

Errare humanum est: What do RFC Errata say about Internet Standards?

Stephen McQuistin*, Mladen Karan[†], Prashant Khare[†], Colin Perkins*, Matthew Purver^{†‡}, Patrick Healey[†],
Ignacio Castro[†], and Gareth Tyson^{†§}

*University of Glasgow, [†]Queen Mary University of London, [‡]Jožef Stefan Institute

[§]Hong Kong University of Science & Technology

sm@smcquistin.uk, {m.karan,p.khare}@qmul.ac.uk, csp@cspkins.org, {m.purver,p.healey,i.castro,g.tyson}@qmul.ac.uk

Abstract—Protocol standards, such as RFCs developed by the IETF, are crucial for the correct operation of the Internet, but many are published containing errors. The RFC Editor allows people to report *errata*, allowing anybody to flag such errors for subsequent correction. This represents an important part of the RFC publication process, and may reveal ways in which standards can be improved. This paper performs the first study of RFC errata reports. We characterize and perform a statistical analysis of the *scale* and *nature* of these errata and explore *who* submits them. Finally, we evaluate the impact, in terms of the number of errata filings, of three different strategies that are designed to improve the standards process. We find that specialist review teams and formal language checkers can reduce the volume of errata filed against standards documents.

I. INTRODUCTION

By ensuring interoperability, protocol standards are crucial for the correct operation of the Internet. An important standards organisation is the Internet Engineering Task Force (IETF). The IETF follows an open and consensus-driven process for developing standards that is effective at identifying solutions that meet technical and operational needs. This process is time-consuming: the median time between the submission of an initial draft and the publication of a standard was 1,170 days in 2020 [11]. It is argued that this laborious process is appropriate, since it results in standards that reflect the needs of a broad heterogeneous community and achieving consensus in such a disparate group takes time. An open question, though, is whether, at the end of this process, the IETF develops standards that are broadly *correct*?

While the protocol standards the IETF publishes in the RFC series¹ are immutable, mistakes can and *do* occur. These errors can be reported to the RFC Editor, and the RFC Editor makes the reports publicly available and coordinates with the IETF on their verification. These *errata* can clarify editorial concerns as well as correct substantive technical errors. Since the presence of significant errors in its specification can undermine the success of a protocol, studying the nature of these errata and their subsequent fixes is important.

To this end, we have analysed the 6,759 errata filed with the RFC Editor between 2001–2022, inclusive, documenting

¹RFC used to stand for *Request For Comments*, but the RFC series of documents has evolved over the past fifty years to become the archival publication venue for Internet standards and other documents [6].

3,288 editorial issues and 3,471 technical issues, and covering 2,240 RFCs (§II). We start by inspecting the characteristics of the errata (§III). We find that a subset of technical areas (*e.g.*, RFCs relating to security protocols) accumulate the majority of errata, with a particular focus on technical concerns. We posit that this is emblematic of the nature of these areas. The trends, however, have evolved across time. For instance, whereas initially filing technical errata was more common, since 2004 far more editorial concerns have been raised. These changes are reflected in other metrics too. We find that there are significantly more unverified errata in recent years, triggering questions over how effective the mechanism currently is. All of these point to the potential of better automated tooling, both during document preparation and errata filing.

We then look at the people filing errata (§IV). While most (66.3%) submit only a single errata, there is a dedicated group (1%) who submit in excess of 20. There are meaningful differences between the people involved, with some targeting individual areas and others spreading errata across many areas. Further, we find some people with a large number of rejected errata, suggesting that the quality of errata varies, potentially creating a burden for the RFC Editor and errata reviewers.

With increasingly complex RFCs, and a significant backlog of unverified errata, there's growing interest in the IETF around mechanisms to improve and sustain the quality of Internet standards. A number of strategies have emerged, and we conclude with a preliminary analysis of three approaches that see active interest in the community (§V). First, we evaluate the impact that cross-area review of documents has on the number of errata filed. While broad cross-area review has long been promoted as a key strength of the IETF process, and likely has other benefits, we do not find a significant correlation between the number of areas represented in the e-mail discussion of a draft and the number of errata that are filed against the resulting RFC. Next, we do find that specialist document review teams can significantly reduce the volume of errata filings, suggesting that *targeted* expert review is useful. Finally, we find that where formal and structured languages are used in documents, these should be supported by automated tools that validate and check them; such tooling reduces the volume of errata filings.

We hope that our findings will be helpful in identifying directions for potential process improvements in the IETF that could reduce RFC errata filings and improve document quality.

II. BACKGROUND AND DATASETS

IETF Protocol Standardisation. Standards development is a collaborative activity. The process begins with a draft document submitted to the IETF, followed by multiple rounds of review, feedback, and discussion. Review takes place in IETF Working Groups (WGs), themselves organised into thematic areas (e.g., security, transport, routing) within the IETF. Once the community agrees a draft is technically ready for publication as an RFC, it is passed to the RFC Editor for copy editing and publication. The RFC Editor publishes RFCs from five publication streams including outputs from the IETF and RFCs from the Internet Research Task Force (IRTF), the Internet Architecture Board (IAB), the Editorial Stream², and independent submissions [7]. In addition, a sixth stream is maintained for legacy RFCs that were published prior to the streams separating in 2007 [9].

Errata. Once published, an RFC cannot be changed. However, anybody can submit errata, highlighting what they perceive to be errors in the document. Since 2000, the RFC Editor has maintained a public database of these errata filings. Errata are separated into *technical* and *editorial*. Technical errata are filed for mistakes in the technical content of the RFC that are likely to result in incorrect, non-conforming, implementations of the standards (for example, *erratum 5996*, which reports a mistake in the pseudocode of an algorithm). Editorial errata include spelling and punctuation errors that do not otherwise impact the technical content (like *erratum 5385*, which reports a grammatical error). There is often some overlap between these two classifications: for example, *erratum 5595* highlights that “*not*” is missing within a specification. This is arguably editorial, but is filed as technical, given that it fundamentally alters the meaning of the text. Errata filed against RFCs published on the IETF stream are checked for correctness by the RFC authors, working group chairs, and area directors. IRTF errata are checked by the authors and the Internet Research Steering Group. The IAB and Independent Submissions Editor check errata for their streams.

Datasets. We analyse the 6,759 errata reports recorded by the RFC Editor from January 2001 through to the end of December 2022 as made available by the RFC Editor.³ Table A1, in the Appendix, lists the fields in the errata database. Of note, in addition to the technical vs. editorial classification mentioned above, is that errata can be in one of four states: *reported*, *verified*, *rejected* (invalid, or significant enough that a new RFC is needed), or *hold for document update* (not a necessary update to this RFC, but worth future consideration).

We combine the errata database with data from a number of other sources. The *RFC index*⁴ provides metadata about RFC publications, including author names, the publication stream, and for RFCs published on the IETF stream, the working group

and area that developed the RFC. The *Datatracker*,⁵ provides information about draft documents, citation data between IETF documents, and comprehensive authorship metadata. Finally, the *e-mail archives* of public IETF lists⁶ allow us to study activity of participants within the IETF community.

Ethical considerations. All data used in this paper was extracted from public archives and APIs of the IETF and RFC Editor. To ensure that our access to these services did not cause operational problems for the IETF, we were in regular contact with the IETF Tools Team and Secretariat, as well as the operators of the Datatracker. We discussed our work with the RFC Editor and IETF leadership (IETF, IAB, and IRTF Chairs, the IETF Executive Director) to ensure that our access falls within their acceptable use policies. Participants in IETF agree to the policies described at <https://www.ietf.org/about/note-well/>, and <https://www.ietf.org/privacy-statement/>. These make explicit provision that information in the Datatracker system will be made public. The RFC Editor follows the same privacy policy.

Code. The code used to produce this paper is available from <https://dx.doi.org/10.5281/zenodo.8008032>.

III. CHARACTERISING ERRATA

Errata over Time. Figure 1 presents the number of errata filed, on average per RFC, since 1969 based on the year of RFC publication. The peak in the number of errata per RFC occurs in 1981. Only 29 RFCs were published that year, but they include major documents such as RFCs 791, 792, and 793 (the original versions of the IP [15], ICMP [14], and TCP [16] standards), with 17, 7, and 47 errata, respectively. These important protocols clearly garnered a great deal of scrutiny and revision. The second highest peak occurs in 2006. In contrast to the previous examples, this has the highest number of RFCs published per year (459), including RFC 4601 [5] that has the most errata (114). Since this second peak, there has been a steady decrease in the number of errata filed. This broadly correlates with the number of RFCs published per year, with Pearson coefficient 0.59 since 2007. Table I lists the top RFCs by errata filing count.

Errata Delay. We next explore how long it takes for errata to be identified and filed. Figure 2 presents a CDF of the number of days between RFC publication and the errata being filed, broken down based on IETF area. We see a wide range of delays. 7.3% of errata are filed within the first 30 days, suggesting that many RFCs are published with issues that could have been identified prior to publication. RFCs from the General (*gen*) area—describing IETF policies and procedures—have the longest delay, with a median of 3,458 days, compared to the Applications and Real-time (*art*) area with a median of 681 days.⁷ Editorial errata are typically filed more quickly, with a median of 987 days, compared to a median of 1,138 days for technical errata.

²The Editorial Stream is newly created [17] and currently contains no RFCs.

³RFC errata are publicly available at <https://www.rfc-editor.org/errata.php>. We make reference to specific filings by *Errata ID*; this can be resolved via the RFC Editor website. We thank the RFC Editor for making the underlying database available to us in machine-readable form for this analysis.

⁴<https://www.rfc-editor.org/rfc-index.xml>

⁵<https://datatracker.ietf.org>

⁶<https://mailarchive.ietf.org>

⁷Errata are filed against RFCs within the *subip* area within a median of 48 days, but this is skewed, with only 19 RFCs being published in that area.

RFC	Title	Year	Area	Filing count
4601	Protocol Independent Multicast - Sparse Mode (PIM-SM): Protocol Specification (Revised)	2006	rtg	114
4880	OpenPGP Message Format	2007	sec	52
793	Transmission Control Protocol	1981	None	47
4634	US Secure Hash Algorithms (SHA and HMAC-SHA)	2006	None	44
5661	Network File System (NFS) Version 4 Minor Version 1 Protocol	2010	tsv	42
1345	Character Mnemonics and Character Sets	1992	app	41
8446	The Transport Layer Security (TLS) Protocol Version 1.3	2018	sec	40
5545	Internet Calendaring and Scheduling Core Object Specification (iCalendar)	2009	app	35
3261	SIP: Session Initiation Protocol	2002	rai	33
5905	Network Time Protocol Version 4: Protocol and Algorithms Specification	2010	int	32

Table I: Top 10 RFCs by errata filing count

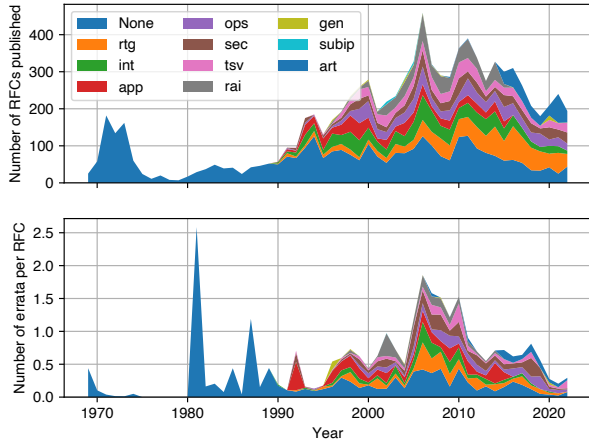


Figure 1: Average errata filed per RFC for each year by RFC publication year, grouped by IETF area (acronyms expanded in Table A2).

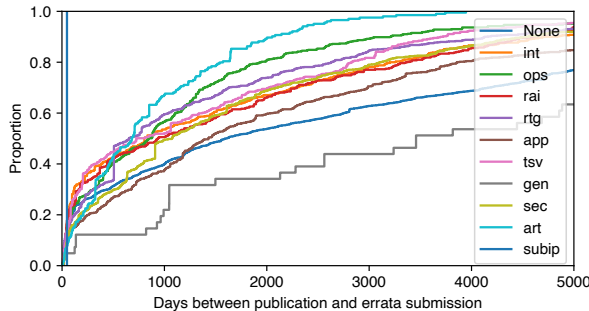


Figure 2: Days from RFC publication and errata filing by IETF area (acronyms expanded in Table A2).

Errata Status. Figure 3 categorises the errata by status and publication year of the RFC to which they relate. The largest share (42.5%) of errata are *verified*: errata that have been confirmed as necessary and accurate. This suggests that many errata are useful to the community. The next largest share (30.3%) are those labelled *hold for document update*. These are errata that are not a necessary update to the RFC, but may be considered on future revisions. For example, *erratum 6278* describes an oversight in RFC 8610 [3]; the solution to this is non-trivial, and so will be considered in the next version of the specification. Of the 930 RFCs that have *hold*

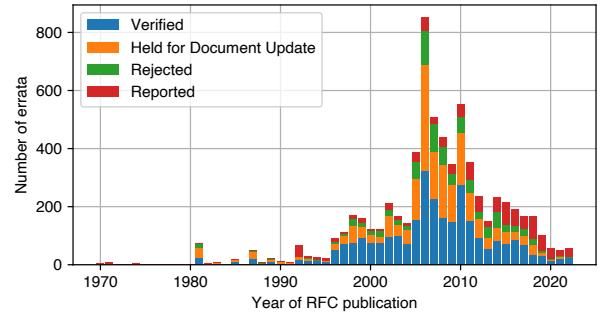


Figure 3: Errata filings by status, by publication year of the RFC.

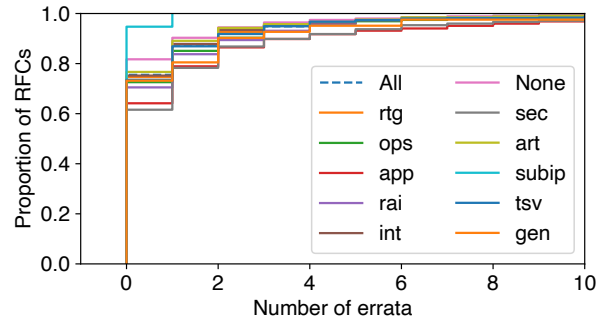


Figure 4: CDF of errata filed per RFC by grouped by IETF area (acronyms expanded in Table A2).

for document update errata filed against them, only 40% have been updated or obsolete by a subsequent RFC. We flag that this may be a cause for concern, or at least a missed opportunity for improvements to standards. The third largest category (13%) is *rejected*, which covers errata that are invalid (like *erratum 6323*, which was rejected because the original text was understood to be correct) or proposes a significant change to the RFC that should be done by publishing a new RFC (like *erratum 5814*, which was rejected for proposing a significant change, rather than reporting an error). Such a large fraction of rejected submissions is unexpected and may flag issues with people’s understanding of the errata process and its place within the wider standardisation process. Finally, 14.2% of errata are *reported* but unverified. Again, we are surprised to see unverified errata from over a decade ago, suggesting the process should be expedited.

Errata per RFC Area, Status, and Stream. Figure 4 shows a CDF of the number of errata filed, per RFC, in each IETF area. Non-IETF RFCs, *e.g.*, IRTF and independent stream RFCs, and legacy IETF RFCs, are labelled as “None”. We confirm errata in standards are common: of the 4,373 standards-track RFCs in our dataset, 32.7% have attracted at least one erratum filing. However, there are three notable outliers. First, RFCs published by the Sub-IP (subip) Area have very few errata, with only 5% of subip RFCs attracting errata filings. This is because this temporary area – established in 2001 and concluded in 2005 – only published 19 RFCs, resulting in a far smaller sample than the other areas. For comparison, the next smallest area, General (gen), published 39 RFCs. *gen* RFCs attract a greater number of errata on average, *vs.* *subip* RFCs, likely due to their broader relevance. Second, we see that both the Application (app) and Security (sec) Areas’ RFCs are more likely to have errata filed for them, with 35.9% of Application and 39% of Security RFCs attracting at least one erratum filing.

Finally, Table A2 details the split between *technical* and *editorial* errata across each area. While there is broadly an even split, there are areas where one type of errata is more dominant. For example, in the Routing (rtg) area, 60.8% of filings are editorial, while in the Applications (app) area, 59% were technical. It remains to determine why this is the case, and, in particular, to establish whether there is something inherent about the RFCs published by these areas that makes them more prone to containing errata, and to containing one type of errata *vs.* another. For example, in the Routing area, structured notation is frequently used to define routing entities; editorial errata are often filed in those definitions. Targeting such areas with improved alternate review procedures may be beneficial.

Tables A3 and A4 further categorise errata by the stream and status, at the time of publication, of each RFC. As expected, the majority of errata are filed against IETF RFCs and *Proposed Standards* since these make up the majority of RFCs that are published. However, there are notable differences in the average number of filings per RFC. *Proposed Standards* (1.01 errata per RFC), *Draft Standards* (1.93), and *Internet Standards* (2.17) attract a far higher number of errata per RFC than *Informational* (0.52) or *Experimental* (0.39) documents. This may be due to the additional readership and attention that standards-track documents receive, and because they are more likely to be the basis for future work and protocol extensions.

Impact of Citations. Figure 5 plots the number of errata filed for an RFC *vs.* the number of citations of that RFC, counting citations from other IETF drafts and RFCs as recorded by the IETF Datatracker. We colour code data points based on the area of the IETF that produced the RFC. The rough trend is that more highly-cited RFCs, which are the basis for later work, tend to have more errata filed for them, with a Pearson coefficient of 0.33. This is intuitive: such RFCs will likely attract more scrutiny from those extending the protocol.

Note, for readability, we remove two outliers from Figure 5. RFC 4601 (related to Protocol Independent Multicast) has 114 errata recorded, yet obtains only 71 citations from other RFCs, while RFC 5741 (specifying the copyright notices and other

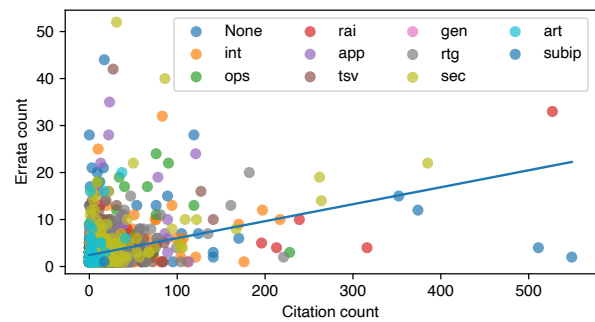


Figure 5: Errata filings by citation count, coloured by IETF area (acronyms expanded in Table A2).

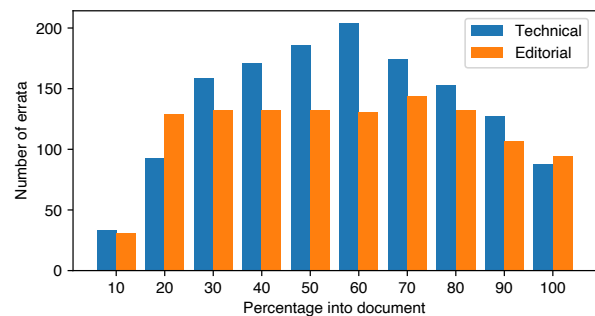


Figure 6: Errata counts by percentile location in document (0 is the beginning; 100 is the end).

boilerplate included in all subsequent RFCs) has 1,628 citations, with only 1 erratum. Such outliers are expected: RFC 4601 is a technical standard, defining a special-purpose, not widely deployed protocol, while RFC 5741 describes aspects of the RFC publication process, and so is much more widely cited. Both RFCs are included in the Pearson coefficient calculation.

Errata Location. Finally, we investigate the location of errata within RFCs. Figure 6 presents the number of errata occurring at each decile of the documents, for the 2,552 filings where accurate location information is available, and after the copyright notice and other boilerplate has been removed. We see that technical errata dominate over editorial in almost all places, except for the very beginning where the *Introduction* is located. Moreover, it shows that the most technical errata are near the middle of the document where the most complex content is. We explore where errata occur, with Figure 7 showing section titles for errata appearing in at least 10 documents. Sections such as the *Introduction* or *References* are dominated by editorial errata while more technical sections, like *IANA Considerations* (*i.e.*, parameter registrations), *Security Considerations* or *Definitions*, have a larger proportion of technical errata. In addition, we see that sections labelled *Appendix* attract a significant proportion of technical errata. While appendices vary in their content, they are widely used to provide pseudocode and test vectors, or to describe algorithms. This suggests that it may be useful to target review efforts on appendices and other dense technical content where errata are more likely.

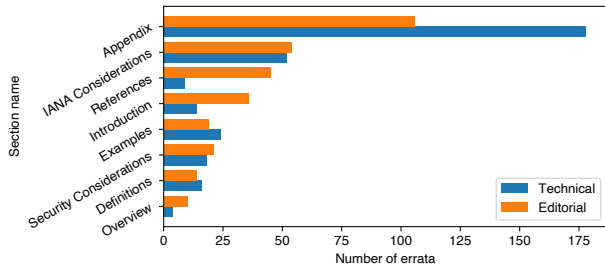


Figure 7: Errata counts by section title for the more frequent section titles.

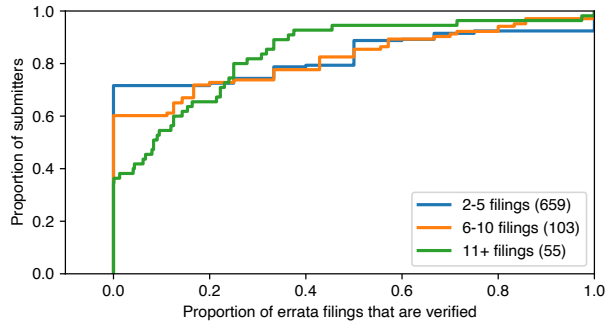


Figure 8: CDF of proportion of errata filings that are verified per submitter.

IV. CHARACTERISING ERRATA SUBMITTERS

Number of Submissions Per Person. The majority (66.3%) of submitters only submit a single errata, comprising 23.7% of all errata filings, with 91.5% of submitters filing fewer than 5 reports. However, we observe a long tail of highly dedicated submitters who report a large number of errata, with 25 people who contribute over 20 errata each.

Of course, not all errata reports are correct or appropriate. Figure 8 presents the proportion of submitters who have errata that are marked as *verified*. We break down submitters based on the number of errata reports they have filed. We find a significant fraction of submitters with a majority of their submissions not marked as verified. On average, 13% of filings are rejected. As shown in Figure 9, this improves with experience, with those submitting more errata (11+), seeing fewer rejections on average (10.8%) than those that submit 2-5 errata (15.1%). 14.2% of errata have yet to be processed at the time of publication. With all filings in the dataset having been made prior to the end of December 2022, this suggests significant scope for improving the quality of submissions and the verification process.

Figure 10 presents the proportion of errata that are technical, rather than editorial, per person reporting. Again, we separate people based on their number of errata submissions. We observe that people who file fewer reports have a larger share of technical errata; those who submit a large number of reports tend to contribute more editorial errata. This may be because editorial errata are easier to identify, often requiring less domain-specific expertise.

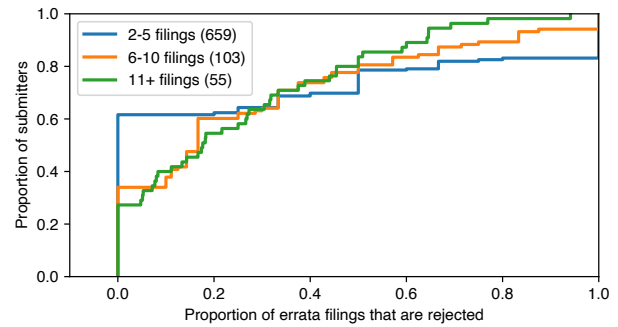


Figure 9: CDF of proportion of errata filings that are rejected per submitter.

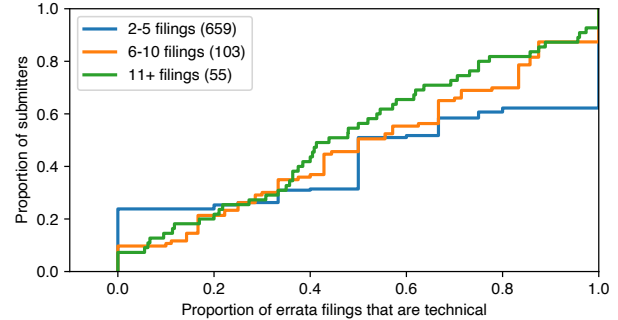


Figure 10: CDF of proportion of errata filings that are technical per submitter.

Top Errata Submitters. For context, we inspect the most active errata submitters, listed in Table A5. The number of submissions per person is highly skewed, with the top submitter (Hoenes) submitting significantly more ($\approx 6.5\times$) errata than the next. The top three submitters (Hoenes, Malykh, and Peasley) collectively submit an average of 63.8 errata per year, representing a significant contribution to the Internet standardisation process. The status of the filings made by these submitters varies substantially. For instance, Hagemeyer has filed 82 errata, of which 56.1% were rejected. Of the submitters who have filed 20 or more errata, 4 had $\geq 20\%$ of their reports rejected. In contrast, others such as Lilly, Newman, and Schaad have $>80\%$ of their errata verified.

We also see that different people exhibit a different emphasis on both IETF area and errata type. Some, such as Børghesen, focus overwhelmingly on technical errata (97.4%) within a single area. In contrast, others such as Malykh, focus on editorial (83%) across many areas. This suggests that these complementary submissions require differing skill sets, and that there is variation in the type and experience of people submitting errata.

Mailing Activity of Errata Submitters. Finally, we check if submitters are active in the areas of the IETF developing RFCs for which they submit the errata. For every errata submitter we analyse their email activity for three years prior to the date of publication of the corresponding RFC. To consider long term shifts, we make these observations in four different periods: 2003-2007, 2008-2012, 2013-2017, and 2018-2022.

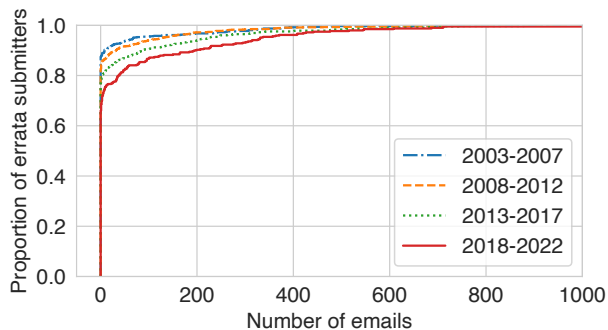


Figure 11: CDF of number of emails sent by errata submitter within the WG of the RFC.

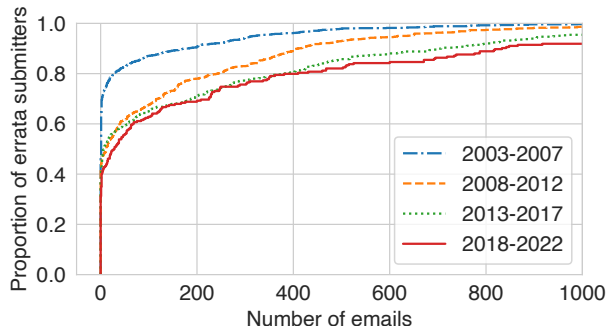


Figure 12: CDF of number of emails sent by errata submitter outside the WG of the RFC.

Figure 11 presents a CDF of the number of emails sent by errata submitters to the working group (WG) that developed the respective RFC. We observe that many submitters are not particularly active in the publishing WG. For example, between 2018–2022, 72.4% of submitters sent fewer than 5 emails within the publishing WG prior to the RFC’s publication. This trend is even more extreme in earlier years: *e.g.*, in 2003–2007 it was 88.9%. Although this suggests errata submitters have been more active in recent years, this is still low. This could indicate that external perspectives are valuable for identifying errata, and that the readership of a published RFC is different to that of a work-in-progress draft.

This is confirmed by Figure 12, which plots the CDF of the number of emails sent by errata submitters to lists *other than* the relevant WG. We observe that errata submitters are far more active within the wider IETF community than in the group that developed the RFC for which they’re making the report. For instance, between 2018–2022, 58.5% of submitters sent 5 or more emails to other IETF lists. While this is skewed by the larger number of lists one could target, it suggests that people from outside the immediate technical group related to RFC preparation tend to submit more errata. This is somewhat intuitive—if errata submitters were active within the working group prior to submission, they would likely have found and discussed the issues prior to publication—but does highlight the importance of ensuring wide cross-WG review occurs.

V. HOW TO REDUCE ERRATA FILINGS?

Having characterised the RFC errata and their submitters, we now consider the success of some of the efforts made to improve document quality in the IETF and reduce errata filings. Specifically we consider 1) the benefits of cross-area review of documents; 2) the use of review teams and directorates to conduct focused expert review of documents prior to publication; and 3) the use of formal and structured languages (*e.g.*, code, domain-specific languages, and mathematical models) in documents. Using errata filing volume as a proxy for document quality, we explore these strategies to identify where they help.

Our exploration of these techniques is preliminary, and intended to set the direction of future work, as well as to help focus efforts within the IETF community.

Cross-area Review. One of the claimed strengths of the IETF process is the breadth of review it provides to drafts before publication. This is facilitated by in-person plenary meetings of the entire community, combined with email-based community-wide discussion and a “last call” period just prior to publication.

To determine if this process is effective at reducing errata, we look at all RFCs published between 2001 and 2020 for which the IETF email archive records the discussion prior to publication. This gives a set of 4356 RFCs. We collect the emails whose subjects mention the draft versions of that RFC and identify the senders. We then look at the email messages sent across all IETF working group mailing lists, and label the senders of those emails based on the IETF area to which they sent the most email. This is naïve approach, and it does not attempt to capture the level of activity within each area. However, it provides a preliminary estimate of the main area of expertise for the reviewers who provided feedback. Thus, each of the 4356 RFCs is associated with a list of areas of expertise for the people who discuss it.

Our first hypothesis is that there is a correlation between the breadth of expertise represented in the discussion prior to publication of an RFC, and the number of errata filed against that RFC. Specifically, we expect that the number of unique IETF areas of expertise in the email discussion relating to the RFC would negatively correlate with the number of errata, *i.e.*, documents that receive broad review attract fewer errata. To test this, we calculate the Pearson correlation coefficient between the number of unique areas covered and the count of errata filings that the RFC attracted within two years of its publication. We find no correlation (0.006), failing to prove our intuition. This is also reflected in Figure 13, which presents a boxplot of the number of errata filed against RFCs for the number of unique areas in the pre-publication email discussion.

Our initial results suggest that cross-area review does not reduce the number of errata filed (despite other potential benefits to the IETF). Our findings could be extended to look at, for example, the depth and nature of the pre-publication e-mail discussion, the wider mix of expertise that participants bring, and how that expertise is identified.

Review Teams. As part of the pre-publication document review process, a number of IETF technical review teams that provide

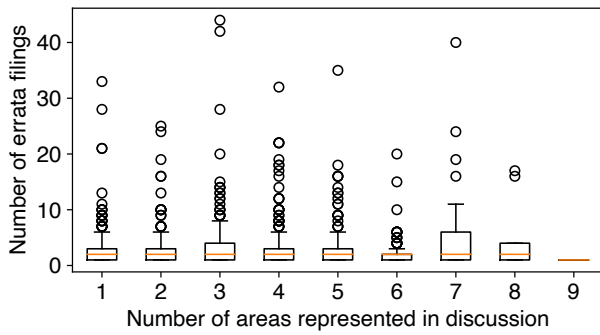


Figure 13: Errata filed against RFCs, by count of unique IETF areas represented in the pre-publication e-mail discussion; RFC 4601 is not shown, as this is an outlier.

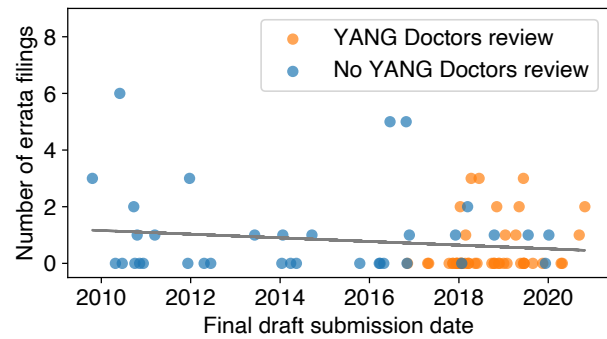


Figure 14: Errata filings against YANG RFCs by submission date of the Internet-Draft that preceded their publication.

in-depth review focused on particular subject areas. Our second hypothesis is that these reviews will be effective at identifying problems with drafts, leading to fewer errata.

One of the longest-established of these review teams is the “YANG Doctors”, who review documents for issues relating to the YANG data modelling language [4]. As an initial case study on the effectiveness of such a review, we consider the 82 RFCs that use YANG. Of these, 38 received a pre-publication review from the YANG Doctors, and 44 did not.

Figure 14 shows the number of errata filed for each RFC that contains YANG, against its final draft submission date, with a line of best fit. There is a reduction in the number of errata submitted against YANG-related RFCs that received an in-depth technical review by the YANG Doctors. Of the 38 RFCs that did not receive a YANG Doctor review, 18 (47.3%) attracted at least 1 errata filing in the two years following their publication. On average, these RFCs attracted 0.97 errata filings. However, of the 44 RFCs that *did* receive a YANG Doctor review, 11 (25%) attracted at least 1 errata filing in the two years following their publication. On average, these RFCs attracted 0.48 errata filings. We further consider the types of errata that are filed against YANG-related RFCs. In the two years after their publication, RFCs that did not receive a YANG Doctor review attract an average of 0.34 editorial and 0.63 technical filings. These rates fall to 0.25 and 0.23 for editorial and technical filings, respectively, highlighting the broad value of targeted reviews in reducing errata filings.

While there are other review teams and directorates to which our preliminary analysis could be extended, these initial results suggest that the IETF may wish to consider more carefully evaluating and expanding its review team activities.

Formal and Structured Languages. Our final hypothesis is that RFCs that include formal and structured languages would benefit from automated checking, reducing errata. For example, RFCs using the ABNF grammars can be automatically processed by an ABNF grammar checker, RFCs that contain a YANG data description can use a YANG syntax checker, and so on. Following on from the analysis of review teams, we study the impact of introducing an automated checker

for YANG models on the RFC errata rate. Support for the automatic checking of YANG modules upon the submission of draft documents was introduced in the IETF Datatracker, starting with v6.16.0, released on 5 March 2016.

Of the 82 RFCs that define YANG modules published between 2001 and 2020 inclusive, 24 had their final draft versions submitted before the YANG checker was released. Thus, these were not subject to the automated check. The remaining 58 RFCs had their final draft version submitted after the YANG checker was released and so were checked. We use this to contrast the efficacy of this formal language checking.

Of the 24 “*pre-checker*” RFCs, 10 (41.7%) attracted one or more errata filing in the two years following their publication; on average, “*pre-checker*” RFCs attract 0.83 errata filings in this time. However, of the 58 “*post-checker*” RFCs, only 19 (32.8%) attracted one or more errata filings in the two years following their publication; an average of 0.65 errata filings. This suggests that the formal checks *do* reduce errata.

It is interesting to recall that, across all standards-track RFCs in our dataset, 32% attract at least one errata filing (§III). This means that the errata rate for pre-checker YANG-module-defining RFCs is actually *higher* than the remaining set of RFCs. We conjecture that, without the assistance of computational checks, the additional complexity of writing formal language specifications can easily introduce errors. This is supported by the rates of different types of errata. On average, pre-checker RFCs have 0.63 technical errata filed against them within two years of publication, while post-checker RFCs have 0.33 technical errata filings. The trend for editorial filings is the opposite, with 0.21 filings on average for pre-checker RFCs, rising to 0.33 filings for post-checker RFCs. While more granular labelling of errata filings is needed, this suggests that the automated tool is effective in removing YANG-related errata, which are likely to be marked as technical. Other hypotheses about the reasons for the larger volume of errata in YANG documents, to be explored in future work, include that errors in formal languages are easier to spot, or that formal definitions attract greater reviewer attention.

VI. RELATED WORK

To the best of our knowledge, we are the first to perform a comprehensive exploration of errata within the IETF. That said, there have been several prior efforts that have focused on other aspects of the IETF. McQuistin *et al.* [11] inspected publication patterns within the IETF, with a particularly focus on understanding what leads to a successful RFC. Similarly, Nikkhah *et al.* [13] statistically explore RFC adoption. Jari Arkko, a former chair of the IETF, maintains a website and tooling that provides various statistics about the IETF [1], including about its documents, authors and their affiliations. BigBang [2] is a Python toolkit for analysing online collaborative communities through mailing list data. Niedermayer *et al.* [12] discuss the challenges of working with large mailing list datasets. Huitema [8] carried out an evaluation of a small set of RFCs, to understand the sources of publication delay. Finally, the IETF reports statistics about authors and publications [10].

VII. CONCLUSIONS

This paper explored the patterns of RFC errata, covering a 22 year period. We characterised the filings themselves (§III), before looking at the people who submitted them (§IV), and finally, evaluating errata-reduction strategies (§V).

We found wide variations in the types of errata that are filed, when, and by whom. 14.2% of the errata in our dataset remains in the *reported* stage, without having been further verified or rejected. This includes errata that were filed more than 5 years ago, and so will likely include errors that have made their way into implementations. We also observed that 7.3% of errata are filed within 30 days of the publication of the RFC they are reported against, suggesting deficiencies in the review phase. We found that targeted, specialist review teams can reduce filings. In contrast, the broader, general cross-area review does not. In addition, we showed that formal languages (like YANG) reduce the presence of errata when accompanied by tooling. We posit that these observations should be considered earlier in the standardization process.

Our findings can form the basis for extensive further work. We hope to extend our preliminary analysis in §V to include other technologies and languages (like ABNF), to show that our findings hold more widely. Our dataset could also be augmented to allow for topics that attract significant volumes of errata to be identified, providing a rationale for the formation of a review team or the use of automated tooling. While further data would be needed to support this, we are hopeful that our findings will be helpful in focusing efforts within the IETF.

ACKNOWLEDGEMENTS

This work is supported by the UK Engineering and Physical Sciences Research Council, under grants EP/S033564/1 and EP/S036075/1. Purver was also supported by the Slovenian Research Agency via research core funding for the programme Knowledge Technologies (P2-0103).

REFERENCES

- [1] Arkko, J.: IETF statistics (2022), <https://www.arkko.com/tools/stats.html>
- [2] Benthall, S.: Testing Generative Models of Online Collaboration with BigBang. In: Proceedings of the 14th Python in Science Conference. pp. 175–181. SciPy Project, Austin, Texas (2015)
- [3] Birkholz, H., Vigano, C., Bormann, C.: Concise Data Definition Language (CDDL): A Notational Convention to Express Concise Binary Object Representation (CBOR) and JSON Data Structures. RFC Editor (Jun 2019), rFC 8610
- [4] Björklund, M.: YANG - A Data Modeling Language for the Network Configuration Protocol (NETCONF). RFC 6020 (Oct 2010). <https://doi.org/10.17487/RFC6020>, <https://www.rfc-editor.org/info/rfc6020>
- [5] Fenner, B., Handley, M., Holbrook, H., Kouvelas, I.: Protocol independent multicast - sparse mode (PIM-SM): Protocol specification (revised). RFC Editor (Aug 2006), RFC 4601
- [6] Flanagan, H.: Fifty years of RFCs. RFC Editor (Dec 2019), RFC 8700
- [7] Housley, R., Daigle, L.L.: The RFC Series and RFC Editor. RFC Editor (Feb 2020), RFC 8729
- [8] Huitema, C.: Evaluation of a sample of RFCs produced in 2018. RFC Editor (Jan 2021), RFC 8963
- [9] IAB, Daigle, L.: The RFC Series and RFC Editor. RFC Editor (Jul 2007), RFC 4844
- [10] IETF: Datatracker stats (2022), <https://datatracker.ietf.org/stats>
- [11] McQuistin, S., Karan, M., Khare, P., Perkins, C., Tyson, G., Purver, M., Healey, P., Iqbal, W., Qadir, J., Castro, I.: Characterising the IETF through the lens of RFC deployment. In: Proceedings of the 21st ACM Internet Measurement Conference. pp. 137–149. ACM, Online (2021)
- [12] Niedermayer, H., Schweltnus, N., Raumer, D., Cordeiro, E., Carle, G.: Information Mining from Public Mailing Lists: A Case Study on IETF Mailing Lists. In: International Conference on Internet Science. pp. 301–309. Springer, Thessaloniki, Greece (2017)
- [13] Nikkhah, M., Mangal, A., Dovrolis, C., Guérin, R.: A statistical exploration of protocol adoption. *IEEE/ACM Transactions on Networking* 25(5), 2858–2871 (2017)
- [14] Postel, J.: Internet Control Message Protocol. RFC Editor (Sep 1981), RFC 792
- [15] Postel, J.: Internet Protocol. RFC Editor (Sep 1981), RFC 791
- [16] Postel, J.: Transmission Control Protocol. RFC Editor (Sep 1981), RFC 793
- [17] Saint-Andre, P.: RFC Editor Model (Version 3). RFC 9280 (Jun 2022). <https://doi.org/10.17487/RFC9280>, <https://www.rfc-editor.org/info/rfc9280>

APPENDIX

This appendix contains supporting data tables as follows. Firstly, Table A1 lists the fields that appear in the errata dataset.

Item	Description
Erratum ID	Unique identifier for the erratum
Doc ID	RFC number
Section	Location of the erratum in the RFC
Status	Reported, Verified, Rejected, or Hold for document update
Type	Technical or Editorial
Original text	Original text from published RFC
Correct text	Corrected text
Notes	Notes from verifier
Submission date	Date erratum reported
Submitter name	Person that reported the erratum
Verifier ID	Unique identifier of the verifier
Verifier name	Person that verified the erratum
Update date	Date the erratum was last modified

Table A1: Contents of Errata Database

Tables A2 through A5 support the results in Section III giving errata statistics by area (Table A2), RFC publication stream (Table A3), and RFC status at publication (Table A4). Finally, Table A5 lists the top submitters (*i.e.*, those that have filed 22 or more errata), with filing statistics.

Area	#	Verified	Held	Rejected	Reported	Technical	Editorial
None	1883	895	505	197	286	930	953
Internet (int)	650	281	223	98	48	342	308
Operations and Management (ops)	558	311	113	67	67	297	261
Real-time Applications and Infrastructure (rai)	457	143	213	48	53	255	202
Security (sec)	888	291	265	115	217	447	441
Routing (rtg)	831	305	378	140	8	326	505
Applications (app)	787	370	175	116	126	464	323
Transport (tsv)	459	188	142	75	54	258	201
General (gen)	41	22	4	5	10	8	33
Applications and Real-Time (art)	204	64	32	16	92	143	61
Sub-IP (subip)	1	1	0	0	0	1	0
All	6759	2871	2050	877	961	3471	3288

Table A2: Errata statistics by area.

Stream	#	Verified	Held	Rejected	Reported	Technical	Editorial
IETF (6619)	5797	2348	1815	798	836	3034	2763
IAB (124)	55	25	13	3	14	23	32
Independent (376)	330	235	32	28	35	172	158
Legacy (1929)	510	226	182	39	63	198	312
IRTF (97)	67	37	8	9	13	44	23
All	6759	2871	2050	877	961	3471	3288

Table A3: Errata statistics by stream; the “*Editorial*” stream has no documents, and is not shown.

Status	#	Verified	Held	Rejected	Reported	Technical	Editorial
Proposed Standard (4084)	4142	1680	1308	555	599	2213	1929
Informational (2894)	1500	719	399	136	246	754	746
Internet Standard (147)	319	118	111	66	24	135	184
Best Current Practice (316)	233	111	53	30	39	80	153
Historic (70)	20	13	4	2	1	9	11
Draft Standard (142)	274	94	93	66	21	150	124
Experimental (563)	221	115	65	20	21	121	100
Unknown (929)	50	21	17	2	10	9	41
All	6759	2871	2050	877	961	3471	3288

Table A4: Errata statistics by status at publication.

Submitter	Filings	Verified	Held for Update	Rejected	Technical	Editorial	Areas	RFCs authored
Alfred Hoenes	1124	403 (35.9%)	602 (53.6%)	118 (10.5%)	459 (40.8%)	665 (59.2%)	9	2
Nikolai Malykh	171	66 (38.6%)	73 (42.7%)	25 (14.6%)	29 (17.0%)	142 (83.0%)	7	0
Maren Peasley	109	0 (0.0%)	93 (85.3%)	16 (14.7%)	45 (41.3%)	64 (58.7%)	1	0
Constantin Hagemeyer	82	7 (8.5%)	29 (35.4%)	46 (56.1%)	36 (43.9%)	46 (56.1%)	3	0
Paul Aitken	78	55 (70.5%)	21 (26.9%)	1 (1.3%)	17 (21.8%)	61 (78.2%)	4	14
Russ Housley	67	53 (79.1%)	5 (7.5%)	1 (1.5%)	49 (73.1%)	18 (26.9%)	6	104
Ivan Panchenko	65	58 (89.2%)	3 (4.6%)	0 (0.0%)	0 (0.0%)	65 (100.0%)	2	0
Julien Élie	55	21 (38.2%)	19 (34.5%)	6 (10.9%)	35 (63.6%)	20 (36.4%)	3	3
Julian Reschke	48	25 (52.1%)	14 (29.2%)	3 (6.2%)	23 (47.9%)	25 (52.1%)	4	30
Mykyta Yevstifeyev	44	7 (15.9%)	26 (59.1%)	11 (25.0%)	5 (11.4%)	39 (88.6%)	4	3
Stéphane Bortzmeyer	42	16 (38.1%)	13 (31.0%)	4 (9.5%)	17 (40.5%)	25 (59.5%)	8	4
Philip Börgesen	38	0 (0.0%)	1 (2.6%)	0 (0.0%)	37 (97.4%)	1 (2.6%)	1	0
Bruce Lilly	36	32 (88.9%)	3 (8.3%)	1 (2.8%)	27 (75.0%)	9 (25.0%)	3	2
Martin Thomson	32	14 (43.8%)	8 (25.0%)	2 (6.2%)	12 (37.5%)	20 (62.5%)	8	36
Ben Smyth	28	0 (0.0%)	0 (0.0%)	0 (0.0%)	10 (35.7%)	18 (64.3%)	1	0
Michael Sweet	26	10 (38.5%)	10 (38.5%)	6 (23.1%)	18 (69.2%)	8 (30.8%)	2	3
Chris Newman	25	20 (80.0%)	3 (12.0%)	0 (0.0%)	24 (96.0%)	1 (4.0%)	4	27
Frank Ellermann	25	15 (60.0%)	8 (32.0%)	2 (8.0%)	10 (40.0%)	15 (60.0%)	7	1
John Klensin	25	11 (44.0%)	5 (20.0%)	3 (12.0%)	13 (52.0%)	12 (48.0%)	5	59
Sean Turner	23	18 (78.3%)	5 (21.7%)	0 (0.0%)	11 (47.8%)	12 (52.2%)	3	53
Jim Schaad	23	20 (87.0%)	1 (4.3%)	1 (4.3%)	22 (95.7%)	1 (4.3%)	3	29
Peter Occil	22	8 (36.4%)	5 (22.7%)	2 (9.1%)	17 (77.3%)	5 (22.7%)	4	0

Table A5: Top errata submitters.