Network Working Group C. Huitema
Internet-Draft Private Octopus Inc.
Intended status: Informational March 4, 2020
Expires: September 5, 2020

Evaluation of a Sample of RFC Produced in 2018
draft-huitema-rfc-eval-project-04

Abstract

This document presents a first attempt at evaluating the recently published RFC. We analyze a set of randomly chosen RFC approved in 2018, looking for history and delays. We also use two randomly chosen sets of RFC published in 2008 and 1998 for comparing delays seen in 2018 to those observed 10 or 20 years ago.
The average RFC in the 2018 sample was produced in 3 years and 4 months, of which 2 years and 10 months were spent in the working group, 3 to 4 months for IETF consensus and IESG review, and 3 to 4 months in RFC production. The main variation in RFC production delays comes from the AUTH-48 phase.

We also measure the number of citations of the chosen RFC using Semantic Scholar, and compare citation counts with what we know about deployment. We show that citation counts indicate academic interest, but correlate only loosely with deployment or usage of the specifications.

The RFCs selected for this survey were chosen at random and represent a small sample of all RFCs produced, and only approximately 10% of the RFCs produced in each of 1998, 2008, and 2018. It is possible that different samples would produce different results. Furthermore, the conclusions drawn from the observations made in this document represent the author's opinions and do not have consensus of the IETF.

Status of This Memo

Table of Contents

1. Introduction

As stated on the organization's web site, "The IETF is a large open international community of network designers, operators, vendors, and researchers concerned with the evolution of the Internet architecture and the smooth operation of the Internet." The specifications produced by the IETF are published in the RFC series, along with independent submissions, research papers and IAB documents. In this memo, the author attempts to evaluate the RFC production process. This is an individual effort, and the author's conclusions presented here did not reach any kind of IETF consensus.

The IETF keeps records of documents and process actions in the IETF data tracker [TRKR]. The IETF data tracker provides information about RFC and drafts, from which we can infer statistics about the production system. We can measure how long it takes to drive a proposition from initial draft to final publication, and how these delays can be split between Working Group discussions, IETF reviews, IESG assessment, RFC Editor delays and final reviews by the authors - or, for independent stream RFCs, draft production, reviews by the Independent Stream Editor, conflict reviews, RFC Editor delays and final reviews. Tracker data is available for all RFCs, not just IETF stream RFCs.

Just measuring production delays may be misleading. If the IETF or the editors of the other series simply rubber-stamped draft proposals and published them, the delays would be short but the quality and impact might suffer. We hope that most of the RFC that are published are useful, but we need a way to measure that usefulness. We try to do that by measuring the number of references of the published RFCs in Semantic Scholar [SSCH], and also by asking the authors of each RFC in the sample whether the protocols and technologies defined in the RFCs were implemented and used on the Internet.

In order to limit the resource required for this study, we selected at random 20 RFC published in 2018, as explained in Section 2.2. The statistical sampling picked both IETF stream and Independent stream documents. For comparison purposes, we also selected at random 20 RFC published in 1998 and 20 published in 2008. Limiting the sample to 20 out of 209 RFCs published in 2018 allows for in depth analysis of each RFC, but readers should be reminded that the this is a small sample. The sample is too small to apply general statistical techniques and quantify specific ratios, and discussions of correlation techniques would be inappropriate. Instead, the purpose is to identify trends, spot issues and document future work.

The information gathered for every RFC in the sample is presented in Section 3. In Section 4 we analyze the production process and the sources of delays, comparing the 2018 sample to the selected samples for 1998 and 2018. In Section 5.1 we present citation counts for the RFC in the samples, and analyze whether citation counts could be used to evaluate the quality of RFC.

The measurement of delays could be automated by processing dates and events recorded in the data tracker. The measurement of published RFC could be complemented by statistics on abandoned drafts, which would measure the efficiency of the IETF triaging process. More instrumentation would help understanding how large delays happen during working group processes. These potential next steps are developed in Section 6.

2. Methodology

The study reported here started with a simple idea: take a sample of RFC, and perform an in-depth analysis of the path from the first presentation of the idea to its publication, while also trying to access the success of the resulting specification. This requires defining the key milestones that we want to track, and drawing a random sample using an unbiased process.

2.1. Defining the Important Milestones

The IETF data tracker records a list of events for each document processed by IETF working groups. This has a high granularity, and also a high variability. Most documents start life as an individual draft, are adopted by a working group, undergo a working group last call, are submitted to the IESG, undergo an IETF last call and an IESG review, get eventually approved by the IESG, and are processed for publication by the RFC Editor, but there are exceptions. Some documents are first submitted to one working group and then moved to another. Some documents are published through the Independent

Stream, and are submitted to the Independent Stream Editor instead of the IESG.

In order to simplify tabulation, we break the delay from between the submission of the first draft and the publication of the RFC in three big components:

- The working group delay, from the first draft to the start of the IETF last call;

- The IETF delay, which lasts from the beginning of the IETF last call to the approval by the IESG, including the reviews by various directorates;

- The RFC production, from approval by the IESG to publication, including the AUTH-48 reviews.

For submissions to the independent stream, we don't have a working group. We consider instead the progression of the individual draft until the adoption by the ISE as the equivalent of the "working group" period, and the delay from adoption by the ISE until submission to the RFC Editor as the equivalent of the IETF delay.

We measure the staring point of the process using the date of submission of the first draft listed on that RFC page in the IETF data tracker. In most case, this first draft is an individual draft that then resubmitted as a working group draft, or maybe resubmitted with a new name as the draft was searching for a home in an IETF working group, or before deciding for submission on the independent stream.

The IETF Data Tracker entries for RFC and drafts do not list working group events like Working Group Last Call. The only intermediate event that we list between the first draft and the submission to the IESG is the Working Group Adoption. For that, we use the date of submission of the version 00 of the draft eventually published as RFC. We use the same definition for drafts submitted to the Independent Stream.

2.2. Selecting a Random Sample of RFC

Basic production mechanisms could be evaluated by processing data from the IETF data tracker, but subjective data requires manual assessment of results, which can be time consuming. Since our resources are limited, we will only perform this analysis for a small sample of RFC, selected at random from the list of RFC approved in 2018. Specifically, we will pick 20 RFC at random between:

o RFC 8307, published in January 2018, and

o RFC 8511, published December 2018.

In order to avoid injecting personal bias in the random selection, we use a random selection process similar to the Nomination Committee selection process defined in [RFC3797]. The process is seeded with the text string "vanitas vanitatum et omnia vanitas", and the results are:

```
Picking 20 numbers between 8307 and 8511,
using MD5(vanitas vanitatum et omnia vanitas)
Rank 1: 8411 -- md5=daba041224a879199b698748808f917d
Rank 2: 8456 -- md5=f5570484d91ada6a672edbdca61d808c
Rank 3: 8446 -- md5=8340e918bb8faf69d197f79c9a58d7b8
Rank 4: 8355 -- md5=19474df74efd9917cf3fe8acce2ac374
Rank 5: 8441 -- md5=5acce2b730f3c24a4a91a5fc1921d1cd
Rank 6: 8324 -- md5=411c11a1cf4c292f83458865599c6921
Rank 7: 8377 -- md5=ac16a89192c0f0727febd35aacbc1f24
Rank 8: 8498 -- md5=bba44f2ba1ab240a1265a82ab71f7e02
Rank 9: 8479 -- md5=1653606b0af95d529a473a8f85ffaea4
Rank 10: 8453 -- md5=0cbfe105667c5a83b027dcfa85062f98
Rank 11: 8429 -- md5=fa51d7738562d990926a0d199fb060b8
Rank 12: 8312 -- md5=96d061523b1a57343356ae7a1e498ca5
Rank 13: 8492 -- md5=1b72b746eb05f79af40ed2bd3faccbe8
Rank 14: 8378 -- md5=645833b936d36cdcc797256518d7c483
Rank 15: 8361 -- md5=2064622c868e410beb0d9c18d0cb522c
Rank 16: 8472 -- md5=ca8a823072a21df011d0ea8b96a6aa47
Rank 17: 8471 -- md5=01b293a7dd0793e6f3297f2a973cd7e3
Rank 18: 8466 -- md5=8e411babe271557fe83bcdececc1643f
Rank 19: 8362 -- md5=8a1ba3efd82856a12b2b35fc5237e1b7
Rank 20: 8468 -- md5=57ae50ee0e1e0708d356d96d116dbfe1
```

When evaluating delays and impact, we will compare the year 2018 to 2008 and 1998, 10 and 20 years ago. To drive this comparison, we pick 20 RFC at random among those published in 2008, and another 20 among those published in 1998. We use the same nomcom-like methodology.

For 2008, we picking random RFC numbers between RFC 5134 (January 2008) and RFC 5405 (December 2008), using the sentence "sed fugit interea fugit irreparabile tempus" as a seed. We actually list here 21 numbers, because the random draw place RFC 5315 in 20th position, but that RFC was never issued. We replace it by RFC 5301, which came in 21st position.

```
Picking 21 numbers between 5134 and 5405,
using MD5(sed fugit interea fugit irreparabile tempus)
Rank 1: 5227 -- md5=61e5b3cd97fda4e75450e93df0744b6d
Rank 2: 5174 -- md5=eed49542394e9d392bcd756ee0f5beed
Rank 3: 5172 -- md5=72055ca953d6a8ee0d60a9fca0b8d738
Rank 4: 5354 -- md5=7a00ef15897479d1d15255159f6d8674
Rank 5: 5195 -- md5=54813be7bb56f48a05af8c894799b51f
Rank 6: 5236 -- md5=153263bbd8c0349501b75a33f3d66f6c
Rank 7: 5348 -- md5=b2d19aa9c1250ef2ddf169045f6e99d7
Rank 8: 5281 -- md5=1c3e61643d46d1da4ba16f0fd7a0aff5
Rank 9: 5186 -- md5=5e87001b830183b1a9427d479a7b0e42
Rank 10: 5326 -- md5=024347839f83d8082549c08bdfa1b43e
Rank 11: 5277 -- md5=049a83016ab08552841c59400480cd9d
Rank 12: 5373 -- md5=a1ce374aaebdacca2e7d6eeff039d970
Rank 13: 5404 -- md5=fb0d6b582a27ce34175e39de33598556
Rank 14: 5329 -- md5=df043ef1f9d42ba12a03a84434d26ead
Rank 15: 5283 -- md5=c40d3f966bc7800d6508d3d82df2371d
Rank 16: 5358 -- md5=6fea5bdb26b19e68befd409a09cb335d
Rank 17: 5142 -- md5=a8844b73287781762e6548fc6f533508
Rank 18: 5271 -- md5=c19eb02984265ecfe4ca076f2c160cfa
Rank 19: 5349 -- md5=33d756f81bf6e40ff344cf6ccaf29f13
Rank 20: 5315 -- md5=d4c30875f88328d72c9f78def2d1dde5
Rank 21: 5301 -- md5=3356419e5560901f0d31309b39d14a80
```

For 1998 we picking random RFC numbers between RFC 2257 (January 1998) and RFC 2479 (December 1998), using the sentence "pulvis et umbra sumus" as a seed.

```
Picking 20 numbers between 2257 and 2479,
using MD5(pulvis et umbra sumus)
Rank 1: 2431 -- md5=6eba444bb3349339fcde7e2be726f0f0
Rank 2: 2381 -- md5=0ce53f69f1a49a49309054af1e9a1d42
Rank 3: 2387 -- md5=53726e77ffeed903d244155d28e30f10
Rank 4: 2348 -- md5=69fcab2d2085555bac9eb0e0f4346523
Rank 5: 2391 -- md5=93358a26a7fce9fa9d61a31cdfdb87b7
Rank 6: 2267 -- md5=652dcab91bd9d5c58f98bdab21b23d80
Rank 7: 2312 -- md5=2505b0bed4af0a00ee663891c6f294d8
Rank 8: 2448 -- md5=6ede4b0dfa935f6ca656a5104b8dc5d0
Rank 9: 2374 -- md5=5312bfe5a9ca45563eb0bf35510d8daa
Rank 10: 2398 -- md5=7748a6deec3860898678f992b0b22792
Rank 11: 2283 -- md5=d73ff67466eb42d971a0e134b9284f83
Rank 12: 2382 -- md5=de1f667ac3e4c64aa529872e08823dff
Rank 13: 2289 -- md5=37773c2569dc25fdd0ab400ea401d5c7
Rank 14: 2282 -- md5=6b3df671a0a0becf9e42d203b59acd08
Rank 15: 2404 -- md5=e1b6819e5355924f456eb79f93beb8fd
Rank 16: 2449 -- md5=4f057df7c226efea773e7013c9c62081
Rank 17: 2317 -- md5=40eee3b536abe4afdb8834a0650c0a04
Rank 18: 2394 -- md5=044f09c53fc9fd1c50fe6bd0c39318e1
Rank 19: 2297 -- md5=78ee7e128436c969c80900fef80c075c
Rank 20: 2323 -- md5=ea6935bbda5f6d97756d3df5c3e2fdfb
```

3. Analysis of 20 Selected RFC

We review each of the RFC listed in Section 2.2 for the year 2018, trying both to answer the known questions and to gather insight for further analyzes. In many cases, the analysis of the data is complemented by direct feedback from the RFC authors.

3.1. 8411

IANA Registration for the Cryptographic Algorithm Object Identifier Range [RFC8411]:

Informational, 5 pages
4 drafts (personal), first 2017-05-08.
Last call announced 2017-10-09
IESG evaluation starts 2017-12-28
Approved 2018-02-26, draft 03
AUTH-48 2018-04-20
AUTH-48 complete 2018-07-17
Published 2018-08-06
IANA action: create table

This RFC was published from the individual draft, which was not resubmitted as a working group draft.

The draft underwent minor copy edit before publication.

Some but not all of the long delay in AUTH-48 is due to clustering with [RFC8410]. MISSREF was cleared on 2018-05-09 and the document re-entered AUTH-48 at once. AUTH-48 lasted over two months after that.

The time after AUTH-48 and before publication (3 weeks) partly overlaps with travel for IETF-102 and is partly due to coordinating the cluster.

3.2. 8456

Benchmarking Methodology for Software-Defined Networking (SDN) Controller Performance [RFC8456]:

Informational, 64 pages
2 personal drafts, 9 WG drafts, first 2015-03-23
WG adoption on 2015-10-18
Last call announced 2018-01-19
IESG evaluation starts 2018-02-27
IESG approved 2018-05-25
AUTH-48 2018-08-31
AUTH-48 complete 2018-10-16
Published 2018-10-30

The draft underwent very extensive copy editing, covering use of articles, turn of phrases, choice of vocabulary. The changes are enough to cause pagination differences. The "diff" tool marks pretty much every page as changed. Some diagrams see change in protocol elements like message names.

According to the author, the experience of producing this draft mirrors a typical one in the Benchmarking Methodologies Working Group (BMWG). There were multiple authors in multiple time zones, which slowed down the AUTH-48 process somewhat, although the AUTH-48 delay of 46 days is only a bit longer than the average draft.

The RFC was part of cluster with [RFC8455].

BMWG publishes informational RFCs centered around benchmarking, and the methodologies in RFC 8456 have been implemented in benchmarking products.

3.3. 8446

The Transport Layer Security (TLS) Protocol Version 1.3 [RFC8446], as the title indicates, defines the new version of the TLS protocol. From the IETF data tracker, we extract the following:

Proposed standard
160 pages
29 WG drafts first 2014-04-17.
Last call announced 2018-02-15
IESG evaluation starts 2018-03-02
Approved 2018-03-21, draft 28
AUTH-48 2018-06-14
AUTH-48 complete 2018-08-10
Published 2018-08-10

This draft started as a WG effort.

The RFC was a major effort in the IETF. Working group members developed and tested several implementations. Researchers analyzed the specifications and performed formal verifications. Deployment tests outlined issues that caused extra work when the specification was almost ready. These complexity largely explains the time spent in the working group.

Comparing the final draft to the published version, we find relatively light copy editing. It includes explaining acronyms on first use, clarifying some definitions standardizing punctiation and capitalization, and spelling out some numbers in text. This generally fall in the category of "style", although some of the clarifications go into message definitions. However, that simple analysis does not explain why the AUTH-48 phase took almost two months.

This document's AUTH-48 process was part of the "Github experiment", which tried to use github pull requests to track the AUTH-48 changes and review comments. The RPC staff had to learn using Github for that process, and this required more work than the usual RFC. Author and AD thoroughly reviewed each proposed edit, accepting some and rejecting some. The concern there was that any change in a complex specification might affect a protocol that was extensively reviewed in the working group, but of course these reviews added time to the AUTH-48 delays.

There are 21 implementations listed in the Wiki of the TLS 1.3 project [TLS13IMP]. It has been deployed on major browsers, and is already used in a large fraction of TLS connections.

3.4. 8355

Resiliency Use Cases in Source Packet Routing in Networking (SPRING) Networks [RFC8355] is an informational RFC. It originated from a use case informational draft that was mostly used for the BOF creating the WG, and then to drive initial work/evolutions from the WG.

Informational, 13 pages.
2 personal drafts (personal), first 2014-01-31. 13 WG drafts.
WG adoption on 2014-05-13
Last call announced 2017-04-20
IESG evaluation starts 2017-05-04, draft 09
Approved 2017-12-19, draft 12
AUTH-48 2018-03-12
AUTH-48 complete 2018-03-27
Published 2018-03-28

Minor set of copy edits, mostly for style.

No implementation of the RFC itself, but the technology behind it such as Segment Routing (architecture RFC 8402, TI-LFA draft-ietf-rtgwg-segment-routing-ti-lfa) is widely implemented and deployment is ongoing.

According to participants in the discussion, the process of adoption of the source packet routing standards was very contentious. The establishment of consensus at both the working group level and the IETF level was difficult and time consuming.

3.5. 8441

Bootstrapping WebSockets with HTTP/2 [RFC8441]

Proposed standard, 8 pages. Updates RFC 6455.
3 personal drafts (personal), first 2017-10-15. 8 WG drafts.
WG adoption on 2017-12-19
Last call announced 2018-05-07, draft 05
IESG evaluation starts 2018-05-29, draft 06
Approved 2018-06-07, draft 07
AUTH-48 2018-08-13
AUTH-48 complete 2018-09-15
Published 2018-09-21
IANA Action: table entries

This RFC defines the support of WebSockets in HTTP/2, which is different from the mechanism defined for HTTP/1.1 in [RFC6455]. The process was relatively straightforward, involving the usual type of discussions, some on details and some on important points.

Comparing final draft and published RFC shows a minor set of copy edit, mostly for style. However, the author recalls a painful process. The RFC includes many charts and graphs that were very difficult to format correctly in the author's production process that involve conversions from markdown to XML, and then from XML to text. The author had to get substantial help from the RFC editor.

There are several implementations, including Firefox and Chrome, making RFC 8441 a very successful specification.

3.6. 8324

DNS Privacy, Authorization, Special Uses, Encoding, Characters, Matching, and Root Structure: Time for Another Look? [RFC8324]. This is an opinion piece on DNS development, published on the Independent Stream.

Informational, 29 pages. Independent stream.
5 personal drafts (personal), first 2017-06-02.
ISE review started 2017-07-10, draft 03
IETF conflict review and IESG review started 2017-10-29
Approved 2017-12-18, draft 04
AUTH-48 2018-01-29, draft 05
AUTH-48 complete 2018-02-26
Published 2018-02-27

This RFC took only 9 months from first draft to publication, which is the shortest in the 2018 sample set. In part, this is because the text was privately circulated and reviewed by ISE designated experts before the first draft was published. The nature of the document is another reason for the short delay. It is an opinion piece, and does not require the same type of consensus building and reviews than a protocol specification.

Comparing the final draft and the published version shows only minor copy edit, mostly for style. According to the author, because this is because he knows how to write in RFC Style with the result that his documents often need a minimum of editing. He also makes sure that the document on which the Production Center starts working already has changes discussed and approved during Last Call and IESG review incorporated rather than expecting the Production Center to operate off of notes about changed to be made.

3.7. 8377

Transparent Interconnection of Lots of Links (TRILL): Multi-Topology [RFC8377]

Proposed standard, 20 pages. Updates RFC 6325, 7177.
3 personal drafts (personal), first 2013-09-03. 7 WG drafts.
WG adoption on 2015-09-01
Last call announced 2018-02-19, draft 05
IESG evaluation starts 2018-03-02, draft 06
Approved 2018-03-12, draft 05
AUTH-48 2018-04-20, draft 06
AUTH-48 complete 2018-07-31
Published 2018-07-31
IANA Table, table entries

Minor set of copy edits, mostly for style, also clarity.

3.8. 8498

A P-Served-User Header Field Parameter for an Originating Call Diversion (CDIV) Session Case in the Session Initiation Protocol (SIP) [RFC8498].

Informational, 15 pages.
5 personal drafts (personal), first 2016-03-21. 9 WG drafts.
WG adoption on 2017-05-15
Last call announced 2018-10-12, draft 05
IESG evaluation starts 2018-11-28, draft 07
Approved 2018-12-10, draft 08
AUTH-48 2019-01-28
AUTH-48 complete 2019-02-13
Published 2019-02-15
IANA Action, table rows added.

Copy edit for style, but also clarification of ambiguous sentences.

3.9. 8479

Storing Validation Parameters in PKCS#8 [RFC8479]

Informational, 8 pages. Independent stream.
5 personal drafts (personal), first 2017-08-08.
ISE review started 2018-12-10, draft 00
IETF conflict review and IESG review started 2018-03-29
Approved 2018-08-20, draft 03
AUTH-48 2018-09-20, draft 04
AUTH-48 complete 2018-09-25
Published 2018-09-26

The goal of the draft was to document what the gnutls implementation was using for storing provably generated RSA keys. This is a short RFC that was published relatively quickly, although discussion

between the author, the Independent Series Editor and the IESG lasted several months. In the initial conflict review, The IESG asked the ISE to not publish this document before IETF working groups had an opportunity to pick up the work. The author met that requirement by a presentation to the SECDISPATCH WG in IETF 102. Since no WG was interested in pickup the work, the document progressed on the Independent Stream.

Very minor set of copy edit, moving some references from normative to informative.

The author is not aware of other implementations than gnutls relying on this RFC.

3.10. 8453

Framework for Abstraction and Control of TE Networks (ACTN) [RFC8453]

Informational, 42 pages.
3 personal drafts, first 2015-06-15. 16 WG drafts.
WG adoption on 2016-07-15
Out of WG 2018-01-26, draft 11
Expert review requested, 2018-02-13
Last call announced 2018-04-16, draft 13
IESG evaluation starts 2018-05-16, draft 14
Approved 2018-06-01, draft 15
AUTH-48 2018-08-13
AUTH-48 complete 2018-08-20
Published 2018-08-20
IANA Action, table rows added.

Minor copy editing.

3.11. 8429

Deprecate Triple-DES (3DES) and RC4 in Kerberos [RFC8429]

BCP, 10 pages.
6 WG drafts, first 2017-05-01.
Last call announced 2017-07-16, draft 03
IESG evaluation starts 2017-08-18, draft 04
Approved 2018-05-25, draft 05
AUTH-48 2018-07-24
AUTH-48 complete 2018-10-31
Published 2018-10-31
IANA Action, table rows added.

This draft started as a working group effort.

This RFC recommends to deprecate two encryption algorithms that are now considered obsolete and possibly broken. The document was sent back to the WG after the first last call, edited, and then there was a second last call. The delay from first draft to working group last call was relatively short, but the number may be misleading. The initial draft was a replacement of a similar draft in the KITTEN working group, which stagnated for some time before the CURDLE working group took up the work. The deprecation of RC4 was somewhat contentious, but the WG had already debated this prior to the production of this draft, and the draft was not delayed by this debate.

Most of the 280 days between IETF LC and IESG approval was because the IESG had to talk about whether this document should obsolete or move to historic RFC 4757, and no one was really actively pushing that discussion for a while.

The 99 days in AUTH-48 are mostly because one of the authors was a sitting AD, and those duties ended up taking precedence over reviewing this document.

Minor copy editing, for style.

The implementation of the draft would be the actual removal of support for 3DES and RC4 in major implementations. This is happening, but very slowly.

3.12. 8312

CUBIC for Fast Long-Distance Networks [RFC8312]

Informational, 18 pages.
2 personal drafts, first 2014-09-01. 8 WG drafts
WG adoption on 2015-06-08
Last call announced 2017-09-18, draft 06
IESG evaluation starts 2017-11-14
Approved 2017-10-04, draft 07
AUTH-48 2018-01-08
AUTH-48 complete 2018-02-07
Published 2018-02-07
IANA Action, table rows added.

Minor copy editing, for style.

The TCP congestion control algorithm Cubic was defined first in 2005, was implemented in Linux soon after, and was implemented in major OSes after that. After some debates from 2015 to 2015, the TCPM working group adopted the draft, with a goal of documenting Cubic in

the RFc series. According to the authors, this was not a high priority effort, as Cubic was already implemented in multiple OSes and documented in research papers. At some point, only one of the authors was actively working on the draft. Ths may explain why another two years was spent progressing the draft after adoption by the WG.

The RFC publication may or may not have triggered further implementations. On the other hand, several OSes picked up bug fixes from the draft and the RFC.

3.13. 8492

Secure Password Ciphersuites for Transport Layer Security (TLS) [RFC8492]

Informational, 40 pages. (Independent Stream)
10 personal drafts, first 2012-09-07. 8 WG drafts
Targeted to ISE stream 2016-08-05
ISE review started 2017-05-10, draft 01
IETF conflict review and IESG review started 2017-09-04
Approved 2017-10-29, draft 04
AUTH-48 2018-10-19, draft 05
AUTH-48 complete 2019-02-19
Published 2019-02-21
IANA Action, table rows added.

This RFC has a complex history. The first individual draft was submitted to the TLS working group on September 7, 2012. It progressed there, and was adopted by the WG after 3 revisions. There were then 8 revisions in the TLS WG, until the WG decided to not progress it. The draft was parked in 2013 by the WG chairs after failing to get consensus in WG last call. The AD finally pulled the plug in 2016, and the draft was then resubmitted to the ISE.

At that point, the author was busy and was treating this RFC with a low priority because, in his words, it would not be a "real RFC". There were problems with the draft that only came up late. In particular, it had to wait for a change in registry policy that only came about with the publication of TLS 1.3, which caused the draft to only be published after RFC 8446, and also required adding references to TLS 1.3. The author also got a very late comment while in AUTH-48 that caused some rewrite. Finally, there was some IANA issue with the extension registry where a similar extension was added by someone else. The draft was changed to just use it.

Changes in AUTH-48 include added reference to TLS 1.3, copy-editing for style, some added requirements, added paragraphs, and changes in algorithms specification.

3.14. 8378

Signal-Free Locator/ID Separation Protocol (LISP) Multicast [RFC8378] is an experimental RFC, defining how to implement Multicast in the LISP architecture.

Experimental, 21 pages.
5 personal drafts, first 2014-02-28. 10 WG drafts
WG adoption on 2015-12-21
Last call announced 2018-02-13, draft 07
IESG evaluation starts 2018-02-28, draft 08
Approved 2018-03-12, draft 09
AUTH-48 2018-04-23
AUTH-48 complete 2018-05-02
Published 2018-05-02

Preparing the RFC took more than 4 years. According to the authors, they were not aggressive pushing it and just let the working group process decide to pace it. They also did implementations during that time.

Minor copy editing, for style.

The RFC was implemented by lispers.net and cisco, and was used in doing IPv6 multicast over IPv4 unicast/multicast at the Olympics in PyeungChang. The plan is to work on a proposedstandard once the experiment concludes.

3.15. 8361

Transparent Interconnection of Lots of Links (TRILL): Centralized Replication for Active-Active Broadcast, Unknown Unicast, and Multicast (BUM) Traffic [RFC8361]

Proposed Standard, 17 pages.
3 personal drafts, first 2013-11-12. 14 WG drafts
WG adoption on 2014-12-16
Last call announced 2017-11-28, draft 10
IESG evaluation starts 2017-12-18, draft 11
Approved 2018-01-29, draft 13
AUTH-48 2018-03-09
AUTH-48 complete 2018-04-09
Published 2018-04-12

According to the authors, the long delays in producing this RFC was due to a slow uptake of the technology in the industry.

Minor copy editing, for style.

There was at least 1 partial implementation.

3.16. 8472

Transport Layer Security (TLS) Extension for Token Binding Protocol Negotiation [RFC8472]

Proposed Standard, 8 pages.
1 personal drafts, 2015-05-29. 15 WG drafts
WG adoption on 2015-09-11
Last call announced 2017-11-13, draft 10
IESG evaluation starts 2018-03-19
Approved 2018-07-20, draft 14
AUTH-48 2018-09-17
AUTH-48 complete 2018-09-25
Published 2018-10-08

This is a pretty simple document, but it took over 3 years from individual draft to RFC. According to the authors,the biggest setbacks occurred at the start: it took a while to find a home for this draft. It was presented in the TLS WG (because it's a TLS extension) and UTA WG (because it has to do with applications using TLS). Then the ADs determined that a new WG was needed, so the authors had to work through the WG creation process, including running a BOF.

Minor copy editing, for style, with the addition of a reference to TLS 1.3.

Perhaps partially due to the delays, some of the implementers lost interest in supporting this RFC.

3.17. 8466

A YANG Data Model for Layer 2 Virtual Private Network (L2VPN) Service Delivery [RFC8466]

Proposed Standard, 158 pages.
5 personal drafts, first 2016-09-01. 11 WG drafts
WG adoption on 2017-02-26
Last call announced 2018-02-21, draft 07
IESG evaluation starts 2018-03-14, draft 08
Approved 2018-06-25, draft 10
AUTH-48 2018-09-17
AUTH-48 complete 2018-10-09
Published 2018-10-12

Copy editing for style and clarity, with also corrections to the yang model.

3.18. 8362

OSPFv3 Link State Advertisement (LSA) Extensibility [RFC8362] is a major extension to the OSPF protocol. It makes OSPFv3 fully extensible.

Proposed Standard, 33 pages.
4 personal drafts, first 2013-02-17. 24 WG drafts
WG adoption on 2013-10-15
Last call announced 2017-12-19, draft 19
IESG evaluation starts 2018-01-18, draft 20
Approved 2018-01-29, draft 23
AUTH-48 2018-03-19
AUTH-48 complete 2018-03-30
Published 2018-04-03

The specification was first submitted as a personal draft in the IPv6 WG, then moved to the OSPF WG. The long delay of producing this RFC is due to the complexity of the problem, and the need to wait for implementations. It is a very important change to OSPF that makes OSPFv3 fully extensible. Since it was a non-backward compatible change, the developers started out with some very complex migration scenarios but ended up with either legacy or extended OSPFv3 LSAs within an OSPFv3 routing domain. The initial attempts to have a hybrid mode of operation with both legacy and extended LSAs also delayed implementation due to the complexity.

Copy editing for style and clarity.

This specification either was or will be implemented by all the router vendors.

3.19. 8468

IPv4, IPv6, and IPv4-IPv6 Coexistence: Updates for the IP Performance Metrics (IPPM) Framework [rfc8468].

Informational, 15 pages.
3 personal drafts, first 2015-08-06. 7 WG drafts
WG adoption on 2016-07-04
Last call announced 2018-04-11, draft 04
IESG evaluation starts 2018-05-24, draft 05
Approved 2018-07-10, draft 06
AUTH-48 2018-09-13
AUTH-48 complete 2018-11-05
Published 2018-11-14

RFC8468 was somehow special in that there was not a technical reason/interest that triggered it, but rather a formal requirement. While writing RFC7312 the IP Performance Metrics working group (IPPM) realized that RFC 2330, the IP Performance Metrics Framework supported IPv4 only and explicitly excluded support for IPv6. Nevertheless, people used the metrics that were defined on top of RFC 2330 (and, therefore, IPv4 only) for IPv6, too. Although the IPPM WG agreed that the work was needed, the interest of IPPM attendees in progressing (and reading/reviewing) the IPv6 draft was limited. Resolving the IPv6 technical part was straight-forward, but subsequently some people asked for a broader scope (topics like header compression, 6lo, etc.) and it took some time to figure out and later on convince people that these topics are out of scope. The group also had to resolve contentious topics, for example how to measure the processing of IPv6 extension headers, which is sometimes non-standard.

The AUTH-48 delay for this draft was longer than average. According to the authors, the main reasons include:

- Work-load and travel caused by busy-work-periods of all co-authors
- Time zone difference between co-authors and editor (at least US, Europe, India, not considering travel)
- Editor proposing and committing some unacceptable modifications that needed to be reverted
- Lengthy discussions on a new document title (required high effort and took a long time, in particular reaching consensus between co-authors and editor was time-consuming and involved the AD)

- Editor correctly identifying some nits (obsoleted personal websites of co-authors) and co-authors attempting to fix them.

The differences between the final draft and the publish RFC show copy editing for style and clarity, but do not account for the back and forth between authors and editors mentioned by the authors.

4. Analysis of Process and Delays

We examine the 20 RFC in the sample, measuring various characteristics such as delay and citation counts, in an attempt to identify patterns in the IETF processes.

4.1. First Draft to RFC Delays

We look at the distribution of delays between the submission of the first draft and the publication of the RFC, using the three big milestones defined in Section 2.1: delay in the Working Group, IETF delay, and publication delay. The following table shows the delays for the 20 RFC in the sample:

RFC	Status	Pages	Overall	WG	IETF	Edit
8411	Info	5	455	154	140	161
8456	Info	64	1317	1033	126	158
8446	PS	160	1576	1400	34	142
8355	Info	13	1517	1175	243	99
8441	PS	8	341	204	31	106
8324	ISE	29	270	38	161	71
8377	PS	8	1792	1630	21	141
8498	Info	15	1061	935	59	67
8479	ISE	8	414	233	144	37
8453	Info	42	1162	1036	46	80
8429	BCP	10	548	76	313	159
8312	Info	18	1255	1113	16	126
8492	ISE	40	2358	1706	172	480
8378	Exp	21	1524	1446	27	51
8361	PS	17	1612	1477	62	73
8472	PS	8	1228	899	249	80
8466	PS	158	771	538	124	109
8362	PS	33	1871	1766	41	64
8468	Info	15	1196	979	90	127
	average	35	1172	939	110	123
	average(not ISE)	35	1202	991	101	109

The average delay from first draft to publication is about 3 years and 2 months, but this varies widely. Excluding the independent

stream submissions, the average delay from start to finish is 3 years and 3 months, of which on average 2 years and 9 months are spent getting consensus in the working group, and 3 to 4 months each for IETF consensus and for RFC production.

The longest delay is found for [RFC8492], 6.5 years from start to finish. This is however a very special case, a draft that was prepared for the TLS working group and failed to reach consensus. After that, it was resubmitted to the ISE, and incurred atypical production delays.

On average, we see that 80% of the delay is incurred in WG processing, 10% in IETF review, and 10% for edition and publication.

For IETF stream RFC, it appears that the delays for informational documents are slightly shorter than those for protocol specifications, maybe six months shorter on average. However, there are lots of differences between individual documents. The delays range from less than a year to more than 5 years for protocol specifications, and from a year and 3 months to a bit more than 4 years for informational documents.

We can compare the delays in the 2018 samples to those observed 10 years ago and 20 years before:

RFC (2008)	Status	Pages	Delay
5326	Exp	54	1584
5348	PS	58	823
5281	Info	51	1308
5354	Exp	23	2315
5227	PS	21	2434
5329	PS	12	1980
5277	PS	35	912
5236	ISE	26	1947
5358	BCP	7	884
5271	Info	22	1066
5195	PS	10	974
5283	PS	12	1096
5186	Info	6	2253
5142	PS	13	1005
5373	PS	24	1249
5404	PS	27	214
5172	PS	7	305
5349	Info	10	1096
5301	PS	6	396
5174	Info	8	427

RFC (1998)	Status	Pages	Delay
2289	PS	25	396
2267	Info	10	unknown
2317	BCP	10	485
2404	PS	7	488
2374	PS	12	289
2449	PS	19	273
2283	PS	9	153
2394	Info	6	365
2348	DS	5	699
2382	Info	30	396
2297	ISE	109	28
2381	PS	43	699
2312	Info	20	365
2387	PS	10	122
2398	Info	15	396
2391	PS	10	122
2431	PS	10	457
2282	Info	14	215
2323	ISE	5	unknown
2448	ISE	7	92

We can compare the median delay, and the delays observed by the fastest and slowest quartiles in the three years:

Year	Fastest 25%	Median	Slowest 25%
2018	604	1179	1522
2008	869	1081	1675
1998	169	365	442

The IETF takes three to four times more times to produce an RFC in 2018 than it did in 1998, but about the same time as it did in 2008. We can get a rough estimate of how this translates in term of "level of attention" per RFC by comparing the number of participants in the IETF meetings of 2018, 2008 and 1998 [IETFCOUNT] to the number of RFC published these years [RFCYEAR].

Year	Nb RFC	Spring P.	Summer P.	Fall	Average P.	Attendees/RFC
2018	208	1235	1078	879	1064	5.1
2008	290	1128	1181	962	1090	3.8
1998	234	1775	2106	1705	1862	9.0

The last column in the table provides the ratio of average number of participants by number of RFC produced. If the IETF was a centralized organization, if all participants and documents were equivalent, this ratio would be the number of participants dedicated to produce an RFC on a given year. This is of course a completely abstract figure because none of the hypotheses above is true, but it still gives a vague indication of the "level of attention" applied to documents. We see that this ratio has increased from 2008 to 2018, as the number of participants was about the same for these two years but the number of published RFC decreased. However, that ratio was much higher in 1998. The IETF had many more participants, and there were probably many more eyes available to review any given draft. If we applied the ratios of 1998, the IETF would be producing 119 documents in 2018 instead of 208.

4.2. Working Group Delays

The largest part of the delays is spent in the working groups, before the draft is submitted to the IESG for IETF review. As mentioned in Section 2.1, the only intermediate milestone that we can extract from

the IETF data tracker is the date at which the document was adopted by the working group, or targeted for independent submission. The breakdown of the delays for the documents in our sample is:

RFC	Status	WG	Until adoption	After adoption
8411	Info	154	0	154
8456	Info	1033	209	824
8446	PS	1400	0	1400
8355	Info	1175	102	1073
8441	PS	204	65	139
8324	ISE	38	0	38
8377	PS	1630	728	902
8498	Info	935	420	515
8479	ISE	233	0	233
8453	Info	1036	396	640
8429	BCP	76	0	76
8312	Info	1113	280	833
8492	ISE	1706	1428	278
8378	Exp	1446	661	785
8361	PS	1477	399	1078
8472	PS	899	105	794
8466	PS	538	178	360
8362	PS	1766	240	1526
8468	Info	979	333	646
	Average	939	292	647

The delay before working group adoption average to a bit less than 10 months, compared to 1 years and a bit more than 9 months for delays after adoption. We see that RFC 8492 stands out, with long delays spent attempting publication through a working group before submission to the independent editor. There are a few documents that started immediately as working group efforts, or were immediately targeted for publication in the independent series. Those documents tend to see short delays, with the exception of RFC 8446 on which the TLS working groupspent a long time working.

4.3. Preparation and Publication Delays

The preparation and publication delays include three components:

- the delay from submission to the RFC Editor to beginning of AUTH-48, during which the document is prepared;
- the AUTH-48 delay, during which authors review and eventually approve the changes proposed by the editors;
- the publication delay, from final agreement by authors and editors to actual publication.

The breakdown of the publication delays for each RFC is shown in the following table.

RFC	Status	Pages	RFC edit	AUTH-48	RFC Pub	Edit(total)
8411	Info	5	53	88	20	161
8456	Info	64	98	46	14	158
8446	PS	160	85	57	0	142
8355	Info	13	83	15	1	99
8441	PS	8	67	33	6	106
8324	ISE	29	42	28	1	71
8377	PS	8	39	102	0	141
8498	Info	15	49	16	2	67
8479	ISE	8	31	5	1	37
8453	Info	42	73	7	0	80
8429	BCP	10	60	99	0	159
8312	Info	18	96	30	0	126
8492	ISE	40	355	123	2	480
8378	Exp	21	42	9	0	51
8361	PS	17	39	31	3	73
8472	PS	8	59	8	13	80
8466	PS	158	84	22	3	109
8362	PS	33	49	11	4	64
8468	Info	15	65	53	9	127
	Average		77.3	41.2	4.2	122.7
-8492	Average		62	37	4	103

On average, the total delay appears to be a bit more than four month, but the average is skewed by the extreme values encountered for [RFC8492]. If we exclude that RFC from the computations, the average delay drops to a just a bit more than 3 months: about 2 months for the preparation, a bit more than one month for the AUTH-48 phase, and 4 days for the publishing.

Of course, these delays vary from RFC to RFC. To try explain the causes of the delay, we compute the correlation factor between the observed delays and 4 plausible explanation factors:

- o The number of pages in the document,
- o The amount of copy edit, as discussed in Section 4.4,
- o Whether or not an IANA action was required.

We find the following values:

Correlation	RFC edit	AUTH-48	Edit(total)
Nb pages	0.50	-0.04	0.21
Copy-Edit	0.42	0.24	0.45
IANA	-0.13	0.26	0.15

None of these indicate strong correlations. The greater number of pages will tend to increase the preparation delay, but it does not appear to impact the AUTH-48 delay at all. The amount of copy editing also tend to increase the the preparation delay somewhat and the AUTH-48 delay a little. The presence or absence of IANA action has very ittle correlation with the delays.

The analysis of RFC 8324 in Section 3.6 explains its short editing delays by the experience of the author. This makes sense: if a document needs less editing, the editing delays would be shorter. This is partially confirmed by the relation between the amount of copy editing and the publication delay.

We also observe that the AUTH-48 delay varies much more than the preparation delay, with a standard deviation of 20 days for AUTH-48 versus 10 days for the preparation delay. In theory, AUTH-48 is just a final verification: the authors receive the document prepared by the RFC production center, and just have to give their approval, or maybe request a last minute correction. The name indicates that this

is expected to last just two days, but in average it lasts more than a month.

We tested a variety of hypotheses that might explain the duration of AUTH-48 by computing the correlation coefficients between various properties of the RFC and the production delays. The results are listed in the following table:

Correlation	RFC production	AUTH-48	RFC Edit(total)
Nb drafts	0.19	-0.30	-0.17
Nb Authors	0.40	-0.04	0.16
WG delay	0.03	-0.16	-0.15

The results show that there is no simple answer. The number of pages, the required amount of copy editing and to a very small extent the number of drafts can help predict the production delay, but there is no obvious predictor for the AUTH-48 delay. In particular, there is no numerical evidence that the number of authors influences the AUTH-48 delay, or that authors who have spent a long time working on the document in the working group somehow spend even longer to answer questions during AUTH-48 - if anything, the numerical results point in the opposite direction.

After asking the authors of the RFC in the sample why the AUTH-48 phase took a long time, and we got three explanations:

1- Some RFC have multiple authors in multiple time zones. This slows down the coordination required for approving changes.

2- Some authors found some of the proposed changes unnecessary or undesirable, and asked that they be reversed. This required long exchanges between authors and editors.

3- Some authors were not giving high priority to AUTH-48 responses.

As mentioned above, we were not able to verify these hypotheses by looking at the data.

4.4. Copy Editing

We can assess the amount of copy editing applied to each published RFC by comparing the text of the draft approved for publication and the text of the RFC. We do expect differences in the "boilerplate"

and in the IANA section, but we will also see differences due to copy editing. Assessing the amount of copy editing is subjective, and we do it using a scale of 1 to 4:

1- Minor editing

2- Editing for style, such as capitalization, hyphens, that versus which, and expending all acronyms at least one.

3- Editing for clarity in addition to style, such as rewriting ambiguous sentences and clarifying use of internal references. For Yang models, that may include model corrections suggested by the verifier.

4- Extensive editing.

The following table shows that with about half of the RFC required editing for style, and the other half at least some editing for clarity.

RFC	Status	Copy Edit
8411	Info	2
8456	Info	4
8446	PS	3
8355	Info	2
8441	PS	2
8324	ISE	2
8377	PS	3
8498	Info	3
8479	ISE	1
8453	Info	2
8429	BCP	2
8312	Info	2
8492	ISE	3
8378	Exp	2
8361	PS	2
8472	PS	2
8466	PS	3
8362	PS	3
8468	Info	3

This method of assessment does not take into account the number of changes proposed by the editors and eventually rejected by the authors, since these changes are not present in either the final draft or the published RFC. It might be possible to get an evaluation of these "phantom changes" from the RFC Production Center.

4.5. Independent Series

Out of 20 randomly selected RFC, 3 were published through the "independent series". One is an independent opinion, another a description of a non-IETF protocol format, and the third was [RFC8492], which is a special case. Apart from this special case, the publication delays were significantly shorter for the Independent Stream than for the IETF stream. This seems to indicate that the Independent Series is functioning as expected.

The authors of these 3 RFC are regular IETF contributors. This observation motivated a secondary analysis of all the RFC published in the "independent" stream in 2018. There are 14 such RFC: 8507, 8494, 8493, 8492, 8483, 8479, 8433, 8409, 8374, 8369, 8367, 8351, 8328 and 8324. (RFC 8367 and 8369 were published on 1 April 2018.) The majority of the documents were published by regular IETF participants, but two of them were not. One describes "The BagIt File Packaging Format (V1.0)" [RFC8493], and the other the "Yeti DNS Testbed" [RFC8483]. They document a data format and a system developed outside the IETF, and illustrate the outreach function of the Independent Stream. In both cases, the authors include one experienced IETF participant, who presumably helped outsiders navigate the publication process.

5. Citation Counts

In this exploration, we want to assess whether citation counts provide a meaningful assessment of the popularity of RFC. We obtain the citation counts through the Semantic Scholar API, using queries of the form:

http://api.semanticscholar.org/
    v1/paper/10.17487/rfc8446?include_unknown_references=true

In these queries, the RFC is uniquely identified by its DOI reference, which is composed of the RFC Series prefix 10.17487 and the rfc identifier. The queries return a series of properties, including a list of citations for the RFC. Based on that list of citations, we compute three numbers:

- The total number of citations
- The number of citations in the year of publication and the year after that
- For the RFC published in 1998 or 2008 that we use for comparison, the number of citations in the years 2018 and 2019.

All the numbers were retrieved on October 6, 2019.

5.1. Citation Numbers

As measured on October 6, 2019, the citation counts for the RFC in our sample set were:

RFC(2018)	Status	Total	2018-2019
8411	Info	1	0
8456	Info	1	1
8446	PS	418	204
8355	Info	3	3
8441	PS	1	1
8324	ISE	0	0
8377	PS	0	0
8498	Info	0	0
8479	ISE	0	0
8453	Info	3	3
8429	BCP	0	0
8312	Info	25	16
8492	ISE	4	4
8378	Exp	1	1
8361	PS	0	0
8472	PS	1	1
8466	PS	0	0
8362	PS	1	1
8468	Info	1	1

The results indicate that [RFC8446] is by far the most cited of the 20 RFC in our sample. This is not surprising, since TLS is a key Internet Protocol. The TLS 1.3 protocol was also the subject of extensive studies by researchers, and thus was mentioned in a number of published papers. Surprisingly, the Semantic Scholar mentions a number of citations that predate the publication date. These are probably citations of the various draft versions of the protocol.

The next most cited RFC in the sample is [RFC8312] which describes the Cubic congestion control algorithm for TCP. That protocol was also the target of a large number of academic publications.The other RFC in the sample only have a small number of citations.

There is probably a small bias when measuring citations at a fixed date. An RFC published in January 2018 would have more time to accrue citations than one published in December. That may be true to some extent, as the second most cited RFC in the set was published in January. However, the effect has to be limited. The most cited RFC was published in August, and the second most cited was published in 2019. (That RFC got an RFC number in 2018, but publication was slowed by long AUTH-48 delays.)

5.2. Comparison to 1998 and 2008

In order to get a baseline, we can look at the number of references for the RFC published in 2008 and 1998. However, we need totake time into account. Documents published a long time ago are expected to have accrued more references. We try to address this by looking at three counts for each document: the overall number of references over the document's lifetime, the number of references obtained in the year following publication, and the number of references observed since 2018:

RFC(2008)	Status	Total	2008-2009	2018-2019
5326	Exp	138	14	15
5348	PS	14	3	0
5281	Info	69	15	7
5354	Exp	17	13	0
5227	PS	19	1	2
5329	PS	24	6	1
5277	PS	32	3	2
5236	ISE	25	5	4
5358	BCP	21	2	0
5271	Info	7	2	0
5195	PS	7	4	2
5283	PS	8	1	0
5186	Info	14	4	2
5142	PS	8	4	0
5373	PS	5	2	0
5404	PS	1	1	0
5172	PS	2	0	0
5349	Info	8	0	2
5301	PS	5	1	0
5174	Info	0	0	0

RFC(1998)	Status	Total	1998-1999	2018-2019
2289	PS	2	0	1
2267	Info	982	5	61
2317	BCP	9	1	2
2404	PS	137	6	1
2374	PS	42	4	0
2449	PS	7	2	0
2283	PS	17	3	2
2394	Info	13	2	1
2348	DS	5	0	0
2382	Info	17	12	0
2297	ISE	36	11	0
2381	PS	39	12	0
2312	Info	14	3	0
2387	PS	4	1	0
2398	Info	17	0	1
2391	PS	31	3	0
2431	PS	3	0	0
2282	Info	8	0	0
2323	ISE	1	0	0
2448	ISE	0	0	0

We can compare the median number of citations and the numbers of citations for the least and most popular quartiles in the three years:

References	Lower 25%	Median	Higher 25%
RFC (2018)	0	1	3
RFC (2008)	6.5	11	21.75
RFC (2008), until 2009	1	2.5	4.5
RFC (2008), 2018 and after	0	0	2
RFC (1998)	4.75	13.5	32.25
RFC (1998), until 1999	0	2	4.25
RFC (1998), 2018 and after	0	0	1

The total numbers shows new documents with fewer citations than the older ones. This can be explained to some degree by the passage of time. If we restrict the analysis to the number of citations accrued in the year of publishing and the year after that, we still see about the same distribution for the three samples.

We also see that the number of references to RFC fades over time. Only the most popular of the RFC produced in 1998 are still cited in 2019.

5.3. Citations Versus Deployments

The following table shows side by side the number of citations as measured in Section 5.1 and the estimation of deployment as indicated in Section 3.

RFC(2018)	Status	Citations	Deployment
8411	Info	1	medium
8456	Info	1	medium
8446	PS	418	high
8355	Info	3	medium
8441	PS	1	high
8324	ISE	0	N/A
8377	PS	0	unknown
8498	Info	0	unknown
8479	ISE	0	one
8453	Info	3	unknown
8429	BCP	0	some
8312	Info	25	high
8492	ISE	4	one
8378	Exp	1	some
8361	PS	0	one
8472	PS	1	medium
8466	PS	0	unknown
8362	PS	1	medium
8468	Info	1	some

From looking at these results, it is fairly obvious that citation counts cannot be used as proxies for the "value" of an RFC. In our sample, the two RFC that have high citation counts were both widely deployed, and can certainly be described as successful, but we also see many RFC that saw significant deployment without garnering a high level of citations.

Citation counts are driven by academic interest, but are only loosely correlated with actual deployment. We saw that [RFC8446] was widely cited in part because the standardization process involved many researchers, and that the high citation count of [RFC8312] is largely due to the academic interest in evaluating congestion control protocols. If we look at previous years, the most cited RFC in the 2008 sample is [RFC5326], an experimental RFC defining security extensions to an experimental delay tolerant transport protocol. This protocol does not carry a significant proportion of Internet traffic, but has been the object of a fair number of academic studies.

The citation process tends to privilege the first expression of a concept. We see that with the most cited RFC in the 1998 set is [RFC2267], an informational RFC defining Network Ingress Filtering that was obsoleted in May 2000 by [RFC2827]. It is still cited frequently in 2018 and 2019, regardless of its formal status in the RFC series. We see the same effect at work with [RFC8441], which garners very few citations although it obsoletes [RFC6455] that has a large number of citations. The same goes for [RFC8468], which is sparsely cited while the [RFC2330] is widely cited. Just counting citations will not indicate whether developers still use an old specification or have adopted the revised RFC.

6. Observations and Next Steps

The goal of this study was to evaluate how the IETF produces documents, from working group processing to RFC production. As shown in Section 4, the average RFC was produced in 3 years and 4 months, which is similar to what was found in the 2008 sample, but more than three times larger than the delays for the 1998 sample.

The Working group process appears to be the main source of delays. Efforts to diminish delays should probably focus there, instead of on the IETF and IESG reviews of the RFC production. For the RFC production phase, most of the variability originates in the AUTH-48 process, which is influenced by a variety of factors such as number of authors or level of engagement of these authors.

Most of the delay is spent in the working group, but the IETF data tracker does not hold much information about what happens inside the working groups. For example, events like Working Group Last Calls are not recorded in the history of the drafts available in the datatracker. Such information would have been interesting. Of course, requiring that information would create an administrative burden, so there is clearly a trade-off between requiring more work from working group chairs and providing better data for process analysis.

The Independent Submission stream operates as expected. The majority of the authors of the independent RFC appear to be in IETF insiders, but there is significant amount of engagement by outside parties.

The analysis of citations in Section 5.1 shows that citation numbers are a very poor indication of the "value" of an RFC. Citation numbers measure the engagement of academic researchers with specific topics, but have little correlation with the level of adoption and deployment of a specific RFC.

This document analyses a small sample of RFC "in depth". This allowed gathering of detailed feedback on the process and the deployments. On the other hand, much of the data on delays is available from the IETF data tracker. It may be worth considering adding an automated reporting of delay metrics in the IETF data tracker.

This document only considers the RFC that were published in a given year. This approach can be criticized as introducing a form of "survivor bias". There are many drafts proposed to the IETF, and only a fraction of them end up being published as RFC. On one hand this is expected, because part of the process is to triage between ideas that can gather consensus and those that don't. On the other hand, we don't know whether that triage is too drastic and discouraged progress on good ideas.

One way to evaluate the triage process would be to look at publication attempts that were abandoned, for example drafts that expired without progressing or being replaced. The sampling methodology could also be used for that purpose. Pick maybe 20 drafts at random, among those abandoned in a target year, and investigate why they were abandoned. Was it because better solutions emerged in the working group? Or maybe because the authors discovered a flaw in their proposal? Or was it because some factional struggle blocked a good idea? Was the idea pursued in a different venue? Hopefully, someone will try this kind of investigation.

7. Security Considerations

This draft does not specify any protocol.

We might want to analyze whether security issues were discovered after publication of specific standards.

8. IANA Considerations

This draft does not require any IANA action.

Peliminary analysis does not indicate that IANA is causing any particular delay in the RFC publication process.

9. Acknowledgements

Many thanks to the authors of the selected RFC who were willing to provide feedback on the process: Michael Ackermann, Zafar Ali, Sarah Banks, Bruno Decraene, Lars Eggert, Nalini Elkins, Joachim Fabini, Dino Farinacci, Clarence Filsfils, Sujay Gupta, Dan Harkins, Vinayak Hegde, Benjamin Kaduk, John Klensin, Acee Lindem, Nikos Mavrogiannopoulos, Patrick McManus, Victor Moreno, Al Morton, Andrei Popov, Eric Rescorla, Michiko Short, Bhuvaneswaran Vengainathan, Lao Weiguo, and Li Yizhou. Many thanks to Adrian Farrel for his useful advice, and to Stephen Farrell and Colin Perkins for their guidance on the use of citations.

10. Informative References

[IETFCOUNT]
IETF, "Past IETF Meetings", 2020, <https://www.ietf.org/how/meetings/past/>.

[RFC2267] Ferguson, P. and D. Senie, "Network Ingress Filtering: Defeating Denial of Service Attacks which employ IP Source Address Spoofing", RFC 2267, DOI 10.17487/RFC2267, January 1998, <https://www.rfc-editor.org/info/rfc2267>.

[RFC2330] Paxson, V., Almes, G., Mahdavi, J., and M. Mathis, "Framework for IP Performance Metrics", RFC 2330, DOI 10.17487/RFC2330, May 1998, <https://www.rfc-editor.org/info/rfc2330>.

[RFC2827] Ferguson, P. and D. Senie, "Network Ingress Filtering: Defeating Denial of Service Attacks which employ IP Source Address Spoofing", BCP 38, RFC 2827, DOI 10.17487/RFC2827, May 2000, <https://www.rfc-editor.org/info/rfc2827>.

[RFC3797] Eastlake 3rd, D., "Publicly Verifiable Nominations Committee (NomCom) Random Selection", RFC 3797, DOI 10.17487/RFC3797, June 2004, <https://www.rfc-editor.org/info/rfc3797>.

[RFC5326] Ramadas, M., Burleigh, S., and S. Farrell, "Licklider Transmission Protocol - Specification", RFC 5326, DOI 10.17487/RFC5326, September 2008, <https://www.rfc-editor.org/info/rfc5326>.

[RFC6455] Fette, I. and A. Melnikov, "The WebSocket Protocol", RFC 6455, DOI 10.17487/RFC6455, December 2011, <https://www.rfc-editor.org/info/rfc6455>.

[RFC8312] Rhee, I., Xu, L., Ha, S., Zimmermann, A., Eggert, L., and R. Scheffenegger, "CUBIC for Fast Long-Distance Networks", RFC 8312, DOI 10.17487/RFC8312, February 2018, <https://www.rfc-editor.org/info/rfc8312>.

[RFC8324] Klensin, J., "DNS Privacy, Authorization, Special Uses, Encoding, Characters, Matching, and Root Structure: Time for Another Look?", RFC 8324, DOI 10.17487/RFC8324, February 2018, <https://www.rfc-editor.org/info/rfc8324>.

[RFC8355] Filsfils, C., Ed., Previdi, S., Ed., Decraene, B., and R. Shakir, "Resiliency Use Cases in Source Packet Routing in Networking (SPRING) Networks", RFC 8355, DOI 10.17487/RFC8355, March 2018, <https://www.rfc-editor.org/info/rfc8355>.

[RFC8361] Hao, W., Li, Y., Durrani, M., Gupta, S., and A. Qu, "Transparent Interconnection of Lots of Links (TRILL): Centralized Replication for Active-Active Broadcast, Unknown Unicast, and Multicast (BUM) Traffic", RFC 8361, DOI 10.17487/RFC8361, April 2018, <https://www.rfc-editor.org/info/rfc8361>.

[RFC8362] Lindem, A., Roy, A., Goethals, D., Reddy Vallem, V., and F. Baker, "OSPFv3 Link State Advertisement (LSA) Extensibility", RFC 8362, DOI 10.17487/RFC8362, April 2018, <https://www.rfc-editor.org/info/rfc8362>.

[RFC8377] Eastlake 3rd, D., Zhang, M., and A. Banerjee, "Transparent Interconnection of Lots of Links (TRILL): Multi-Topology", RFC 8377, DOI 10.17487/RFC8377, July 2018, <https://www.rfc-editor.org/info/rfc8377>.

[RFC8378] Moreno, V. and D. Farinacci, "Signal-Free Locator/ID Separation Protocol (LISP) Multicast", RFC 8378, DOI 10.17487/RFC8378, May 2018, <https://www.rfc-editor.org/info/rfc8378>.

[RFC8410] Josefsson, S. and J. Schaad, "Algorithm Identifiers for Ed25519, Ed448, X25519, and X448 for Use in the Internet X.509 Public Key Infrastructure", RFC 8410, DOI 10.17487/RFC8410, August 2018, <https://www.rfc-editor.org/info/rfc8410>.

[RFC8411] Schaad, J. and R. Andrews, "IANA Registration for the Cryptographic Algorithm Object Identifier Range", RFC 8411, DOI 10.17487/RFC8411, August 2018, <https://www.rfc-editor.org/info/rfc8411>.

[RFC8429] Kaduk, B. and M. Short, "Deprecate Triple-DES (3DES) and RC4 in Kerberos", BCP 218, RFC 8429, DOI 10.17487/RFC8429, October 2018, <https://www.rfc-editor.org/info/rfc8429>.

[RFC8441] McManus, P., "Bootstrapping WebSockets with HTTP/2", RFC 8441, DOI 10.17487/RFC8441, September 2018, <https://www.rfc-editor.org/info/rfc8441>.

[RFC8446] Rescorla, E., "The Transport Layer Security (TLS) Protocol Version 1.3", RFC 8446, DOI 10.17487/RFC8446, August 2018, <https://www.rfc-editor.org/info/rfc8446>.

[RFC8453] Ceccarelli, D., Ed. and Y. Lee, Ed., "Framework for Abstraction and Control of TE Networks (ACTN)", RFC 8453, DOI 10.17487/RFC8453, August 2018, <https://www.rfc-editor.org/info/rfc8453>.

[RFC8455] Bhuvaneswaran, V., Basil, A., Tassinari, M., Manral, V., and S. Banks, "Terminology for Benchmarking Software-Defined Networking (SDN) Controller Performance", RFC 8455, DOI 10.17487/RFC8455, October 2018, <https://www.rfc-editor.org/info/rfc8455>.

[RFC8456] Bhuvaneswaran, V., Basil, A., Tassinari, M., Manral, V., and S. Banks, "Benchmarking Methodology for Software-Defined Networking (SDN) Controller Performance", RFC 8456, DOI 10.17487/RFC8456, October 2018, <https://www.rfc-editor.org/info/rfc8456>.

[RFC8466] Wen, B., Fioccola, G., Ed., Xie, C., and L. Jalil, "A YANG Data Model for Layer 2 Virtual Private Network (L2VPN) Service Delivery", RFC 8466, DOI 10.17487/RFC8466, October 2018, <https://www.rfc-editor.org/info/rfc8466>.

[rfc8468] Morton, A., Fabini, J., Elkins, N., Ackermann, M., and V. Hegde, "IPv4, IPv6, and IPv4-IPv6 Coexistence: Updates for the IP Performance Metrics (IPPM) Framework", RFC 8468, DOI 10.17487/RFC8468, November 2018, <https://www.rfc-editor.org/info/rfc8468>.

[RFC8468] Morton, A., Fabini, J., Elkins, N., Ackermann, M., and V. Hegde, "IPv4, IPv6, and IPv4-IPv6 Coexistence: Updates for the IP Performance Metrics (IPPM) Framework", RFC 8468, DOI 10.17487/RFC8468, November 2018, <https://www.rfc-editor.org/info/rfc8468>.

[RFC8472] Popov, A., Ed., Nystroem, M., and D. Balfanz, "Transport Layer Security (TLS) Extension for Token Binding Protocol Negotiation", RFC 8472, DOI 10.17487/RFC8472, October 2018, <https://www.rfc-editor.org/info/rfc8472>.

[RFC8479] Mavrogiannopoulos, N., "Storing Validation Parameters in PKCS#8", RFC 8479, DOI 10.17487/RFC8479, September 2018, <https://www.rfc-editor.org/info/rfc8479>.

[RFC8483] Song, L., Ed., Liu, D., Vixie, P., Kato, A., and S. Kerr, "Yeti DNS Testbed", RFC 8483, DOI 10.17487/RFC8483, October 2018, <https://www.rfc-editor.org/info/rfc8483>.

[RFC8492] Harkins, D., Ed., "Secure Password Ciphersuites for Transport Layer Security (TLS)", RFC 8492, DOI 10.17487/RFC8492, February 2019, <https://www.rfc-editor.org/info/rfc8492>.

[RFC8493] Kunze, J., Littman, J., Madden, E., Scancella, J., and C. Adams, "The BagIt File Packaging Format (V1.0)", RFC 8493, DOI 10.17487/RFC8493, October 2018, <https://www.rfc-editor.org/info/rfc8493>.

[RFC8498] Mohali, M., "A P-Served-User Header Field Parameter for an Originating Call Diversion (CDIV) Session Case in the Session Initiation Protocol (SIP)", RFC 8498, DOI 10.17487/RFC8498, February 2019, <https://www.rfc-editor.org/info/rfc8498>.

[RFCYEAR] RFC Editor, "Number of RFC Published per YEAR", 2020, <https://www.rfc-editor.org/rfcs-per-year/>.

[SSCH] Allen Institute for AI, "Semantic Scholar", 2020, <https://www.semanticscholar.org/>.

[TLS13IMP]
TLS WG, "TLS 1.3 Implementations", 2020, <https://github.com/tlswg/tlswg-wiki/blob/master/IMPLEMENTATIONS.md>.

[TRKR] IETF, "IETF Data Tracker", 2020, <https://datatracker.ietf.org/>.

Author's Address

Christian Huitema
Private Octopus Inc.
427 Golfcourse Rd
Friday Harbor WA 98250
U.S.A

Email: huitema@huitema.net