Internet Engineering Task Force
INTERNET-DRAFT
draft-kohler-dccp-ccid3-drops-00.ps
Expires: 30 April 2007

E. Kohler
UCLA

30 October 2006

Datagram Congestion Control Protocol (DCCP) Congestion Control ID 3 Dropped Packets Option

Status of This Memo

Abstract

This document describes the Dropped Packets option, a mechanism for reporting the number of lost and marked packets per loss interval in the Datagram Congestion Control Protocol (DCCP)'s Congestion Control ID 3, TCP-Friendly Rate Control. This option may be useful for applications that need to know precisely how many packets are being dropped.

Table of Contents

List of Tables

1. Introduction

The Datagram Congestion Control Protocol (DCCP) [RFC4340] allows the use of several distinct congestion control mechanisms. One of these, Congestion Control Identifier 3 [RFC4342], specifies the use of TCP-Friendly Rate Control (TFRC) [RFC3448]. The core information reported by CCID 3 receivers is a list of recent loss intervals, where a loss interval begins with a lost or ECN-marked data packet; continues with at most one round-trip time's worth of packets that may or may not be lost or marked; and completes with an arbitrarily long series of non-dropped, non-marked data packets. Loss intervals model the congestion behavior of TCP NewReno senders, which reduce their sending rate at most once per window of data packets. Consequently, the number of packets lost in a loss interval is not important for either TCP's or TFRC's congestion response. CCID 3's Loss Intervals option reports the length of each loss interval's lossy part, not the number of packets that were actually lost or marked in that lossy part.

Nevertheless, applications and experimental variants of TFRC, such as the Small Packet variant, may be interested in the number of packets lost or marked in a loss interval, over and above the length of the loss interval in packets. This document specifies the Dropped Packets option, a CCID 3-specific option that reports this information. Together with the existing Loss Intervals option, the Dropped Packets option allows CCID 3 senders to discover exactly how many packets were dropped from each loss interval.

The mechanisms in this document do not change existing CCID 3 congestion response behavior. CCID 3's congestion response still depends entirely on loss interval lengths, not the number of packets dropped per loss interval. Most CCID 3 senders will therefore ignore the contents of any Dropped Packets options they receive. Sending applications may, however, be interested in Dropped Packets information.

2. Conventions

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].

All multi-byte numerical quantities in CCID 3, such as arguments to options, are transmitted in network byte order (most significant byte first).

A DCCP half-connection consists of the application data sent by one endpoint and the corresponding acknowledgements sent by the other endpoint. The terms "HC-Sender" and "HC-Receiver" denote the endpoints sending application data and acknowledgements, respectively. Since CCIDs apply at the level of half-connections, we abbreviate HC-Sender to "sender" and HC-Receiver to "receiver" in this document. See [RFC4340] for more discussion.

For simplicity, we say that senders send DCCP-Data packets and receivers send DCCP-Ack packets. Both of these categories are meant to include DCCP-DataAck packets.

The phrases "ECN-marked" and "marked" refer to packets marked ECN Congestion Experienced unless otherwise noted.

3. Options and Features

This document defines a single CCID 3-specific option, Dropped Packets.

Type	Option Length	Meaning	DCCP-Data?	Section Reference
195	variable	Dropped Packets	N	3.1

Table 1: Additional DCCP CCID 3 Options

The "DCCP-Data?" column indicates that Dropped Packets MUST be ignored when it occurs on a DCCP-Data packet.

A CCID 3-specific feature governing the use of the Dropped Packets option is also defined.

Number	Meaning	Rec'n Rule	Initial Value	Req'd	Section Reference
195	Send Dropped Packets	SP	0	N	3.2

Table 2: Additional DCCP CCID 3 Feature Numbers

The column meanings are described in [RFC4340], Table 4. "Rec'n Rule" defines the feature's reconciliation rule, where "SP" means server-priority. "Req'd" specifies whether every CCID 3 implementation MUST understand a feature; Send Dropped Packets is optional, in that it behaves like an extension ([RFC4340], Section 15).

3.1. Dropped Packets Option

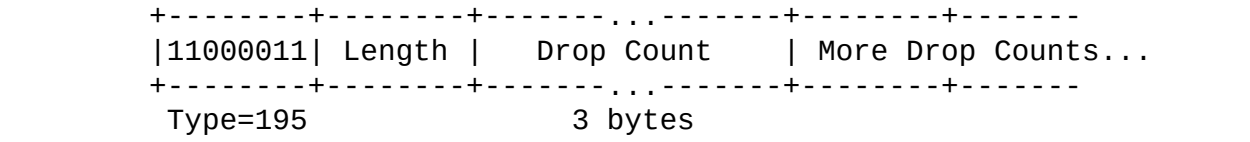

The receiver reports the number of lost or marked packets in its recently observed loss intervals using a Dropped Packets option.

The Dropped Packets option contains information about one to 84 consecutive loss intervals, always including the most recent loss interval. As with CCID 3's Loss Intervals option, intervals are listed in reverse chronological order. Should more than 84 loss intervals need to be reported, multiple Dropped Packets options can be sent; the second option begins where the first left off, and so forth.

One Drop Count is specified per loss interval. Drop Count is a 24-bit number that equals the number of packets lost or received ECN-marked during the corresponding loss interval. By definition, this number MUST NOT exceed the corresponding loss interval's Loss Length.

Dropped Packets options SHOULD be sent in tandem with corresponding Loss Intervals options. Consider a CCID 3 receiver that is reporting Dropped Packets information. When this receiver sends a feedback packet containing information about the N most recent loss intervals (packaged in one or more Loss Intervals options), it SHOULD include on the same feedback packet one or more Dropped Packets options covering exactly those N loss intervals. CCID 3 senders MUST ignore Drop Counts information for loss intervals not

covered by a Loss Intervals option on the same feedback packet. Conversely, a CCID 3 sender might want to interpolate Drop Counts information for a loss interval not covered by any Dropped Packets options; such a sender SHOULD use the corresponding loss interval's Loss Length as its Drop Count.

Each loss interval's Drop Count MUST by definition be less than or equal to its Loss Length. A Drop Count that exceeds the corresponding Loss Length MUST be ignored.

3.1.1. Example

Consider the following sequence of packets, where "−" represents a safely delivered packet and "*" represents a lost or marked packet. This sequence is repeated from [RFC4342].

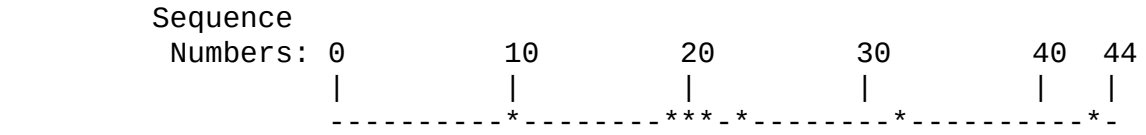

Assuming that packet 43 was lost, not marked, this sequence might be divided into loss intervals as follows:

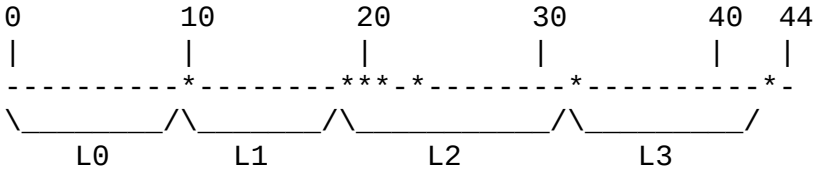

A Loss Intervals option sent on a packet with Acknowledgement Number 44 to acknowledge this set of loss intervals might contain the bytes 193,39,2, 0,0,10, 128,0,1, 0,0,10, 0,0,8, 0,0,5, 0,0,10, 0,0,8, 0,0,1, 0,0,8, 0,0,10, 128,0,0, 0,0,15; for interpretation of this option, see [RFC4342]. A Dropped Packets option sent in tandem on this packet would contain the bytes 195,14, 0,0,1, 0,0,4, 0,0,1, 0,0,0. This is interpreted as follows.

195 The Dropped Packets option number.

14 The length of the option, including option type and length bytes. This option contains information about (14 − 2)/3 = 4 loss intervals. Note that the two most recent sequence numbers are not yet part of any loss interval -- the Loss Intervals option includes them in its Skip Length -- and are thus not included in the Dropped Packets option.

0,0,1
: These bytes define the Drop Count for L3, which is 1. As required, the Drop Count is less than or equal to L3's Loss Length, which is also 1.

0,0,4
: The Drop Count for L2 is 4.

0,0,1
: The Drop Count for L1 is 1.

0,0,0
: Finally, the Drop Count for L0 is 0.

3.2. Send Dropped Packets Feature

The Send Dropped Packets feature lets CCID 3 endpoints negotiate whether the receiver MUST provide Dropped Packets options on its acknowledgements. DCCP A sends a "Change R(Send Dropped Packets, 1)" option to ask DCCP B to send Dropped Packets options as part of its acknowledgement traffic.

Send Dropped Packets has feature number 195 and is server-priority. It takes one-byte Boolean values. DCCP B MUST send Dropped Packets options on its acknowledgements when Send Dropped Packets/B is one, although it MAY send Dropped Packets options even when Send Dropped Packets/B is zero. Values of two or more are reserved. A CCID 3 half-connection starts with Send Dropped Packets equal to zero.

4. Security Considerations

The Dropped Packets option does not affect the existing security considerations for DCCP CCID 3, which have been discussed in [RFC4340] and [RFC4342]. For instance, the Dropped Packets option neither helps nor hinders the existing CCID 3 mechanisms for ECN Nonce verification.

5. IANA Considerations

This specification allocates two values in namespaces managed by IANA. Specifically, the value 195 is allocated from the DCCP CCID 3-specific option type registry for the Dropped Packets option (Table 1), and the value 195 is allocated from the DCCP CCID 3-specific feature number registry for the Send Dropped Packets feature (Table 2).

6. Thanks

Thanks to Sally Floyd for inspiring this document.

Normative References

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

[RFC3448] Handley, M., Floyd, S., Padhye, J., and J. Widmer, "TCP Friendly Rate Control (TFRC): Protocol Specification", RFC 3448, January 2003.

[RFC4340] Kohler, E., Handley, M., and S. Floyd, "Datagram Congestion Control Protocol (DCCP)", RFC 4340, March 2006.

[RFC4342] Floyd, S., E. Kohler, and J. Padhye, "Profile for Datagram Congestion Control Protocol (DCCP) Congestion Control ID 3: TCP-Friendly Rate Control (TFRC)", RFC 4342, March 2006.

Authors' Addresses

Eddie Kohler
4531C Boelter Hall
UCLA Computer Science Department
Los Angeles, CA 90095
USA

EMail: kohler@cs.ucla.edu

Full Copyright Statement

Intellectual Property

Acknowledgement

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