

Internet Architecture Board (IAB)
Request for Comments: 8700
Updates: 2555, 5540
Category: Informational
ISSN: 2070-1721

H. Flanagan, Ed.
RFC Editor
December 2019

Fifty Years of RFCs

Abstract

This RFC marks the fiftieth anniversary for the RFC Series. It includes both retrospective material from individuals involved at key inflection points as well as a review of the current state of affairs. It concludes with thoughts on possibilities for the next fifty years for the Series. This document updates the perspectives offered in RFCs 2555 and 5540.

Status of This Memo

This document is not an Internet Standards Track specification; it is published for informational purposes.

This document is a product of the Internet Architecture Board (IAB) and represents information that the IAB has deemed valuable to provide for permanent record. It represents the consensus of the Internet Architecture Board (IAB). Documents approved for publication by the IAB are not candidates for any level of Internet Standard; see Section 2 of RFC 7841.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <https://www.rfc-editor.org/info/rfc8700>.

Copyright Notice

Copyright (c) 2019 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document.

Table of Contents

1. Introduction
2. Key Moments in RFC History
3. Perspectives
 - 3.1. The Origins of RFCs - by Stephen D. Crocker
 - 3.2. The RFC Management and Editing Team - by Vint Cerf
 - 3.3. Formalizing the RFC Editor Model - by Leslie Daigle
 - 3.4. The Continuation, or Creation, of a Stream - by Nevil Brownlee
 - 3.5. A View from inside the RFC Editor - by Sandy Ginoza
4. The Next Fifty Years of RFCs
 - 4.1. Preservation
 - 4.2. Evolution of the RFC Format
 - 4.3. Stream Structure
5. Conclusion

6. IANA Considerations
7. Security Considerations
8. Informative References
IAB Members at the Time of Approval
Acknowledgements
Contributors
Author's Address

1. Introduction

The RFC Series began in April 1969 with the publication of "Host Software" by Steve Crocker. The early RFCs were, in fact, requests for comments on ideas and proposals; the goal was to start conversations rather than to create an archival record of a standard or best practice. This goal changed over time, as the formality of the publication process evolved and the community consuming the material grew. Today, over 8500 RFCs have been published, ranging across best practice guidance, experimental protocols, informational material, and, of course, Internet standards. Material is accepted for publication through the IETF, the IAB, the IRTF, and the Independent Submissions streams, each of which have clear processes on how drafts are submitted and potentially approved for publication as an RFC. Ultimately, the goal of the RFC Series is to provide a canonical source for the material published by the RFC Editor and to support the preservation of that material in perpetuity.

The RFC Editor as a role came a few years after the first RFC was published. The actual date the term "RFC Editor" was first used is unknown, but it was formalized by [RFC0902] in July 1984; Jon Postel, the first RFC Editor, defined the role by his actions and later by defining the initial processes surrounding the publication of RFCs. What is certain is that the goal of the RFC Editor is to produce documents that are readable, clear, consistent, and reasonably uniform, and that the archival record of what has been published is maintained.

Change does come to the Series, albeit slowly. First, we saw the distribution method change from postal mail to FTP and then to email. RFCs could not be distributed electronically in the beginning, as the means to do that distribution would not be defined until years after the first RFC was "published". Not all early RFCs were even created electronically; some were written out by hand or on a typewriter. Eventually, the process for creating RFCs became more structured; authors were provided guidance on how to write an RFC. The editorial effort went from Steve Crocker to a more official model with a designated editor, Jon Postel, and later to a team of five to seven individuals. The actual editing and publishing work split from the service for registration of protocol code points. The whole RFC Editor structure was reviewed [RFC4844], refined [RFC5620], and refined again [RFC6635]. And, in the last few years, the process to change the format of the RFC documents themselves has started [RFC7990].

This is evolution; and the Series will continue to be adapted in order to meet the needs and expectations of the implementers, operators, historians, and community of authors that uses the RFC Series. These changes will always be balanced against the core mission of the Series: to maintain a strong, stable, archival record of technical specifications, protocols, and other information relevant to the Advanced Research Projects Agency Network (ARPANET) and Internet networking communities.

There is more to the history of the RFC Series than can be covered in this document. Readers interested in earlier perspectives may find the following RFCs of particular interest. These RFCs focus on the enormous contributions of Jon Postel, Czar of Socket Numbers

[RFC0433] and first RFC Editor:

- * [RFC2441]"Working with Jon, Tribute delivered at UCLA, October 30, 1998"
- * [RFC2555]"30 Years of RFCs"
- * [RFC5540]"40 Years of RFCs"

In this document, the history of the Series is viewed through the eyes of several individuals who have been a part of shaping it. Narratives of this nature offer a limited perspective on events; there are almost certainly other viewpoints, memories, and perspectives on events that are equally valid and would reflect a different history. So, while these retrospectives are enormously valuable and provide an insight to events of the day, they are just one lens on the history of the RFC Series.

Steve Crocker, author of [RFC0001], offers his thoughts on how and why the Series began. Leslie Daigle, a major influence in the development of the RFC Editor model, offers her thoughts on the change of the RFC Editor to a stronger, contracted function. Nevil Brownlee, Independent Submissions Editor from 2010 through February 2018, shares his view on the clarification of the Independent Stream (IS) and its transition upon the retirement of Bob Braden from the position. As the current RFC Series Editor, I will put my thoughts in on the most recent changes in formalizing the digital preservation of the Series, the process to modernize the format while respecting the need for stability, and my thoughts on the next fifty years of RFCs.

This document updates the perspectives offered in [RFC2555] and [RFC5540].

2. Key Moments in RFC History

Marker	Date	Event
[RFC0001]	April 1969	First RFC published
[RFC0114]	April 1971	First distribution of RFCs over the network
[RFC0433]	December 1972	First mention of the Czar of Socket Numbers and the proposal for a formal registry
[RFC0690]	June 1975	Relationship starts between the Information Sciences Institute (ISI) and the RFC Editor (judging by Jon Postel's affiliation change)
[RFC0748]	April 1977	First April 1st RFC published
[IETF1]	January 1986	First Internet Engineering Task Force (IETF) meeting
[IAB-19880712]	July 1988	IAB approved the creation of an Internet-Draft series
[RFC1122][RFC1123]	December	First major effort to review

	1988	key specifications and write applicability statements
[RFC1083]	October 1989	Three-stage standards process first defined
[RFC1150]	March 1990	FYI sub-series started
[RFC1311]	March 1992	STD sub-series started
[RFC1818]	August 1995	BCP sub-series started
[RFC-ONLINE]	approx. 1998	RFC Online Project to restore early RFCs that were "lost" started
[RFC2441]	15 October 1998	Jon Postel's death
[RFC4844]	July 2007	RFC Series administrative structure documented
[RFC4846]	July 2007	Independent Submission document stream is formalized
[RFC5620]	August 2009	RFC Editor organization officially established as RFC Series Editor, Independent Submission Editor, RFC Production Center, and RFC Publisher
[ISI-to-AMS]	October 2009	Transition of RFC Production Center and RFC Publisher starts from Information Sciences Institute (ISI) to Association Management Solutions (AMS)
[RFC5540]	January 2010	Bob Braden retires from RFC Editor
[RFC5743]	December 2009	Internet Research Task Force document stream formalized
[RFC-ONLINE]	approx. 2010	RFC Online Project to restore early RFCs that were "lost" finished
[RFC6360]	August 2011	FYI sub-series ended
[RFC6410]	October 2011	Two-stage standards process formalized
[RFC6635]	June 2012	Updated responsibilities of RFC Series allocated to RFC Series Editor, RFC Production Center, and RFC Publisher

[RFC6949]	May 2013	RFC format change project started
[RFC8153]	April 2017	RFCs no longer printed to paper upon publication

Table 1: Key Moments in RFC History

3. Perspectives

3.1. The Origins of RFCs - by Stephen D. Crocker

(This is a revision of material included more than 30 years ago in [RFC1000].)

The Internet community now includes millions of nodes and billions of users. It owes its beginning to the ARPANET, which was once but a gleam in the eyes of J. C. R. Licklider, Bob Taylor, and Larry Roberts of the Advanced Research Projects Agency (ARPA). While much of the development proceeded according to plan, the initial design of the protocols and the creation of the RFCs was largely accidental.

The procurement of the ARPANET was initiated in the summer of 1968; remember Vietnam, flower children, etc.? There had been prior experiments at various ARPA sites to link together computer systems, but this was the first version to explore packet switching as a core part of the communication strategy. ("ARPA" didn't become "DARPA" (Defense Advanced Research Projects Agency) until 1972. It briefly changed back to ARPA in 1993 and then back again to DARPA.) The government's Request for Quotations (RFQ) called for four packet-switching devices, called Interface Message Processors ("IMPs"), to be delivered to four sites in the western part of the United States: University of California, Los Angeles (UCLA); SRI International (Stanford Research Institute) in Menlo Park, CA; University of California, Santa Barbara (UCSB); and the University of Utah in Salt Lake City. These sites were running a Scientific Data Systems (SDS) Sigma 7, an SDS 940, an IBM 360/75, and a DEC PDP-10, respectively. These machines not only had different operating systems, but even details like character sets and byte sizes varied. Other sites would have further variations.

The focus was on the basic movement of data. The precise use of the ARPANET was not spelled out in advance, thus requiring the research community to take some initiative. To stimulate this process, a meeting was called in August 1968 with representatives from the selected sites, chaired by Elmer Shapiro from SRI. Based on Shapiro's notes from that meeting, the attendees were Dave Hopper and Jeff Rulifson from SRI; Glen Culler and Gordon Buck from Santa Barbara; R. Stephenson, C. Stephen Carr, and W. Boam from Utah; Vint Cerf and me from UCLA; and a few others from potential future sites.

That first meeting was seminal. We had lots of questions. How would IMPs and "hosts" (I think that was the first time I was exposed to that term) be connected? What would hosts say to each other? What applications would be supported? The only concrete answers were remote login as a replacement for dial-up, telephone-based interactive terminal access, and file transfer, but we knew the vision had to be larger. We found ourselves imagining all kinds of possibilities: interactive graphics, cooperating processes, automatic database query, electronic mail, etc., but no one knew where to begin. We weren't sure whether there was really room to think hard about these problems; surely someone senior and in charge, likely from the East, would be along by and by to bring the word. But we did come to one conclusion: we ought to meet again. Over the next

several months, we met at each of our sites, thereby setting the precedent for regular face-to-face meetings. We also instantly felt the irony. This new network was supposed to make it possible to work together at a distance, and the first thing we did was schedule a significant amount of travel.

Over the next several months, a small, fairly consistent set of graduate students and staff members from the first four sites met. We used the term Network Working Group (NWG) to designate ourselves. This was the same term Elmer Shapiro had used when he convened our first meeting, although it had been used until that point to refer to the principal investigators and ARPA personnel: senior people who had been planning the network. Our group was junior and disjointed from the prior group, except, of course, that each of us worked for one of the principal investigators.

The first few meetings were quite tenuous, primarily because we weren't sure how narrow or expansive our goals should be. We had no official charter or leadership, and it remained unclear, at least to me, whether someone or some group would show up with the official authority and responsibility to take over the problems we were dealing with. Without clear definition of what the host-IMP interface would look like, or even a precise definition of what functions the IMP would provide, we focused on broader ideas. We envisioned the possibility of application-specific protocols, with code downloaded to user sites, and we took a crack at designing a language to support this. The first version was known as DEL, for "Decode-Encode Language" and a later version was called NIL, for "Network Interchange Language".

In late 1968, Bolt Beranek and Newman (BBN) in Cambridge, MA won the contract for the IMPs and began work in January 1969. A few of us flew to Boston in the middle of February to meet the BBN crew. The BBN folks, led by Frank Heart, included Bob Kahn, Severo Ornstein, Ben Barker, Will Crowther, Bernie Cosell, and Dave Walden. They were organized, professional, and focused. Their first concern was how to meet their contract schedule of delivering the first IMP to UCLA at the beginning of September and how to get bits to flow quickly and reliably. The details of the host-IMP interface were not yet firm; the specification came a few months later as BBN Report 1822. In particular, BBN didn't take over our protocol design process, nor did any other source of authority appear. Thus, we doggedly continued debating and designing the protocols.

A month later, our small NWG met in Utah. As the meeting came toward an end, it became clear to us that we should start writing down our discussions. We had accumulated a few notes on the design of DEL and other matters, and we decided to put them together in a set of notes. We assigned writing chores to each of us, and I took on the additional task of organizing the notes. Though I initiated the RFCs, my role was far less than an editor. Each of the RFCs were numbered in sequence. The only rule I imposed was the note had to be complete before I assigned a number because I wanted to minimize the number of holes in the sequence.

I tried a couple of times to write a note on how the notes would be organized, but I found myself full of trepidation. Would these notes look as if we were asserting authority we didn't have? Would we unintentionally offend whomever the official protocol designers were? Finally, unable to sleep, I wrote a few humble words. The basic ground rules were that anyone could say anything and that nothing was official. And to emphasize the point, I used Bill Duvall's suggestion and labeled the notes "Request for Comments". I never dreamed these notes would eventually be distributed through the very medium we were discussing in these notes: talk about Sorcerer's Apprentice! (See [APPRENTICE].)

After BBN distributed the specification for the IMP hardware and software interface to the initial ARPANET sites, our attention shifted to low-level matters. The ambitious ideas for automatic downloading of code evaporated. It would be several years before ideas like mobile code, remote procedure calls, ActiveX, JAVA, and Representational State Transfer (RESTful) interfaces appeared.

Over the spring and summer of that year, we grappled with the detailed problems of protocol design. Although we had a vision of the vast potential for intercomputer communication, designing usable protocols was another matter. We knew a custom hardware interface and a custom software addition in the operating system was going to be required for anything we designed, and we anticipated these would pose some difficulty at each of the sites. We looked for existing abstractions to use. It would have been convenient if we could have made the network simply look like a regular device, e.g., a tape drive, but we knew that wouldn't do. The essence of this network was peer-to-peer cooperation among the machines and the processes running inside them, not a central machine controlling dependent devices. We settled on a virtual bit stream layer as the basic building block for the protocols; but even back then, we knew that some applications like voice might need to avoid that layer of software. (Why a virtual bit stream instead of a virtual byte stream? Because each computer had its own notion of how many bits were in a byte. 8-bit bytes didn't become standard until a few years later.)

Over the next two years, we developed, exchanged, and implemented ideas. I took a leave from UCLA in June 1971 to spend time working at ARPA. Jon Postel took over the care and feeding of the RFCs, evolving the process and adding collaborators over the next twenty-seven years.

The rapid growth of the network and the working group also led to a large pile of RFCs. When the 100th RFC was in sight, Peggy Karp at the MIT Research Establishment (MITRE) took on the task of indexing them. That seemed like a large task then, and we could have hardly anticipated seeing more than 1000 RFCs several years later and the evolution toward Internet-Drafts yet later.

When we first started working on the protocols, the network did not exist. Except for our occasional face-to-face meetings, RFCs were our only means of communication. In [RFC0003], I set the bar as low as possible:

The content of a NWG note may be any thought, suggestion, etc. related to the HOST software or other aspect of the network. Notes are encouraged to be timely rather than polished. Philosophical positions without examples or other specifics, specific suggestions or implementation techniques without introductory or background explication, and explicit questions without any attempted answers are all acceptable. The minimum length for a NWG note is one sentence.

These standards (or lack of them) are stated explicitly for two reasons. First, there is a tendency to view a written statement as ipso facto authoritative, and we hope to promote the exchange and discussion of considerably less than authoritative ideas. Second, there is a natural hesitancy to publish something unpolished, and we hope to ease this inhibition.

Making the RFCs informal was not only a way of encouraging participation; it was also important in making the communication effective. One of the early participants said he was having trouble writing and sending an RFC because his institution wanted to subject

them to publication review. These are not "publications", I declared, and the problem went away. Another small detail, handled instinctively and without debate, was the distribution model. Each institution was required to send a copy directly to each of the other handful of participating institutions. Each institution handled internal copies and distribution itself. Submission to a central point for redistribution was not required so as to minimize delays. SRI's Network Information Center, however, did maintain a central repository of everything and provided an invaluable record.

We didn't intentionally set out to challenge the existing standards organizations, but our natural mode of operation yielded some striking results. The RFCs are open in two important respects: anyone can write one for free and anyone can get them for free. At the time, virtually everyone in the ARPANET community was sponsored by the government, so there was little competition and no need to use documents as a way of raising money. Of course, as soon as we had email working on the ARPANET, we distributed RFCs electronically. When the ARPANET became just a portion of the Internet, this distribution process became worldwide. The effect of this openness is often overlooked; even now, students and young professionals all over the world have been able to download RFCs, learn about the technology within, and in turn, build the most amazing software. (They are also a fantastic resource for historians.)

Where will it end? The ARPANET begat the Internet, and the underlying technology transitioned from the original host-host protocol to TCP/IP. But, the superstructure of protocol layers, community-driven protocol design, and RFCs continued. Through the changes in physical-layer technology, resulting in speed increases from thousands to billions of bits per second, and similarly from thousands to billions of users, this superstructure, including the RFCs, has continued to serve the community. All of the computers have changed, as have all of the transmission lines, but the RFCs march on. Maybe I'll write a few words for RFC 10,000.

Quite obviously, the circumstances have changed. Email and other media are most often used for the immediate exchange of inchoate thoughts. Internet-Drafts are the means for exchanging substantial, albeit sometimes speculative, content, while RFCs are reserved for fully polished, reviewed, edited, and approved specifications. Comments to RFCs are not requested, although usage-related discussions and other commentary on mailing lists often take place nonetheless. Rather than bemoan the change, I take it as a remarkable example of adaptation. RFCs continue to serve the protocol development community. Indeed, they are the bedrock of a very vibrant and productive process that has fueled and guided the Internet revolution.

3.2. The RFC Management and Editing Team - by Vint Cerf

As Steve Crocker mentions in Section 3.1, Jon Postel assumed the role of RFC manager in 1971 when Steve left UCLA for ARPA. Jon took on this role in addition to his subsequent "numbers Czar" responsibilities. Initially, his focus was largely on assigning RFC numbers to aspiring writers, but with time, and as the standardization of the ARPANET and Internet protocols continued apace, he began to serve in an editorial capacity. Moreover, as an accomplished software engineer, he had opinions about technical content in addition to writing style and did not hesitate to exercise editorial discretion as would-be published authors presented their offerings for his scrutiny. As the load increased, he recruited additional "volunteer" talent, most notably Joyce K. Reynolds, a fellow researcher at USC/ISI. Over the ensuing years, he also drafted Robert (Bob) Braden onto the team, and when Jon unexpectedly passed away in October 1998 (see [RFC2468]), Joyce and Bob undertook

carrying on with the RFC work in his stead, adding Sandy Ginoza to the team. During the period when Jon and Joyce worked closely together, Joyce would challenge me to tell which edits had been made by Jon and which by her. I found this impossible, so aligned were they in their editorial sensibilities. Sadly, three of these tireless Internauts have passed on, and we have only the product of their joint work and Sandy Ginoza's and others' corporate memory by which to recall history.

3.3. Formalizing the RFC Editor Model - by Leslie Daigle

I was the chair of the Internet Architecture Board, the board responsible for the general oversight of the RFC Series, at a particular inflection point in the evolution of all Internet technology institutions. To understand what we did, and why we had to, let me first paint a broader picture of the arc of these institutions.

Like many others who were in decision-making roles in the mid '00s, I wasn't present when the Internet was born. The lore passed down to me was that, out of the group of talented researchers that developed the core specifications and established the direction of the Internet, different individuals stepped up to take on roles necessary to keep the process of specification development organized and open. As the work of specification expanded, those individuals were generally supported by organizations that carried on in the same spirit. This was mostly Jon Postel, managing the allocation and assignment of names and numbers, as well as working as the editor of RFCs, but there were also individuals and institutions supporting the IETF's Secretariat function. By the late 20th century, even this model was wearing thin; the support functions were growing, and organizations didn't have the ability to donate even more resources to run them. In some cases (IANA), there was significant industry and international dependence on the function and its neutrality.

The IETF, too, had grown in size, stature, and commercial reliance. This system of institutional pieces "flying in formation" was not providing the kind of contractual regularity or integrated development that the IETF needed. People who hadn't been there as the institutions developed, including IETF decision makers, didn't innately understand why things "had to be the way they were" and were frustrated when trying to get individual systems updated for new requirements as well as better integrated across the spectrum of activities.

Internet engineering had expanded beyond the point of being supportable by a loosely coupled set of organizations of people who had been there since the beginning and knew each other well. New forms of governance were needed along with a rationalized funding model. The IANA function was absorbed into a purpose-built international not-for-profit organization. The IETF stepped up to manage its own organizational destiny, creating the IETF Administrative Support Activity (IASA), and the Secretariat became one of its contracted functions.

This left the RFC Editor function as an independent effort supported by the Internet Society.

That independent nature was necessary for the historic role of the RFC Series in considering all technical contributions. But, at that inflection point in the Series' history, it needed a new governance and funding model, just as the other organizations supporting Internet technical specification had. Also, the IETF leadership had some concerns it felt needed to be addressed in its own technical publication stream. While the RFC Series had been established before there was an IETF, and had historically continued to have documents

in it that didn't originate from the IETF, the IETF was its largest and most organized contributor. There was no particular organization of independent contributors. Equally, the funding for the RFC Editor was at that point coming from the Internet Society in the guise of "support for the IETF". For people who hadn't been involved with the institution from the outset, it was pretty easy to perceive the RFC Series uniquely as the IETF's publication series. So, the challenge was to identify and address the IETF's issues, along with governance and funding, without sacrificing the fundamental nature of the RFC Series as a broader-than-IETF publication series.

To give a sense of the kind of tensions that were prevalent, let me share that the one phrase that stuck in my mind from those discussions was "push to publish". There were those in IETF leadership who felt that it would significantly reduce costs and improve timeliness if an RFC could be published by, literally, pushing a button on a web interface the moment it was approved by the IESG. It would also, they argued, remove the specification issues being introduced by copy editors that were hired as occasional workers to help with improving publication rates but who weren't necessarily up to speed on terms of art in technical specifications. (There were some pretty egregious examples of copy editors introducing changes that significantly changed the technical meaning of the text that I forbear from citing here; let's just say it wasn't strictly a problem of Internet engineers getting uptight about their cheese being moved.) While "push to publish" would have addressed those issues, it would not have addressed the loss of clarity from the many significant text improvements copy editors successfully introduced, or the fact that not all RFCs are approved by the IESG.

Institutionally, it was clear that the target was to have the RFC Editor function governance within the reach of the Internet technical community (as opposed to any particular private organization) without tying it specifically to the IETF. That was reasonably achievable by ensuring that the resultant pieces were established under the oversight of the IAB (which is, itself, independent of the IETF even as it is supported by the IASA organization).

The IETF worked on a document outlining functional requirements for its technical specification publication. This could have been useful for establishing its own series, but it also was helpful in establishing awareness of the challenges in document publishing (it always looks easy when you haven't thought about it) and also in laying the groundwork for dialogue with the RFC Editor. The requirements document was published as [RFC4714] as an Informational RFC that stands today to provide guidance in the editing processes surrounding IETF publications.

There was still, however, a certain lack of clarity about responsibilities for making decisions and changes in the RFC Series itself. To that end, I and the IAB worked with the various involved parties to produce [RFC4844]. That document captured the RFC Series mission (for a purpose greater than IETF technical specification publication) as well as the roles and responsibilities of the parties involved. The RFC Editor is responsible for ensuring the implementation of the mission. The IAB continues to have oversight responsibilities, including policy oversight, which it could act on by changing the person (organization) in the role of RFC Editor. At the same time, operational oversight was migrated to the IASA support function of the IETF (and IAB).

The discussions, and the resulting publication of [RFC4844], allowed greater visibility into and commitment to the RFC Series as a general Internet publication. It also meant that subsequent adjustments could be made as requirements evolved; the responsible parties are clearly identified.

3.4. The Continuation, or Creation, of a Stream - by Nevil Brownlee

Arguably starting in 2006 with [RFC4714], the IAB and the IETF community spent some time in the mid-'00s evolving the structure of the RFC Series. This work included defining how those groups that published into the RFC Series (initially including the IETF, the IAB [RFC4845], and the Independent Submissions Stream [RFC4846], and later growing to include the IRTF [RFC5743]) would handle approving documents to be published as RFCs. In 2009, the IAB published "RFC Editor Model (Version 1)" [RFC5620]. In this model, a new role was created within the RFC Editor: the RFC Series Editor (RSE). This individual would oversee RFC publishing and development while leaving the process for approving documents for publication outside his or her mandate. While, arguably, this was a role long filled by people like Jon Postel, Bob Braden, and Joyce Reynolds, [RFC5620] saw the role of RFC Series Editor defined in such a way as to distinctly separate it from that of the Independent Submissions Editor (ISE).

Before 2009, the RFC Editor could accept "Independent" submissions from individuals and, if they were judged significant, publish them as RFCs; the Independent Stream was set up to continue that function. From February 2010 through February 2018, I was the ISE. After reading [RFC4846], I went on to develop the Independent Stream (IS).

First, I spent several days at the RFC Production Center at the Information Sciences Institute (ISI) in Marina Del Ray with the RFC Editor (Bob Braden), Sandy Ginoza, and Alice Hagens so as to learn how RFCs were actually edited and published. All RFCs reach the Production Center as Internet-Drafts; they are copy edited until the edited version can be approved by its authors (AUTH48). At any stage, authors can check their draft's status via the RFC Editor website.

For the Independent Submissions, Bob kept a journal (a simple ASCII file) of his interactions with authors for every draft, indexed by the draft name. Bob also entered the Independent drafts into the RFC Editor database so that authors could track their draft's status. After my few days with his team at ISI, he handed that journal (covering about 30 drafts) over to me and said, "Now it's over to you!"

I began by following in Bob's footsteps, maintaining a journal and tracking each draft's status in the RFC Editor database. My first consideration was that every serious Internet-Draft submitted needed several careful reviews. At that time, if the ISE knew of suitable reviewers, he or she could simply ask them. Otherwise, if the draft related to an IETF or IRTF Working Group, the ISE could ask Working Group Chairs or Area Directors to suggest reviewers. The Independent Submissions Editorial Board (Ed Board) was another place the ISE could request reviewers from. My experience with reviewers was that most of those I approached were happy to help.

Most drafts were straightforward, but there were some that needed extra attention. Often, a draft requested IANA code points, and for that, IANA was always quick to offer help and support. Code points in some IANA registries require Expert Review [RFC8126]; sometimes the interactions with Expert Reviewers took quite a long time! Again, sometimes a draft seemed to fit better in the IETF Stream; for these, I would suggest that the draft authors try to find an Area Director to sponsor their work as an individual submission to the IETF Stream.

After my first few years as ISE, the IETF Tools Team developed the Datatracker [DATATRACKER] to show draft status and perform all the "housekeeping" tasks for all of the streams. At that stage, I

switched to using the Datatracker rather than the RFC Editor database.

Once a draft has been reviewed and the authors have revised it in dialogue with their reviewers, the ISE must submit that draft to the IESG for an "IESG Review" [RFC5742]. Overall, each IS draft benefited from discussions (which were usually simple) with my Ed Board and the IESG. A (very) few were somewhat controversial; for those, I was able to work with the IESG to negotiate a suitable "IESG Statement" and/or an "ISE Statement" to make it clear why the ISE published the draft.

One rather special part of the Independent Stream is the April 1st RFCs. These are humorous RFCs that have no formal review and approval process. The authors must send them directly to the ISE or the RFC Editor. Only a few of them can be published each year, and each is reviewed by the ISE and the RSE. Bob Braden's criteria for April 1st drafts were:

- * They must relate to the Internet (like all drafts).
- * Their readers should reach the end of page two before realizing it is an April 1st RFC.
- * They must actually be funny!

April 1st RFCs have a large following, and feedback from the Internet community on April 1st of each year has been enthusiastic and quick!

159 RFCs were published in the Independent Stream during my eight years as ISE. Over those eight years, I worked with most of their authors and often met with them at IETF meetings. For me, that was a very rewarding experience, so I thank all those that contributed. During those eight years, I also worked with most of the IESG members, who all also gave me a lot of helpful interaction. Lastly, I've always enjoyed working with the RSE and all the staff of the RFC Production Center. The IETF (as a whole) is very fortunate to have such an effective team of talented professional staff.

3.5. A View from inside the RFC Editor - by Sandy Ginoza

When I joined ISI, shortly after Jon Postel passed away, the RFC Editor model as we know it today (as defined in [RFC5620] and as obsoleted by [RFC6548] and [RFC6635]) did not exist. The RFC Editor functioned as one unit: there was no RSE, Production Center, Publisher, or Independent Submissions Editor. All of these roles were performed by the "RFC Editor", which was comprised of four individuals: Bob Braden, Joyce Reynolds, a part-time student programmer, and me.

Bob provided high-level guidance and reviewed Independent Submissions. While Bob was a researcher in "Div 7" (Networking) at ISI, ostensibly, the percentage of time he had for the RFC Editor was 10%, but he invested much more time to keep the Series running. He pitched in where he could, especially when processing times were getting longer; at one point, he even NROFFed a couple of RFCs-to-be.

Joyce was a full-time ISI employee. However, while continuing to ensure RFCs were published, she was also serving as a User Services Area Director and a keynote speaker about the Internet, and she was also temporarily on loan to IANA for 50% of her time while IANA was getting established after separating from ISI. The student programmer performed programming tasks as requested and was, at the time, responsible for parsing MIBs.

I was a full-time staffer and had to quickly learn the ropes so RFCs

would continue to be published. My primary tasks were to manage the publication queue, format and prepare documents for Joyce's review, carry out AUTH48 once Joyce completed her review, and publish, index, and archive the RFCs (both soft and hard copies).

The workload increased significantly over the next few years. As the workload increased, the RFC Editor reacted and slowly grew their staff over time. To understand the team growth, let's first take a look at the publication rates throughout history. The table below shows average annual publication rates during 5-year periods.

Years	Avg. Pubs per Year
1969 - 1972	80
1973 - 1977	55
1978 - 1982	20
1983 - 1987	39
1988 - 1992	69
1993 - 1997	171
1998 - 2002	237
2003 - 2007	325
2008 - 2012	333
2013 - 2017	295

Table 2: Annual Publication Rates

There were significant jumps in the publication rates in the '90s and onward, with the number of publications almost doubling between 1993 and 2007. The annual submission count surpassed the 300 mark for the first time in 2004 and reached an all-time high of 385 in 2011. The submission rate did not drop below 300 until 2016 (284).

As the submissions grew, the RFC Editor experienced growing pains. Processing times began to increase as the existing staff was unable to keep up with the expanding queue size. In an attempt to reduce the training hump and to avoid permanently hiring staff in case the submission burst was a fluke, ISI brought on temporary copy editors; this way, the staff could easily be resized as needed. However, as Leslie noted, this didn't work very well. The effects of the experiment would be lasting, as this led to a form of the process we have now, where the RFC Editor asks more questions during AUTH/AUTH48 and technical changes require approval from the relevant Area Directors or stream managers, depending on the document stream. These changes added to the workload and extended publication times; many often now jokingly refer to AUTH48 as the authors' "48 days", "48 weeks", etc.

In addition to the increase in document submissions, we engaged in tools testing and went through several editorial process changes. Because of the lesson learned with temporary copy editors, our team grew to be more permanent. While we added other editors in between, two additions are of particular interest, as they experienced much of the RFC Editor's growing pains, helped work us out of a backlogged state, shaped the RFC Editor function, and are still with the team today: Alice Russo joined the team in 2005 and Megan Ferguson joined

us in 2007.

With the understanding that the record-breaking number of submissions was not an anomaly, we made significant upgrades to the infrastructure of the RFC Editor function to facilitate document tracking and reporting. For example, the illustrious "black binder" (an actual 3-ring binder used to track RFC number assignment), a manually edited HTML file for the queue page, and a Rube Goldberg set of text files and scripts that created queue statistics, all were eventually replaced; an errata system was proposed and implemented; and XML became a newly accepted source file.

In 2009, [RFC5620] was published, introducing the initial version of the RFC Editor model we have now. While it was published in 2009, it did not go into effect until 2010, when the RFC Editor project as I knew it was disbanded and divvied up into four pieces: RFC Series Editor (RSE), Independent Submissions Editor (ISE), RFC Production Center (RPC), and the Publisher function. In addition, the RFC Series Advisory Group (RSAG) was created to "provide expert, informed guidance (chiefly, to the RSE) in matters affecting the RFC Series operation and development" [RSAG].

In 2010, the RPC and Publisher contracts were awarded to Association Management Solutions (AMS). There, we started with three existing team members (Alice Russo, Megan Ferguson, and me), and we were pleased to be joined by Lynne Bartholomew and Rebecca VanRheenen, new colleagues to anchor us in the AMS office.

I was wary of this model and was especially worried about the hole Bob Braden's departure would create. Luckily for us, Bob Braden provided wise counsel and insight during the transition (and beyond). He gave the staff transitioning to AMS particularly helpful parting words, "keep the RFCs coming", and that is what we did.

AMS embraced the RFC Series and helped us quickly get set up on new servers. The RFC Production Center and Publisher were now part of the AMS family and it was all hands on deck to make sure the transition went smoothly to minimize the impact on document processing.

Our focus during transition was to 1) keep the trains running; that is, we wanted to get ourselves up and running with minimal down time, and 2) work with the Transitional RSE (a role that concluded before the transition ended), the ISE (Nevil Brownlee), RSAG, and the IETF Administrative Director (IAD) to better understand and implement the newly defined RFC Editor model.

Though some portions of the transition were challenging and lasted longer than expected, the Acting RSE (Olaf Kolkman) officially handed the reins over to the new RSE (Heather Flanagan) in 2012. She had to jump in, learn the RFC Editor and IETF culture, and work through a backlog of issues that had been left unattended.

Two of the backlogged issues were so old that they were ones someone had asked me about at my first IETF meeting: When was the RFC Editor going to allow non-ASCII characters in RFCs? When would the RFC Editor adopt a more modern publication format?

At that time, while we understood the desire to move toward supporting a broader range of character sets and having more-modern outputs, we also routinely received emails from individuals requesting that we send them plaintext files (instead of pointing them to the website) because their Internet access was limited. We also regularly received complaints from users of <<https://www.rfc-editor.org>> whenever something on the site didn't work correctly with their older browsers. In short, we could not advance without leaving

a large number of users behind.

However, we now find ourselves on the precipice of change. The next few years promise to be exciting for the RFC Series as we transition from publishing plaintext, ASCII-only files to publishing multiple file formats (XML, HTML, PDF/A-3, and TXT) that allow both non-ASCII characters and SVG art.

Interestingly enough, I find that the RFC Editor has been in an almost constant state of change since I joined the team, even though the goal of the RFC Editor remains the same: to produce archival quality RFCs in a timely manner that are easily accessible for future generations.

4. The Next Fifty Years of RFCs

As Steve Crocker mentioned, the Series began with goals of communication over formality and openness over structure. As the Internet has grown and become a pervasive, global construct, we still aim for openness and communication, but recognize that for protocols and other information to support interoperability, there must be points of stability to build from. Everyone, from small-time app developers to multi-billion dollar companies, is on the same footing. Anyone should be able to look back at a point in time and understand what was done and why.

While the informality has given way to increased structure, the openness and solid foundation that the Series provides must continue. With that in mind, what does the future hold for the next fifty years of RFCs?

4.1. Preservation

The RFC Editor exists to edit, publish, and maintain an archive of documents published in the RFC Series. A proper digital archive, however, is more than just saving RFCs to disk and making sure the disks are backed up; the field of digital preservation has grown and transformed into an industry in and of itself. "Digital Preservation Considerations for the RFC Series" [RFC8153] reviews what a digital archive means today and describes ways to support the archive into the future. It also recommends ways for the RFC Editor to take advantage of those organizations that specialize in this field.

The future of digital preservation, as far as the RFC Series is concerned, will mean both finding new partners that can absorb and archive RFCs into a public, maintained digital archive and reviewing the RFC format to ensure that the published documents are archivable according to whatever the industry best practice is over time.

4.2. Evolution of the RFC Format

RFCs have been digital documents since very early in the days of the Series. While not always published in US-ASCII, that format has been the canonical format for decades. The fact that this format has lasted through so much evolution and change is remarkable.

Unfortunately, the US-ASCII format does not extend enough to meet the expectations and requirements of many users today. The entire field of online document presentation, consumption, and preservation has, in some cases, only been invented years after the first RFC was published. While it can be (and has been) argued that those newer fields and their tools have not had a chance to stand the test of time, the RFC Series Editor (in consultation with the community) started a concerted effort in 2012 to bring the RFC Series into alignment with a new array of possibilities for preservation and display.

Information on the RFC format project and the initial reasoning and requirements for the changes underway can be found in [RFC7990]. With the advent of these changes, the door has been opened to consider further changes in the future as the specifications for archiving digital material evolves, and as the expectation of web development advances.

4.3. Stream Structure

In the eyes of many, particularly within the IETF, the RFC Series is synonymous with the IETF. While the Series itself predates the IETF by eighteen years, over time, the IETF has become the source of the majority of documents submitted for publication to the RFC Editor. The policies developed for IETF Stream drafts tend to apply across all four document streams, and publication-related tools tend to focus on the IETF as the primary audience for their use. It is difficult for people to see how, or even why, there is a distinction between the Series and the IETF.

We are in the midst of that question now more than ever. What is the future of the Series? If people cannot tell where the IETF ends and the Series starts, should we consider this an artificial distinction and declare them to be the same entity?

Ultimately, this will be something the community decides, and conversations are underway to consider the ramifications of possible changes.

5. Conclusion

As the Internet evolves, expectations and possibilities evolve, too. Over the next fifty years, the Series will continue to demonstrate a balance between the need to stay true to the original mission of publication and preservation, while also staying relevant to the needs of the authors and consumers of RFCs. The tension in balancing those needs rests on the RFC Editor and the community to resolve. We will not run short of challenges.

6. IANA Considerations

This document has no IANA actions.

7. Security Considerations

This document has no security considerations.

8. Informative References

[APPRENTICE]

Wikipedia, "The Sorcerer's Apprentice", December 2019, <https://en.wikipedia.org/w/index.php?title=The_Sorcerer%27s_Apprentice&oldid=925824658>.

[DATATRACKER]

Internet Engineering Task Force, "IETF Datatracker", <<https://datatracker.ietf.org>>.

[IAB-19880712]

IAB, "IAB Minutes 1988-07-12", July 1988, <<https://www.iab.org/documents/minutes/minutes-1988/iab-minutes-1988-07-12/>>.

[IETF1]

The MITRE Corporation, "Proceedings of the 16-17 January 1986 DARPA Gateway Algorithms and Data Structures Task Force", IETF 1, January 1986,

<<https://www.ietf.org/old/2009/proceedings/prior29/IETF01.pdf>>.

[ISI-to-AMS]

IETF Administrative Support Activity (IASA), "RFC Production Center Agreement between Association Management Solutions, LLC and The Internet Society", October 2009, <<https://iaoc.ietf.org/documents/AMS-RPC-Public-Final-2009.pdf>>.

[RFC-ONLINE]

RFC Editor, "History of RFC Online Project", 2000, <<https://www.rfc-editor.org/rfc-online-2000.html>>.

[RFC0001] Crocker, S., "Host Software", RFC 1, DOI 10.17487/RFC0001, April 1969, <<https://www.rfc-editor.org/info/rfc1>>.

[RFC0003] Crocker, S., "Documentation conventions", RFC 3, DOI 10.17487/RFC0003, April 1969, <<https://www.rfc-editor.org/info/rfc3>>.

[RFC0114] Bhushan, A., "File Transfer Protocol", RFC 114, DOI 10.17487/RFC0114, April 1971, <<https://www.rfc-editor.org/info/rfc114>>.

[RFC0433] Postel, J., "Socket number list", RFC 433, DOI 10.17487/RFC0433, December 1972, <<https://www.rfc-editor.org/info/rfc433>>.

[RFC0690] Postel, J., "Comments on the proposed Host/IMP Protocol changes", RFC 690, DOI 10.17487/RFC0690, June 1975, <<https://www.rfc-editor.org/info/rfc690>>.

[RFC0748] Crispin, M., "Telnet randomly-lose option", RFC 748, DOI 10.17487/RFC0748, April 1978, <<https://www.rfc-editor.org/info/rfc748>>.

[RFC0902] Reynolds, J. and J. Postel, "ARPA Internet Protocol policy", RFC 902, DOI 10.17487/RFC0902, July 1984, <<https://www.rfc-editor.org/info/rfc902>>.

[RFC1000] Reynolds, J. and J. Postel, "Request For Comments reference guide", RFC 1000, DOI 10.17487/RFC1000, August 1987, <<https://www.rfc-editor.org/info/rfc1000>>.

[RFC1083] Defense Advanced Research Projects Agency and Internet Activities Board, "IAB official protocol standards", RFC 1083, DOI 10.17487/RFC1083, December 1988, <<https://www.rfc-editor.org/info/rfc1083>>.

[RFC1122] Braden, R., Ed., "Requirements for Internet Hosts - Communication Layers", STD 3, RFC 1122, DOI 10.17487/RFC1122, October 1989, <<https://www.rfc-editor.org/info/rfc1122>>.

[RFC1123] Braden, R., Ed., "Requirements for Internet Hosts - Application and Support", STD 3, RFC 1123, DOI 10.17487/RFC1123, October 1989, <<https://www.rfc-editor.org/info/rfc1123>>.

[RFC1150] Malkin, G. and J. Reynolds, "FYI on FYI: Introduction to the FYI Notes", RFC 1150, DOI 10.17487/RFC1150, March 1990, <<https://www.rfc-editor.org/info/rfc1150>>.

[RFC1311] Postel, J., "Introduction to the STD Notes", RFC 1311, DOI 10.17487/RFC1311, March 1992,

<<https://www.rfc-editor.org/info/rfc1311>>.

- [RFC1818] Postel, J., Li, T., and Y. Rekhter, "Best Current Practices", RFC 1818, DOI 10.17487/RFC1818, August 1995, <<https://www.rfc-editor.org/info/rfc1818>>.
- [RFC2441] Cohen, D., "Working with Jon, Tribute delivered at UCLA, October 30, 1998", RFC 2441, DOI 10.17487/RFC2441, November 1998, <<https://www.rfc-editor.org/info/rfc2441>>.
- [RFC2468] Cerf, V., "I REMEMBER IANA", RFC 2468, DOI 10.17487/RFC2468, October 1998, <<https://www.rfc-editor.org/info/rfc2468>>.
- [RFC2555] Editor, RFC. and et. al., "30 Years of RFCs", RFC 2555, DOI 10.17487/RFC2555, April 1999, <<https://www.rfc-editor.org/info/rfc2555>>.
- [RFC4714] Mankin, A. and S. Hayes, "Requirements for IETF Technical Publication Service", RFC 4714, DOI 10.17487/RFC4714, October 2006, <<https://www.rfc-editor.org/info/rfc4714>>.
- [RFC4844] Daigle, L., Ed. and Internet Architecture Board, "The RFC Series and RFC Editor", RFC 4844, DOI 10.17487/RFC4844, July 2007, <<https://www.rfc-editor.org/info/rfc4844>>.
- [RFC4845] Daigle, L., Ed. and Internet Architecture Board, "Process for Publication of IAB RFCs", RFC 4845, DOI 10.17487/RFC4845, July 2007, <<https://www.rfc-editor.org/info/rfc4845>>.
- [RFC4846] Klensin, J., Ed. and D. Thaler, Ed., "Independent Submissions to the RFC Editor", RFC 4846, DOI 10.17487/RFC4846, July 2007, <<https://www.rfc-editor.org/info/rfc4846>>.
- [RFC5540] Editor, RFC., "40 Years of RFCs", RFC 5540, DOI 10.17487/RFC5540, April 2009, <<https://www.rfc-editor.org/info/rfc5540>>.
- [RFC5620] Kolkman, O., Ed. and IAB, "RFC Editor Model (Version 1)", RFC 5620, DOI 10.17487/RFC5620, August 2009, <<https://www.rfc-editor.org/info/rfc5620>>.
- [RFC5742] Alvestrand, H. and R. Housley, "IESG Procedures for Handling of Independent and IRTF Stream Submissions", BCP 92, RFC 5742, DOI 10.17487/RFC5742, December 2009, <<https://www.rfc-editor.org/info/rfc5742>>.
- [RFC5743] Falk, A., "Definition of an Internet Research Task Force (IRTF) Document Stream", RFC 5743, DOI 10.17487/RFC5743, December 2009, <<https://www.rfc-editor.org/info/rfc5743>>.
- [RFC6360] Housley, R., "Conclusion of FYI RFC Sub-Series", RFC 6360, DOI 10.17487/RFC6360, August 2011, <<https://www.rfc-editor.org/info/rfc6360>>.
- [RFC6410] Housley, R., Crocker, D., and E. Burger, "Reducing the Standards Track to Two Maturity Levels", BCP 9, RFC 6410, DOI 10.17487/RFC6410, October 2011, <<https://www.rfc-editor.org/info/rfc6410>>.
- [RFC6548] Brownlee, N., Ed. and IAB, "Independent Submission Editor Model", RFC 6548, DOI 10.17487/RFC6548, June 2012, <<https://www.rfc-editor.org/info/rfc6548>>.

- [RFC6635] Kolkman, O., Ed., Halpern, J., Ed., and IAB, "RFC Editor Model (Version 2)", RFC 6635, DOI 10.17487/RFC6635, June 2012, <<https://www.rfc-editor.org/info/rfc6635>>.
- [RFC6949] Flanagan, H. and N. Brownlee, "RFC Series Format Requirements and Future Development", RFC 6949, DOI 10.17487/RFC6949, May 2013, <<https://www.rfc-editor.org/info/rfc6949>>.
- [RFC7990] Flanagan, H., "RFC Format Framework", RFC 7990, DOI 10.17487/RFC7990, December 2016, <<https://www.rfc-editor.org/info/rfc7990>>.
- [RFC8126] Cotton, M., Leiba, B., and T. Narten, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 8126, DOI 10.17487/RFC8126, June 2017, <<https://www.rfc-editor.org/info/rfc8126>>.
- [RFC8153] Flanagan, H., "Digital Preservation Considerations for the RFC Series", RFC 8153, DOI 10.17487/RFC8153, April 2017, <<https://www.rfc-editor.org/info/rfc8153>>.
- [RSAG] RFC Editor, "RFC Series Advisory Group", <<https://www.rfc-editor.org/about/rsag/>>.

IAB Members at the Time of Approval

Jari Arkko
Alissa Cooper
Stephen Farrell
Wes Hardaker
Ted Hardie
Christian Huitema
Zhenbin Li
Erik Nordmark
Mark Nottingham
Melinda Shore
Jeff Tantsura
Martin Thomson
Brian Trammell

Acknowledgements

Many thanks to John Klensin for his feedback and insights on the history of the Series, as someone who directly engaged and influenced many of the key individuals involved in developing the RFC Series.

Additional thanks to members of the RFC Series Advisory group and the Independent Submissions Editorial Board, in particular, Scott Bradner, Brian Carpenter, and Adrian Farrel, for their early reviews and input into the sequence of key moments in the history of the Series.

Contributors

Many thanks to Steve Crocker, Vint Cerf, Leslie Daigle, Nevil Brownlee, and Sandy Ginoza for their perspectives on the Series and their ongoing support.

Author's Address

Heather Flanagan (editor)
RFC Editor

Email: rse@rfc-editor.org
URI: <https://orcid.org/0000-0002-2647-2220>

