

Internet Architecture Board (IAB)
Request for Comments: 8153
Category: Informational
ISSN: 2070-1721

H. Flanagan
RFC Editor
April 2017

Digital Preservation Considerations for the RFC Series

Abstract

The RFC Editor is both the publisher and the archivist for the RFC Series. This document applies specifically to the archivist role of the RFC Editor. It provides guidance on when and how to preserve RFCs and describes the tools required to view or re-create RFCs as necessary. This document also highlights gaps in the current process and suggests compromises to balance cost with best practice.

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 a candidate 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 <http://www.rfc-editor.org/info/rfc8153>.

Copyright Notice

Copyright (c) 2017 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 (<http://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
1.1. Terminology	4
1.2. Life Cycle of Digital Preservation	4
2. Updating Policy and Procedure	5
2.1. Acquisition of Documents	6
2.2. Ingestion of Documents	6
2.3. Metadata and Document Registration	7
2.4. Normalization and Standardization of Canonical File Structure and Format	9
2.4.1. 'Best Effort' Data Retention	10
2.4.2. Single Format for Archival Purposes	11
2.4.3. Holistic Archiving of the Computing Environment	12
2.5. Transformation/Migration to Current Publication Formats ...	12
2.6. System Parameters	13
2.7. Financial Impact	13
3. Recommendations	14
4. Summary	15
5. IANA Considerations	15
6. Security Considerations	15
7. Informative References	16
IAB Members at the Time of Approval	18
Author's Address	18

1. Introduction

The RFC Editor is both the publisher and the archivist for the RFC Series, a series of technical specifications and policy documents that includes foundational Internet standards [RFC6635] [RFC-SERIES]. The goal of the RFC Editor is to produce clear, consistent, and readable documents for the Internet community. Over time, the RFC Editor will use as many modern features, such as hyperlinks and content markup, within the document as necessary to convey the information the authors intended for their audience. As the archivist, however, the main goal is to preserve both the information described and the documents themselves for the indefinite future. To meet both of these goals, the RFC Editor must find the necessary balance between the publication needs of today and the archival needs of tomorrow, while acknowledging a finite set of resources to complete both aspects of the RFC Editor function.

While many files are created during the editing process, this document focuses on the archival needs of the Internet-Drafts (I-Ds) that were approved for publication and the RFCs that resulted from these I-Ds; I-Ds before they are approved for publication by the appropriate stream-approving body are out of scope.

To summarize, the key areas of tension between the roles of publisher and archivist are:

- o the desire of the publisher to meet the needs expressed by authors who want to use the latest technology (e.g., vector graphics, live links, and a rich set of metadata) within their documents; and
- o the desire of the archivist to support only the simplest format for documents possible -- currently held by the Series to be plain-text, ASCII-only documents -- so that the tools needed to view the documents are equally simple and resistant to changes in technology, resulting in a set of documents that will be easier to archive for at least the next several decades, if not centuries.

Through most of the history of the RFC Series, the file format for RFCs has been plain text with an ASCII-only character set. This choice offered the simplest format likely to remain available to the largest number of consumers and the format most likely to be resistant to changes in technology over time. Increasingly, however, consumers and authors are requesting additional features that would allow for easy reading on a wider array of devices while retaining all the metadata authors intended in their documents. In 2013, RFC 6949 ("RFC Series Format Requirements and Future Development") captured the high-level requirements for the Series; the fundamental issue was that plain-text, ASCII-only documents no longer meet the needs of the communities interested in using and producing RFCs [RFC6949].

The assertion that plain-text, ASCII-only documents no longer meet the needs of the community suggests that the simple archival process maintained by the RFC Editor is also no longer sufficient. More complex tools and file formats require a more complex process to ensure that RFCs can be read and rendered far into the future. This document describes the considerations that must inform any changes in policy and procedure, and it describes a model for the RFC Series to follow when additional formats beyond plain-text, ASCII-only RFCs are published. The functional model that provides the framework for the archival process described in this document was derived from the ISO Open Archival Information System (OAIS) reference model, defined in "Space data and information transfer systems -- Open archival information system (OAIS) -- Reference model" [ISO14721].

1.1. Terminology

Acquisition: The point at which a document is accepted by the RFC Editor for future inclusion into the archive.

Ingestion: The point at which a digital object is assigned all necessary metadata to describe the object and its contents and is added to the archive.

Bitstream preservation: The process of storing and maintaining digital objects over time, ensuring that there is no loss or corruption of the bits making up those objects.

Content preservation: The retention of the ability to read, listen, or watch a digital file in perpetuity. Content preservation is not about the bits being stored; it is about being able to access and present those bits to the user.

1.2. Life Cycle of Digital Preservation

The basic process for preserving digital information has been described by a variety of organizations. From the Life cycle Information For E-Literature (LIFE) project [LIFE] in the United Kingdom to the ongoing digital preservation work in the U.S. Library of Congress [USLOC], the basic digital preservation process is straightforward. Documents are acquired and processed, metadata is recorded, physical media is refreshed, and content is regularly checked to see if it is still accessible by interested parties. Complexities arise when one considers the need to preserve both the bits of the digital objects themselves and the tools with which to express those bits in an environment that experiences rapid changes in technology.

For most of the existence of the RFC Series, the digital preservation process has been fairly simple, focusing on bitstream preservation and relying on paper copies of digital files.

The current archival process for the RFC Series is as follows:

1. **Acquisition:** The RFC Editor database is updated to indicate an I-D has been approved for publication. At this point, the document is taken through the editorial process on the way to publication [RFC-PUB].
2. **Ingestion:** The RFC is added to the archive at the time of publication.

3. Metadata creation: The details regarding an RFC, including RFC number, author, title, abstract, etc., are created at time of publication. Additional metadata in the form of status and errata can be added or changed at any time, following the process of the originating document stream.
4. Bitstream preservation: This part of the process is handled as part of the IT system administration; all servers, disks, and backup technology are refreshed on a regular cycle.
5. Content preservation: All RFCs since January 2010 have been printed out on standard office paper at time of publication, and the electronic files have been preserved on disk and in backups with no particular focus on preserving the entire computing environment used to create the electronic documents. Most RFCs prior to January 2010 are also available on paper, but there are gaps in the record and issues of ownership around the paper copies before that date.

When the format for RFCs transitions from plain-text, ASCII-only files to an XML format with multiple outputs, the overall archival process will become more complex. Additional metadata and some (or possibly all) of the computing environment may need to be added to the archive.

2. Updating Policy and Procedure

RFCs are created and published as digital objects. Unlike paper-based publications, a digital collection requires a focus on retaining the details of the technology as well as retaining the object itself. Specifically, a digital archive needs to:

- o consider the inherent instability of digital media,
- o plan for a relatively short path to technological obsolescence,
- o schedule regular media updates,
- o apply predefined criteria for technology evaluation, and
- o ensure the continued authenticity and integrity of documents through any changes in technology.

As the custodian and canonical source of RFCs and associated errata, the RFC Editor must consider how to ensure the availability and integrity of this document series far into the future and determine whether the focus must be on bitstream preservation, content preservation, or both.

The RFC Editor has several advantages in acting as the digital archivist for the Series. Since the RFC Editor is the publisher as well as the archivist, the RFC Editor controls the format of the material and the process for adding that material to an archive and can add any additional metadata considered necessary. External material, while a major consideration for more general archives, is no longer accepted by the RFC Editor. (See "Internet Archaeology: Documents from Early History" [RFC-HISTORY] for the list of non-RFC digital objects held by the RFC Editor.)

This document describes several different preservation models that may fit the needs of the Series and raises several points for community consideration. Specifically, this document covers information on:

- o Acquisition of documents
- o Ingestion of documents
- o Metadata and document registration
- o Normalization and standardization of canonical file structure and format
- o Transformation/migration to current publication formats
- o Content and computing environment preservation
- o System parameters
- o Financial impact

2.1. Acquisition of Documents

The acquisition process for documents intended for the archive starts with the submission of an approved I-D for publication. During the editorial process, information such as the document metadata is finalized prior to publication. However, the initial I-D as submitted and the RFC produced from it do not formally enter the archive until the time of publication, which is considered the point of ingestion from an archival perspective.

2.2. Ingestion of Documents

Once an RFC is published, the canonical format is considered immutable. At this point, the RFC Production Center, one of the internal roles within the RFC Editor, assigns the document metadata that an archivist needs to identify the unique object.

In the case of RFCs, the metadata assigned to a document at the time of publication includes:

- o the RFC number
- o ISSN
- o publication date
- o Digital Object Identifier (DOI)

Additional metadata, such as author name, is assigned earlier in the document creation process, but it is subject to change up to the point of publication. More information on metadata is available in Section 2.3 ("Metadata and Document Registration").

In terms of deciding what to accept in the archive -- a major question for most archives and yet a simple one for the RFC Series -- the RFC Editor accepts documents that are approved for publication by the approving body of one of the document streams: the IETF, IAB, IRTF, or Independent Submission streams [RFC7841]. Each document stream has defined processes on when and how I-Ds are approved and submitted to the RFC Editor for publication. The RFC Editor does not select documents for publication and archiving; the RFC Editor edits and publishes documents approved for publication by the document streams.

The RFC Editor holds no copyright on I-Ds or RFCs. As per the IETF Trust Legal Provisions [TLP], the copyright for RFCs is held by the authors and the IETF Trust. At any point in time, the current entities providing RFC Editor services must be able to release the archive of RFCs to the IETF Trust.

Note: The RFC Editor is currently only responsible for RFCs; any associated datasets or other research data is not considered within the RFC Editor's mandate at this time; therefore, no consideration to the archival requirements of such datasets is covered in this document.

2.3. Metadata and Document Registration

Metadata is data about data. In the field of digital archiving, this is the data that clearly identifies every aspect of a document, from its identifier (i.e., the RFC number and the I-D draft string) to the size and file format of the document and more. Metadata is stored in a central registry that records information on exactly what is being

preserved and where it is located, information on authenticity and provenance, and details on the hardware and/or software needed to view or create the documents.

The RFC Editor maintains this registry in the form of a database that includes all metadata available for documents being edited and for published RFCs. This database feeds the search engine on the RFC Editor website and the info pages available for every RFC (e.g., <http://www.rfc-editor.org/info/rfc####>).

Following is the current list of metadata presented in the RFC info pages:

- o RFC number
- o Canonical URI
- o Title
- o Status
- o Updates (if applicable)
- o Updated by (if applicable)
- o Obsoletes (if applicable)
- o Obsoleted by (if applicable)
- o Authors
- o Stream
- o Abstract
- o Content-Type
- o Character Set
- o ISSN
- o Publication date
- o Digital Object Identifier (DOI)

The following metadata will be added in the future:

- o Publication format URIs

Info pages also include links to errata, IPR searches, and both plain-text and XML citation files.

In terms of best practice, all documents used as normative references within an RFC would also be stored in the archive. While this is done automatically when the normative reference is another RFC (the usual case), retaining a copy of third-party documents is considered out of scope for the RFC Editor. As the digital archive industry stabilizes, services such as Perma.cc [PERMACC] may be a reasonable compromise. These services provide a permanent URI and image capture of online documents, with a goal of buffering against URI and online availability changes.

2.4. Normalization and Standardization of Canonical File Structure and Format

The normalization process is perhaps the most technically critical part of digital archiving. The purpose is content preservation -- making sure the data accepted for archiving are in the most stable and easily accessed formats possible for the long-term future and require the least amount of re-engineering and emulation of environments in order to view the document in the future. Normalization is about enabling long-term access to the information within a document.

Over the history of the RFC Series, documents have been submitted for publication in a variety of formats, including paper for the earliest RFCs. Today, the majority of RFCs are available in both a canonical plain-text format and PDF format. For exceptions, see the RFC Online Project [RFC-ONLINE].

Currently, all RFCs are printed out to paper and stored at time of publication. This has been a reasonable backup plan for several decades. With few of the features one might expect from a digital document format (such as links, metadata within the document, and line drawings), plain-text files do not lose much, if any, information when printed out to paper. However, as the published formats change (see RFC 6949), printing to paper provides less value as much of the metadata that is an intrinsic yet invisible part of the rendered document will be lost in such printing. With that in mind, the focus needs to change to preserving the new file formats electronically.

While each RFC today is printed to paper and all electronic versions stored on multiple hard drives, no particular effort is made to ensure copies of the software used to render or read the canonical

plain-text RFC are also archived. The RFC Editor has several choices on how to adapt to the need to archive a more complex set of data and follow best practice as defined by the digital archive community:

- o a simplified bitstream preservation model that focuses on standard "best effort" data-retention practices, which rely on backups, upgrades, and regular equipment change to preserve the data. This model assumes that emulators may be built when needed if the formats used go out of common use (a significant part of the model currently followed by the RFC Editor).
- o a content preservation model that focuses on one publication format as the version most likely to be viewable and provide all necessary metadata in the future. This is a viable option considering that PDF/A-3 [PDF], one of the intended publication formats, was designed for this type of archiving.
- o a complex bitstream and content preservation model that focuses on archiving the canonical XML and the entire computing environment required to create, view and render all outputs from that file. This is the "best practice" from an archivist's perspective.

Those options are listed in order of least to greatest complexity and expense. More detail on each option is described below.

2.4.1. 'Best Effort' Data Retention

When dealing with very simple data structures such as plain-text, ASCII-only files, the experience of the RFC Series suggests that for the last few decades, hardware and operating system changes have had minimal impact on the document files being stored. While a complete failure of an operating system migration corrupted the dataset in the past, that situation represents a somewhat different problem than the tools themselves changing such that plain-text files are not easily read with existing technology. Given that the basic plain-text format and ASCII encoding remain in common use, the standard protections against file corruption and data loss, such as disk mirroring, off-site backups, and periodic restoration testing, will continue to provide access to the entirety of the RFC Series for the foreseeable future. As has been pointed out, both in this document and in broader community discussion, that is not sufficient for complex formats such as XML, HTML, PDF, or other proprietary formats offered by today's large IT companies. The risk of technological change resulting in the file formats mentioned being deprecated or changed without backwards compatibility is fairly high when looking decades or centuries into the future.

It is recommended that this model of archiving the RFC Series cease to be the primary model after the plain-text, ASCII-only format is no longer the canonical format. Best effort data retention is a necessary but not sufficient level of effort for preserving a digital archive. For more guidance on how to define best effort data retention, the section on "Media and Formats, Summary Recommendations" in the 2009 version of the Digital Preservation Handbook [DPC2009] provides useful and concrete information.

2.4.2. Single Format for Archival Purposes

If preserving the information described by a document, rather than the document itself, is the primary purpose of an archive, then focusing efforts on a single file format is a reasonable option. Some well-supported archival tooling projects follow this route, such as Archivemata [ARCHIVEMATICA]. By selecting a feature-rich yet fundamentally stable file format for documents, an organization may avoid expensive whole-environment reconstruction in order to view the document. The PDF/A formats were designed to be an archival format for electronic documents, and PDF/A-3 is one of the options intended for publication as the RFC Series moves from a plain-text canonical format to an XML canonical format with multiple publication formats. A PDF/A-3 file can be produced that embeds the XML from which the PDF/A-3 file was created; this allows for both original and rendered document validation if one has the correct tools available to see the source of the PDF/A-3 file [RFC7995]. The XML is not otherwise visible when viewing the PDF/A-3 file through typical PDF reader software.

When looking at the need to archive RFCs in a resource-limited environment, a content-preservation-only model has merit, but it is not without risks. First, PDF/A-3 will not be the canonical format; it is intended to be one of the rendered outputs. It may contain rendering bugs that were not intended to be in the document. Second, while the various PDF/A formats were designed to be archival, they have not been put to the test of time to determine if they will actually live up to the design goals.

This is a valid option to consider, but the risks, priorities, and costs must be discussed by the community before a decision is made to follow this path. The best option may be to combine this with one of the other methods of archiving described in this document to help minimize both risk and cost.

2.4.3. Holistic Archiving of the Computing Environment

Preserving everything published by the RFC Editor in order to have a permanent record of information, standards, and best practice is arguably the whole point of being an archival series. One can argue that it is not only about the information described in an RFC, it is also about supporting Intellectual Property Rights (IPR) and retaining the history of the Internet. In following this model, however, one must consider the complexity of the archival environment as matching, and possibly exceeding, the complexity of the file formats being preserved.

Consider a future where XML has been obsoleted for half a century, HTML5 was a format used three to four human generations ago, and PDF/A-3 is no longer supported by any existing company's reading software. For RFCs that were produced with XML as their canonical format, an archive must not only hold the data, it must also hold the entire computing environment that allows the data to be rendered and viewed. Operating systems and hardware on which those OSs can run, each major version of each piece of software used or relied upon during the publication of an RFC, browsers and readers for HTML, PDF, and any other publication format must be preserved in some fashion. This is considered best practice when archiving digital documents. This is also the most expensive method, and the cost only increases over time as more and more instances of the computing environment must be preserved over the lifetime of the Series.

This is a valid option to consider, but the sheer scope of resources required suggests that this must be discussed by the community before a decision is made. Pursuing this may require an entirely different paradigm for the RFC Editor from what has been considered in the past; expanding the scope and resources for the RFC Editor, finding a third party to take over the responsibilities of archiving, or some other option may be necessary.

2.5. Transformation/Migration to Current Publication Formats

Because normalization is a complex subject, it is important to consider how to mitigate the risk of failure of the normalization process.

The RFC Editor is responsible for making RFCs available to the Internet community. The canonical version of an RFC does not change once published; any formats officially rendered from the canonical version, however, may change. One way to mitigate the need to preserve the entire computing environment for an RFC, including web browsers and PDF readers, would be to take advantage of the non-canonical nature of the publication formats and re-render them from

the canonical source at the point that browser or reader technology has changed sufficiently to make RFCs largely unavailable to 'modern' tools.

For example, the RFC Editor may develop the practice of annually reviewing the tools needed to view the publication formats created by the RFC Editor to determine whether or not the current common and popular reader technologies (i.e., web browsers, PDF viewers, e-readers) can view the existing publication formats. During that review, the RFC Editor would work with the community to determine if the current publication formats meet the needs of the community and whether any should be retired or added to improve the availability of information to the community at that time.

2.6. System Parameters

While the industry best practice on the backup and restoration of data is not sufficient as a long-term archival solution, it is still a necessary part of keeping the Series available now and into the future. In the past, nearly 800 RFCs had to be manually transcribed from paper back to electronic format due to a failed server migration and insufficient backups.

The underlying servers hosting the tools, database, RFCs, and errata are the physical link in the archival environment. While such systems cannot and should not remain static and unchanging, there must be clear documentation regarding the environment, in particular, the storage, backups, and recovery processes for all RFC-related material. The documentation must include information on the refresh cycle for the physical storage and backup media and describe a regular cycle of data restoration and/or migration testing.

2.7. Financial Impact

Having a policy regarding digital archiving provides input into the budget process. The main costs associated with digital archives come from the complexity and quantity of the material being archived, as described in Section 2.4 on normalization.

Estimating potential costs and providing figures are outside of the scope of this document, but it should be noted that costs are a major factor when determining what level of archival practice an organization will follow.

For more information on potential business plans and cost modeling for digital preservation, see the "Business cases, benefits, costs, and impact" section of the Digital Preservation Handbook [DPC].

3. Recommendations

Given the need to balance cost and complexity with retention of information for historic, legal, and informational purposes, preservation efforts should focus on the XML canonical format files, the PDF/A-3 format files, the xml2rfc tool and its documentation, and at least two PDF reader applications capable of extracting the embedded XML. Care should be taken that the software being included in this archive has a provision for free copies for backup or archival purposes. All other formats and the overall computing environment should be stored as described in "best effort" data retention (Section 2.4.1), which should in turn be described in the appropriate vendor contract for the RFC Publisher.

Particular preservation efforts should be made by:

- o choosing a format designed for archiving RFCs (PDF/A-3 as indicated by [RFC7995])
- o embedding the canonical XML format within the PDF/A-3 file for RFCs
- o retaining a copy of the plain-text or XML file submitted for approved I-Ds
- o retaining all major versions of the tools and their associated documentation used to acquire and ingest an RFC
- o retaining the final XML file as well as the PDF/A-3 file with the embedded XML
- o retaining at least two software reader applications to ensure the PDF/A-3 and XML files can be viewed in the future
- o partnering with other digital archives around the world to mirror copies of the target data

In order to control costs and focus the archiving effort on the entire content of an RFC, including the metadata and other features embedded within each RFC published in more than just plain text, printing each RFC to paper upon publication is no longer reasonable. Proper data storage and mirrored copies of RFCs provide more efficient and effective copies in case of catastrophic failure of the existing archive of material.

Particular focus should be given to finding partners that specialize in digital preservation to ingest RFCs. Ideally, they will ingest all material associated with an RFC, including all metadata, digital

signatures, and the approved I-D that was submitted to the RFC Editor. The possibilities and options should be discussed with each archival partner; at minimum, they must ingest copies of RFCs as they are published, with the basic metadata associated with each document.

Preservation efforts should be reviewed and validated through a biennial audit that will verify that the targeted content and all its associated metadata can be read with existing tools. The full process from acquisition to ingestion should be reviewed to ensure that best current practice is being followed from the perspective of the digital archive community. Since the overall model for the digital archive maintained by the RFC Editor follows the OAIS reference model, the associated audit guidelines should also be followed. While the RFC Editor does not seek to be recognized as 'OAIS-compliant' at this time, use of the ISO standard "Space data and information transfer systems -- Audit and certification of trustworthy digital repositories" [ISO16363] would provide a solid, accepted method for structuring an audit for this digital archive.

4. Summary

The RFC Series is worth archiving. It contains the history of the early Internet, as well as some of the key standards for Internet technology and best practice today. Who knows what the community will create in the future? There are many ways to preserve the Series, from relying on preservation of the bits, to focusing on a single file format, to preserving the entire computing environment. Each possibility, or permutations of them, involves risks and requires varying levels of resources. The goal of this document is to describe the possibilities and associated risks so that the community can come to an informed decision regarding what it is willing to see supported far into the future.

5. IANA Considerations

This document does not require any IANA actions.

6. Security Considerations

This document assumes that the origination of RFCs via the RFC Editor is secure and trusted. With that assumption, the activities discussed in this document do not affect the security of the Internet.

7. Informative References

- [ARCHIVEMATICA] "Archivemata", <https://www.archivemata.org/wiki/Main_Page>.
- [DPC] Digital Preservation Coalition, "Digital Preservation Handbook", 2015, <<http://dpconline.org/handbook>>.
- [DPC2009] Digital Preservation Coalition, "Digital Preservation Handbook", 2009, <<http://www.dpconline.org/docman/digital-preservation-handbook/304-digital-preservation-handbook-media-and-formats>>.
- [ISO14721] International Organization for Standardization, "Space data and information transfer systems -- Open archival information system (OAIS) -- Reference model", ISO 14721:2012, 2012.
- [ISO16363] International Organization for Standardization, "Space data and information transfer systems -- Audit and certification of trustworthy digital repositories", ISO 16363:2012, 2012.
- [LIFE] Hole, B., "LIFE^3: Predictive Costing of Digital Preservation", July 2010, <http://www.life.ac.uk/3/docs/Hole_pasig_v1.pdf>.
- [PDF] International Organization for Standardization, "Document management -- Electronic document file format for long-term preservation -- Part 3: Use of ISO 32000-1 with support for embedded files (PDF/A-3)", ISO 19005-3:2012, 2012.
- [PERMACC] "Perma.cc", <<http://perma.cc/>>.
- [RFC-HISTORY] RFC Editor, "Internet Archaeology: Documents from Early History", <<http://www.rfc-editor.org/history.html>>.
- [RFC-ONLINE] RFC Editor, "History of RFC Online Project", <<http://www.rfc-editor.org/rfc-online-2000.html>>.
- [RFC-PUB] RFC Editor, "Publication Process", <<http://www.rfc-editor.org/pubprocess.html>>.

- [RFC-SERIES] RFC Editor, "About Us",
<<http://www.rfc-editor.org/RFCoverview.html>>.
- [RFC6635] Kolkman, O., Ed., Halpern, J., Ed., and IAB, "RFC Editor Model (Version 2)", RFC 6635, DOI 10.17487/RFC6635, June 2012, <<http://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, <<http://www.rfc-editor.org/info/rfc6949>>.
- [RFC7841] Halpern, J., Ed., Daigle, L., Ed., and O. Kolkman, Ed., "RFC Streams, Headers, and Boilerplates", RFC 7841, DOI 10.17487/RFC7841, May 2016, <<http://www.rfc-editor.org/info/rfc7841>>.
- [RFC7995] Hansen, T., Ed., Masinter, L., and M. Hardy, "PDF Format for RFCs", RFC 7995, DOI 10.17487/RFC7995, December 2016, <<http://www.rfc-editor.org/info/rfc7995>>.
- [TLP] IETF Trust, "Trust Legal Provisions (TLP)", <<https://trustee.ietf.org/trust-legal-provisions.html>>.
- [USLOC] LeFurgy, B., "Life Cycle Models for Digital Stewardship", February 2012, <<http://blogs.loc.gov/digitalpreservation/2012/02/life-cycle-models-for-digital-stewardship/>>.

IAB Members at the Time of Approval

The IAB members at the time this document was approved were (in alphabetical order):

Jari Arkko
Ralph Droms
Ted Hardie
Joe Hildebrand
Lee Howard
Erik Nordmark
Robert Sparks
Andrew Sullivan
Dave Thaler
Martin Thomson
Brian Trammell
Suzanne Woolf

Author's Address

Heather Flanagan
RFC Editor

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