RPC Self-Evaluation for 2021

The core function of the RFC Production Center (RPC) is to edit and publish high quality RFCs in a timely manner. In 2021, the RPC published 240 RFCs and met the production times defined in the Service Level Agreement (SLA).   At the same time, members of the team also engaged in other activities aimed at improving our process, responding to the changing needs of the community, and evolving the Series.

This evaluation will discuss document throughput, highlight engagement that falls outside of day-to-day operations, review challenges faced by the RPC, and describe the work plan for 2022.

Throughput and Quality

In 2021, the RPC published 240 RFCs (6387 pages) and successfully met the Service Level Agreement for the first time since transitioning to v3 XML in September 2019 (see table 1). This is a 14% increase in the number of documents published in 2020, and a 23% increase in page counts. This means the documents published in 2021 were, on average, slightly longer than those published in 2020 (see table 2).

SLA statistics from Q4 2019 to Q4 2021

	2019	2020				2021			
	Q4	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4
Submissions									
Pages	1413	2423	929	1084	935	2601	1310	1513	930
Docs	49	64	40	42	34	67	49	41	30
PGTE	1576	2537	904	814	1134	2655	1403	2011	1163
Publications									
Pages	878	1383	890	1825	1114	2117	1859	1507	904
Docs	40	71	34	53	52	89	61	54	36
Docs met SLA	19	7	9	0	17	49	48	40	33
SLA tier	Tier 2*	Tier 2	Tier 2*	Tier 1	Tier 1	Tier 2	Tier 2*	Tier 2	Tier 2*
SLA	✖	✖	✖	✖	✖	✔	✔	✔	✔

Table 1. SLA Table since Transition to v3 XML

Published	2020	2021
RFCs	210	240
Pages	5212	6387
Pages/RFC	25	27

Table 2. RFC Publications in 2020 and 2021

A significant portion of the documents were published in Q1. Of the 89 documents published, 47 of those were part of Cluster 238 (WebRTC); an additional 3 documents were part of Cluster 391, a related cluster that was requested for simultaneous publication. There was a significant effort to coordinate a cluster of that size and publish all of the documents at the same time.

The following table summarizes submissions and publications by stream for 2021.

Streams	Submissions	Publications
IAB	3	4
IETF (WG + Individual)	167	224
Independent	13	8
IRTF	4	4
Total	187	240

Table 3. 2021 Submissions and Publications by Stream

93% of the RFCs published in 2021 are a product of the IETF; 3% originate with the Independent Stream; and the IAB and IRTF account for 2% each.

In addition to the 187 I-Ds entering the RFC Editor queue, more were released from MISSREF into the EDIT queue (see "About the Queue" for state definitions). While there were 41 I-Ds in MISSREF to start 2021, there were only 14 by the end of the year.

Since the transition to v3 XML, the process of creating publication-ready v3 XML files has been subsumed by the RPC.   In most cases, this means converting v2 XML to v3 XML; in some cases, this means using id2xml to create an XML file from plain text, the output of which often requires a heavy formatting lift. Table 4 gives a summary of the source files submitted by authors for the RFCs published in 2021.

Source Files	2020	2021
XML v2	172	190
XML v3	25	39
No XML	13	11

Table 4. Summary of Author-Submitted Source Files

There was a 10% increase in the number of v3 XML files provided by authors, and a 15% decrease in files with no author-submitted XML source file.

Our goal is not only to publish RFCs, but to publish high quality RFCs and provide excellent customer service throughout the process. The following figure shows that 97% of respondents to a post-publication survey indicated that the RPC improved their drafts. 37% said the drafts were significantly improved, and 60% said the drafts were slightly improved. Two respondents indicated that the RPC made their documents slightly worse, and a handful selected "Meh" on the rating scale. (Note: This data was provided to the RPC by John Levine, the Temporary RFC Series Project Manager, who conducted the survey.)
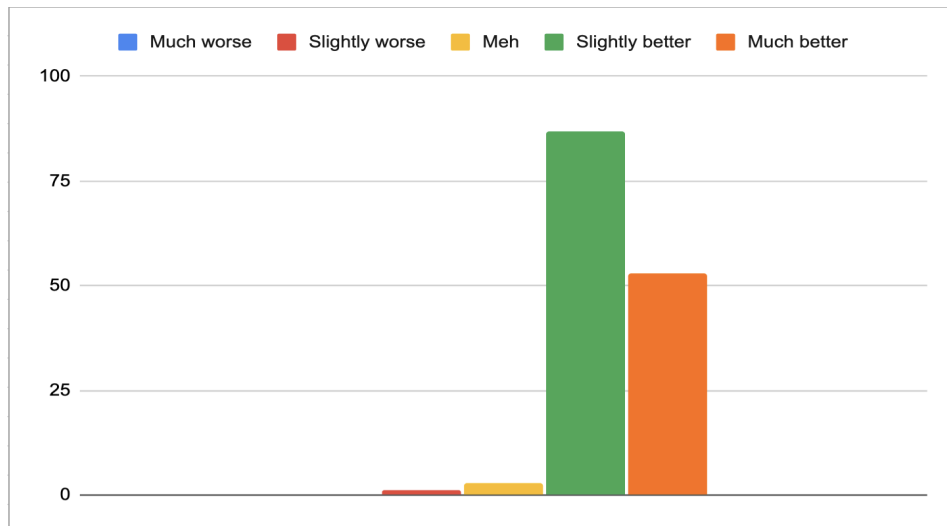
Many respondents expressed an appreciation for the work of the RPC and suggested that improvements to the tracking of edits and issues during AUTH48 are needed, and three different tools were suggested (MS Word, Google Docs, and GitHub).   For example, these are a few of the comments received via the survey:

> I find the RFC Editor's technical knowledge, editorial insight, and professionalism to be outstanding and have been very happy to work with them.

> I always appreciate the work that the RPC puts in, and this was no exception. Everything went smoothly as we got through this 160-page document.

Would like: communicate these needed changes through emails are error-prone and inefficient. I wish we can address the edits over google documents or similar online editing tools.

Would like: Provide diffs in word file with questions as comments. Current diff tool is not very good.

Having the XML in github would save a lot of busywork on the side of the authors when the RPC makes changes.

One author, who chose to remain anonymous, indicated that the RPC introduced errors in their document. It is unclear to us whether there were issues with the XML tagging, the edits, or both. The AUTH48 process is designed to correct any such issues if they arise during the editing process, and we welcome any feedback.

In addition to the survey, comments received directly from authors during AUTH48 are an indication of the quality of edits and service provided by the RPC. Below are some comments received from authors during 2021.

▨▨▨▨▨▨▨▨▨▨

Thank you for your kind consideration of my proposals. I've reviewed updates you've suggested and I agree with you, they make the text clearer.

▨▨▨▨▨▨▨▨

First, thank you for the careful work by the RFC Editor. This has clearly required expertise and effort and made the specification much better. Thank you.

▨▨▨▨▨▨▨▨▨▨

Thank YOU for such a thorough proof reading!

Process Improvements

Process improvement activities included sharing feedback with Ops Area ADs regarding YANG checks that could be implemented by the YANG review team and exploring ways to improve the AUTH48 experience for authors and ADs.

YANG Reviews: We suggested improvements to the YANG I-D review process to the three ADs working on updates to the document shepherd templates, including an expanded section on the use of YANG models. This includes items the RPC has been asked to do over the years, but it's not clear that YANG authors are aware of these recommended practices. One of the ADs recommended we work with the appropriate parties to have the [Yang Review Tools wiki page](#) updated.

Errata System: We participated in discussion and reviewed a proposal from Area Directors about evolving the way errata are handled within the IESG. The major change to the process was for the RPC to review and approve Editorial errata. The outcome is better quality files for readers of RFCs – specifically, the Verified errata show up in the RFCs with in-line errata. We established a process to review and approve typos and grammatical errors and (when needed) pass along other items to the Stream Managers for further review. This new process began in February, and so far the RPC has verified 19 Editorial errata. We raised this for discussion with the other stream managers, and they chose to adopt this new handling as well.

AUTH48: We had fruitful discussions with the IESG regarding transparency during AUTH48, with short-term and long-term requirements identified. Will further explore a short-term solution in which a public mailing list to capture AUTH48 discussions will be set up.

To facilitate the AUTH48 process for clusters and to improve transparency, the RPC introduced cluster-specific mailing lists (links to the mailing list archives: C238, C310, C430, C431). This allowed authors to view and track the various documents in the cluster as they completed AUTH48. We continue to experiment with mailing lists for clusters in which the documents seem tightly integrated, and have different author sets and/or span multiple areas. This may be overtaken by the short-term solution mentioned above.

Using GitHub during AUTH48: Exploring the use of GitHub for handling AUTH48 issues and edits:

- We vastly improved our processes for experimenting with GitHub for AUTH48:
    - We updated procedures to identify good candidates for AUTH48 experiments, reviewing author use of GitHub features for their documents and asking the authors and ADs if they would like to participate when their documents entered the queue.
    - We sought expertise from Barbara Stark, an author of RFC 8874 "Working Group GitHub Usage Guidance", and Paul Hoffman, an author of RFC 8875 "Working Group GitHub Administration", and invited them to observe and provide feedback on the RFC9131 AUTH48 experiment, which we then incorporated into our procedures.
    - We created a template for AUTH48 repos. This template contains a README providing instructions to authors on using GitHub features for AUTH48. This template can be continually improved and used to set up a new AUTH48 repo with each new experiment.
    - We documented on our internal wiki the procedures for identifying and contacting candidates and setting up the repo for AUTH48, and provided instructions and links on how to use GitHub features in order to allow more editors to participate in these experiments.
- GitHub experiments:
    - RFC 8829 (an ad hoc use of the authors' GitHub repo requested by one of the authors during AUTH48):
        - Summary of AUTH48 issue tracking

- ■ [Feedback from participants](#)
    - ○ RFC 9131 (a planned experiment in which the RPC set up the repo and invited the author and AD to participate):
        - ■ https://github.com/rfc-editor/rfc9131-AUTH48
        - ■ [Feedback from participants](#)
    - ○ RFC 9069 (the second planned experiment, currently on-going, which uses a repo built from the AUTH48 template):
        - ■ https://github.com/rfc-editor/rfc9069-AUTH48/

Community Support and Engagement

The RFC Editor typically meets with the stream managers 3-4 times a year. This allows the RFC Editor to discuss and collaborate on series-wide efforts as well as discuss any stream-specific issues. These meetings were suspended when the RSE departed at the end of 2019. We restarted these meetings, which was especially helpful for discussing the changes in errata handling and adoption of [NISTR 8366](#) (details below).

We provided office hours in Gather at each of the 3 virtual IETFs so that participants could ask questions about the publication process.

To improve transparency for the community and per the IETF Executive Director's request, we reviewed and redacted text from the RFC self reviews (2012 - 2020) so that they can be posted publicly.

We served as a hub for RFCXML-related questions by raising topics with the CMT, adding tickets to the issuetracker, and updating documentation. Some resources written by the RPC were used as input to the new central resource site for I-D authors (https://authors.ietf.org). We also provided feedback on a v3 template provided by authors.ietf.org.

We provided feedback to the IESG on an early draft of I-D nits / Guidelines for authors, the contents of which were added to [authors.ietf.org](#).

We updated the [online portion of the style guide](#) to include guidance about Inclusive Language. In addition, the team drafted a more significant update which will be posted in Q1 2022.

We have been closely tracking the RFC Editor Model v3 (rfced-future) discussion on the rfc-future mailing list and in meeting sessions and program calls. Each update to [draft-iab-rfcefdp-rfced-model](#) has been reviewed to ensure the RPC could implement the model as defined and highlight concerns, if any. In preparation for the pending transition, we have identified action items such as updates to our database and code, errata system, and webpages and have discussed them with stream managers and the IETF Executive Director.

The RPC participates in various community mailing lists in order to monitor RFC-Editor-related issues and also to serve as a resource where possible. These mailing lists include tools-development, tools-discuss, rfc-interest, rfc-markdown, emo-dir, xml-sg-cmt, and xml2rfc.

Adoption of Inclusive Language Policy

As described in the [IESG Statement on Inclusive Language](#), the IETF adopted [NISTR 8366](#), which provides guidance to NIST staff on inclusive language in documents. The RPC coordinated the adoption with the stream managers, and subsequently updated the RFC Style Guide, making it a series-wide recommendation (see [Web Portion of the Style Guide](#) for details). The RPC adjusted their processes to check for any of the "potentially biased language" in I-Ds prior to publication.

Note: [NISTR 8366](#) says:

> Table 1 was extracted in part to show examples of sentences with potentially biased language along with suggestions for unbiased alternatives. The full table can be found on the [NIST website](#).

The RPC checks for words included in the "full table" on the NIST website and discusses whether updates are needed with the authors and stream managers. As discussed with the stream managers, the RPC may request that authors review such language, but the authors together with the steam managers ultimately decide whether any changes are required.

Infrastructure Improvements

We performed a significant amount of testing in preparation for the OS upgrade of the RPC server (from OpenSuse 42.3(14.3) to OpenSuse 15.3), which was completed in October. This included major updates to essential software packages including Python and PHP. We tested all RPC workflows and tools beforehand on a separate server. This required close coordination with the AMS IT Director and RPC programmer.

Challenges

The RPC routinely reviews our processes and feedback received from authors to identify areas for possible improvements. These are some challenges the RPC faced in 2021.

AUTH48 process and source file management: A number of authors who seemingly keep their document source in version control repositories regularly request updates/reversions in the formatting during AUTH48 to eliminate "unnecessary" diffs. Author requests during AUTH48 include the editing of line breaks, spaces, target names, and anchor names in the XML in order

to make diffs cleaner and round-tripping between XML and markdown easier for the authors. Sometimes the authors make the changes themselves or offer to make them, other times they ask the RPC to make them. These requests can add length to the AUTH48 process.

As noted by one author in their post-publication survey:

> For markdown authors, it is cumbersome to port RPC edits back into the markdown sources, and sometimes it's impossible, because the XML gets tweaked in ways that markdown doesn't allow.

Using the RFCXML vocabulary: During AUTH48, some authors raise questions about the best practices for v3 XML formatting. For example, some don't like how <tt> is displayed or how <blockquote> is used. There is a documentation gap in that some of the formatting guidance resulting from discussion with various parties (e.g., tool developer, RSE, Temporary RSE, XML and Style Guide Change Management Team) has not been recorded and made available to the community.

Queue management and internal flexibility: Ongoing is the difficulty associated with managing clusters and the bursty nature of submissions. During 2021, almost 50% of the documents were part of a cluster, where cluster is as described on "What is a cluster?" On average, about a third of the documents published in the last few years have been part of a cluster, so this is an almost 20% increase. We speculate that the increase is due to the number of documents released from MISSREF. Note that the level of difficulty in editing clusters and the level of coordination increases as the size of a cluster increases. We continue to internally experiment to facilitate editorial work and communication across clustered documents.

Document submissions are bursty, and the RPC often suffers delays when work accumulates in one state. For example, if 50 documents are submitted in one month, there is a heavy load in EDIT. As documents move through the queue, there may be a buildup in subsequent states. To help address this, the RPC focused on cross-training in 2021. A couple of editors are now performing more senior-level reviews, which means that a greater portion of the editorial team can swing between document production states as needed.

We are internally tracking time spent on formatting and content editing in order to determine formatting overhead and process bottlenecks. Note that format and content updates are often intermingled, and the editors touch format at various times throughout the process. We'll be analyzing the time-tracking data to determine the level of uncertainty in the data and to improve the time-tracking process.

Errata system in need of evolution: There is also a known issue with the errata system in that updates to support verifiers are needed and the bot deterrent on the submission form isn't working as well as it should be.

Looking Ahead

2022 is going to be a year of change for the RPC, and the RFC community as a whole. In addition to editing and publishing RFCs, the RPC will also be engaged in the following:

XML v3 Formatting: As noted above, there is a documentation gap regarding best practices for v3 XML formatting. The RPC is discussing how to provide this guidance.

RFC Editor Model v3: We will continue to prepare for the upcoming transition to RFC Editor Model v3, which includes updating our system and code base to include the Editorial Stream, participating in RSWG and RSAB discussions, and updating web content to reflect the new model. We'll be adjusting to the new model as it's implemented.

Website Review: We inventoried the rfc-editor.org website in 2021 and started a content review to ensure the pages are up to date and accurate. Completion of the current review is a priority for 2022, and we will be reviewing the web content at regular intervals moving forward.

AUTH48: There is a need to update the AUTH48 process to facilitate issue tracking and diff reviews, and to increase transparency. We will continue to experiment with GitHub during AUTH48 and there are other potential AUTH48 process experiments on the horizon.
We will also further explore a short-term solution with the IESG to make the AUTH48 process more transparent.

RPC Tools: After discussion with the IETF Executive Director, the RPC will be evaluating our entire toolset and processes along with the IETF LLC Project Manager in 2022.

Other Updates: We will be reviewing the rest of the comments from the post-publication survey closely to assess where additional process review is needed, and whether there are any tool requirements that may feed into the RPC tools project mentioned above.

The RPC has recovered from the v3 transition, meeting the SLA and achieving a 97% approval rating while handling an increase in the number of documents, document page counts, and number of documents in clusters with the same number of personnel. In addition to editing and publishing, RPC now reviews and approves editorial errata, is experimenting with AUTH48 processes, and has increased its community engagement.

We look forward to collaborating with the RSWG and streams to advance the Series, and to working with the community to improve our processes to facilitate interactions.