

6LoWPAN
Internet-Draft
Intended status: Standards Track
Expires: September 24, 2009

P. Thubert, Ed.
Cisco
March 23, 2009

LoWPAN simple fragment Recovery
draft-thubert-6lowpan-simple-fragment-recovery-03

Status of this Memo

This Internet-Draft is submitted to IETF in full conformance with the provisions of [BCP 78](#) and [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/1id-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 September 24, 2009.

Copyright Notice

Copyright (c) 2009 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](#) and the IETF Trust's Legal Provisions Relating to IETF Documents in effect on the date of publication of this document (<http://trustee.ietf.org/license-info>). Please review these documents carefully, as they describe your rights and restrictions with respect to this document.

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.

Table of Contents

1.	Introduction	3
2.	Terminology	3
3.	Rationale	4
4.	Requirements	5
5.	Overview	6
6.	New Dispatch types and headers	7
6.1.	Recoverable Fragment Dispatch type and Header	7
6.2.	Fragment Acknowledgement Dispatch type and Header	8
7.	Outstanding Fragments Control	8
8.	Security Considerations	10
9.	IANA Considerations	10
10.	Acknowledgments	10
11.	References	10
11.1.	Normative References	10
11.2.	Informative References	10
	Author's Address	11

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 is encouraged to read [[RFC4963](#)] and follow the references for more information. That experience led to the definition of the Path MTU discovery [[RFC1191](#)] protocol that limits 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 September 24, 2009

[Page 3]

Internet-Draft

LoWPAN simple fragment Recovery

March 2009

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

Number of fragments

The recovery mechanism must support highly fragmented packets, with a maximum of 32 fragments per packet.

Minimum acknowledgement overhead

Because the radio is half duplex, and because of silent time spent in the various medium access mechanisms, an acknowledgement consumes roughly as many resources as data fragment.

The recovery mechanism should be able to acknowledge multiple fragments in a single message.

Controlled latency

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

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 congestion control

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

The recovery mechanism should provide means for controlling the number of fragments in transit over the LoWPAN.

Backward compatibility

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

[5.](#) Overview

Considering that a multi-hop LoWPAN can be a very sensitive environment due to the limited queueing capabilities of a large population of its nodes, this draft recommends a simple and conservative approach to congestion control, based on TCP congestion avoidance.

Congestion on the forward path is assumed in case of packet loss, and packet loss is assumed upon time out.

Congestion on the forward path can also be indicated by an Explicit Congestion Notification (ECN) mechanism. Though whether and how ECN [[RFC3168](#)] is carried out over the LoWPAN is out of scope, this draft provides a way for the destination endpoint to echo an ECN indication

back to the source endpoint in an acknowledgement message as represented in Figure 3 in [Section 6.2](#).

From the standpoint of a source LoWPAN endpoint, an outstanding fragment is a fragment that was sent but for which no explicit acknowledgement was received yet. This means that the fragment might be on the way, received but not yet acknowledged, or the acknowledgement might be on the way back. It is also possible that either the fragment or the acknowledgement was lost on the way.

Because a meshed LoWPAN might deliver frames out of order, it is virtually impossible to differentiate these situations. In other words, from the sender standpoint, all outstanding fragments might still be in the network and contribute to its congestion. There is an assumption, though, that after a certain amount of time, a frame is either received or lost, so it is not causing congestion anymore. This amount of time can be estimated based on the round trip delay between the LoWPAN endpoints. The method detailed in [\[RFC2988\]](#) is recommended for that computation.

The reader is encouraged to read through "Congestion Control Principles" [\[RFC2914\]](#). Additionally [\[RFC2309\]](#) and [\[RFC2581\]](#) provide deeper information on why this mechanism is needed and how TCP handles Congestion Control. Basically, the goal here is to manage the amount of fragments present in the network; this is achieved by to reducing the number of outstanding fragments over a congested path by throttling the sources.

[Section 7](#) describes how the sender decides how many fragments are (re)sent before an acknowledgement is required, and how the sender adapts that number to the network conditions.

[6.](#) New Dispatch types and headers

This specification extends "Transmission of IPv6 Packets over IEEE 802.15.4 Networks" [\[RFC4944\]](#) with 4 new dispatch types, for Recoverable Fragments (RFRAG) headers with or without Acknowledgement Request, and for the Acknowledgement back, with or without ECN Echo.

Pattern	Header Type
11 101000	RFRAG - Recoverable Fragment
11 101001	RFRAG-AR - RFRAG with Ack Request
11 101010	RFRAG-ACK - RFRAG Acknowledgement
11 101011	RFRAG-AEC - RFRAG Ack with ECN Echo

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

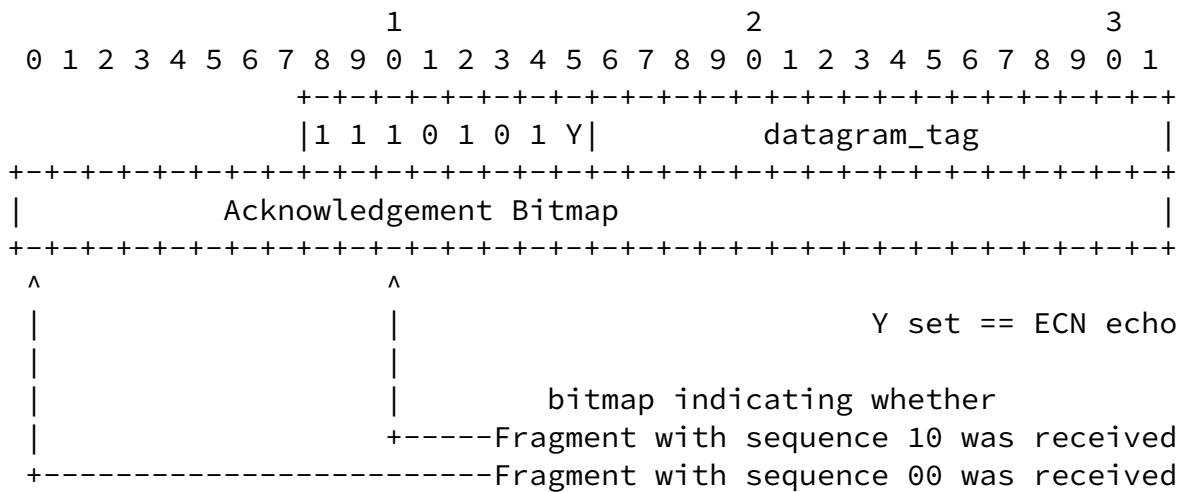


Figure 3: Fragment Acknowledgement Dispatch type and Header

Y bit

When set, the sender indicates that at least one of the acknowledged fragments was received with an Explicit Congestion Notification, indicating that the path followed by the fragments is subject to congestion.

Acknowledgement Bitmap

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. Outstanding Fragments Control

A mechanism based on TCP congestion avoidance dictates the maximum number of outstanding fragments.

The maximum number of outstanding fragments for a given packet toward a given LoWPAN endpoint is initially set to a configured value,

unless recent history indicates otherwise.

Each time that maximum number of fragments is fully acknowledged, that number can be incremented by 1. ECN echo and packet loss cause the number to be divided by 2.

The sender transfers a controlled number of fragments and flags the last fragment of a series with an acknowledgement request.

The sender arms a timer to cover the fragment that carries the Acknowledgement request. Upon time out, the sender assumes that all the fragments on the way are received or lost. It divides the maximum number of outstanding fragments by 2 and resets the number of outstanding fragments to 0.

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 with sequence n was received, otherwise the bit is set to 0.

The receiver MAY issue unsolicited acknowledgements. An unsolicited acknowledgement enables the sender endpoint to resume sending if it had reached its maximum number of outstanding fragments. Note that acknowledgements might consume precious resources so the use of unsolicited acknowledgements should be configurable and not enabled by default.

The receiver MUST acknowledge a fragment with the acknowledgement request bit set. If any fragment immediately preceding an acknowledgement request is still missing, the receiver MAY intentionally delay its acknowledgement to allow in-transit fragments to arrive. This mechanism might defeat the round trip delay computation so it should be configurable and not enabled by default.

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.

The process must complete within an acceptable time that is within the boundaries of upper layer retries. Additional work is required to define how this is achieved. When the source endpoint decides

that a packet should be dropped and the fragmentation process cancelled, it sends a pseudo fragment with the datagram_offset,

sequence and datagram_size all set to zero, and no data. Upon reception of this message, the receiver should clean up all resources for the packet associated to the datagram_tag.

8. Security Considerations

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

9. IANA Considerations

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

10. Acknowledgments

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

11. References

11.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [RFC2988] Paxson, V. and M. Allman, "Computing TCP's Retransmission Timer", [RFC 2988](#), November 2000.
- [RFC4944] Montenegro, G., Kushalnagar, N., Hui, J., and D. Culler, "Transmission of IPv6 Packets over IEEE 802.15.4 Networks", [RFC 4944](#), September 2007.

11.2. Informative References

[I-D.ietf-tsvwg-udp-guidelines]

Eggert, L. and G. Fairhurst, "Unicast UDP Usage Guidelines for Application Designers",
[draft-ietf-tsvwg-udp-guidelines-11](#) (work in progress),
October 2008.

[I-D.mathis-frag-harmful]

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

Thubert

Expires September 24, 2009

[Page 10]

Internet-Draft

LoWPAN simple fragment Recovery

March 2009

July 2004.

[RFC1191] Mogul, J. and S. Deering, "Path MTU discovery", [RFC 1191](#),
November 1990.

[RFC2309] Braden, B., Clark, D., Crowcroft, J., Davie, B., Deering, S., Estrin, D., Floyd, S., Jacobson, V., Minshall, G., Partridge, C., Peterson, L., Ramakrishnan, K., Shenker, S., Wroclawski, J., and L. Zhang, "Recommendations on Queue Management and Congestion Avoidance in the Internet", [RFC 2309](#), April 1998.

[RFC2581] Allman, M., Paxson, V., and W. Stevens, "TCP Congestion Control", [RFC 2581](#), April 1999.

[RFC2914] Floyd, S., "Congestion Control Principles", [BCP 41](#),
[RFC 2914](#), September 2000.

[RFC3168] Ramakrishnan, K., Floyd, S., and D. Black, "The Addition of Explicit Congestion Notification (ECN) to IP",
[RFC 3168](#), September 2001.

[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.

[RFC4963] Heffner, J., Mathis, M., and B. Chandler, "IPv4 Reassembly Errors at High Data Rates", [RFC 4963](#), July 2007.

Author's Address

Pascal Thubert (editor)
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