RPC Self-Evaluation for 2022

Although the RFC Production Center's (RPC's) responsibilities changed during 2022 (as described in RFC 9280), the core function remains to edit and publish high quality RFCs in a timely manner. This review highlights how the RPC performed on their core mission, discusses other activities such as adapting to RFC Editor Model v3 and experimenting with GitHub, and describes the challenges faced by the RPC. It also provides a preview into the work ahead for 2023.

Document Throughput and Editorial Quality

Throughput

At the end of 2019, queue processing times grew extraordinarily long as the RPC and community transitioned to RFCXML. As the RPC adapted to the new vocabulary and process changes, the processing times decreased significantly over the following two years. As shown in Table 1, the number of RFCs being published in fewer than 6 weeks is steadily increasing, while the number of documents with an RFC Editor time of greater than 12 weeks is down to 14% from 74% in 2020.

The table below shows the number of documents and page counts moving into and out of the queue over the last few years. It also provides some details about their processing times.

Year	2018	2019	2020	2021	2022
Docs to EDIT (PGTE)	235 (6663)	231 (7420)	180 (5389)	188 (7232)	203 (5936)
RFC Pubs (pages)	208 (5631)	180 (5152)	209 (5212)	240 (6386)	194 (5811)
RET <= 6 wks	46%	40%	12%	29%	39%
6 < RET <= 12 wks	52%	47%	14%	45%	48%
RET > 12 wks	2%	13%	74%	26%	14%
Median RET	6.4	9	16.6	9	7.3
Clusters	34%	41%	33%	49%	44%

Table 1. Document Production

> Note: the publication number for 2021 is unusually high because a 47-document cluster (C238) was published. The number of documents moving into edit is not elevated, as the documents were approved over time, with a large number being released into the edit queue at the end of 2019 (in the midst of transition to v3 XML).

While processing times continue to improve, the RPC is not meeting the currently defined SLA. The SLA has not yet been updated to reflect the additional workload required by the transition to

v3 XML or the additional responsibilities created by RFC Editor Model version 3. In discussion with the LLC, there is agreement that the SLA needs to be redefined. In addition, the RPC was asked to "focus resources on the transformational projects rather than trying to bring performance in line with the SLA" (e.g., tools modernization, document complexity, and RFC Editor Model v3 implementation).

	2021				2022			
	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4
Submissions								
Pages	2685	1310	1513	930	1983	1590	1396	1707
Docs	68	49	41	30	60	49	48	46
PGTE	2655	1403	2011	1163	1590	1212	1435	1699
Publications								
Pages	2117	1858	1507	904	1416	1715	1603	1077
Docs	89	61	54	36	44	58	56	37
Docs met SLA	49	48	46	33	15	20	18	21
SLA tier	Tier 2	Tier 2*	Tier 2	Tier 2*	Tier 1	Tier 1	Tier 1	Tier 1
SLA	✔	✔	✔	✔	✘	✘	✘	✘

Table 2. Production Times and the Service Level Agreement

Document Sources

There has been a steady increase in the number of v3 XML source files being submitted by authors. In 2022, there was a 27% increase in the number of v3 XML files submitted over the previous year. Over the last three years, about 6% of the submitted Internet-Drafts did not have an XML file, so id2xml was used to create them. As seen in Figure 1, id2xml was used at a fairly consistent rate over the last three years.
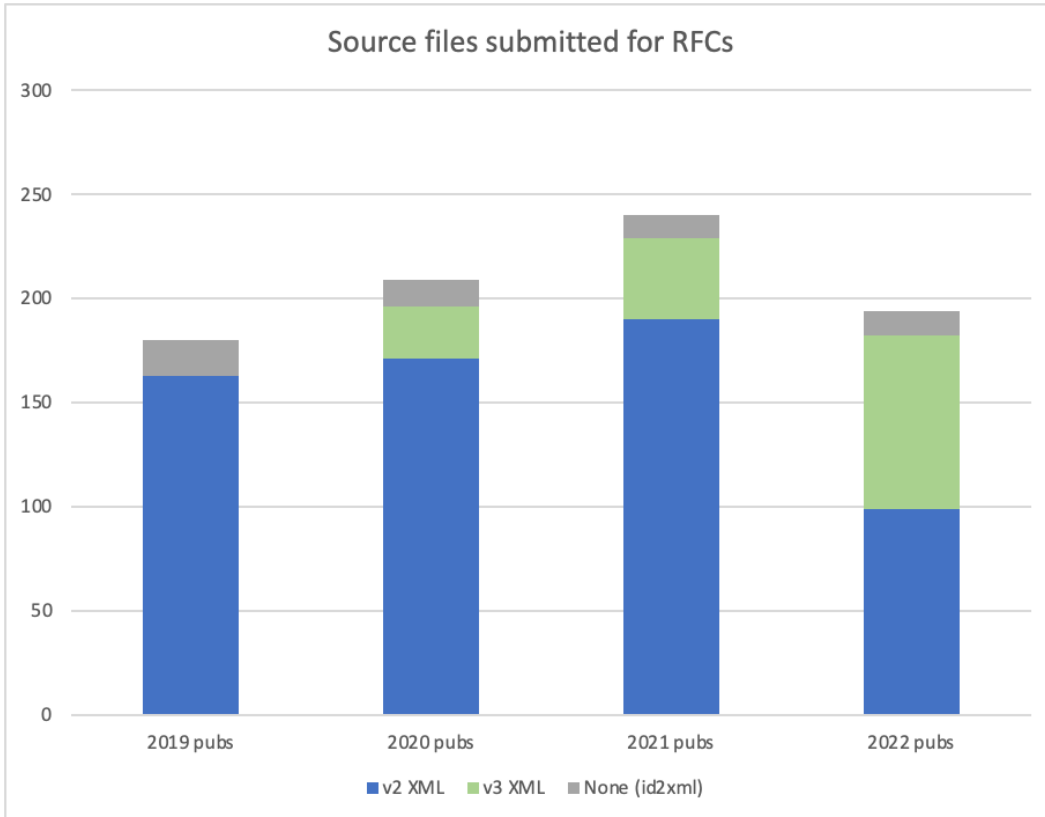
Figure 1. RFC Source Files Submitted by the Authors

With authors becoming more familiar with RFCXML and with more central documentation about RFCXML and its features (i.e., authors.ietf.org), we are seeing a slight increase in formatting complexity.

The following figure shows that 83% of the RFCs originated within the IETF (includes IETF working group and individual submissions). 12% were Independent Submissions, which is unusually high; over the last five years, 7% on average were from the Independent Submissions Stream. The Independent Submissions Editor (ISE) indicated this year was higher than normal because it was a transitional year for the ISE.

2022 RFCs by Stream and Status

Figure 2. RFCs by Stream of Origin and Status

Author Feedback

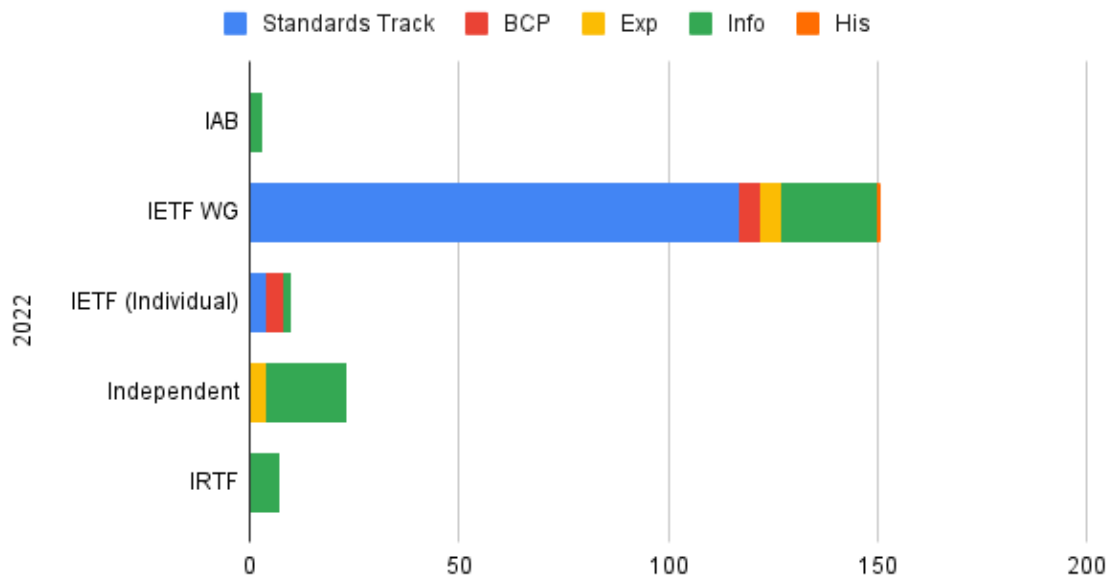There is a 94% overall satisfaction rate per author feedback via the post-publication surveys, with 48% indicating the editors made their documents "significantly better." This is a 12% increase in authors reporting "significantly better" from the previous year. 5% indicated the editors did not positively or negatively affect their documents, while 1% indicated that editors had a negative impact on their document. The author suggested the RPC should "[a]void nitpicking and changing author's writing style."   28% of the respondents expressed appreciation for the work of the editors, for example:

> This was my first RFC Editor experience, and the quality of the changes/suggestions and the related discussion was outstanding. Easily the best text-review cycle I've been through.

> The RFC Editor was diligent, thoughtful, and helpful (for example, in correcting the format of the tables and adding a missing reference), which was much appreciated.

> Thank you so much!!!
> It always is lovely to see what editorial suggestions the editor comes up with. A second set of eyes on a document greatly helps.

> I'm impressed with the care in the editing and the intelligence of the questions -- the editors clearly made the document better!

Note that there was a 21% response rate to the post-publication survey, with a total of 130 respondents.

2021: Did the editing service provided by the RFC Editor improve the quality of your document?

- significantly better
- slightly better
- no difference
- slightly worse
- significantly worse

1%
3%
36%
60%
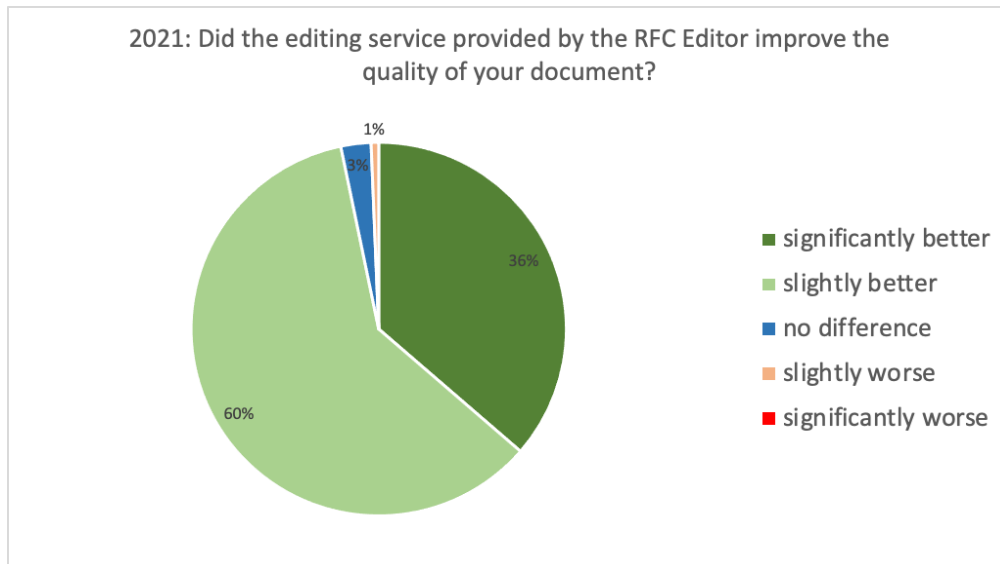
Figure 3. 2021 Post-publication Survey Results

2022: Did the editing service provided by the RFC Editor improve the quality of your document?

- significantly better
- slightly better
- no difference
- slightly worse
- significantly worse

1% 0%
5%
48%
46%
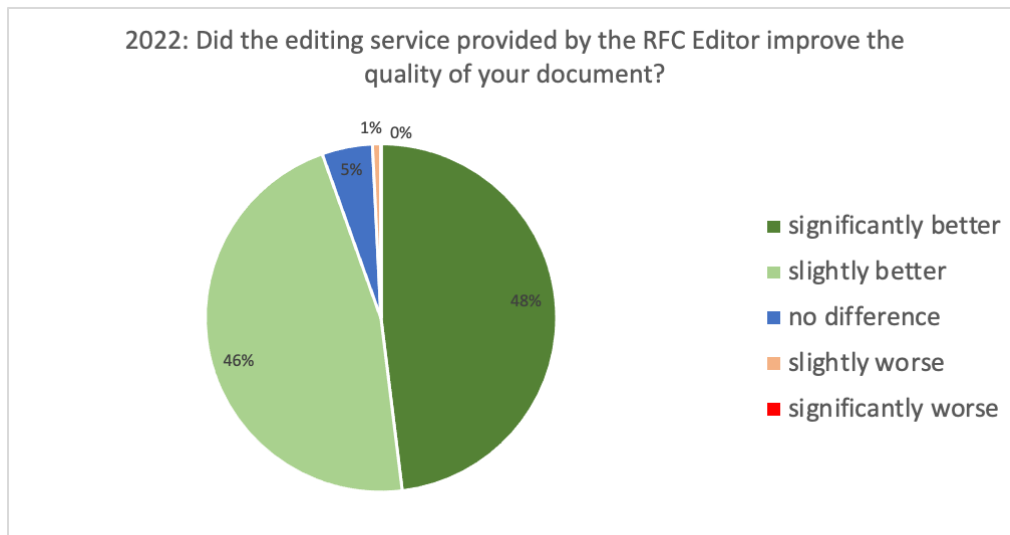
Figure 4. 2022 Post-publication Survey Results

In addition to author and document-shepherd feedback, we received the following from the ISE and IRTF Chair regarding the RPC's performance of its primary responsibilities, for example
- formatting, editing, and publishing RFCs in a timely manner
- liaising, training, and communicating with the community
- facilitating and experimenting with the publication process
- adopting new policies/guidance and v3 of the RFC Editor model

Eliot Lear, as the ISE:
Last year, the ISE presented to the RPC a substantial number of documents, compared to an average year. Independent documents tend to need additional care because they get less vetting than others. These documents were processed efficiently and with great attention to detail by the editors, and I think the results have been quite good.

We had healthy discussions in the RSAB.

Next year we will see some new challenges, including boiler plate changes, and proposals to handle formal language definitions. I am confident that between the RPC and the tooling team, we will find satisfactory results here as well.

In addition, Eliot sent the following in Dec 2022:
It's been a year of learning for me, and part of the joy of this position as been to see the incredibly professional way in which you all make the ISE and the IETF look good (in as much as that is possible)

Colin Perkins, as the IRTF Chair:
From the IRTF side, I'm very happy with the RFC Production Centre performance. Everything seems to be running smoothly, work is proceeding in a timely manner with no unexpected problems, and you're being responsive to queries from both the IRTF chair and document authors. Requests for experimentation and special treatment are being handled appropriately. The transition to the V3 process does not seem to have affected the RFC publication process, and I have no complaints about how the RPC handled the transition.

Process Improvements

auth48archive Mailing List

An [archive of AUTH48 discussions](#) between editors, authors, and stream managers was launched 31 May 2022 with the [creation of the auth48archive mailing list](#). Feedback from stream managers has been positive.

GitHub Experiments

The RPC hosted a workshop to discuss the use of GitHub during AUTH48 at IETF 114. Many participants provided feedback, which has informed our more recent GitHub experiments. [Notes](#) for the meeting can be found on our wiki.

GitHub experiments:

Document	Date of pub	GitHub Repo	Source files
RFC 9069	Feb 2022	[Repo](#)	XML
RFC 9245	Jun 2022	[Repo](#)	Markdown
RFC 9345	In progress	[Repo](#)	Markdown

In addition, draft-ietf-sipcore-multiple-reasons (4 pages) and draft-irtf-cfrg-hash-to-curve-16 (175 pages) will be edited in Markdown and AUTH48 will take place in GitHub.

After the workshop at IETF 114, we intended to work with XML and GitHub for our next experiments. However, authors requested that we experiment with Markdown. In particular, more than a handful of responses to the post-publication survey desired a process using a GitHub(like) environment and some requested that the RPC accept and edit Markdown files.

We have taken on the task of experimenting with editing in Markdown and using GitHub for AUTH48. However, it has been challenging, as we are learning Markdown and experimenting with new-to-us features in GitHub as we go. In addition, there is a surprising learning curve associated with user expectations for GitHub that sometimes slows the AUTH48 process.

We are using these experiments to create/update procedures for using GitHub and Markdown, and continue training our staff on these experimental processes.

Document Complexity

The RPC is analyzing data about the degree of complexity associated with Internet–Drafts in an attempt to determine better predictive information about items such as level of effort to edit a document, when a document will move to AUTH48, and overall staffing resources. We are now capturing more data about each document as Internet-Drafts progress through the editorial process. As part of this analysis, in addition to document length and cluster coordination, we are looking at things such as:
- the author-submitted format and formatting workload to get into publication-ready shape
- the required updates to IANA-related text
- whether code needs to be formatted and checked; is there related boilerplate that needs to be checked
- the number of clarifying questions sent to the authors

This is an ongoing effort as we collect more data in 2023. When the database is updated, we will be collecting new kinds of data, such as how long a document waits for an editor to become available to work on it.

Improvements to the Current Toolchain

Programming efforts during 2022 included improvements to external data and internal workflows, as well as necessary updates for preparing to publish RFCs on the new Editorial stream. Highlights include:
- For improving the metadata provided via JSON files, the RPC made fixes to errata.json as well as put a better system in place to refresh the RFC-specific JSON files (which contain the metadata of each RFC).
- For improving data used to notify relevant parties about their documents, the RPC started using the Datatracker API to maintain WG-specific data.
- For mitigating junk reports on the errata submission form, the RPC deployed a Captcha. There has been a slight reduction in junk reports, where on average:
  - Preceding the change, 11 junk reports per month (n=6).
  - After the change, 8 junk reports per month (n=8), which is 19% of submitted reports.
- For the addition of the new Editorial Stream, we made updates to many aspects of the current system, so that when RFCs from the new stream are published, they will have the features as any other RFC would. This included updates to:
  - display (various locations on rfc-editor.org, including rfc-index.xml)
  - email generation
  - errata system
  - statistics & metrics site

Community Support & Engagement

Website Updates

Style Guide

The RPC team updated the online portion of the Style Guide in August 2022 and let the community know via an [announcement](#) to the rfc-interest mailing list.

General Content

The RPC team held a two-day, face-to-face meeting in which the team updated many aspects of the existing content of the rfc-editor.org website.

A page that captures RPC formatting practices and also provides guidance on how XML elements and attributes are used within a document was added to the website: https://www.rfc-editor.org/pubprocess/how-we-update/

RSAB

In 2022, we saw the inauguration of the RFC Series Advisory Board (RSAB), as defined in RFC 9280. The group met in person for the first time during IETF 114 and met again during IETF 115. The RFC Editor's about page was updated to reflect the model defined in RFC 9280, and the Datatracker points to the RSAB group. Below we highlight the topics discussed with RSAB and the resulting actions.

Headers and Boilerplate

Per RFC 9280, the RPC now maintains information about the RFC headers and boilerplate. The RPC drafted a webpage to replace the existing IAB page.

RPAT

The RFC Production Center Advisory Team (RPAT) was created in mid-September to advise on operational practices and issues affecting production. The team has advised on many topics; see the list archive.

The use of <u> (related to RFC-to-be 9290)

draft-ietf-core-problem-details (now RFC 9290) included a Hewbrew string, which posed a new and interesting challenge, both with XML tagging and reviewing the document via different applications to understand how the RTL-read language is displayed. The output was not clear when using the <u> element. As such, in addition to discussing alternatives with the authors, the RPC asked RSAB to consider whether non-ASCII characters can be used in running text without the using the <u> element.

After discussion with the RPAT and RSAB, the RSAB ultimately announced:

> After discussion, the RSAB concluded that section 3.1 of RFC 7997 describes the current policy that covers the case in question. That section permits non-ASCII within <t> without the use of <u> "[w]here the use of non-ASCII characters is purely part of an example and not otherwise required for correct protocol operation" [RFC7997]. The current practice by the RPC is manifested in tooling limitations that are more strict than the policy in RFC7997 and only allow the use of <u> for non-ASCII chars within <t>.

The Online Portion of the Style Guide was updated to provide guidance about the use of non-ASCII characters.

RSWG

The RFC Series Working Group, specified by RFC 9280, launched in July 2022, and the RPC has been participating in discussions (see the [mail archive](#)).

RSCE

As of 1 September, the RPC is working closely with Alexis Rossi, the RFC Series Consulting Editor (RSCE), on items such as Archive agreements, archiving out links, and the overhaul of the [RFC Editor website](#).

Tools Modernization

In January 2022, the IETF Administration LLC [solicited bids](#) for the 'RFC Production Center Tools Modernisation BA/PM services' RFP, which requested a business analyst/project manager to enable the RPC to make informed decisions on modernizing their toolchain, document their requirements, assist with working with vendors to deliver a new toolchain, and to manage the implementation.
In preparation for replacing the rfc-editor.org database and upgrading editing tools, the existing tools chain and database was documented, and potential paths forward were researched and documented.

The contract [was awarded](#) to Association Management Solutions LLC (AMS), with Jean Mahoney taking on the position of business analyst and project manager in July 2022. (Previous to joining the RPC in 2019, Jean was a feature engineer and project manager at a large software company, where she managed customer and engineering requirements and software team deliverables.)

Since taking on the BA/PM role, Jean has documented the current database, tool chain, and the processes for editing and publishing documents, while gathering requirements for improving the systems and researching potential paths forward. This year, she will oversee the upgrade of the editing software and the migration of rfc-editor.org database to using datatracker technology.

Challenges

During 2022, the RPC continued to navigate the new normal of waning pandemic life, with our hybrid attendance at the first of the hybrid IETF meetings (113 in Vienna) as well as gathering in person as a team in October 2022. We encountered various challenges described in more detail below.

Because new-staff training is time intensive, turnover has been a challenge during 2022. A new editor was brought on in May to replace an editor that parted ways with the RPC earlier in the year. At the same time, Jean Mahoney shifted to her new role as the Tools Modernization

Project Manager, so another editor was brought on to help recover the loss of her editorial contributions. Onboarding two editors simultaneously has presented some benefits and challenges.   A primary challenge has been losing production of two senior editors at the same time. Training slows the production queue as newer editors receive training and senior editors coach the incoming editors.   In addition, another editor parted ways with the RFC later in 2022, so a search is underway for a replacement.

The RPC reorganized their team to account for the new responsibilities defined by RFC 9280 and the team is still adjusting to version 3 of the RFC Editor Model. This requires more coordination, planning, meetings, discussion, and consideration about with whom things should be discussed and in which venue.

Since the RFC format change in 2019, authors have not made much use of specific features – including the use of non-ASCII characters (outside the Latin blocks) and SVG ([examples of feature usage](#)).   The few submitted documents that have made use of these features have required detailed review, troubleshooting, and discussion with various parties (e.g., RPAT, RSAB).

We uncovered some unexpected issues with the file outputs that require rerendering of the PDF files. Guidelines for rerendering in this case are documented in the [RPC wiki](#). The PDF of RFC 9019 was rerendered in December 2022 so that Figures 2 and 3 are not truncated. The PDF of RFC 8743 also is a candidate for rerendering due to the same issue of a truncated figure (2 final lines missing from Figure 22). Its rerendering is on hold because the rerendered PDF would be a deterioration (with poor page breaking and some blank pages) compared to the current PDF.

Another issue was the leakage of fonts outside the accepted font families into the PDF files. The RPC was previously unaware that the fonts within the PDF needed to be checked prior to publication. Since discovering this in October 2022, the RPC has been checking the fonts used prior to publication.

The citation library site, [bib.ietf.org](#), was launched this summer, and the RPC has been working with the Tools team to improve the data that is served. The RPC prepared to transition to relying on [bib.ietf.org](#) in the first half of 2022, but ongoing issues with the data delayed the transition. We have since switched to using bib.ietf.org to reference RFCs and Internet-Drafts. Oftentimes, author info or role information is incorrect or missing. As such, the full reference entries (i.e., not using xi:include) may be included in the RFCs. When this occurs, the entries are checked again during publication and manually updated to ensure the information is current (for example, version info, publication date, etc). An added challenge has been reopening tickets when regressive behavior appears as other bugs are fixed. We are reporting issues and tracking their progress via [GitHub Issues](#).

Looking Ahead

The coming year will be revolutionary as the RPC will be upending its toolset, which includes the database, all of the scripts that support the editorial and production process, and scripts that provide data to users and authors (e.g., queue, RFC search). The future environment provides an opportunity to evolve the internal editing process. The team will be examining the recommended editing tools and reorganizing to optimize efficiency and quality. These improvements will take time to implement, and we foresee short-term slowdowns to document throughput while these transformations take place.

In response to requests from the community, the team will also continue experimenting with use of GitHub and Markdown; documenting procedures and training for both will be ongoing.   We intend to follow up the GitHub workshop held at IETF 114, ideally at IETF 117. We would like to further refine our experiments and better understand how to optimize the process for all authors, not just those that are already expert GitHub users. The processes will continue to advance as we gain more experience with the new toolset and can better understand how and where GitHub effectively fits into RPC operations.

The RPC will continue to track RSWG discussion and provide input as needed. This will help us understand community issues and the direction of possible upcoming policies. Some priorities will be set by the RSWG, for example, if action is needed once the v3 XML is finalized or once a formal definition of "archival" is derived, and, down the road, if versioning is adopted. The RPC will remain flexible enough to shift priorities as needed.

As shown in Table 1, processing times have continued to decrease each year since the transition to v3 XML, and the post-publication surveys indicate the editing service is improving the RFCs, with nearly half of the responses indicating that the RFCs have been "significantly" improved. In addition to editing and publishing RFCs, much of 2022 was spent implementing the new RFC Editor Model and paving the way to evolve the suite of editing and data management tools. We expect 2023 will be a year of grand change, as we anticipate the pieces of the model settling in and advancing their work items and we will be adopting new editing tools and adapting our processes throughout the year. We are looking forward to an exciting and demanding 2023!