

**Digital Preservation Considerations for the RFC Series
draft-flanagan-rfc-preservation-02**

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 the tools required to view or re-create RFCs as necessary. This document also highlights where gaps are in the current process, and where compromises are suggested to balance cost with ideal best practice.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of [BCP 78](#) and [BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on June 5, 2015.

Copyright Notice

Copyright (c) 2014 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. Code Components extracted from this document must include Simplified BSD License text as described in [Section 4.e](#) of

the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

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. Ingest 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](#) [11](#)
- 2.5. Transformation/migration to current publication formats . [12](#)
- [2.6. System Parameters](#) [13](#)
- [2.7. Financial Planning](#) [13](#)
- [3. Recommendations](#) [14](#)
- [4. Summary](#) [15](#)
- [5. IANA Considerations](#) [15](#)
- [6. Security Considerations](#) [15](#)
- [7. Draft Change Log](#) [15](#)
- [7.1. -01 to -02](#) [15](#)
- [7.2. -00 to -01](#) [15](#)
- [8. Informative References](#) [15](#)
- Author's Address [17](#)

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](#)] [[RFCSERIES](#)]. As the publisher of these documents, the goal is to produce clear, consistent, and readable documents for the community using 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.

Flanagan

Expires June 5, 2015

[Page 2]

While many files are created during the publication process, this document focuses on the archival needs of RFCs and the Internet-Drafts (I-Ds) that are approved for publication; 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 the authors who want to use the latest technology within their documents, such as vector graphics, live links, and a rich set of metadata;
- o the desire of the archivist to support only the simplest format for documents possible--currently held by the Series to be ASCII-only plain-text--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 one 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 and retain all the metadata an author intended in their document. In 2013, [RFC 6949](#), "RFC Series Format Requirements and Future Development," captured the high level requirements for the Series; the fundamental issue being that the plain-text, ASCII-only documents no longer met 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 in turn suggests that the simple archive process maintained by the RFC Editor is also no longer sufficient. More complex tools and file formats require a more complex process to make sure that RFCs can still be read and rendered far into the future. This document describes the considerations that must inform any changes in policy and procedure, and describes a model for the RFC Series to follow when additional formats beyond the ASCII-only, plain-text 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" [[IS014721](#)].

1.1. Terminology

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

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

Bit stream 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. It 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 in the United Kingdom, to the ongoing digital preservation work in the U.S. Library of Congress, the basic digital preservation process is straightforward [[LIFE](#)] [[USLOC](#)]. 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. The 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 bit stream preservation and relying on paper copies of digital files.

The archival process for the RFC Series is as follows:

1. Acquisition: The RFC Editor database is updated to indicate an Internet Draft (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. Ingest: 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. Bit stream 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 are printed out on paper at time of publication, and the electronic files preserved on disk and in backups with no particular focus on preserving the entire computing environment used to create the electronic documents.

When the format for RFCs changes from plain-text, ASCII-encoded files, the archival process overall 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 RFCs 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 bit stream 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, the process for adding those materials to an archive, and

can add any additional metadata considered necessary. External materials, while a major consideration for more general archives, are no longer accepted by the RFC Editor. (See "Internet Archaeology: Documents from Early History" for the list of non-RFC digital objects held by the RFC Editor [[RFC-HISTORY](#)].)

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

- o Acquisition of documents
- o Ingest 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 are finalized prior to publication. The initial I-D as submitted and the RFC produced from it do not formally enter the archive, however, until the time of publication, which is considered the point of ingest from an archival perspective.

2.2. Ingest 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 an archivist needs to identify the unique object.

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

- o the RFC number

- o ISSN
- o publication date
- o Digital Object Identifier (DOI) --future

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 "Metadata and document registration."

The publication of an RFC--the point at which responsibility for the document moves to the RFC Publisher, another internal role within the RFC Editor--starts the formal archival process for the documents. At that time, the canonical document should be digitally signed. Information regarding the signatures and how to verify them must be made available on the RFC Editor website.

In terms of deciding what to accept in the archive--a major question for most archives, and yet simple for the RFC Series--the RFC Editor accepts documents that are approved for publication by the stream approving body of one of the document streams: the IETF, IAB, IRTF, or Independent Submissions streams [[RFC5741](#)]. 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 as directed by the document streams.

The RFC Editor holds no copyright on I-Ds or RFCs. As per the IETF Trust Legal Provisions, the copyright for RFCs is held by the authors and the IETF Trust [[TLP](#)]. 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 data sets or other research data is not considered within the RFC Editor's mandate at this time and 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, the I-D draft string) to the size and file format of the document and more. Metadata is stored in a central registry that stores information on what exactly is being preserved, 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 engaged in the final editing and publication process. 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####>).

Current list of metadata presented in the RFC Info pages

- o RFC number
- o Canonical URI
- o Title
- o Status
- o Updates
- o Authors
- o Stream
- o Abstract
- o Content-Type
- o Character Set
- o ISSN
- o Publication date

Metadata to be added in the future

- o Digital Object Identifier (DOI)
- o Publication format URIs

Info pages also include links to: errata, IPR searches, 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 may be a reasonable compromise [[PERMACC](#)]. Those 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 parts of digital archiving. The purpose here 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, requiring 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 in the earliest RFCs. Today, the majority of RFCs are available in both a canonical plain-text format and PDF format. For exceptions to this list, 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 (including links, metadata within the document, or line drawings), plain-text files do not lose much, if any, information when printed out to paper. As the published formats change (see [RFC 6949](#)), however, 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 on 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 a more complex set of data to archive and follow best practice as defined by the digital archive community:

- o a simplified bit stream preservation model that focuses on "best effort" standard data retention practices, which rely on backups, upgrades, and regular equipment change to preserve the data, and assuming that emulators may be built when needed if the formats

used go out of common use (a significant part of the existing model);

- o a content preservation model that focuses on one publication format as a version most likely to be viewable and provide all necessary metadata in the future (a viable option considering the fact that PDF/A-3--one of the intended publication formats--was designed for this type of archiving) [[PDF](#)];
- o a complex bit stream 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 (the "best practice" when looking at this 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 in the past had corrupted the data set, 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 when one moves into more 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 at a future of decades or centuries.

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 latest version of the Digital Preservation Handbook provides useful, concrete information [[DPC](#)].

2.4.2. Single format for archival purposes

If one ascribes to the idea that 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 Archivemetica https://www.archivemetica.org/wiki/Main_Page . 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, which in turn 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 [[draft-hansen-rfc-use-of-pdf](#)].

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, but 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, it has not been put to the test of time to determine if will actual live up to its design goals.

It 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 through 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 no longer supported by any existing company's reading software. In order 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. It is also the most expensive, 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 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 than 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

Noting that normalization is a complex subject, it is important to consider what to do 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 a practice of starting an annual review of the tools needed to view the publication formats created by the RFC Editor, and 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 archive 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 Planning

Having a digital archive policy 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 the section on Normalization. To quote the Digital Protection Conservancy Handbook:

The complexity of the material submitted and number of objects acquired generally has more impact on costs than the total storage size. The type and variety of formats accepted into the repository will also affect cost, because for example proprietary formats are likely to be more difficult and expensive to manage in the long term. It may be possible to reduce costs by limiting the formats the repository will accept, or transforming material into a standard common format. This can be done to reduce the number of file types and possibly reducing the storage size. However, it is also necessary to realise that due to storage redundancies required for back up each gigabyte of deposited data requires more than one gigabyte of disk space in repository storage. --
<http://www.dpconline.org/advice/preservationhandbook/institutional-strategies/costs-and-business-modelling>

Estimating potential costs and providing figures it 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.

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, the PDF/A-3 format, the xml2rfc tool and its documentation, and at least one PDF reader application. All other formats and the overall computing environment should be stored as described in "best effort" data retention, 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)
- o embedding the canonical XML format within the PDF/A-3 file for RFCs
- o adding a digital signature and checksum for the canonical XML and the PDF/A-3 files
- 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 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 upon publication to paper is no longer reasonable. Proper data storage and mirrored copies of RFCs provides more efficient and effective copies in case of catastrophic failure of the existing archive of material.

Preservation efforts should be reviewed and validated through a bi-annual 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 ingest should be reviewed to ensure that best current practice is being followed from a digital archive community perspective. Since the overall model for the RFC Editor-maintained digital archive follows the OAIS Reference model, the

associated audit guidelines should be followed. While the RFC Editor does not seek to be recognized as 'OAIS-compliant' at this time, use of the ISO standard, "Audit and Certification of Trustworthy Digital Repositories," would provide a solid, accepted method for structuring an audit for this digital archive [[IS016363](#)].

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 the permutations from them, involves risks and 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 they are willing to see supported far into the future.

5. IANA Considerations

None

6. Security Considerations

TBD

7. Draft Change Log

To be removed before publication

7.1. -01 to -02

Updated text where appropriate to indicate approved I-Ds should also be targeted for archiving

7.2. -00 to -01

Recommendations: added the requirement to archive reader software, and to stop printing out to paper

8. Informative References

[[draft-hansen-rfc-use-of-pdf](#)]

Hansen, T., Masinter, L., and M. Hardy, "PDF for an RFC Series Output Document Format", [draft-hansen-rfc-use-of-pdf-02](#), July 2014.

- [DPC] DigitalPreservationCoalition, "Digital Preservation Handbook", 2012,
<<http://www.dpconline.org/advice/preservationhandbook>>.
- [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 16363:2011 , 2011.
- [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, "'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.
- [PERMACC] "Perma.CC", n.d., <<http://perma.cc/>>.
- [RFC-HISTORY] RFC Editor, "Internet Archaeology: Documents from Early History", n.d., <<http://www.rfc-editor.org/history.html>>.
- [RFC-ONLINE] RFC Editor, "History of RFC Online Project", n.d.,
<<http://www.rfc-editor.org/rfc-online-2000.html>>.
- [RFC-PUB] RFC Editor, "RFC Editor Publication Process", n.d.,
<<http://www.rfc-editor.org/pubprocess.html>>.
- [RFCSERIES] RFC Editor, "Overview of RFC Document Series", n.d.,
<<http://www.rfc-editor.org/RFCoverview.html>>.
- [TLP] IETF Trust, "IETF Trust Legal Provisions", n.d.,
<<http://trustee.ietf.org/docs/IETF-Trust-License-Policy.pdf>>.

- [USLOC] Library of Congress, "Life Cycle Models for Digital Stewardship", n.d., <<http://blogs.loc.gov/digitalpreservation/2012/02/life-cycle-models-for-digital-stewardship/>>.
- [RFC5741] Daigle, L., Kolkman, O., and IAB, "RFC Streams, Headers, and Boilerplates", [RFC 5741](#), December 2009.
- [RFC6635] Kolkman, O., Halpern, J., and IAB, "RFC Editor Model (Version 2)", [RFC 6635](#), June 2012.
- [RFC6949] Flanagan, H. and N. Brownlee, "RFC Series Format Requirements and Future Development", [RFC 6949](#), May 2013.

Author's Address

Heather Flanagan
RFC Editor

Email: rse@rfc-editor.org

