

6LoWPAN
Internet-Draft
Intended status: Standards Track
Expires: August 10, 2008

P. Thubert
Cisco
February 7, 2008

LoWPAN simple fragment Recovery
draft-thubert-lowpan-simple-fragment-recovery-00

Status of this Memo

By submitting this Internet-Draft, each author represents that any applicable patent or other IPR claims of which he or she is aware have been or will be disclosed, and any of which he or she becomes aware will be disclosed, in accordance with [Section 6 of BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/lid-abstracts.txt>.

The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>.

This Internet-Draft will expire on August 10, 2008.

Copyright Notice

Copyright (C) The IETF Trust (2008).

Abstract

Considering that 6LoWPAN packets can be as large as 2K bytes and that an 802.15.4 frame with security will carry in the order of 80 bytes of effective payload, a packet might end up fragmented into as many as 25 fragments at the 6LoWPAN shim layer. If a single one of those fragments is lost in transmission, all fragments must be resent, further contributing to the congestion that might have caused the initial packet loss. This draft introduces a simple protocol to recover individual fragments between 6LoWPAN endpoints.

Internet-Draft

LoWPAN simple fragment Recovery

February 2008

Table of Contents

| | | |
|-----------------------|---|--------------------|
| 1. | Introduction | 3 |
| 2. | Terminology | 3 |
| 3. | Rationale | 4 |
| 4. | Requirements | 4 |
| 5. | Overview | 5 |
| 6. | New Dispatch types and headers | 6 |
| 6.1. | Recoverable Fragment Dispatch type and Header | 7 |
| 6.2. | Fragment Acknowledgement Dispatch type and Header | 7 |
| 7. | Security Considerations | 8 |
| 8. | IANA Considerations | 8 |
| 9. | Acknowledgments | 8 |
| 10. | References | 8 |
| 10.1. | Normative References | 8 |
| 10.2. | Informative References | 8 |
| | Author's Address | 9 |
| | Intellectual Property and Copyright Statements | 10 |

1. Introduction

Considering that 6LoWPAN packets can be as large as 2K bytes and that a 802.15.4 frame with security will carry in the order of 80 bytes of effective payload, a packet might be fragmented into about 25 fragments at the 6LoWPAN shim layer. This level of fragmentation is much higher than that traditionally experienced over the Internet with IPv4 fragments. At the same time, the use of radios increases the probability of transmission loss and Mesh-Under techniques compound that risk over multiple hops.

Past experience with fragmentation has shown that missassociated or lost fragments can lead to poor network behaviour and, eventually, trouble at application layer. The reader might start his research from [[I-D.mathis-frag-harmful](#)] and follow the references. That experience led to the definition of the Path MTU discovery [[RFC1191](#)] protocol that avoids fragmentation over the Internet.

An end-to-end fragment recovery mechanism might be a good complement to a hop-by-hop MAC level recovery with a limited number of retries. This draft introduces a simple protocol to recover individual fragments between 6LoWPAN endpoints. Specifically in the case of UDP, valuable additional information can be found in UDP Usage Guidelines for Application Designers [[I-D.ietf-tsvwg-udp-guidelines](#)].

2. Terminology

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 [[RFC2119](#)].

Readers are expected to be familiar with all the terms and concepts that are discussed in "IPv6 over Low-Power Wireless Personal Area Networks (6LoWPANs): Overview, Assumptions, Problem Statement, and Goals" [[RFC4919](#)] and "Transmission of IPv6 Packets over IEEE 802.15.4

Networks" [[RFC4944](#)].

ERP

Error Recovery Procedure.

LoWPAN endpoints

The LoWPAN nodes in charge of generating or expanding a 6LoWPAN header from/to a full IPv6 packet. The LoWPAN endpoints are the points where fragmentation and reassembly take place.

Thubert

Expires August 10, 2008

[Page 3]

Internet-Draft

LoWPAN simple fragment Recovery

February 2008

[3.](#) Rationale

There are a number of usages for large packets in Wireless Sensor Networks. Such usages may not be the most typical or represent the largest amount of traffic over the LoWPAN; however, the associated functionality can be critical enough to justify extra care for ensuring effective transport of large packets across the LoWPAN.

The list of those usages includes:

Towards the LoWPAN node:

Packages of Commands: A number of commands or a full configuration can be packaged as a single message to ensure consistency and enable atomic execution or complete roll back. Until such commands are fully received and interpreted, the intended operation will not take effect.

Firmware update: For example, a new version of the LoWPAN node software is downloaded from a system manager over unicast or multicast services. Such a reflashing operation typically involves updating a large number of similar 6LoWPAN nodes over a relatively short period of time.

From the LoWPAN node:

Waveform captures: A number of consecutive samples are measured at a high rate for a short time and then transferred from a sensor to a gateway or an edge server as a single large report.

Large data packets: Rich data types might require more than one fragment.

Uncontrolled firmware download or waveform upload can easily result in a massive increase of the traffic and saturate the network. When a fragment is lost in transmission, all fragments are resent, further contributing to the congestion that caused the initial loss, and potentially leading to congestion collapse.

This saturation may lead to excessive radio interference, or random early discard (leaky bucket) in relaying nodes. Additional queueing and memory congestion may result while waiting for a low power next hop to emerge from its sleeping state.

[4.](#) Requirements

This paper proposes a method to recover individual fragments between

Thubert

Expires August 10, 2008

[Page 4]

Internet-Draft

LoWPAN simple fragment Recovery

February 2008

LoWPAN endpoints. The method is designed to fit the following requirements of a LoWPAN (with or without a Mesh-Under routing protocol):

Controlled latency

The ERP mechanism must succeed or give up within the time boundary imposed by the recovery process of the Upper Layer Protocols.

Minimum acknowledgement overhead

Because the radio is inherently half duplex, an acknowledgement consumes roughly as many resources as the fragment itself.

Support for out-of-order fragment delivery

A Mesh-Under load balancing mechanism such as the ISA100 Data Link Layer can introduce out-of-sequence packets. The recovery mechanism must account for packets that appear lost but are actually only delayed over a different path.

Optional flow control

The aggregation of multiple concurrent flows may lead to the saturation of the radio network and congestion collapse.

Backward compatibility

A node that implements this draft should be able to communicate with a node that implements [[RFC4944](#)]. The current draft assumes that compatibility information about the remote LoWPAN endpoint is obtained by external means.

[5.](#) Overview

The fragmentation/reassembly of a packet must complete within an acceptable overall latency, otherwise the packet expires and must be dropped. This latency must be smaller than Upper Layer Protocol retry values, and smaller than expiration period of the information transported.

The sender transfers a controlled number of fragments (possibly all of them) and flags the last fragment of a series with an Acknowledgement request.

The sender sets a retry timer for the fragment that carries the Acknowledgement request. That fragment is retransmitted individually

upon time out. This is repeated until an Acknowledgement comes back or the packet expires.

Upon receipt of an Acknowledgement request, the receiver responds with an Acknowledgement containing a bitmap that indicates which fragments were actually received. The bitmap is a 32bit DWORD, which accommodates up to 32 fragments and is sufficient for the 6LoWPAN MTU. For all n in $[0..31]$, bit n is set to 1 in the bitmap to indicate that fragment n was received, otherwise the bit is set to 0. If any fragment immediately preceding the acknowledgement request is missing, the receiver MAY intentionally delay its response to allow in-transit fragments to arrive.

The sender has either one or no Acknowledgement pending. An Acknowledgement that is not expected or does not acknowledge the

pending sequence in the bitmap is a duplicate and is ignored.

When a valid Acknowledgement is received, the sender resumes sending fragments and the process is repeated until all fragments are acknowledged or the packet expires.

Fragments are sent in a round robin fashion: the sender sends all the fragments for a first time before it retries any lost fragment; lost fragments are retried in sequence, oldest first. This mechanism enables the receiver to acknowledge fragments that were delayed in the network before they are actually retried.

It is up to the sender to decide how many fragments are (re)sent before an acknowledgement is received, and the sender can adapt that number to the network conditions. This way, the number of outstanding fragments can be used as a flow control mechanism to protect the network.

6. New Dispatch types and headers

This specification extends "Transmission of IPv6 Packets over IEEE 802.15.4 Networks" [[RFC4944](#)] with 3 new dispatch types, for Recoverable Fragments (RFRAG) headers with or without Acknowledgement Request, and for the Acknowledgement back.

| Pattern | Header Type |
|-----------|---|
| 11 101000 | RFRAG - Recoverable Fragment |
| 11 101001 | RFRAG-AR - RFRAG with Acknowledgement Req |
| 11 101010 | RFRAG-ACK - RFRAG Acknowledgement |

Figure 1: Additional Dispatch Value Bit Patterns

In the following sections, the semantics of "datagram_tag," "datagram_offset" and "datagram_size" and the reassembly process are unchanged from [[RFC4944](#)] [Section 5.3](#). "Fragmentation Type and

Header."

6.1. Recoverable Fragment Dispatch type and Header




Figure 2: Recoverable Fragment Dispatch type and Header

X bit

When set, the sender requires an Acknowledgement from the receiver

Sequence

The sequence number of the fragment. Fragments are numbered [0..N] where N is in [0..31].

6.2. Fragment Acknowledgement Dispatch type and Header

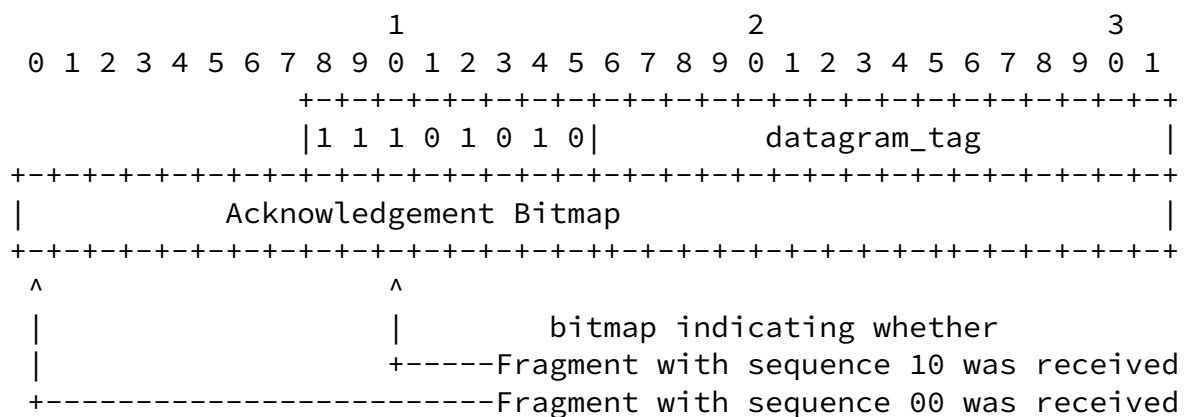


Figure 3: Fragment Acknowledgement Dispatch type and Header

Each bit in the Bitmap refers to a particular fragment: bit n set indicates that fragment with sequence n was received, for n in [0..31]. All zeroes means that the fragment was dropped because it corresponds to an obsolete datagram_tag. This happens if the packet was already reassembled and passed to the network upper layer, or the packet expired and was dropped.

7. Security Considerations

The process of recovering fragments does not appear to create any opening for new threat.

8. IANA Considerations

Need extensions for formats defined in "Transmission of IPv6 Packets over IEEE 802.15.4 Networks" [[RFC4944](#)]. ? Is that IANA ?.

9. Acknowledgments

The author wishes to thank Jay Werb, Christos Polyzois, Soumitri Kolavennu and Harry Courtice for their contribution and review.

10. References

10.1. Normative References

- [RFC1191] Mogul, J. and S. Deering, "Path MTU discovery", [RFC 1191](#), November 1990.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [RFC4944] Montenegro, G., Kushalnagar, N., Hui, J., and D. Culler, "Transmission of IPv6 Packets over IEEE 802.15.4 Networks", [RFC 4944](#), September 2007.

10.2. Informative References

- [I-D.ietf-tsvwg-udp-guidelines]
Eggert, L. and G. Fairhurst, "UDP Usage Guidelines for Application Designers", [draft-ietf-tsvwg-udp-guidelines-04](#) (work in progress), November 2007.

[I-D.mathis-frag-harmful]

Mathis, M., "Fragmentation Considered Very Harmful",
[draft-mathis-frag-harmful-00](#) (work in progress),
July 2004.

[RFC4919] Kushalnagar, N., Montenegro, G., and C. Schumacher, "IPv6
over Low-Power Wireless Personal Area Networks (6LoWPANs):
Overview, Assumptions, Problem Statement, and Goals",
[RFC 4919](#), August 2007.

Author's Address

Pascal Thubert
Cisco Systems
Village d'Entreprises Green Side
400, Avenue de Roumanille
Batiment T3
Biot - Sophia Antipolis 06410
FRANCE

Phone: +33 4 97 23 26 34
Email: pthubert@cisco.com

Internet-Draft

LoWPAN simple fragment Recovery

February 2008

Full Copyright Statement

Copyright (C) The IETF Trust (2008).

This document is subject to the rights, licenses and restrictions contained in [BCP 78](#), and except as set forth therein, the authors retain all their rights.

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY, THE IETF TRUST AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Intellectual Property

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in [BCP 78](#) and [BCP 79](#).

Copies of IPR disclosures made to the IETF Secretariat and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the IETF on-line IPR repository at <http://www.ietf.org/ipr>.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at

ietf-ipr@ietf.org.

Acknowledgment

Funding for the RFC Editor function is provided by the IETF Administrative Support Activity (IASA).

Thubert

Expires August 10, 2008

[Page 10]