

DCCP Working Group
Internet-Draft
Intended status: Experimental
Expires: November 31, 2008

Gorry Fairhurst
Arjuna Sathiaselalan
University of Aberdeen

Intended status: Experimental

June 1, 2008

Quick-Start for Datagram Congestion Control Protocol (DCCP)
draft-fairhurst-tsvwg-dccp-qs-03.txt

Status of this Draft

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/lid-abstracts.html>

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

This Internet-Draft will expire on November 31, 2008.

Abstract

This document specifies the use of the Quick-Start mechanism by the Datagram Congestion Control Protocol (DCCP). DCCP is a transport protocol that allows the transmission of congestion-controlled, unreliable datagrams. DCCP is intended for applications such as streaming media, Internet telephony, and on-line games. In DCCP, an application has a choice of congestion control mechanisms, each specified by a Congestion Control Identifier (CCID). This document specifies general procedures applicable to all DCCP CCIDs and specific procedures for the use of Quick-Start with DCCP CCID-2 and CCID-3. Quick-Start enables a DCCP sender to cooperate with any Quick-Start routers along the end-to-end path to determine an allowed sending rate at the start and, at times, in the middle of a DCCP connection (e.g., after an idle or application-limited period). The present specification is provided for use in controlled environments, and not as a mechanism that would be intended or

appropriate for ubiquitous deployment in the global Internet.

Expires November 2008

[page 1]

Table of Contents

1. Introduction
2. Quick-Start for DCCP
 - 2.1 Sending a Quick-Start Request for a DCCP flow
 - 2.2 The Quick-Start Interval
 - 2.3 Receiving a Quick-Start Request for a DCCP flow
 - 2.3.1 The Quick-Start Response Option
 - 2.4 Receiving a Quick-Start Response
 - 2.5 Procedure when no response to a Quick-Start Request
 - 2.6 Procedure when a Quick-Start Packet is dropped
 - 2.7 Interactions with Mobility and Signalled Path Changes
 - 2.8 Interactions with Path MTU Discovery
 - 2.9 Interactions with Middle boxes
3. Mechanisms for Specific CCIDs
 - 3.1 Quick-Start for CCID-2
 - 3.1.1 The Quick-Start Request for CCID-2
 - 3.1.2 Sending a Quick-Start Response with CCID-2
 - 3.1.3 Using the Quick-Start Response with CCID-2
 - 3.1.4 Reported Loss during Quick-Start Mod
 - 3.1.5 CCID-2 Feedback Traffic on the Reverse Path
 - 3.2 Quick-Start for CCID-3
 - 3.2.1 The Quick-Start Request for CCID-3
 - 3.2.2 Sending a Quick-Start Response with CCID-3
 - 3.2.3 Using the Quick-Start Response with CCID-3
 - 3.2.4 The Quick-Start Validation Phase
 - 3.2.5 Reported Loss during Quick-Start Mode or Validation Phase
 - 3.2.6 An Example Quick-Start Scenario with CCID-3
 - 3.2.7 CCID-3 Feedback Traffic on the Reverse Path
 - 3.3 Quick-Start for CCID-4
 - 3.3.1 The Quick-Start Request for CCID-4
 - 3.3.2 Sending a Quick-Start Response with CCID-4
 - 3.3.3 Using the Quick-Start Response with CCID-4
 - 3.3.4 CCID-4 Feedback Traffic on the Reverse Path
4. Discussion of Issues
 - 4.1 Over-run and Quick-Start Validation
 - 4.2 Experimental Status
5. IANA Considerations
6. Acknowledgments
7. Security Considerations
8. References
 - 8.1 Normative References
 - 8.2 Informative References
9. Authors' Addresses
10. IPR Notices
 - 10.1 Intellectual Property Statement
 - 10.2 Disclaimer of Validity
11. Copyright Statement

Expires November 2008

[Page 2]

1. Introduction

The Datagram Congestion Control Protocol (DCCP) [[RFC4340](#)] is a transport protocol for congestion-controlled, unreliable datagrams, intended for applications such as streaming media, Internet telephony, and on-line games.

In DCCP, an application has a choice of congestion control mechanisms, each specified by a Congestion Control Identifier (CCID) [[RFC4340](#)]. There are general procedures applicable to all DCCP CCIDs that are described in [Section 2](#), and details that relate to how individual CCIDs should operate, which are described in [Section 3](#). This separation of CCID-specific and DCCP general functions is in the spirit of the modular approach adopted by DCCP.

Quick-Start [[RFC4782](#)] is an Experimental mechanism for transport protocols specified for use in controlled environments. The current specification of this mechanism is not intended or appropriate for ubiquitous deployment in the global Internet.

Quick-Start is designed for use between end hosts within the same network or on Internet paths that include IP routers. It works in cooperation with any routers, allowing a sender to determine an allowed sending rate at the start and at times in the middle of a data transfer (e.g., after an idle or application-limited period).

This document assumes that the reader is familiar with [RFC4782](#) [[RFC4782](#)], which specifies the use of Quick-Start with IP and with TCP. [Section 7 of RFC4782](#) also provides guidelines for the use of Quick-Start with other transport protocols, including DCCP. This document answers some of the issues that were raised by [RFC4782](#) and provides a definition of how Quick-Start must be used with DCCP.

In using Quick-Start, the sending DCCP end host indicates the desired sending rate in bytes per second, using a Quick-Start option in the IP header of a DCCP packet. Each Quick-Start capable router along the path could, in turn, either approve the requested rate, reduce the requested rate, or indicate that the Quick-Start Request is not approved.

If the Quick-Start Request is approved by all the routers along the path, then the DCCP receiver returns an appropriate Quick-Start Response. On receipt of this, the sending end host can send at up to the approved rate for one round-trip time. Subsequent transmissions will be governed by the default CCID congestion control mechanisms for the connection. If the Quick-Start Request is not approved, then the sender must use the default congestion control mechanisms.

Expires November 2008

[Page 3]

2. Quick-Start for DCCP

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](#) [[RFC2119](#)].

Unless otherwise specified, DCCP end hosts follow the procedures specified in [Section 4 of \[RFC4782\]](#), following the use specified for Quick-Start with TCP.

2.1 Sending a Quick-Start Request for a DCCP flow

A DCCP sender MAY use a Quick-Start Request during the start of a connection, when the sender would prefer to have a larger initial rate than allowed by standard mechanisms (e.g. [RFC-to-be-3448.bis] or [[RFC3390](#)]).

A Quick-Start Request MAY be also used once a DCCP flow is connected (in the middle of a DCCP flow). In standard operation, DCCP CCIDs can constrain the sending rate (or window) to less than that desired (e.g. when an application increases the rate at which it wishes to send). A DCCP sender that has data to send after an idle period or application-limited period (i.e. where the sender transmits at less than the allowed sending rate) can send a Quick-Start Request using the procedures defined in [Section 3](#).

Quick-Start Requests will be more effective if the Quick-Start Rate is not larger than necessary. Each requested Quick-Start Rate that has been approved, but was not fully utilized, takes away from the bandwidth pool maintained by Quick-Start routers that would be otherwise available for granting successive requests [[RFC4782](#)].

In contrast to most TCP applications, many DCCP applications have the notion of a natural media rate that they wish to achieve. For example, during the initial connection, a host may request a Quick-Start rate equal to the media rate of the application.

When sending a Quick-Start Request, the DCCP sender SHOULD send the Quick-Start Request using a packet that requires an acknowledgement, such as a DCCP-Request, DCCP-Response, or DCCP-Data.

2.2 The Quick-Start Interval

Excessive use of the Quick-Start mechanism is undesirable. This document therefore introduces the concept of the Quick-Start Interval. End hosts therefore MUST NOT make a subsequent Quick-Start Request within a period specified as the Quick-Start Interval.

When a connection is established, the Quick-Start Interval is

initialized to a value of 6 seconds, and the previous Quick-Start Interval is set to 0 seconds. This value is chosen to be sufficiently large to prevent excessive router processing over typical Internet paths.

Whenever a Quick-Start Request is sent, the Quick-Start Interval is then recalculated as:

$$\text{Quick-Start Interval} = \max(6 \text{ seconds}, \text{QSPrev_Interval} * 2, 4 * \text{RTT})$$

where the QSPrev_Interval is the value of the previous Quick-Start Interval.

Each unsuccessful Quick-Start Request, therefore results in the Quick-Start Interval being doubled (resulting in an exponential back-off). The maximum time the sender can back-off is 64 seconds, after which a sender ceases using Quick-Start and MUST NOT send any further Quick-Start Requests for the remainder of the DCCP connection.

Whenever a Quick-Start Request is approved, the previous Quick-Start Interval and QSPrev_Interval are reset to their initial value.

2.3 Receiving a Quick-Start Request for a DCCP flow

The procedure for processing a received Quick-Start Request is normatively defined in [\[RFC4782\]](#), and summarised in this paragraph. An end host that receives an IP packet containing a Quick-Start Request passes the Quick-Start Request, along with the value in the IP TTL field, to the receiving DCCP layer. If the receiving host is willing to permit the Quick-Start Request, it SHOULD respond immediately by sending a packet that carries the Quick-Start Response option in the DCCP header of the corresponding feedback packet (e.g. using a DCCP-Ack packet or in a DCCP-DataAck packet).

The Rate Request in the Quick-Start Response option is set to the received value of the Rate Request in the Quick-Start option or to a lower value if the DCCP receiver is only willing to allow a lower Rate Request. Where information is available (e.g. knowledge of the local layer 2 interface speed), a QS receiver SHOULD verify that the received rate does not exceed its expected receive link capacity. The TTL Diff in the Quick-Start Response is set to the difference between the IP TTL value and the Quick-Start TTL value. The Quick-Start Nonce in the Response is set to the received value of the Quick-Start Nonce in the Quick-Start option.

If an end host receives an IP packet with a Quick-Start Request with a request rate of zero, then this host SHOULD NOT send a Quick-Start Response [\[RFC4782\]](#).

The Quick-Start Response MUST NOT be resent if it is lost in the network [\[RFC4782\]](#). Packet loss could be an indication of congestion on the return path, in which case it is better not to approve the

Quick-Start Request.

2.3.1 The Quick-Start Response Option

Expires November 2008

[Page 5]

The Quick-Start Response message must be carried by the transport protocol using Quick-Start. This section defines a DCCP Header option used to carry the Quick-Start response. This header option is REQUIRED for end hosts to utilise the Quick-Start mechanism with DCCP flows. The format resembles that defined for TCP [[RFC4782](#)].

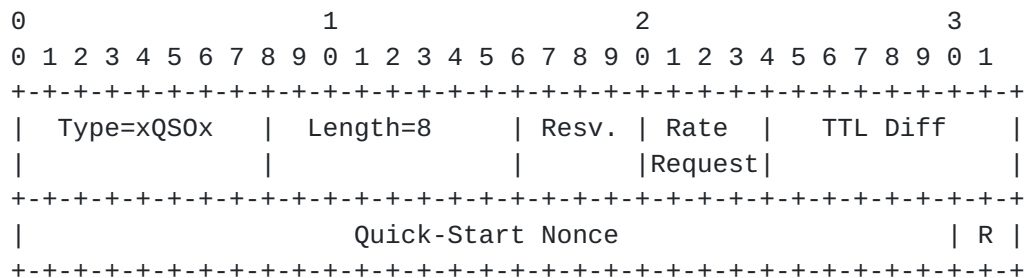


Figure 1. The Quick-Start Response option.

IANA ACTION, PLEASE REPLACE xQSOx with the assigned value in the figure above.###

IANA ACTION, PLEASE ALSO REPLACE xQSOx with the assigned value in the paragraph below.###

The first byte of the Quick-Start Response option contains the option kind, identifying the DCCP option (xQSOx).

The second byte of the Quick-Start Response option contains the option length in bytes. The length field MUST be set to 8 bytes.

The third byte of the Quick-Start Response option contains a four-bit Reserved field, and the four-bit allowed Rate Request, formatted as in the IP Quick-Start Rate Request option [[RFC4782](#)].

The fourth byte of the DCCP Quick-Start Response option contains the TTL Diff. The TTL Diff contains the difference between the IP TTL and Quick-Start TTL fields in the received Quick-Start Request packet, as calculated in [[RFC4782](#)].

Bytes 5-8 of the DCCP option contain the 30-bit Quick-Start Nonce and a two-bit Reserved field.

2.4 Receiving a Quick-Start Response

Reception of a Quick-Start Response packet results in the sender entering the Quick-Start Mode. The procedure following reception of a Quick-Start Response packet is CCID-specific and described in [Section 3](#).

2.5 Procedure when no response to a Quick-Start Request

As in TCP, if a Quick-Start Request is dropped (i.e., the Request or Response is not delivered by the network) the DCCP sender MUST

Expires November 2008

[Page 6]

revert to the congestion control mechanisms it would have used if the Quick-Start Request had not been approved. The connection is not permitted to send a subsequent Quick-Start Request before expiry of the current Quick-Start Interval ([section 2.1](#)).

[2.6](#) Procedure when a Quick-Start Packet is dropped

While the sender is in the Quick-Start Mode, all sent packets are known as Quick-Start Packets [[RFC4782](#)]. Loss of a Quick-Start Packet is an indication of potential network congestion. The behaviour of a DCCP sender following the loss of a Quick-Start Packet is specific to a particular CCID (see [section 3](#)).

[2.7](#) Interactions with Mobility and Signalled Path Changes

The use of Quick-Start may assist hosts in determining when it is appropriate to increase their rate following an explicitly signalled change of the network path. Senders must ensure this does not generate an excessive rate of Quick-Start Requests by using the method below.

A sender that has explicit information that the network path has changed (e.g. a mobile IP binding update [[RFC3344](#)], [[RFC3775](#)]) MAY reset the Quick-Start Interval and QSPrev_Interval to their initial values (specified in [section 2.1](#)).

The sender MAY also send a Quick-Start Request to determine a new safe transmission rate, but must observe the following rules:

- . It MUST NOT send a Quick-Start Request within a period less than the initial Quick-Start Interval (i.e., 6 seconds) since it previously sent a Quick-Start Request. That is, it must wait for at least a period of 6 seconds after the previous request, before sending a new Quick-Start Request.
- . If it has not sent a Quick-Start Request within the previous 6 seconds, it SHOULD defer sending a Quick-Start Request for a randomly chosen period between 0 and 6 seconds. The random period should be statistically independent between different hosts and between different connections on the same host. This delay is to mitigate the effect on router load of synchronised responses by multiple connections in response to a path change that effects multiple connections.

[2.8](#) Interactions with Path MTU Discovery

DCCP implementations are encouraged to support Path MTU Discovery (PMTUD) when applications are able to use a DCCP packet size that exceeds the default Path MTU [[RFC4340](#)], [[RFC4821](#)]. Quick-Start

Requests SHOULD NOT be sent with packets that are used as a PMTUD Probe Packet, since these packets could be lost in the network increasing the probability of loss. It may therefore be preferable to separately negotiate the PMTU and the use of Quick-Start.

The DCCP protocol is datagram-based and therefore the size of the segments that are sent is a function of application behaviour as well as being constrained by the largest supported Path MTU.

2.9 Interactions with Middle boxes

A Quick-Start Request is carried in an IP packet option [[RFC4782](#)]. Interactions with network devices (middleboxes) that inspect or modify IP options could therefore lead to discard, ICMP error, or DCCP-Reset when attempting to forward packets carrying a Quick-Start Request.

If a DCCP sender sends a DCCP-Request that also carries a Quick-Start Request, and does not receive a DCCP-Response to the packet, the DCCP sender SHOULD resend the DCCP-Request packet without including a Quick-Start Request.

Similarly, if a DCCP sender receives a DCCP-Reset in response to a DCCP-Request packet that also carries a Quick-Start Request, then the DCCP sender SHOULD resend DCCP-Request packet without the Quick-Start Request.

The DCCP sender then ceases to use the Quick-Start Mechanism for the remainder of the connection.

A DCCP sender that uses a Quick-Start Request within an established connection, and does not receive a response will treat this as non-approval of the request. Successive unsuccessful attempts will result in an exponential increase in the Quick-Start Interval ([section 2.1](#)). If this grows to a value exceeding 64 seconds the DCCP sender ceases to use the Quick-Start Mechanism for the remainder of the connection.

3. Mechanisms for Specific CCIDs

This section specifies the use of Quick-Start with DCCP CCID-2, CCID-3, and CCID-4.

3.1 Quick-Start for CCID-2

This section describes the Quick-Start mechanism to be used with DCCP CCID-2 [[RFC4341](#)]. CCID-2 uses a TCP-like congestion control mechanism.

3.1.1 The Quick-Start Request for CCID-2

A Quick-Start Request MAY be sent to allow the sender to determine

if it is safe to use a larger initial cwnd. This permits a faster start-up of a new DCCP CCID-2 flow.

A Quick-Start Request MAY also be sent for an established connection to request a higher sending rate after an idle period or application-limited period (described in [section 2.1](#)). This allows a receiver to use a larger cwnd than allowed with standard operation.

A Quick-Start Request SHOULD NOT be sent within a period less than the Quick-Start Interval following a previous Quick-Start Request ([section 2.2](#)), unless it is sent as a result of a network path change ([section 2.7](#)). A Quick-Start Request that follows a loss or congestion event MUST NOT request a Quick-Start rate that exceeds the largest congestion window achieved by the CCID-2 connection since the last packet drop (translated to a sending rate).

3.1.2 Sending a Quick-Start Response with CCID-2

A receiver processing a Quick-Start Request uses the method described in [Section 2.2](#). On receipt of a Quick-Start Request, the receiver MUST send a Quick-Start Response (even if a receiver is constrained by the ACK Ratio).

3.1.3 Using the Quick-Start Response with CCID-2

On receipt of a valid Quick-Start Response option, the sender MUST send a Quick-Start Approved option [[RFC4782](#)] as an IP option using the first Quick-Start Packet or send this as an option using a DCCP control packet if there are no DCCP-Data packets pending transmission.

If the approved Quick-Start rate is less than current sending rate, the sender does not enter the Quick-Start Mode, and continues using the procedure defined in CCID-2.

If the approved Quick-Start rate at the sender exceeds the current sending rate, the sender enters the Quick-Start Mode and continues in the Quick-Start Mode for a maximum period of 1 RTT. While in the Quick-Start Mode, all DCCP packets that it sends are known as Quick-Start Packets.

The sender sets its Quick-Start cwnd (QS_cwnd) as follows:

$$QS_cwnd = (R * T) / (s + H) \quad (1)$$

where R is the Rate Request in bytes per second, s is the packet size, and H is the estimated DCCP/IP header size in bytes (e.g., 32 bytes for DCCP layered directly over IPv4).

A CCID-2 sender MAY then increase its cwnd to the QS_cwnd. The cwnd should not be reduced (i.e., a QS_cwnd lower than cwnd should be ignored, since the CCID-2 congestion control method already permits

this rate). CCID-2 is not a rate-paced protocol. Therefore, if the QS_cwnd is used, the sending host MUST implement a suitable method to pace the rate at which the Quick-Start Packets are sent until it

receives an ACK for a packet sent during the Quick-Start Mode [[RFC4782](#)]. The sending host SHOULD also record the previous cwnd and note that the new cwnd has been determined by Quick-Start, rather by other means (e.g. by setting a flag to indicate that it is in Quick-Start Mode).

When the sender receives the first ACK to a packet sent in the Quick-Start Mode, the sender MUST reduce the cwnd to the actual flight size (the current amount of unacknowledged data sent) [[RFC4782](#)].

[3.1.4](#) Reported Loss during Quick-Start Mode

A sender in the Quick-Start Mode or Validation Phase that detects congestion (e.g. receives a feedback packet that reports new packet loss or a packet with a congestion marking), MUST immediately leave the Quick-Start Mode or Validation Phase and enter the congestion avoidance phase [[RFC4341](#)].

[3.1.5](#) CCID-2 Feedback Traffic on the Reverse Path

A CCID-2 receiver sends feedback for groups of received packets [[RFC4341](#)]. Approval of a higher transmission rate using Quick-Start will increase control traffic on the reverse path. A return path that becomes congested could have a transient negative impact on other traffic flows sharing the return link. The lower rate of feedback will then limit the achievable rate in the forward direction.

[3.2](#) Quick-Start for CCID-3

This section describes the Quick-Start mechanism to be used with DCCP CCID-3 [[RFC4342](#)]. The rate-based congestion control mechanism used by CCID-3, leads to specific issues that are addressed by Quick-Start in this section, and include the introduction of a Quick-Start Validation Phase.

[3.2.1](#) The Quick-Start Request for CCID-3

A Quick-Start Request MAY be sent to allow the sender to determine if it is safe to use a larger initial sending rate. This permits a faster start-up of a new DCCP flow.

A Quick-Start Request MAY also be sent to request a higher sending rate after an idle period (in which the nofeedback timer expires [[RFC-to-be-3448bis](#)]) or an application-limited period (described in [section 2.1](#)). This allows a receiver to increase the sending rate faster than allowed with standard operation (i.e. faster than twice the rate reported by a CCID-3 receiver in the most recent feedback

message).

The requested rate specified in a Quick-Start Request must consider the current loss event rate (if any), either from calculation at the

Expires November 2008

[Page 10]

sender or from feedback received from the receiver. CCID-3 requires that a sender must not send more than the upper bound dictated by the loss event rate. This rate offers a safe response in the presence of expected congestion.

3.2.2 Sending a Quick-Start Response with CCID-3

When processing a received Quick-Start Request, the receiver uses the method described in [Section 2.3](#). In addition, if a CCID-3 receiver uses the window counter to send periodic feedback messages, then the receiver sets its local variable `last_counter` to the value of the window counter reported by the segment containing the Quick-Start Request. The next feedback message would then be sent when the `window_counter` is greater or equal to `last_counter + 4`. If the CCID-3 receiver uses a feedback timer to send period feedback messages, then the DCCP receiver **MUST** reset the CCID-3 feedback timer, causing the feedback to be sent as soon as possible. This helps to align the timing of feedback to the start and end of the period in which Quick-Start Packets are sent, and will normally result in feedback at a time that is approximately the end of the period when Quick-Start Packets are received.

A Quick-Start Request **SHOULD NOT** be sent within a period less than the Quick-Start Interval following a previous Quick-Start Request ([section 2.2](#)), unless it is sent as a result of a network path change ([section 2.7](#)).

3.2.3 Using the Quick-Start Response with CCID-3

On receipt of a valid Quick-Start Response option, the sender enters the Quick-Start Mode. The sender **MUST** send a Quick-Start Approved option [[RFC4782](#)] as an IP option using the first Quick-Start Packet or send this as an option using a DCCP control packet if there are no DCCP-Data packets pending transmission.

If the approved Quick-Start rate is less than current sending rate, the sender does not enter the Quick-Start Mode, and continues using the procedure defined in CCID-3. While in the Quick-Start Mode, all DCCP packets that it sends are known as Quick-Start Packets.

If the approved Quick-Start rate exceeds the current sending rate, the sender enters the Quick-Start Mode and continues in the Quick-Start Mode for a maximum period of 1 RTT. The sender sets its Quick-Start sending rate (`QS_sendrate`) as follows:

$$QS_sendrate = R * s / (s + H) \quad (2)$$

where `R` the Rate Request in bytes per second, `s` is the packet size, and `H` the estimated DCCP/IP header size in bytes (e.g., 32 bytes). A CCID-3 host **MAY** then increase its sending rate (`sendrate`) to the

QS_sendrate. The rate should not be reduced.

CCID-3 is a rate-paced protocol. Therefore, if the QS_sendrate is used, the sending host MUST pace the rate at which the Quick-Start

Packets are sent over the next RTT. The sending host SHOULD also record the previous congestion-controlled rate and note that the new rate has been determined by Quick-Start rather by other means (e.g. by setting a flag to indicate that it is in Quick-Start Mode).

The sender exits the Quick-Start Mode after either:

- * Receipt of a feedback packet acknowledging one or more Quick-Start Packets,
- * A period of 1 RTT after receipt of a Quick-Start Response, or
- * Detection of a loss or congestion event (see [Section 3.2.5](#)).

3.2.4 The Quick-Start Validation Phase

After transmitting a set of Quick-Start Packets (and providing that no loss or ECN marking is reported), the sender enters the Quick-Start Validation Phase. This phase persists for a period during which the sender seeks to affirm that the capacity used by the Quick-Start Packets did not introduce congestion. (This phase is introduced, because unlike TCP (and CCID-2), CCID-3 does not receive frequent feedback that would indicate the congestion state of the forward path). While in the Quick-Start Validation Phase, the sender is tentatively permitted to continue sending at the `QS_sendrate`. On conclusion of the Validation Phase, the sender expects to receive assurance that it may safely use the current rate.

A sender that receives feedback that reports a loss or congestion event MUST follow the procedures described in [Section 3.2.5](#).

The sender SHOULD exit the Quick-Start Validation Phase on receipt of feedback that acknowledges all packets sent in the Quick-Start Mode (i.e. all Quick-Start Packets). It MUST exit this phase on expiry of the Quick-Start validation time. The Quick-Start Validation Phase is limited to the Quick-Start Validation Time (a maximum of 1.5 RTTs).

A sender that completes the Quick-Start Validation phase with no reported packet loss or congestion, stops using the `QS_sendrate` and MUST recalculate a suitable sending rate using the standard congestion control mechanisms [[RFC4342](#)]. For example, if the DCCP sender was in slow-start prior to the Quick-Start Request, and no packets were lost or marked since that time, then the sender continues in slow-start after exiting Quick-Start Mode until the sender sees a packet loss, or congestion is reported.

If no feedback is received within the Quick-Start Validation Phase, the sender MUST return to the minimum of the original rate (at the start of the Quick-Start Mode) and one half of the `QS_sendrate`.

3.2.5 Reported Loss during Quick-Start Mode or Validation Phase

Expires November 2008


[Page 12]


```
X = max(min(X_calc, min(2*X_recv, 2* QS_recv-rate)), s/t_mbi);
```

```
X = max(min(X_calc, min(recv_limit, 2* QS_recv-rate)),
        s/t_mbi);
```

When [RFC4342](#) adopts [\[RFC-to-be-3448bis\]](#), the send rate, X would also be calculated using the above formula from [\[RFC-to-be-3448bis\]](#).

	DCCP Sender	DCCP Receiver
Quick-Start Request/Response	<pre> +-----+ Quick-Start Request --> <-- Quick-Start Response Quick-Start Approve --> +-----+ </pre>	
Quick-Start Mode	<pre> +-----+ Quick-Start Packets --> <-- Feedback from Receiver +-----+ </pre>	
Quick-Start Validation Phase	<pre> +-----+ Packets --> <-- Feedback from Receiver +-----+ </pre>	



Expires November 2008

[Page 13]

Figure 2. The Quick-Start Mode and Validation Phase.

Figure 2 shows an example use of Quick-Start with CCID-3.

[3.2.7](#) CCID-3 Feedback Traffic on the Reverse Path

A CCID-3 receiver sends feedback at least once each RTT [[RFC4342](#)]. Use of Quick-Start is therefore not expected to significantly increase control traffic on the reverse path.

[3.3](#) Quick-Start for CCID-4

This section describes the Quick-Start mechanism to be used with DCCP CCID-4 [[ID.CCID4](#)]. CCID-4 is similar to CCID-3 except that a sender using CCID-4 is limited to a maximum of 100 packets/second. The Quick-Start procedure defined below therefore resembles that for CCID-3.

[3.3.1](#) The Quick-Start Request for CCID-4

The procedure for sending a Quick-Start Request using CCID-4 is the same as for CCID-3, defined in [section 3.2.1](#). In addition, the requested rate MUST be less than or equal to the equivalent of a sending rate of 100 packet per second [[ID.CCID4](#)].

[3.3.2](#) Sending a Quick-Start Response with CCID-4

This procedure is the same as for CCID-3, defined in [section 3.2.2](#).

[3.3.3](#) Using the Quick-Start Response with CCID-4

This procedure is the same as for CCID-3, defined in sections [3.2.3](#), [3.2.4](#), and [3.2.5](#), except that the congestion control procedures defined in [[ID.CCID4](#)] and used in place of those defined in [[RFC4342](#)]

[3.3.4](#) CCID-4 Feedback Traffic on the Reverse Path

A CCID-4 receiver sends feedback at least once each RTT (defined in [[RFC4342](#)]). Use of Quick-Start is therefore not expected to significantly increase control traffic on the reverse path.

[4](#). Discussion of Issues

The considerations for using Quick-Start with DCCP are not significantly different to those for Quick-Start with TCP.

[4.1](#) Over-run and Quick-Start Validation

CCID-3 raises an issue in that a sender using Quick-Start may continue to use the rate specified by a Quick-Start Response for a period that exceeds one path round trip time (i.e., that which TCP

would have used). This over-run is a result of the less frequent feedback interval used by TFRC (i.e., CCID-3 and CCID-4 provide feedback once per RTT, rather than once for a few packets). In the method specified by this document, the Quick-Start Validation Time bounds this over-run to be not more than an additional 1.5 RTTs.

The currently selected method is chosen as a compromise that reflects the need to terminate quickly following the loss of a feedback packet, and the need to allow sufficient time for end host and router processing as well as the different perceptions of the path RTT held at the sender and receiver. Any reported loss or congestion results in immediate action without waiting for completion of the Quick-Start Validation period.

4.2 Experimental Status

There are many cases in which Quick-Start Requests would not be approved [[RFC4782](#)]. These include communication over paths containing routers, IP tunnels, MPLS paths, and the like, that do not support Quick-Start. These cases also include paths with routers or middleboxes that drop packets containing IP options. Quick-Start Requests could be difficult to approve over paths that include multi-access layer-two networks.

Transient effects could arise when the transport protocol packets associated with a connection are multiplexed over multiple parallel (sometimes known as alternative) link or network-layer paths, and Quick-Start is used, since it will be effective on only one of the paths, but could lead to increased traffic on all paths.

A CCID-2 sender using Quick-Start can increase the control traffic on the reverse path, which could have a transient negative impact on other traffic flows sharing the return link ([section 3.1.5](#)). The lower rate of feedback will then limit the achievable rate in the forward direction.

[RFC4782] also describes environments where the Quick-Start mechanism could fail with false positives, with the sender incorrectly assuming that the Quick-Start Request had been approved by all of the routers along the path. As a result of these concerns, and as a result of the difficulties and seeming absence of motivation for routers, such as core routers, to deploy Quick-Start, Quick-Start has been proposed as a mechanism that could be of use in controlled environments, and not as a mechanism that would be intended or appropriate for ubiquitous deployment in the global Internet.

Further experimentation would be required to confirm the deployment

of Quick-Start and to investigate performance issues that may arise, prior to any recommendation for use over the general Internet.

5. IANA Considerations

Expires November 2008

[Page 15]

This document requires IANA involvement for the assignment of a DCCP Option Type from the DCCP Option Types Registry. This Option is applicable to all CCIDs and is known as the "Quick-Start Response" Option and is defined in [Section 2.2.1](#). It specifies a length value in the format used for options numbered 32-128.

[6. Acknowledgments](#)

The author gratefully acknowledges the previous work by Sally Floyd to identify issues that impact Quick-Start for DCCP, and her comments to improve this document. We also acknowledge comments and corrections from Pasi Sarolahti, Mark Allman and others in the IETF DCCP WG.

[7. Security Considerations](#)

Security issues are discussed in [\[RFC4782\]](#). Middlebox deployment issues are also highlighted in [section 2.7](#). No new security issues are raised within this document.

[8. References](#)

[8.1 Normative References](#)

[RFC2119] Bradner, S., "Key Words for Use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), 1997.

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

[RFC4341] Floyd, S. and E. Kohler, "Profile for Datagram Congestion Control Protocol (DCCP) Congestion Control ID 2: TCP-like Congestion Control", [RFC 4341](#), March 2006.

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

[RFC4782] Floyd, S., Allman, M., Jain, A., and P. Sarolahti, "Quick-Start for TCP and IP", [RFC 4782](#), January 2007.

[ID.CCID4] Floyd, S., Kohler, E., "Profile for Datagram Congestion Control Protocol (DCCP) Congestion ID 4: TCP-Friendly Rate Control for Small Packets (TFRC-SP)", IETF Work In Progress, 2007.

[RFC-to-be-3448bis] Floyd, S., Padhye, J., Widmer, J., "TCP Friendly Rate Control (TFRC): Protocol Specification", IETF Work In Progress, 2008.

Expires November 2008

[Page 16]

8.2 Informative References

[RFC3344] Perkins, C., Ed., "IP Mobility Support for IPv4", [RFC 3344](#), August 2002.

[RFC3775] Johnson, D., Perkins, C., and J. Arkko, "Mobility Support in IPv6", [RFC 3775](#), June 2004.

[RFC3390] Allman, M., Floyd, S., Partridge, C., "Increasing TCP's Initial Window", RFC 3390, October 2002.

[RFC4821] Mathis, M. and J. Heffner, "Packetization Layer Path MTU Discovery", [RFC 4821](#), March 2007.

9. Authors' Addresses

Godred Fairhurst
School of Engineering
University of Aberdeen
Aberdeen, AB24 3UE
UK
Email: gorry@erg.abdn.ac.uk
Web: <http://www.erg.abdn.ac.uk/users/gorry>

Arjuna Sathiaselan
School of Engineering
University of Aberdeen
Aberdeen, AB24 3UE
UK
Email: arjuna@erg.abdn.ac.uk
Web: <http://www.erg.abdn.ac.uk/users/arjuna>

10. IPR Notices

10.1 Intellectual Property Statement

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

Expires November 2008

[Page 17]

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.

10.2 Disclaimer of Validity

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.

11. 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.

Acknowledgment

Funding for the RFC Editor function is currently provided by the Internet Society.

Expires November 2008

[Page 18]

[RFC EDITOR NOTE:

This section must be deleted prior to publication]

DOCUMENT HISTORY

Individual Draft 00

This is the first presentation of this document.

Individual Draft 01

This includes fixes for NiTs (thanks Pasi)

It also includes a note on initial rates in 2.1

All mention of packet loss now qualified with loss/congestion.

It adds supports for CCID-2.

It also defines the Quick-Start Interval as a way of controlling the rate at which hosts may issue Quick-Start requests.

Individual Draft 02 - Draft intended for more general review

Resolution of many minor outstanding editorial issues.

Includes feedback on a longer Quick-Start period from Mark Allman.

Includes new section on the interaction with middleboxes.

CCID-2 and CCID-3 text now use the same style.

Added description for CCID-4, based on CCID-3.

Added clarification of PMTUD interaction.

Reorganised to create a section on the QS Interval

Rewritten sections on what to do after loss/congestion

Clarified path change triggers (e.g. from mobility binding updates)

There are no currently known remaining issues to be addressed.

Individual Draft 03

This includes fixes for NiTs, especially to shorten some parts of text.

It includes some additional clarification based on the progress of [RFC3448](#).bis.

Replaced reference to Faster Restart.

Change to paragraph on mobility usage.

[END of RFC EDITOR NOTE]

Expires November 2008

[Page 19]