Internet Engineering Task Force
INTERNET-DRAFT
draft-ietf-dccp-ccid3-thin-01.txt
Expires: January 2005

Eddie Kohler
UCLA

18 July 2004

DCCP CCID 3-Thin

Status of this Memo

This document is an Internet-Draft.

By submitting this Internet-Draft, we certify that any applicable patent or other IPR claims of which we are aware have been disclosed, or will be disclosed, and any of which we become aware will be disclosed, in accordance with RFC 3668 (BCP 79).

By submitting this Internet-Draft, we accept the provisions of Section 3 of RFC 3667 (BCP 78).

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 a "work in progress."

The list of current Internet-Drafts can be accessed at http://www.ietf.org/1id-abstracts.html

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

Copyright Notice

Abstract

This document describes the Thin variant of the Datagram Congestion Control Protocol (DCCP) Congestion Control Identifier 3, TCP-

Friendly Rate Control (TFRC). The Thin variant is more restricted than CCID 3; it limits allowable options, acceptable feature values, and so forth. CCID 3-Thin packets are still valid DCCP CCID 3 packets. CCID 3-Thin was designed for small clients where a full DCCP implementation would be too expensive.

Table of Contents

1. Introduction

The Datagram Congestion Control Protocol (DCCP) [DCCP] implements a congestion-controlled stream of unreliable unicast datagrams. DCCP uses Congestion Control Identifiers, or CCIDs, to determine the congestion control mechanism in use on a half-connection. (A half-connection might consist of data packets sent from DCCP A to DCCP B, plus acknowledgements sent from DCCP B to DCCP A. DCCP A is the HC-Sender, and DCCP B the HC-Receiver, for this half-connection. In this document, I abbreviate HC-Sender and HC-Receiver as "sender" and "receiver", respectively. These terms are defined more fully in [DCCP].) DCCP CCID 3, TCP-Friendly Rate Control (TFRC) [CCID 3 PROFILE], defines a receiver-based congestion control mechanism that provides a TCP-friendly send rate, while minimizing abrupt rate changes [RFC 3448].

This document describes the Thin variant of DCCP CCID 3. All CCID 3-Thin packets are valid CCID 3 packets, but the converse is not true: the Thin variant restricts the options that may be sent, the values that features may take, and so forth. CCID 3-Thin was designed for small clients and servers where a full DCCP implementation would be too expensive.

Note that this version of this document is more a proof-of-concept than a final proposal. The right set of restrictions for CCID 3-Thin should be subject to further discussion. For example, a final CCID 3-Thin might support ECN.

This document assumes familiarity with both DCCP proper [DCCP] and the CCID 3 profile [CCID 3 PROFILE].

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 [RFC 2119].

3. The CCID 3-Thin Option

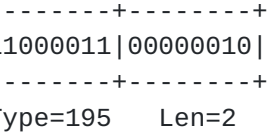

The CCID 3-Thin option is sent on the initial packets in a CCID 3 connection to indicate that the endpoints agree to be bound by the CCID 3-Thin restrictions. It is a CCID-specific option, which means that the CCID feature must be set to 3 before the CCID 3-Thin option

is processed; this may happen simply by sending CCID 3-Thin after Mandatory options that negotiate the CCID.

CCID 3-Thin options will generally be preceded by a Mandatory option (Type). This will force the receiving DCCP to reset the connection if it cannot abide by the CCID 3-Thin restrictions, below.

4. CCID 3-Thin Restrictions

A CCID 3-Thin connection agrees to follow these restrictions.

- Both half-connections use CCID 3 (their CCID features have value 3).
- All features are fixed to the values listed below, whether or not their values are negotiated. No feature negotiation options will be sent after the CCID 3-Thin handshake completes.
- Both Allow Short Seqnos features have value 1 (allow short sequence numbers).
- Both Sequence Window features have value 100 (the default).
- Both ECN Capable features have value 0 (ECN incapable). All packets are sent as ECN-incapable.
- Both Ack Ratio features have value 0 (Ack Ratio not in use).
- Both Send Ack Vector features have value 0 (do not send Ack Vectors; the default). No Ack Vector options will be sent. CCID 3-Thin applications that want information about which packets have arrived will have to send application-level acknowledgements.
- Both Send NDP Count features have value 0 (do not send NDP Count options; the default). No NDP Count options will be sent.
- Both Minimum Checksum Coverage features have value 0 (require full Checksum Coverage; the default). No packets with less-than-full Checksum Coverage will be sent.
- Both Check Data Checksum features have value 0 (do not necessarily check Data Checksum options; the default). No Data Checksum options will be sent.
- Both Send Loss Event Rate features have value 1 (send Loss Event Rate options; the default).

- Both Send Loss Intervals features have value 0 (do not send Loss Intervals options; the default). No Loss Intervals options will be sent.

- Only the following types of packets can be sent:

  1. DCCP-Request packets without data.

  2. DCCP-Response packets without data and without an Init Cookie option.

  3. DCCP-Data packets with no options.

  4a. DCCP-Ack packets with no options.

  4b. DCCP-Ack packets with 20 bytes of option. The options MUST be exactly as follows: Elapsed Time with length 4; Padding; Padding; Receive Rate with length 6; Padding; Padding; and Loss Event Rate with length 6.

  4c. A fat client responding to a thin server may follow a DCCP-Response with a DCCP-Ack that contains other options; see Section 5.4.

  5. DCCP-Close packets with no options.

  6. DCCP-Reset packets with no options.

  7. DCCP-Sync packets with no options.

  8. DCCP-SyncAck packets with no options.

- Packets that would normally be reported with Data Dropped are instead treated as network losses (they increase the reported Loss Event Rate).

- Following from the earlier restrictions, neither DCCP will ever send Slow Receiver, Init Cookie, Ack Vector, Data Dropped (except on a DCCP-Response), Timestamp, Timestamp Echo, Identification, Challenge, Payload Checksum, or Loss Intervals options, or DCCP-DataAck or DCCP-CloseReq packets.

A CCID 3-Thin connection is a normal DCCP connection, and MUST follow the DCCP [DCCP] and CCID 3 [CCID 3 PROFILE] specifications, except that if any of the above restrictions is violated, the receiving DCCP SHOULD reset the connection with Reason 195, Thin Violation.

5. Connection Establishment

This section describes how CCID 3-Thin connections are initiated. The terminology "thin" means that a client or server can only communicate with CCID 3-Thin restricted partners, and agrees to abide by those restrictions itself. "Fat" clients and/or servers can communicate with any DCCP.

5.1. Thin Client Initiates Connection

A thin client initiating a connection to any server MUST send a DCCP-Request packet containing only the following options in this exact order: Mandatory, Change L(CCID, 3), Mandatory, Change R(CCID, 3), Mandatory, CCID 3-Thin, Padding, Padding, Padding.

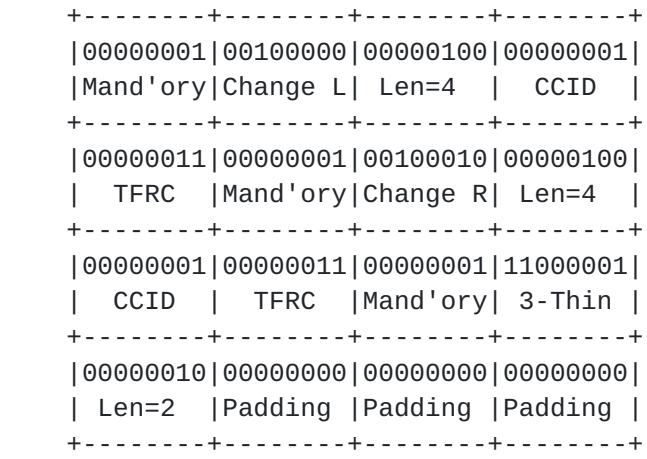

The DCCP-Request MUST NOT contain data.

5.2. Server Responds to Thin Client

A server responding to this packet sequence will first change both CCID features to 3 (TFRC). If TFRC is not acceptable to the server, it MUST reset the connection, since the Change options are marked as Mandatory. The server will then process the Mandatory CCID 3-Thin option. If the server agrees to the CCID 3-Thin restrictions above, it MUST send a DCCP-Response packet containing only the following options, in this order: Confirm R(CCID 3), Confirm L(CCID, 3), Mandatory, CCID 3-Thin, and Padding.

```
+--------+--------+--------+--------+
|00100011|00000100|00000001|00000011|
|ConfirmR| Len=4  |  CCID  |  TFRC  |
+--------+--------+--------+--------+
|00100001|00000100|00000001|00000011|
|ConfirmL| Len=4  |  CCID  |  TFRC  |
+--------+--------+--------+--------+
|00000001|11000001|00000010|00000000|
|Mand'ory| 3-Thin | Len=2  |Padding |
+--------+--------+--------+--------+
```

The server SHOULD reset the connection if the DCCP-Request does not follow the CCID 3-Thin restrictions -- for example, if the CCID 3-Thin option is followed by other options.

If the server does not understand the CCID 3-Thin option, or does not agree to the Thin restrictions, it MUST reset the connection with Reason set to Option Error (because of the Mandatory option).

5.3. Thin Server Responds to Fat Client

A thin server might be contacted by a fat client, which might send options and feature values on its DCCP-Request that do not follow the CCID 3-Thin restrictions. The server MUST process these options, but not necessarily as a conventional DCCP would. In particular:

- The server MUST process options immediately preceded by a Mandatory option. If a Mandatory option does not fit the CCID 3-Thin restrictions, the server MUST reset the connection with Reason set to Option Error. If a Mandatory option does fit the restrictions, the server SHOULD respond as required; for example, when sent a Mandatory Change option with acceptable values, the server SHOULD respond with the appropriate Confirm option. An extremely restricted server MAY instead choose to reset the connection (with Reason set to Option Error) whenever any Mandatory option is received.

- The server SHOULD process non-Mandatory options as well. For example, when sent a Change L(CCID, 2) option, the server SHOULD reset the connection with Reason set to Fruitless Negotiation; but when sent a Change L(CCID, 1 2 3) option, the server SHOULD respond with a Confirm R(CCID, 3) option. An extremely restricted server MAY instead ignore all non-mandatory options; the fat client will perform the necessary settings (or reset the connection) when it receives the server's Mandatory CCID 3-Thin option.

- If the DCCP-Request contains data, the server's DCCP-Response MUST contain a Data Dropped option with Drop Code 0 (data dropped due to protocol constraints).

- Any DCCP-Response MUST contain Mandatory feature negotiation options that set both CCIDs to 3, followed by a Mandatory CCID 3-Thin option. The feature negotation options might be Mandatory Change L(CCID, 3) and Mandatory Change R(CCID, 3) options, as in the thin-client case; or if the server processed one or more Change(CCID) options on the DCCP-Request, they might be Confirm options.

5.4. Fat Client Responds to Thin Server's Response

A fat client receiving a DCCP-Response containing a Mandatory CCID 3-Thin option MUST either reset the connection, if it cannot abide by the CCID 3-Thin restrictions, or respond with a DCCP-Ack. This DCCP-Ack MUST complete any feature negotiations initiated by the DCCP-Response, and MUST also contain a CCID 3-Thin option (either Mandatory or not).

The thin server MUST be prepared to handle such a DCCP-Ack, but it MAY ignore the feature negotiation options on that Ack. Because the server sent a Mandatory CCID 3-Thin option on its DCCP-Response, it can assume that any ensuing DCCP-Ack abides by the CCID 3-Thin restrictions.

6. Security Considerations

Security considerations for DCCP have been discussed in [DCCP]; security considerations for TFRC have been discussed in [RFC 3448]; and security considerations for DCCP CCID 3 have been discussed in [CCID 3 PROFILE]. CCID 3-Thin is not as secure against spoofed feedback, misbehaving receivers, and DOS attacks as straight DCCP with CCID 3. In particular, CCID 3-Thin does not support ECN, so the ECN Nonce-based verification mechanisms of CCID 3 are not available to it.

7. IANA Considerations

The CCID 3-Thin specification allocates two values from IANA-administered registries:

- The CCID 3-specific option 195, for CCID 3-Thin.

- The CCID 3-specific Reset Reason 195, for Thin Violation.

8. Thanks

Thanks to Tom Phelan for his DCCP-Lite draft, which showed what he thought could be elided from full CCID 3.

Normative References

[CCID 2 PROFILE] S. Floyd and E. Kohler. Profile for DCCP Congestion Control ID 2: TCP-like Congestion Control, draft-ietf-dccp-ccid2-06.txt, work in progress.

[CCID 3 PROFILE] S. Floyd, E. Kohler, and J. Padhye. Profile for DCCP Congestion Control ID 3: TCP-Friendly Rate Control, draft-ietf-dccp-ccid3-06.txt, work in progress.

[DCCP] E. Kohler, M. Handley, S. Floyd, and J. Padhye. Datagram Congestion Control Protocol, draft-ietf-dccp-spec-07.txt, work in progress.

[RFC 2119] S. Bradner. Key Words For Use in RFCs to Indicate Requirement Levels. RFC 2119.

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

Authors' Addresses

Eddie Kohler <kohler@cs.ucla.edu>
4531C Boelter Hall
UCLA Computer Science Department
Los Angeles, CA 90095
USA

Full Copyright Statement

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.