

August 2020
Geoff Huston

The Making of an RFC in today's IETF

I'm a co-author (or is that "co-editor" in today's politically correct environment?) of an Internet Draft that is closing in for publication as an RFC. It has gone through the Full Monty of the current IETF standardization process, including the steps of document review for Working Group adoption, further cycles of Working Group review including a Working Group Last Call, various Area Reviews, including the Operations Area directorate (OPSDIR), the General Area (GENART), Routing Area (RTGDIR), Internet Area Directorate (INTDIR) and the Transport Area (TSVART). The document has also been reviewed by the Internet Engineering Steering Group (IESG), including various individual AD comments, most of whom arrived at the rather non-committal (or should that be "wimpy"?) ballot position of "No Objection". And finally, a review by the RFC Editor team. It's still not a published RFC as there is the final call from the RFC Editor to the authors in the so-called AUTH48 (a 48-hour last chance review, which bears the distinction of being anything but 48 hours) as a final chance to amend the document. All up, the document preparation process for this particular document has taken two years and five months and the draft has been revised 22 times. Perhaps you might believe that this 29-month ageing process makes the output somehow better. You might think that the result of this truly exhaustive document review process is some bright shiny truth that is stated with precision and clarity. We collectively appear to believe that if we simply follow the procedure the output will be invariably a work of high quality. But that is not necessarily so.

There are some current 1,866 internet drafts in the IETF's internet drafts repository. On average, each current draft document has already been revised 5 or 6 times (the average is 5.44 revisions per draft). Each internet draft document in the current draft repository has been in development for an average of 2 years and 17 days. The draft that has been sitting in the draft repository for the longest time is draft-kunze-ark, which has been an Internet Draft for 19 years 1 month and 30 days, and has been revised a total of 24 times in this period. The distribution of draft ages of current drafts is shown in Figure 1.

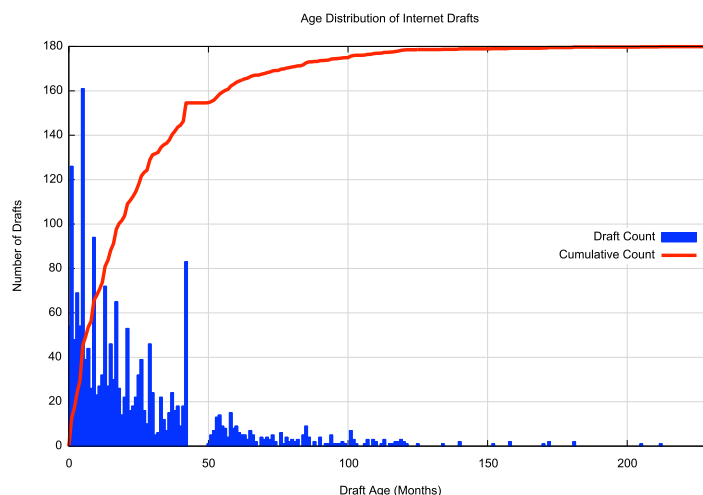


Figure 1 – Age Distribution of internet drafts

What's wrong with taking an average of two years for the IETF process to take a draft of an idea and figure out whether to publish it as an RFC or abandon it? Perhaps that's perfectly reasonable for a staid and venerable twentieth century standards body that submits working drafts to ponderously slow review. But that's not what we wanted from the IETF back in the 1990's. We wanted the IETF to reflect a mode of operation that preferred code over text and emphasised the interoperation of running code as the primary purpose of IETF-published specifications. The intent of an IETF specification was to allow a coder to create a working implementation that would interoperate with all other implementations of the same protocol. A quote attributed to Dave Mills (the man behind the NTP time protocol) from

2004 embodies this purpose: "Internet standards tended to be those written for implementers. International standards were written as documents to be obeyed" ¹ The intended values of the IETF process were to be open, agile, practical, realistic and useful. This was deliberately phrased to be a counterpoint to incumbent model of international technical standards of the time, encompassed under the umbrella of the Open Systems Interconnect (OSI) protocol suite. The OSI effort was summarily dismissed in IETF circles at the time as nothing more than paperware about vapourware, while the IETF attempted to base their effort on running code and a collection of openly available software. Documentation was seen as an aid to interoperability, not a prescriptive statement of required compliance.

In retrospect, this characterisation of the IETF is best seen as the naive idealism of an organisation in its youth. With the waning of the OSI effort the Internet was no longer the challenger but the incumbent and the IETF was the keeper of the keys of this technology. What was obvious in retrospect was that the specifications, in form of RFCs, would change their role and rather than being a commentary on working software where the software was ultimately the reference point, the specification itself became the reference point. Specifications, in the form of RFCs, transformed from being descriptive to being prescriptive. As a natural consequence of this change the IETF started to obsess about the accuracy, precision and clarity of its specifications. And such an obsession is not necessarily helpful.

A look over the errata list maintained by the RFC Editor ² reveals some 3,061 errata records, of which 1,436 have been verified, 534 are reported but not yet verified nor rejected, 605 held for document update and 486 have been rejected. The RFC numbering sequence is up to 8,881 these days so that's an average erratum reporting rate of 1 in 3 RFCs. In some ways that might appear to be a low rate, but considering the intensity of document review that the IETF indulges in the question still needs to be asked: Does all this structured review activity really produce helpful outcomes?

Let's go back to the Internet draft that I referred to above. This particular document was first drafted in April; 2018. In an earlier generation of the IETF the draft would have been published in a month or two after the initial draft with the implicit commentary that "well, this is what we think today" and its publication would've invited others to comment and post revisions, updates or operational experience. But not today. These days we've invented a baroque process for review of this document that has involved 29 months, 22 revisions of the document, and numerous reviews, as already noted. And the result is that the revision of the document (revision 17) that passed RFC Editor review is still one that contains errors, terminology issues and clumsy elisions. Given that RFCs are intended to be prescriptive these days, then doubtless we should fix the document. But to correct these problems looks like it may require more than a change in punctuation or grammatical style. It could require a change in the underlying content of the document. And that would involve returning the document back into the IETF's document review process for yet another round of Working Group review, cross area review, IETF last call review, IESG review and RFC Editor review. It's perhaps not surprising that, as individual volunteers who are not getting paid by the hour, few of us are enthusiastically willing to head back into the grinder for another cycle of this process.

What's the alternative? Another option, and a painless one at that, is to just let through and allow it to be published. Or we could excise the offending text and make minor changes to what's left and hope that is a small enough change to be considered as "editorial" in nature and avoid a tedious second round of reviews. The document has all the prerequisite approvals, so doing absolutely nothing, or as little as possible, certainly represents a path of least pain. The result is uncomfortable. Another RFC is published that contains precisely those issues that this extensive and massively ornate review process was meant to catch.

If this is a failure of the IETF document review process, then the IETF has two options, as I see it:

- The IETF could add even more reviews by even more volunteers. This option seems head down a path that expresses the forlorn hope that by involving ever increasing proportions of the world's population in reviewing a draft document before publication, the process will eventually manage to catch every last issue and a bright shiny truth will emerge!
- The other option is to strip out all the cruft of extensive reviews and publish quickly. Yes, it won't catch all of the errors, omissions, elisions and terminology errors, but the current process isn't performing very well on that score either. In the case of this particular draft we would've saved up to 22 document revisions, a good part of the 28 months of elapsed time and the cumulative sum of hundreds of hours of volunteer review time in the process and we probably would've ended up with a document of comparable quality to the initial draft.

¹ "Rough Consensus and Running Code' and the Internet-OSI Standards War", Andrew Russell, <https://www2.cs.duke.edu/courses/common/compsci092/papers/govern/consensus.pdf>

² https://www.rfc-editor.org/errata_search.php?rec_status=15&errata_type=2&presentation=table

Ever since Edsger Dijkstra's letter "Go To Statement Considered Harmful" was published in 1968 ³, the term "considered harmful" and its variants are seen in Computer Science circles as a form of censure. The internet draft I'm referring to here is one titled "IP Fragmentation Considered Fragile" ⁴. However, it's perhaps not IP fragmentation that is my real concern here. I suspect that it's the baroque IETF document review process itself that could be considered harmful. It's a process that can be forbidding, tediously slow, prone to unpredictable outcomes and subject to uncomfortable compromises.

Can we improve this situation?

Adding even more reviews to the process really won't help, as it's not the IETF's review process itself that is the root cause of the problem here. It's a symptom, but not the cause. Lurking underneath is the more subtle change in the role of RFCs over the past thirty years or so from being descriptive to being prescriptive, and the desire to promote the notion that RFCs are fixed immutable documents. Maybe a related issue is that after seeing off the OSI effort and assuming the mantle of the incumbent standards body for the data communications industry, the IETF has strived to be an outstanding exemplar of a twentieth century paper-based prescriptive standards publication body. And it's unclear to me how this stance is truly helpful for a twenty first century Internet.

Perhaps the IETF's earlier descriptive approach, where running interoperable code was ultimately the arbiter of what it meant to conform to an Internet standard, was actually a more useful form of standard. After all, it's running code that really should matter, not loads of ascii text, however reviewed and polished that text may be!

And what about that draft on IP Fragmentation? The draft document takes 28 pages of ascii text to convey a very simple message: If you are planning to use packet fragmentation in the public IPv4 Internet then you should be aware that it's not all that reliable. You should have a Plan B ready for those situations where packet fragmentation causes silent but consistent packet drop. For the IPv6 public Internet the problem of fragmented packets being discarded is far worse than IPv4. You should make your Plan B your Plan A!

Postscript

It appears that my assumption that the IETF RFC process is even slightly concerned about the quality of the resultant documents is sadly misplaced, at least in the view of some of the IESG Area Directors. According to one IESG member the individuals named as authors on RFCs are merely editors of Working Group consensus documents. This Area Director's view is that the Working Group consensus process owns the content.

The implication of this view is clear. When a draft makes it to AUTH48 last call then the document is at that point sacrosanct, and any changes to the document at this point are unsanctioned. I feel various levels of dismay that there are members of the IESG who, as the managers of the IETF document process, apparently value blind and absolute adherence to the process itself over and above any consideration as to the quality of its outcomes.

Is this my problem? As an author I may well feel an obligation to care a whole lot about accuracy, clarity and precision in published work attributed in whole or in part to me, but as a mere document editor with no direct responsibility for its content then my care factor is a whole lot less. So, it's not my problem I guess.

However, it is someone's problem. It's a problem for folk who choose to rely on these RFCs, who have to work through with incredible care in figuring out what's useful information and what's fanciful BS. Not everything you read in an RFC, no matter how polished the text and references, is necessarily a complete and accurate description of the situation. I wish them luck in sorting out the difference between useful truths and a collection of lazy errors!

³ "go to to statement considered harmful", Edsger W. Dijkstra, Letters to the Editor, Communications of the ACM Vol. 11, No. 3, March 1968. (<https://cacm.acm.org/magazines/1968/3>)

⁴ "Fragmentation Considered Fragile", internet draft, Ron Bonica et al., <https://datatracker.ietf.org/doc/draft-ietf-intarea-frag-fragile/>. The draft's title is a play on an earlier 1987 paper "Fragmentation Considered Harmful" by Christopher Kent and Jeffrey Mogul (<https://www.hpl.hp.com/techreports/Compaq-DEC/WRL-87-3.pdf>)

Disclaimer

The above views do not necessarily represent the views or positions of the Asia Pacific Network Information Centre.

Author

Geoff Huston AM, B.Sc., M.Sc., is the Chief Scientist at APNIC, the Regional Internet Registry serving the Asia Pacific region.

www.potaroo.net