

Internet Engineering Task Force
Internet-Draft
Intended status: Best Current Practice
Expires: December 2007

S. Floyd
M. Allman
ICIR / ICSI
June 2007

Specifying New Congestion Control Algorithms
draft-ietf-tsvwg-cc-alt-04.txt

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>.

Copyright Notice

Copyright (C) The IETF Trust (2007).

Abstract

The IETF's standard congestion control schemes have been widely shown to be inadequate for various environments (e.g., high-speed networks). Recent research has yielded many alternate congestion control schemes that significantly differ from the IETF's congestion control principles. Using these new congestion control schemes in the global Internet has possible ramifications to both the traffic using the new congestion control and to traffic using the currently standardized congestion control. Therefore, the IETF must proceed with caution when dealing with alternate congestion control proposals. The goal of this document is to provide guidance for considering alternate congestion control algorithms within the IETF.

TO BE DELETED BY THE RFC EDITOR UPON PUBLICATION:

Changes from [draft-ietf-tsvwg-cc-alt-03.txt](#):

- * Minor rewordings in response to IESG review.

Expires: December 2007

[Page 1]

[draft-ietf-tsvwg-cc-alt-04.txt](#)

June 2007

Changes from [draft-ietf-tsvwg-cc-alt-02.txt](#):

- * Removed references from abstract.
- * Added a note that we are focused on documents produced within the IETF (i.e., these are not guidelines that the IRTF or the RFC Editor would necessarily have to follow).
- * Added a list of 'difficult environments' the IETF has thought about in the past (even while admitting that an exhaustive list of 'difficult environments' is impossible to produce).
- * Removed [section 5](#) (conclusions). Some felt that it was redundant and not needed.
- * Made a few of the references normative.
- * Various small wording tweaks.

Changes from [draft-ietf-tsvwg-cc-alt-01.txt](#):

- * Very minor wording tweaks gathered during WGLC.

Changes from [draft-ietf-tsvwg-cc-alt-00.txt](#):

- * Added text to the introduction to clarify the relationship of this document and [RFC 2914](#). In addition, added a requirement (0) in [section 3](#) that says new congestion control schemes that significantly diverge from the principles in [RFC 2914](#) must explain this divergence.

Changes from [draft-floyd-tsvwg-cc-alt-00.txt](#):

- * Changed the name to [draft-ietf-tsvwg-cc-alt-00.txt](#).
- * Added a sentence about robustness with various queueing algorithms in the routers, especially both RED and DropTail. Suggestion from Jitendra Padhye.

- * Added a sentence about robustness with the routers, middleboxes, and such deployed in the current Internet. Concern taken from a talk by Henry Sanders.
- * Add a section about minimum requirements necessary for approval for deployment in the global Internet. Suggestion by Jitendra Padhye.
- * Added more examples to guideline 3 about difficult environments, and added that TCP performance in difficult environments is still an active research topic. Suggestion from Doug Leith.
- * Added citations to examples of discussions of these issues in Experimental RFCs 3649 and 4782.

Expires: December 2007

[Page 2]

[draft-ietf-tsvwg-cc-alt-04.txt](#)

June 2007

- * Added examples of high speed TCP proposals. Suggestion from Bob Braden.
- * Changed the fairness bullets to better reflect that new congestion controllers are expected to assess the impact to standard congestion controlled flows---without commenting on how that assessment should be done. From discussions with bob Briscoe.
- * Made numerous editing changes suggested by Gorrry Fairhurst.

Changes from [draft-floyd-cc-alt-00.txt](#):

- * Changed the name to [draft-floyd-tsvwg-cc-alt-00.txt](#).
- * Added a bullet about incremental deployment. Feedback from Colin Perkins
- * Clarified the fairness section; this section is not saying that strict TCP-friendliness is a requirement.
- * Clarified that as an alternative to Full Backoff, a flow could stop sending when the packet drop rate is above a certain threshold.
- * Clarified that the Full Backoff bullet does not require that different flows with different round-trip times use the same criteria about when they should back off to one packet per round-trip time or less.

- * Added a paragraph about Informational RFCs.
- * Added a bullet about response to transient events, including routing events or moving from a private to a shared network.

END OF NOTES TO BE DELETED.

1. Introduction

This document provides guidelines for the IETF to use when evaluating suggested congestion control algorithms that significantly differ from the general congestion control principles outlined in [[RFC2914](#)]. The guidance is intended to be useful to authors proposing alternate congestion control and for the IETF community when evaluating whether a proposal is appropriate for publication in the RFC series.

The guidelines in this document are intended to be consistent with the congestion control principles from [[RFC2914](#)] of preventing congestion collapse, considering fairness, and optimizing the flow's own performance in terms of throughput, delay, and loss. [[RFC2914](#)] also discusses the goal of avoiding a congestion control 'arms race' among competing transport protocols.

Expires: December 2007

[Page 3]

[draft-ietf-tsvwg-cc-alt-04.txt](#)

June 2007

This document does not give hard-and-fast requirements for an appropriate congestion control scheme. Rather, the document provides a set of criteria that should be considered and weighed by the IETF in the context of each proposal. The high-order criteria for any new proposal is that a serious scientific study of the pros and cons of the proposal needs to have been done such that the IETF has a well rounded set of information to consider.

After initial studies, we encourage authors to write a specification of their proposals for publication in the RFC series to allow others to concretely understand and investigate the wealth of proposals in this space.

2. Document Status

Following the lead of HighSpeed TCP [[RFC3649](#)], alternate congestion control algorithms are expected to be published as "Experimental" RFCs until such time that the community better understands the

solution space. Traditionally, the meaning of "Experimental" status has varied in its use and interpretation. As part of this document we define two classes of congestion control proposals that can be published with the "Experimental" status. The first class includes algorithms that are judged to be safe to deploy for best-effort traffic in the global Internet and further investigated in that environment. The second class includes algorithms that, while promising, are not deemed safe enough for widespread deployment as best-effort traffic on the Internet, but are being specified to facilitate investigations in simulation, testbeds, or controlled environments. The second class can also include algorithms where the IETF does not yet have sufficient understanding to decide if the algorithm is or is not safe for deployment on the Internet.

Each alternate congestion control algorithm published is required to include a statement in the abstract indicating whether or not the proposal is considered safe for use on the Internet. Each alternate congestion control algorithm published is also required to include a statement in the abstract describing environments where the protocol is not recommended for deployment. There may be environments where the protocol is deemed **safe** for use, but still is not **recommended** for use because it does not perform well for the user.

As examples of such statements, [[RFC3649](#)] specifying HighSpeed TCP includes a statement in the abstract stating that the proposal is Experimental, but may be deployed in the current Internet. In contrast, the Quick-Start document [[RFC4782](#)] includes a paragraph in the abstract stating the the mechanism is only being proposed for controlled environments. The abstract specifies environments where the Quick-Start request could give false positives (and therefore would be unsafe to deploy). The abstract also specifies environments where packets containing the Quick-Start request could be dropped in the network; in such an environment, Quick-Start would not be unsafe to deploy, but deployment would still not be recommended because it could cause unnecessary delays for the

connections attempting to use Quick-Start.

For authors of alternate congestion control schemes who are not ready to bring their congestion control mechanisms to the IETF for standardization (either as Experimental or as Proposed Standard), one possibility would be to submit an internet-draft that documents the alternate congestion control mechanism for the benefit of the IETF and IRTF communities. This is particularly encouraged in order to get algorithm specifications widely disseminated to facilitate

further research. Such an internet-draft could be submitted to be considered as an Informational RFC, as a first step in the process towards standardization. Such a document would also be expected to carry an explicit warning against using the scheme in the global Internet.

Note: we are not changing RFC publication process for non-IETF produced documents (e.g., those from the IRTF or independent RFC-Editor submissions). However, we would hope the guidelines in this document inform the IESG as they consider whether to add a note to such documents.

3. Guidelines

As noted above, authors are expected to do a well-rounded evaluation of the pros and cons of proposals brought to the IETF. The following are guidelines to help authors and the IETF community. Concerns that fall outside the scope of these guidelines are certainly possible; these guidelines should not be considered as an all-encompassing check-list.

(0) Differences with Congestion Control Principles [[RFC2914](#)]

Proposed congestion control mechanisms should include a clear explanation of the deviations from [[RFC2914](#)].

(1) Impact on Standard TCP, SCTP [[RFC2960](#)], and DCCP [[RFC4340](#)].

Proposed congestion control mechanisms should be evaluated when competing with standard IETF congestion control [[RFC2581](#),[RFC2960](#),[RFC4340](#)]. Alternate congestion controllers that have a significantly negative impact on traffic using standard congestion control may be suspect and this aspect should be part of the community's decision making with regards to the suitability of the alternate congestion control mechanism.

We note that this bullet is not a requirement for strict TCP-friendliness as a prerequisite for an alternate congestion control mechanism to advance to Experimental. As an example, HighSpeed TCP is a congestion control mechanism that is Experimental, but that is not TCP-friendly in all environments. We also note that this guideline does not constrain the fairness offered for non-best-effort traffic.

As an example from an Experimental RFC, fairness with standard TCP is discussed in Sections 4 and 6 of [RFC3649] (HighSpeed TCP) and using spare capacity is discussed in Sections 6, 11.1, and 12 of [RFC3649].

(2) Difficult Environments.

The proposed algorithms should be assessed in difficult environments such as paths containing wireless links. Characteristics of wireless environments are discussed in [RFC3819] and in Section 16 of [Tools]. Other difficult environments can include those with multipath routing within a connection. We note that there is still much to be desired in terms of the performance of TCP in some of these difficult environments. For congestion control mechanisms with explicit feedback from routers, difficult environments can include paths with non-IP queues at layer-two, IP tunnels, and the like. A minimum goal for experimental mechanisms proposed for widespread deployment in the Internet should be that they do not perform significantly worse than TCP in these environments.

While it is impossible to enumerate all possible "difficult environments", we note that the IETF has previously grappled with paths with long delays [RFC2488], high delay bandwidth products [RFC3649], high packet corruption rates [RFC3155], packet reordering [RFC4653] and significantly slow links [RFC3150]. Aspects of alternate congestion control that impact networks with these characteristics should be detailed.

As an example from an Experimental RFC, performance in difficult environments is discussed in Sections 6, 9.2, and 10.2 of [RFC4782] (Quick-Start).

(3) Investigating a Range of Environments.

Similar to the last criteria, proposed alternate congestion controllers should be assessed in a range of environments. For instance, proposals should be investigated across a range of bandwidths, round-trip times, levels of traffic on the reverse path, and levels of statistical multiplexing at the congested link. Similarly, proposals should be investigated for robust performance with different queueing mechanisms in the routers, especially Random Early Detection (RED) [FJ03] and Drop-Tail. This evaluation is often not included in the internet-draft itself, but in related papers cited in the draft.

A particularly important aspect of evaluating a proposal for standardization is in understanding where the algorithm breaks down. Therefore, particular attention should be paid to characterizing the areas where the proposed mechanism does not

perform well.

As an example from an Experimental RFC, performance in a range of environments is discussed in [Section 12 of \[RFC3649\]](#)

Expires: December 2007

[Page 6]

[draft-ietf-tsvwg-cc-alt-04.txt](#)

June 2007

(HighSpeed TCP) and [Section 9.7 of \[RFC4782\]](#) (Quick-Start).

(4) Protection Against Congestion Collapse.

The alternate congestion control mechanism should either stop sending when the packet drop rate exceeds some threshold [\[RFC3714\]](#), or should include some notion of "full backoff". For "full backoff", at some point the algorithm would reduce the sending rate to one packet per round-trip time and then exponentially backoff the time between single packet transmissions if congestion persists. Exactly when either "full backoff" or a pause in sending comes into play will be algorithm-specific. However, as discussed in [\[RFC2914\]](#), this requirement is crucial to protect the network in times of extreme congestion.

If "full backoff" is used, this bullet does not require that the full backoff mechanism must be identical to that of TCP [\[RFC2988\]](#). As an example, this bullet does not preclude full backoff mechanisms that would give flows with different round-trip times comparable bandwidth during backoff.

(5) Fairness within the Alternate Congestion Control Algorithm.

In environments with multiple competing flows all using the same alternate congestion control algorithm, the proposal should explore how bandwidth is shared among the competing flows.

(6) Performance with Misbehaving Nodes and Outside Attackers.

The proposal should explore how the alternate congestion control mechanism performs with misbehaving senders, receivers, or routers. In addition, the proposal should explore how the alternate congestion control mechanism performs with outside attackers. This can be particularly important for congestion control mechanisms that involve explicit feedback from routers along the path.

As an example from an Experimental RFC, performance with misbehaving nodes and outside attackers is discussed in Sections

9.4, 9.5, and 9.6 of [[RFC4782](#)] (Quick-Start). This includes discussion of misbehaving senders and receivers; collusion between misbehaving routers; misbehaving middleboxes; and the potential use of Quick-Start to attack routers or to tie up available Quick-Start bandwidth.

(7) Responses to Sudden or Transient Events.

The proposal should consider how the alternate congestion control mechanism would perform in the presence of transient events such as sudden congestion, a routing change, or a mobility event. Routing changes, link disconnections, intermittent link connectivity, and mobility are discussed in more detail in Section 17 of [[Tools](#)].

Expires: December 2007

[Page 7]

[draft-ietf-tsvwg-cc-alt-04.txt](#)

June 2007

As an example from an Experimental RFC, response to transient events is discussed in [Section 9.2 of \[RFC4782\]](#) (Quick-Start).

(8) Incremental Deployment.

The proposal should discuss whether the alternate congestion control mechanism allows for incremental deployment in the targeted environment. For a mechanism targeted for deployment in the current Internet, it would be helpful for the proposal to discuss what is known (if anything) about the correct operation of the mechanism with some of the equipment installed in the current Internet, e.g., routers, transparent proxies, WAN optimizers, intrusion detection systems, home routers, and the like.

As a similar concern, if the alternate congestion control mechanism is intended only for specific environments (and not the global Internet), the proposal should consider how this intention is to be carried out. The community will have to address the question of whether the scope can be enforced by simply stating the restrictions or whether additional protocol mechanisms are required to enforce the scoping. The answer will necessarily depend on the change being proposed.

As an example from an Experimental RFC, deployment issues are discussed in Sections [10.3](#) and [10.4](#) of [[RFC4782](#)] (Quick-Start).

[4.](#) Minimum Requirements

This section suggests minimum requirements for a document to be approved as Experimental with approval for widespread deployment in the global Internet.

The minimum requirements for approval for widespread deployment in the global Internet include the following guidelines (1) on assessing the impact on standard congestion control, (3) on investigation of the proposed mechanism in a range of environments, guideline (4) on protection against congestion collapse and guideline (8), discussing whether the mechanism allows for incremental deployment.

For other guidelines, i.e., (2), (5), (6), and (7), the author must perform the suggested evaluations and provide recommended analysis. Evidence that the proposed mechanism has significantly more problems than those of TCP should be a cause for concern in approval for widespread deployment in the global Internet.

5. Security Considerations

This document does not represent a change to any aspect of the TCP/IP protocol suite and therefore does not directly impact Internet security. The implementation of various facets of the Internet's current congestion control algorithms do have security

Expires: December 2007

[Page 8]

[draft-ietf-tsvwg-cc-alt-04.txt](#)

June 2007

implications (e.g., as outlined in [[RFC2581](#)]). Alternate congestion control schemes should be mindful of such pitfalls, as well, and should examine any potential security issues that may arise.

6. IANA Considerations

This document does not require any IANA action.

Acknowledgments

Discussions with Lars Eggert and Aaron Falk seeded this document. Thanks to Bob Briscoe, Gorrry Fairhurst, Doug Leith, Jitendra Padhye, Colin Perkins, Pekka Savola, members of TSVWG, and participants at the TCP Workshop at Microsoft Research for feedback and contributions. This document also draws from [[Metrics](#)].

Normative References

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

[RFC2914] S. Floyd, Congestion Control Principles, [RFC 2914](#), Best Current Practice, September 2000.

[RFC2960] Stewart, R., Xie, Q., Morneault, K., Sharp, C., Schwarzbauer, H., Taylor, T., Rytina, I., Kalla, M., Zhang, L., and V. Paxson, Stream Control Transmission Protocol, [RFC 2960](#), October 2000.

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

Informative References

[FJ03] Floyd, S., and Jacobson, V., Random Early Detection Gateways for Congestion Avoidance, IEEE/ACM Transactions on Networking, V.1 N.4, August 1993.

[Metrics] S. Floyd, Metrics for the Evaluation of Congestion Control Mechanisms. Internet-draft [draft-irtf-tmrg-metrics-07](#), work in progress, February 2007.

[RFC2488] M. Allman, D. Glover, and L. Sanchez. Enhancing TCP Over Satellite Channels using Standard Mechanisms. [RFC 2488](#). January 1999.

[RFC2988] Vern Paxson, Mark Allman. Computing TCP's Retransmission Timer, November 2000. [RFC 2988](#).

[RFC3150] S. Dawkins, G. Montenegro, M. Kojo, V. Magret, End-to-end Performance Implications of Slow Links, [RFC 3150](#), July 2001.

[RFC3155] S. Dawkins, G. Montenegro, M. Kojo, V. Magret, N. Vaidya, End-to-end Performance Implications of Links with Errors, [RFC 3155](#),

Expires: December 2007

[Page 9]

[draft-ietf-tsvwg-cc-alt-04.txt](#)

June 2007

August 2001.

[RFC3649] S. Floyd, HighSpeed TCP for Large Congestion Windows, [RFC 3649](#), September 2003.

[RFC3714] S. Floyd and J. Kempf, IAB Concerns Regarding Congestion Control for Voice Traffic in the Internet, [RFC 3714](#), March 2004.

[RFC3819] P. Karn, C. Bormann, G. Fairhurst, D. Grossman, R. Ludwig, J. Mahdavi, G. Montenegro, J. Touch, and L. Wood, Advice for Internet

Subnetwork Designers, [RFC 3819](#), July 2004

[RFC4653] Sumitha Bhandarkar, A. L. Narasimha Reddy, Mark Allman, Ethan Blanton, Improving the Robustness of TCP to Non-Congestion Events, [RFC 4653](#), August 2006.

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

[Tools] S. Floyd and E. Kohler, Tools for the Evaluation of Simulation and Testbed Scenarios, Internet-draft [draft-irtf-tmrg-tools-03.txt](#), work in progress, December 2006.

Authors' Addresses

Sally Floyd
ICIR (ICSI Center for Internet Research)
1947 Center Street, Suite 600
Berkeley, CA 94704-1198
Phone: +1 (510) 666-2989
Email: floyd at icir.org
URL: <http://www.icir.org/floyd/>

Mark Allman
ICSI Center for Internet Research
1947 Center Street, Suite 600
Berkeley, CA 94704-1198
Phone: (440) 235-1792
Email: mallman at icir.org
URL: <http://www.icir.org/mallman/>

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

Expires: December 2007

[Page 10]

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.

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.

Copyright Statement

Copyright (C) The IETF Trust (2007). 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: December 2007

[Page 11]