofer FKIE
Intended status: Experimental August 11, 2014
Expires: February 12, 2015

Optimized flooding for NHDP Dual Stack routers
draft-rogge-manet-nhdp-dualstack-optimization-01

Abstract

This document specifies an optimization for the flooding of RFC5444 control traffic with NHDP in dualstack deployments.

Table of Contents

1. Introduction

While a lot of MANETs have been running as pure IPv4 networks in the past, networks with both IPv4 and IPv6 support are much more important today.

It is possible to run instances of OLSRv2 and NHDP on the same router, but this introduces unnecessary overhead to the network. This document describes a way to reduce the overhead of a Dual Stack MANET while keeping backward compatibility to MANET routers without this capability, including routers that run two separated instances of the routing protocol for both IPv4 and IPv6.

2. Terminology

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

The terminology introduced in [RFC5444], [RFC7181] and [RFC6130], including the terms "packet", "message" and "TLV" are to be interpreted as described therein.

Additionally, this document uses the following terminology and notational conventions:

RFC5444 IPv4 packet - a RFC5444 packet that is transported within an IPv4 UDP packet.

RFC5444 IPv6 packet - a RFC5444 packet that is transported within an IPv6 UDP packet.

3. Applicability Statement

The Dual Stack optimization described in this document is applicable for all combined IPv4 and IPv6 deployments of RFC5444 based routing protocols that share a single combined implementation for both IP address types. It is also applicable for deployments of IPv4 and IPv6 implementations on the same router that can communication between each other over a local connection.

4. Dual Stack Optimization Rational

RFC5444 based routing protocols can aggregate messages in UDP packets to reduce the number of media access and the overhead introduced by IP and MAC header.

This specification allows routers to aggregate messages with different address length (e.g. IPv4 and IPv6 based messages) in a single UDP packet, which allows for a further reduction of the number of media access and overhead.

5. Dual Stack Optimization Functioning & Overview

This specification uses an additional TLV inside the IPv6 [RFC6130] HELLO messages to signal the corresponding IPv4 originator address of the same router. This allows the router to determine which neighbors are Dual Stack capable and which IPv4/IPv6 originator address pair belong to each other.

Whenever a [RFC5444] message with Hoplimit field larger than 1 is created or forwarded, the router counts the number of IPv4-only, IPv6-only and Dual Stack neighbors in its Link Set Tuples on each interface. If the interface has at least one IPv4-only neighbor, all IPv4 messages must be forwarded in RFC5444 IPv4 packets. If the interface has at least one IPv6-only neighbor, all IPv6 messages must be forwarded in RFC5444 IPv6 packets. Other messages can be

forwarded in any UDP packet, the protocol prefers IPv4 UDP packets because of the lower IP header overhead.

6. Data Structures

This specification extends the Link Set Tuples of the Interface Information Base, as defined in [RFC6130] section 7.1, by the following additional elements for each link tuple when being used with this metric:

L_DualStack_Originator is the originator address of the Dual Stack partner router instance of the link.

This field is only used for IPv6 Link Set Tuples.

6.1. Initial Values

When generating a new tuple in the Link Set, as defined in [RFC6130] section 12.5 bullet 3, the values of the elements specified in Section 6 are set as follows:

- o L_DualStack_Originator := UNDEFINED.

7. Packets and Messages

7.1. Definitions

For the purpose of this section, note the following definitions:

- o IPv4ONLY(if): number of IPv4 Link Set Tuples for the interface "if" that have no IPv6 Link Set Tuple for the same interface with L_DualStack_Originator set to their IPv4 originator address.

- o IPv6ONLY(if): number of IPv6 Link Set Tuples for the interface "if" that have no IPv4 Link Set Tuple for the same interface with the IPv4 originator address set to the L_DualStack_Originator element.

- o DUALSTACK(if): number of IPv6 Link Set Tuples for the interface "if" that have a IPv4 Link Set Tuples for the same interface with the IPv4 originator address set to the L_DualStack_Originator element.

- o hoplimit: the value of the Hop Limit header field of a RFC5444 messsage (as defined in [RFC5444] Section 5.2), UNDEFINED if the message has no Hop Limit field.

7.2. Requirements

This protocol requires the router to be able to receive and process incoming [RFC5444] messages both with address length 4 and 16, regardless of the IP address family of the UDP packet.

[RFC5444] messages that have no Hoplimit field or a Hoplimit field with value 1, e.g. [RFC6130] HELLO messages are never sent in RFC5444 packets within UDP packets which don't match the address length of the message.

This specification also requires [RFC6130] HELLO messages with an unique originator address, e.g. as described in [RFC7181].

7.3. NHDP message generation

For each generated [RFC6130] HELLO message with address length 16, the following steps have to be followed:

1. Add a message TLV of type ADDRESS with type extension ADDR_ORIGINATOR and length 4 to the HELLO message with the IPv4 originator address of the local router as the value.

7.4. NHDP message processing

For each incoming [RFC6130] HELLO message with an address length of 16 (IPv6), additional processing MUST be carried out after the packet messages have been processed as specified in [RFC6130] and [RFC7181].

The router MUST update the Link Set Tuple corresponding to the originator of the packet:

- o If the message contains an ADDRESS TLV with type extension ORIGINATOR and length 4:
  - * L_DualStack_Originator := tlvvalue.
- o Otherwise:
  - * L_DualStack_Originator := UNDEFINED.

7.5. Dualstack RFC5444 Message Aggregation

The following process decides if a RFC5444 message should be sent within an IPv4 or IPv6 RFC5444 packet on an interface. Each message is only sent once on an interface.

For each [RFC5444] IPv4 message that is ready to be put into a RFC5444 packet on the interface 'if', the following steps need to be followed:

o If hoplimit == UNDEFINED or hoplimit == 1 or DUALSTACK(if) == 0 or IPv4ONLY(if) != 0 or IPv6ONLY(if) == 0

   * Put the message into a RFC5444 IPv4 packet.

   Otherwise

   * Put the message into a RFC5444 IPv6 packet.

For each [RFC5444] message with address length 16 (IPv6) that is ready to be put into a RFC5444 packet on the interface 'if', the following steps need to be followed:

o If hoplimit == UNDEFINED or hoplimit == 1 or DUALSTACK(if) == 0 or IPv6ONLY(if) != 0

   * Put the message into an RFC5444 IPv6 packet.

   Otherwise

   * Put the message into an RFC5444 IPv4 packet.

7.6. TLVs

This specification defines one Message TLV.

Note that, following [RFC5444] and network byte order, bits in an octet are numbered from 0 (most significant) to 7 (least significant).

7.6.1. Message TLVs

The ADDRESS TLV is used in [RFC5444] messages to transport addresses with a different address length than the message address block.

Type	Type Extension	Value
ADDRESS	ADDR_ORIGINATOR (0)	Originator Address of a router.

Table 1: ADDRESS TLV Definition

8. IANA Considerations

This specification defines one Message TLV Type, which have been allocated from the "Message TLV Types" registry of [RFC5444].

8.1. Expert Review: Evaluation Guidelines

For the registries where an Expert Review is required, the designated expert SHOULD take the same general recommendations into consideration as are specified by [RFC5444].

8.2. Message TLV Types

This specification defines one Message TLV Type, which has been allocated from the "Message TLV Types" namespace defined in [RFC5444]. IANA has made allocations in the 0-127 range for this type. The new Type Extension registries have been created with assignment as specified in Table 2.

Name	Type	Type Extension	Description	Allocation Policy
ADDRESS	TBD	ADDR_ORIGINATOR (0)	Originator address of a router.	
ADDRESS	TBD	1-255	unassigned	Expert Review

Table 2: Message TLV Type Assignment: ADDRESS TLV

Type extensions indicated as Expert Review SHOULD be allocated as described in [RFC5444], based on Expert Review as defined in [RFC5226].

9. Security Considerations

[RFC6130] HELLO messages with address length 16 could announce an IPv4 originator address that does belong to a different router, which could lead to database inconsistencies. A router implementing this specification might want to include consistency checks so that the mapping between IPv4 and IPv6 Link Set Tuples is strictly one-to-one.

10. Acknowledgements

This effort/activity is supported by the European Community Framework Program 7 within the Future Internet Research and Experimentation Initiative (FIRE), Community Networks Testbed for the Future Internet ([CONFINE]), contract FP7-288535.

11. References

11.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", RFC 2119, BCP 14, March 1997.

[RFC5226] Nartan, T. and H. Alverstrand, "Guidelines for Writing an IANA Considerations Section in RFCs", RFC 2119, BCP 14, May 2008.

[RFC5444] Clausen, T., Dearlove, C., Dean, J., and C. Adjih, "Generalized Mobile Ad Hoc Network (MANET) Packet/Message Format", RFC 5444, February 2009.

[RFC6130] Clausen, T., Dearlove, C., and J. Dean, "Mobile Ad Hoc Network (MANET) Neighborhood Discovery Protocol (NHDP)", RFC 6130, April 2011.

[RFC7181] Clausen, T., Jacquet, P., and C. Dearlove, "The Optimized Link State Routing Protocol version 2", RFC 7181, March 2013.

11.2. Informative References

[CONFINE] Braem, B., Blondia, C., Barz, C., Rogge, H., Freitag, F., Navarro, L., Bonicioli, J., Papathanasiou, S., Escrich, P., Vinas, R., Kaplan, A., Neumann, A., Balaguer, I., Tatum, B., and M. Matson, "A case for research with and on community networks", July 2013, <http://www.confine-project.eu>.

Appendix A. Dual Stack routing implementations

While traditional routing protocol implementations often handle IPv4 and IPv6 in completely separated instances or even programs, this optimization requires some coordination and communication between these two parts.

If both IPv4 and IPv4 are handled with the same executable, the implementation of this dual stack optimization should be easy to do.

The routing program needs two new multiplexer parts that allow generating and processing RFC5444 messages within RFC5444 packets of different address lengths, one in the RFC5444 parser and one in the RFC5444 packet aggregation. The multiplexer for outgoing messages needs access to both the IPv4 and IPv6 NHDP Link Set.

If IPv4 and IPv6 are handled in different programs the implementation will be more difficult. To implement this dual stack optimization, both programs would need to communicate over an internal connection, either a local network socket or a pipe. The network protocol running over this connection would need to allow sending RFC5444 messages between each instance and accessing each others Link Set database.

Author's Address

Henning Rogge
Fraunhofer FKIE
Fraunhofer Strasse 20
53343 Wachtberg
Germany

Phone: +49 228 9434 961
Email: henning.rogge@fkie.fraunhofer.de
URI: http://www.fkie.fraunhofer.de