

Network Working Group
Internet-Draft
Intended status: Informational
Expires: January 14, 2009

S. Krishnan
Ericsson
P. Eronen
Nokia
E. Gray
Ericsson
July 13, 2008

Guidelines to authors and reviewers regarding the IETF review process
draft-krishnan-review-process-01

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>.

This Internet-Draft will expire on January 14, 2009.

Abstract

This document describes the authors' experiences with the IETF review process and provides guidelines to draft authors and reviewers on how to effectively participate in it. This document does not define the IETF review process itself.

Internet-Draft

IETF Review Guidelines

July 2008

Table of Contents

1.	Introduction	3
2.	Sources of reviews	3
3.	Guidelines for authors	4
3.1.	Who responds to reviews	4
3.2.	Before reading the review	4
3.3.	Responding to reviews	4
3.3.1.	Timeframe for response	5
3.3.2.	Contents of a response	5
3.4.	Keeping track of the issues	7
3.5.	New version of the draft	7
4.	Guidelines for reviewers	8
4.1.	Level of review	8
4.2.	Recipients of the review	9
4.3.	Classification of the issues	10
4.4.	Prioritization of the issues	10
4.5.	Reviewing changes	11
5.	Acknowledgements	12
6.	IANA Considerations	12
7.	Security Considerations	12
8.	Normative References	12
	Authors' Addresses	12
	Intellectual Property and Copyright Statements	14

1. Introduction

The Tao of the IETF [[RFC4677](#)] provides an informal overview of the inner workings of the IETF. This document is intended to augment the information available in [[RFC4677](#)] and is mainly concerned about reviews.

An Internet Draft's life cycle begins when the author(s) submit the document as a personal draft; it may become a Working Group draft, and usually ends when the draft is published as an RFC, or is abandoned and eventually expires. During its lifetime, a draft is likely to receive review comments from a number people.

In our experience, inefficient handling of review comments is one big source of delay and frustration in the IETF standards process. The goal of this document is to help document authors in dealing with reviews in an effective manner.

2. Sources of reviews

An internet draft gets reviewed at different stages of its lifecycle by different sets of people. When the draft is in the initial stages it gets reviewed mostly by the members of the working group that the draft targets. But as the draft progresses, it gets reviewed by a much more varied set of people with differing fields of expertise. These reviews are performed with a very different point of view than the wg reviews and are likely to bring out cross-area issues. The different sources from which the author can receive reviews are

- o Working Group participants: Reviews from working group participants are not just reviews, but also an important method of participating in the working group. This document deals primarily with late reviews though some of the concepts may be applicable to early reviews (such as reviews from WG participants) as well.
- o Other IETF participants -- especially during IETF last call

- o Review from the area specific review teams. Some review teams try to review all documents -- these include Gen-ART (General Area Review Team) assisting the General Area Director, and SecDir (Security Directorate) assisting the Security Area Directors. Other review teams are more specialized: for example, MIB Doctors review only MIB documents.
- o Reviews from IESG members: after IETF last call,
- o Reviews from the IAB

[3.](#) Guidelines for authors

[3.1.](#) Who responds to reviews

There are several persons who can respond to a received review. If the document is an individual document, it is customary for the author to respond to the review. If the document has several authors, they should co-ordinate among themselves before responding. Otherwise the reviewer may receive conflicting responses from different authors. If the document is a working group document, ideally the WG chairs or the document shepherd should coordinate things, since the author and reviewer cannot unilaterally decide to make substantial changes. The coordination effort (by WG chairs and/or document shepard) may be more or less active, depending on the working group.

[3.2.](#) Before reading the review

Receiving a long review, or a critical review for your favorite document can be disheartening. Before reading the review -- or at the very least, before starting to reply -- it's good to remind yourself about a crucial difference between Internet Draft reviews, and say, reviews about books or movies.

A good movie review attempts to give a balanced view: if the reviewer liked the movie, the review will say good things about it. However, a typical review of an Internet Draft will comment only the parts that the reviewer did not like and wants to see changed.

The reviewer may well think that your work is excellent quality, and a valuable contribution to Internet engineering -- but it is not common to say this in a review. Thus, when you start reading a review, imagine that the review started with a positive comment -- after all, the reviewer considered this important enough to review it!

Hint for reviewers: it helps if you start your review with a positive overall opinion. Although it may not seem so, there's a big difference between saying "The following things have to be fixed:" and "This is a well-written and valuable document; however, I'd like to propose some minor improvements for the following things:"

[3.3.](#) Responding to reviews

Reviewers spend a non-trivial amount of time in performing a review, and reviews also play an important part in how IETF determines consensus. Thus, the single most important rule for handling reviews is the following: every review merits a response.

Krishnan, et al.

Expires January 14, 2009

[Page 4]

Internet-Draft

IETF Review Guidelines

July 2008

Even if the review said "everything looks ok", saying "thank you for reviewing this" is part of common courtesy. If the reviewer raised issues, a response is deserved, irrespective whether the authors agree with the reviewer or not.

[3.3.1.](#) Timeframe for response

The authors are expected to respond to the reviews within a reasonable amount of time. What constitutes a reasonable amount of time is left to the judgement of the persons involved. e.g. If a document is up for IESG review on a certain date, it may be useful to respond before that date. It is possible that the authors may be unavailable to respond to a review within a given timeline since they may be sick/on vacation/busy with dayjob etc. In this case, the document shepherd (if any) can respond to the review and follow up with the authors at a later time. If the review is not responded to in a timely fashion, then other folks are likely to assume that there is a real problem that needs to be addressed, rather than finding out that the reviewer simply misread the document (reviewers make mistakes too).

[3.3.2.](#) Contents of a response

[3.3.2.1.](#) The initial response

The initial response that the author sends might be as simple as a single mail that acknowledges the receipt of the review and the intent of the author to look into the contents of the review. It is often helpful if the responding author(s) can also provide some indication of when a subsequent response will be provided to the comments/issues contained in the review -- assuming that a response cannot be provided at once. This sort of response can be helpful in telling the difference between a review that has been forgotten, and one that is being actively worked on. e.g. If the author is busy with his/her day job or on vacation, the document shepherd can ask some other working group member to evaluate the review.

[3.3.2.2.](#) The substantial response

In the substantial response the author tries to address the concerns raised in the review. There are several ways to do this:

- o Accept the concern as valid and propose a solution
- o Accept the concern as valid and propose a timeline for a solution
- o Accept the concern as valid and solicit text from the reviewer for a solution.

- o Ask for clarifying information to better understand the concern
- o Explain why the concern is invalid (out of scope, already covered in the document, misunderstanding, etc.)
- o Other responses that may apply in special cases
- o Any combination of the above

The substantial response needs to be sent at least to the reviewer, the relevant working groups and the mailing list of the reviewing body(if any). It can also be sent to the other recipients of the original review.

[3.3.2.2.1.](#) Common issues with the substantial response

This section tries to document some common mistakes that authors make while responding to a review.

- o If the reviewer has not understood some part of the draft, a very common response is to explain the topic in email. However, often that explanation should really be in the draft itself, not just in emails to the authors, so that future readers will also understand the text.
- o Soliciting text from the reviewer should not be used as a default response. It can be helpful, if the authors, even after asking for more information, do not understand the reviewer's concern. It can also be used as a technique for reducing the number of iterations required to arrive at text that is acceptable to the reviewer, or to help in producing review comments that are direct, with a specific focus toward some resolution. Note that, if the author understands the concern, it may be best if an author proposes text -- since text proposed by an author may be less likely to introduce stylistic discontinuities, or result in other issues that an author (who is likely to be more familiar with the history of the draft) might avoid.
- o The reviewer should not be pushy while trying to suggest new text. Instead of writing "Here is the fix for my problem" that imposes the reviewer's will on the author, he or she might write "Here is my proposal for one way to fix this issue; there may of course be other ways to do so. Please consult with the WG and decide on the text to adopt".
- o Simply accepting the concern as valid is a response that is often seen, but is really good enough only when the solution is obvious (fixing a typo, etc.). For complex comments, saying "I'll change this in the next draft version" may not be enough, since it it may make it difficult for the reviewer to verify the fix, or may result in a fix that is not acceptable to the reviewer.

[3.4.](#) Keeping track of the issues

When the author makes an attempt in good faith to resolve all the issues raised by the reviewer, it is possible that some of the issues are left unresolved in the revised version of the draft. In order to prevent this from happening, it is useful to keep a systematic record of the issues and the associated resolutions. One common way of doing this is by using an issue tracker. The IETF tools team

provides an issue tracker on request to any of the working groups and the chairs can add authorized users for the tracker. When the issues are raised by the reviewer, the author can open the issue on the issue tracker and close it when it is resolved. The author can also add the suggested resolution text into the tracker. This way both the author and the reviewer can keep tabs on the change process without missing out any issues.

For a document in Working Group Last Call or IETF wide Last Call it is considered good practice for the WG chair or the document shepherd to ensure that someone posts a summary of all the comments received during the last call, and current proposal for handling them. If the WG chair or shepherd does not do this, the task of producing a summary falls to the author, or authors. Note that - if no such summary is produced - it will be very difficult for working group members (or the IETF at large) to determine that the changes made were in line with the changes discussed and agreed to. This summary could include links to mailing list archives, or if an issue tracker is used, issue numbers or links to "tickets".

[3.5.](#) New version of the draft

The details of dealing with response resolutions may vary slightly from working group to working group, but -- generally -- the following approaches are used at the indicated stages:

- o Prior to acceptance as a WG draft, Author(s) are free to publish a new version with or without including specific resolutions - bearing in mind that acceptance by a WG is only likely if a) the text is essentially headed in the right direction and b) the WG is inclined to believe the current Author(s) are the right ones for the job.
- o Once accepted as a WG draft, and through completion of a final WG Last Call, the Author(s) need confirmation (typically from discussions on the working group mailing list, or at meetings and subsequently confirmed on the mailing list) of all substantive changes prior to publishing a modified version. Prior to last call, the necessary confirmation may simply be a general direction proposed and accepted at an IETF meeting and not refuted in minutes published. Often prior to last call, the Author(s) may

explicitly solicit help from WG members who've expressed concerns

about specific portions of draft in an effort to reach convergence in a subsequent IETF WG meeting. Once Last Call has started, each substantive change will typically require explicit confirmation on the WG mailing list.

- o Once the document has entered IESG processing (post WG Last Call), new versions should not be posted unless the document shepherd or (responsible) AD explicitly says so. Typically, this will occur only after the responsible individual is certain that the Authors have addressed all of the outstanding comments - both substantive and editorial.
- o At any point after WG Last Call, the IESG may decide that all remaining comments may be addressed via a note to the RFC Editor. Typically, this would occur if all remaining comments are relatively minor, strictly editorial or of a type that would best be dealt with by the RFC Editor in any case.

[4.](#) Guidelines for reviewers

[4.1.](#) Level of review

The level of review performed on the draft varies greatly based on the source of the review and on the stage of the document lifecycle. In the earlier stages of the lifecycle, the document is usually reviewed by the members of the relevant working group(s) only, although the working group chairs and (usually) ADs should take at least a brief look at the draft. Let's call these reviews "early reviews".

At this point the comments on the draft are deeply technical and mainly remove obvious errors from the draft. It is also a good idea for working group members, working group chairs and/or ADs, to review early drafts for "meta issues" (such as Internet Architecture impact, special security exposures, manageability, scale, consistency with chartered goals, etc.) that may be a problem later and result in potentially going back to the drawing board. Reviewer(s) and author(s) are usually from similar backgrounds and are able to understand the underlying technologies and jargon in a consistent fashion.

The reviews that occur after the draft has been progressed from the working group have a very different focus with respect to the contents of the draft. They explore broader aspects of the draft (such as security, operations, readability, etc.) that are not usually on the top of the list of priorities of the authors and the early reviewers - at least during most reviews. They also explore how the mechanisms proposed in the document fit into the Internet

architecture as a whole and what detrimental effects they will have, if any. Note that these reviews are where it will come out if the very earliest reviews did not consider where the drafts were going because it is at this point that a poorly directed draft is very likely to stall. Let's call these "late reviews".

The person performing the review should consider this when the review is being performed and concentrate their efforts on the right parts of the document.

[4.2.](#) Recipients of the review

The list of recipients of the review is tricky to get right. The main idea is to make sure all the relevant people receive the review. The recipient list is determined mainly by the following factors.

- o The timeframe of the review (early vs. late)
- o The contents of the review (editorial vs. technical)

Early reviews are usually performed by active participants of a working group. The preferred destination for these reviews is the working group mailing list since it can be reasonably assumed that the persons interested in the document are subscribed to the mailing list. This applies for both technical and editorial issues. Alternately editorial issues can be resolved using a private mail to the author(s).

Late reviews are usually performed by persons who are not active participants of the working group and who usually review the draft from a different point of view than the working group. If the contents of the review are mainly editorial in nature, the reviews can be sent just to the authors, the working group chair(s), the document shepherd(s). If the review is of a more technical nature it is considered polite to include the working group mailing list and/or the IETF discussion list. As it is not reasonable to assume that the reviewer will subscribe to the working group mailing list just for discussing this issue, the working group chair(s) need to make sure that this review will get past any moderation controls imposed on non-subscribers by the working group mailing list.

[4.3.](#) Classification of the issues

While writing a review, the reviewer needs to distinguish between different types of comments. On a really high level review comments may be classified into two types

- o Actionable comments: These comments contain statements that can be either proved or disproved. e.g. "The protocol described in this document will not work in networks with lossy link layers". This creates the opportunity for the reviewer and working group to discuss the issue in terms that can be tested. In order to ask the working group to do so, a reviewer should be able to back his/her statement with a reasonably clear, worked set of reasons for the assertion. A working group can go about disproving the statement, and if it fails to do so, it can either make the necessary technical changes or write up an applicability statement that documents the cases where the protocol works well enough to be useful.
- o Observations: These comments are not testable statements but merely opinion/commentary of the reviewer. e.g. " I think this protocol should run on top of TCP rather than on top of UDP because it needs reliable transport". This kind of comments may be very valuable as well, but they should not blindly overrule the consensus of the working group. These should be treated as issues that the wg and the author need to evaluate carefully.

Reviewers should be sensitive to the difference between their personal opinions (and preferences) and issues that will affect the correct operation or interoperation of the documents under review.

[4.4.](#) Prioritization of the issues

A review often results in a list of issues or requested clarifications. The reviewer may choose to list them in the order they appear in the document. For long drafts, where structure and content maturity is fairly well established, this approach is easier to follow. Alternatively (and possibly preferably, depending on the maturity of the draft and other factors), the reviewer can use another classification scheme where the issues are grouped together

by importance and/or potential impact on the draft. This makes it easier for eventual recipients of the review to attach the appropriate priority for resolving the issues. At early stages of draft maturity, this approach can be critical, because resolving some issues may impact on appropriate resolution alternatives for other issues. If a review is sent on behalf of a reviewing body (such as might be the case with a review reported by a liaison statement), this also helps the beneficiary of the review in taking a position on the document. For example, A reviewer might classify issues into the

following categories

- o Major
- o Minor
- o Editorial
- o Nits

The purpose of such a prioritization scheme might be to indicate the level of discomfort the reviewing entity will feel with the results of any resolution that does not address the concerns at each level of priority.

It is almost always a good idea to separate at least the editorial comments (and NITs) from those impacting on the substance of the draft at any stage of draft maturity. A key reason for doing this is that this allows most people who will read the review to concentrate on the portions of the review that impact on the substance of the document, primarily to ensure that they do not have fundamental issues with the review comments or proposed resolutions. Note that this is not a license to put forward omission of key words (such as "not") as minor editorial comments (or NITs).

If comments are grouped in any way, however, the grouping may very well be "of the essence" of the comments. In other words, if the responder does not agree with the distinctions made by the reviewer, they should be clear about this. This is, in part, to ensure that something that a reviewer thought to be minor -- but which an author felt to be fundamental -- receives the right amount of attention on reading by others.

[4.5.](#) Reviewing changes

Once the authors submit a new revision of the draft, the reviewer can look over the new revision to see if the agreed changes have been made to the draft. If the reviewer finds out that some changes have not been made, he/she can alert the authors of this fact. There are tools available that can make the reviewer's life easier in this regard. The rfcdiff tool can be used to identify the changes made in the latest version of the draft. The reviewer can just look over these changes instead of rereading the entire draft. The reviewers can also keep track of issues using the issue tracker used by the authors(if one was used). If the reviewer is satisfied with the changes, he/she can send out a mail to the original recipients of the review mentioning this. If not, a new generation of the review cycle is started and the steps described earlier are redone.

Krishnan, et al.

Expires January 14, 2009

[Page 11]

Internet-Draft

IETF Review Guidelines

July 2008

[5.](#) Acknowledgements

The authors would like to thank Bernard Aboba, Thomas Narten, Frank Ellermann, Ted Hardie, Paul Hoffman, Thomas Goldbeck-Lowe, Scott Brim, Joel Halpern, Brian Carpenter, Iain Calder, Dan Romascanu and SM for their contributions to this document.

[6.](#) IANA Considerations

This document does not require any action from the IANA.

[7.](#) Security Considerations

This document does not create any new security issues.

[8.](#) Normative References

- [RFC4677] Hoffman, P. and S. Harris, "The Tao of IETF - A Novice's Guide to the Internet Engineering Task Force", [RFC 4677](#), September 2006.

Authors' Addresses

Suresh Krishnan
Ericsson
8400 Decarie Blvd.
Town of Mount Royal, QC
Canada

Phone: +1 514 345 7900 x42871
Email: suresh.krishnan@ericsson.com

Pasi Eronen
Nokia Research Center
P.O. Box 407
FI-00045 Nokia Group
Finland

Email: pasi.eronen@nokia.com

Krishnan, et al.

Expires January 14, 2009

[Page 12]

Internet-Draft

IETF Review Guidelines

July 2008

Eric Gray
Ericsson
900 Chelmsford Street
Lowell, MA
USA

Email: Eric.Gray@Ericsson.com

Full 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.

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 DISCLAIMS 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.

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.