Network Working Group
Internet-Draft
Intended status: Informational
Expires: January 17, 2019

H. Levkowetz
Elf Tools AB

July 16, 2018

Implementation notes for RFC 7991, "The 'xml2rfc' Version 3 Vocabulary"

draft-levkowetz-xml2rfc-v3-implementation-notes-00

Abstract

This memo documents issues and observations found while implementing RFC 7991. Individual notes are organised into separate sections, depending on their characters.

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 https://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 January 17, 2019.

Table of Contents

1. Introduction

Implementation of tool support for [RFC7991] and related specifications has been done during 2017 and 2018, split in the following individual parts, all implemented as individual modes of the python-based xml2rfc processor [XML2RFC]:

* An XML converter from vocabulary version 2 [RFC7749] to version 3 [RFC7991]

* A Normalisation processor, "PrepTool", [RFC7997]

* An XML to plain text converter [RFC7994] for the version 3 vocabulary

* An XML to html converter [RFC7992] for the version 3 vocabulary (pending as of 08 Jul 2018)

* A HTML to PDF converter [RFC7995] for the version 3 vocabulary (pending as of 08 Jul 2018)

During the implementation work, a number of issues with the specification has been found (this was expected at the outset by all parties) and a number of observations has been made about limitations of the specification and vocabulary version 3 schema, and also limitations in the specification of the work to be done.

The purpose of this memo is to collect those issues and observations in one place.

2. Fitness for Purpose

The introduction to [RFC7991] states:

> "This document defines the "xml2rfc" version 3 vocabulary: an XML-based language used for writing RFCs and Internet-Drafts. It is heavily derived from the version 2 vocabulary that is also under discussion. This document obsoletes the v2 grammar described in RFC 7749."

However, an unstated assumption seems to have been that the new tools and formatters would be used primarily to produce HTML output, in order to transition to publication of renderings of RFCs in more modern formats than plain-text ASCII.

This is a reasonable and worthwhile goal, but as a result, the schema as specified in [RFC7991] has some drawbacks compared with the version 2 vocabulary when used to produce Internet-Drafts in the text format common within the IETF (Internet Engineering Task Force) at this time.

2.1. Degraded Table of Contents

Lack of pagination has little impact on direct online readability, but when comparing the output of the new text formatter with the old one, one aspect leaps out: Since there is no pagination, the table of contents simply lists the section headers to a certain depth, without any accompanying page numbers. This makes a surprising difference in how useful the table of contents is in getting an initial feel for the document. The at-a-glance information which lets a reader know if this is a document of 10 pages or 100 is simply lacking.

Recommendation: Add support for pagination in a future version of the text formatter.

2.2. Justification of Tables and Artwork

The version 3 schema deprecates the previously available 'align' attribute for artwork and tables, and the PrepTool will remove these attributes if used. This makes a previous feature that was appreciated by some authors unavailable. In the text formatter, the effect is simply to make all tables and artwork left-aligned, which may not be the most readable and polished output, but for the HTML formatter it also potentially removes the option of letting text flow around smaller artwork and tables in a controlled way.

Recommendation: Make the 'align' attribute for artwork and tables available again. (The current text formatter code already has

support for the 'align' attribute for these elements; but since the attribute is stripped away by the PrepTool, the code is never invoked.)

2.3. RFC Publication Date Policy

The specification [RFC7998] says that an error should be generated if a <date> specification is found with missing elements; but the RFC Editor publishes documents (except for April 1st RFCs) with only year and month, no day of month. The specification disallows this, and in effect makes it impossible for the RFC Editor to publish documents according to the current policy regarding publication date format.

Recommendation: Revert to to the old behaviour, where the tool in RFC mode would issue a date with or without day depending on whether the <date> element had a day attribute or not.

3. Issues with the Schema

3.1. RFC 7991

3.1.1. In Section 2.5.5, "name" Attribute

"A filename suitable for the contents (such as for extraction to a local file)."

Given the existing use of "name" on seriesInfo, this attribute name has a semantic dissonance.

Recommendation: Deprecate "name" for use on <artwork> and <sourcecode>, and instead use "file", which for <sourcecode> will be explicitly rendered, as established as best current practice for YANG modules (see for instance RFC 6087 [RFC6087])

3.1.2. In Section 2.20, <dl>

The current specification says:

"The "hanging" attribute defines whether or not the term appears on the same line as the definition. hanging="true" indicates that the term is to the left of the definition, while hanging="false" indicates that the term will be on a separate line."

This does not match established typographic terminology. In typographic terminology, "hanging indent" describes the case where the indentation of the second and subsequent lines of a paragraph is greater than the indentation of the first line. Whether the definition in a definition list starts on the first line or not has

nothing to do with the presence of hanging indent; our definition lists will _always_ have hanging indent.

The 'hanging' attribute also describes something different from what the term has been used to describe in the version 2 vocabulary. This will be confusing to users.

A more descriptive name for the attribute we're talking about would be 'start-definition-on-first-line', but that's unwieldy. Maybe 'newline="false"' to start the definition on the first line, or something like 'definition-start="first"'?

Recommendation: Change this to a different term that is more descriptive and does not use typographically incorrect terminology.

3.1.3. New Section 2.20.4, "indent" Attribute

The deprecation of the "hangIndent" attribute on <list> leaves no opportunity to control the size of the hanging indent. In some definition lists, it is desirable to have a wide indentation, in order to clearly show the terms, in other cases it is more important to allow for a larger text volume than the width of the terms would allow.

Recommendation: Add an "indent" attribute on <dl> to control the size of the hanging indent.

3.1.4. In Section 2.29,

3.1.4.1. Unordered lists with arbitrary symbols

When <li> is used with <ul empty="true">, the rendering is under-specified (the specification say 'no label will be show", but doesn't say whether list indentation (leading white-space) should be eliminated or not.

If the intention is to make it possible to render unordered lists with arbitrary symbols, chosen on a per-list-item basis, the current attributes of <li> are insufficient to indent and line-wrap list items properly with <ul empty='true'>.

It is not possible, for instance, to use <ul> lists to generate XML for a table of content, since if the with of the bullet (the section number, in this case) is unknown, the proper indentation and line wrapping cannot be determined.

Recommendation: Add an explicit "bullet" attribute to support this use case.

3.1.4.2. Mixed Content Model

The mixed content model for <li> --- either text and inline elements like sub, sup, bcp14, _or_ <t>, <ul>, <figure> etc, is non-intuitive and may be hard for users to keep straight.

Recommendation: Consider simplifying the schema by requiring that text and inline elements always are placed within a <t> element.

This would apply also to other elements that today have alternative content models: <blockquote>, <dd>, <td>, and <th>.

3.1.5. In Section 2.32, <name>

So the <name> element can contain text or <tt>, and <tt> can contain other markup like <sub> and <sup> etc., but why cannot <name> contain <sup> etc. directly?

3.1.6. In Section 2.42, <references>

The v3 schema cannot properly model multiple reference subsections contained within one numbered section. The v2 formatter handled this by silently inserting a containing section, but with the introduction of the preptool, which in theory should produce a master file from which various formatters would produce equivalent results, this becomes troublesome, as the automatic insertion of a container section is specified for the html formatter, in section 9.8. of RFC 7992, but not for the text formatter. It would be much better to make the prepped xml explicitly show exactly what should be rendered, and not rely on formatters silently insert elements.

Recommendation: Update the schema to make it possible for <references> to contain <references>, and have the prepped xml explicitly show both the encapsulating section and the subsections. The current preptool implementation does this.

3.1.7. In Section 2.45.1, "category" Attribute

Changing the "category" attribute of <rfc> to a name value in an additional <seriesInfo> makes it much harder than it needs to be to look it up. It also makes the semantics of <seriesInfo> less clear.

Recommendation: Remove this, and keep the "category" attribute on <rfc>

3.1.8. In Section 2.53.3 and 2.53.4.

3.1.8.1. Unnecessary limitation on where the "keepWithNext" attribute can be used

Why keepWithNext only on <t>? It would be very natural to expect to be able to say keepWithNext for 2 tables, or 2 figures, or 2 lists?

Recommendation: Permit keepWithNext on all elements that can be siblings to <t>.

3.1.8.2. Violation of KISS and DRY principles

keepWithNext on one element is equivalent with keepWithPrevious on the following element, provided the following element can have a keepWithPrevious attribute. Providing both violates both KISS and DRY.

Recommendation: Keep only one of these two attributes, preferably keepWithNext.

3.1.9. In Section 2.63.2, <ul> "empty" attribute

In v2, this results in a list using space as the bullet, thus each list entry is indented as with other bullet symbols. However, this leaves no way to get list entries with arbitrary text that are not indented, in order to produce lists such as that used in Table of Content and Index.

The current implementation introduces a new attribute "bare" with the possible values "false" | "true" to signal this. This works, but is maybe clumsier than necessary.

3.1.10. In Section 3.4.2, "hangIndent" Attribute

"Deprecated. Use <dl> instead."

This causes capability loss. The "hangIndent" attribute not only signalled that hanging indent should be used, but also gave the size of the indent. No equivalent control has been provided for the <dl> element in the version 3 vocabulary.

3.1.11. In Appendix C. Relax NG schema

The "colspan" attribute is given a default value of "0", this should be "1". "0" is not otherwise defined in the text, and the only reasonable interpretation would be to hide the cell (make it occupy zero columns).

The "rowspan" attribute is given a default value of "0", this should be "1". "0" is not otherwise defined in the text, and the only reasonable interpretation would be to hide the cell (make it occupy zero rows).

3.1.12. Use of the term 'counter'.

The classical meaning of this term is a a monotonically increasing sequence of integers, globally unique or unique within a context. In this document, it is instead meant to indicate section, table, figure numbers, which for sections are not plain counters.

To make more interesting, in other contexts in the document, the notation "-nnn", which also would normally indicate a dash followed by digits, i.e., a counter, is also re-interpreted to include section numbers; strings of numbers including embedded period signs. This is bad terminology.

Recommendation: Instead of "counter", use "number" as the attribute value, and explicitly say "Section number, Figure number, Table number or ordered list labels" in the description. Use "-n.n" instead of "-nnn".

3.2. RFC 7998

3.2.1. In Section 5.2.6, Attribute Default Value Insertion

The <seriesInfo> "stream" attribute has a default value of "IETF". The effect of setting default values after the XInclude processing is to set stream="IETF" on all reference <seriesInfo> which don't have a stream set. This is probably not right.

The current implementation removes the default value for the "stream" attribute from the schema.

3.2.2. In Section 5.4.2.1, Compare <rfc> "submissionType" and <seriesInfo> "stream".

It doesn't seem like a good fit to have tag attributes that all have to be set to the same value. This is not DRY, and unnecessarily introduces the possibility of conflict, as a result of multiple <seriesInfo> elements being permitted (Relevant to the v3 schema, not the preptool).

3.2.3. In Section 5.4.6, "pn" Numbering.

The list of elements that are given p- or paragraph tags is severely limited, and since the presence of a pn= attribute is required in order to make internal <xref> instances work, this limits the elements to which it is possible to reference with html fragment identifiers. Why?
Why is <dt> and <li> present, but not <ol>, <dl>, ?

The current implementation adds p- numbering to <list>, <dl>, <dd>, <ol>, <ul>, which all are allowed to have pn= attributes according to the schema.

4. Non-Schema Issues

4.1. RFC 7991

4.1.1. In Section 2.17, <date>

4.1.1.1. Current Date Requirement

> "When the prep tool is used to create Internet-Drafts, it will reject a submitted Internet-Draft that has a <date> element in the boilerplate for itself that is anything other than today."

It is not up to the format definition to set policy for acceptance or rejection of draft submissions. The matter is more complex than the text assumes, see for instance datatracker issue #2422. In addition to being inappropriate, this text also quietly changes policy from +/- 3 days to +/- 0 days, without saying that it updates RFC 4228 [RFC4228], which is the current specification of permissible dates in draft submissions. Finally, enforcing this would cause _a lot_ of grief and problems.

This specification item has been ignored in the implementation.

4.1.1.2. Date Specification in References

> "Bibliographic references: In dates in <reference> elements, the date information can have prose text for the month or year. For example, vague dates (year="ca. 2000"), date ranges (year="2012-2013"), non-specific months (month="Second quarter"), and so on are allowed."

The text regarding prose text for month and year in bibliographic references is not workable. How should month and year be combined? Some bibliographic references may have date text which requires year first, others year last, and so on. Mixing the described fuzziness

into the otherwise strict year, month, date format makes little sense when the result of combining the year, month and date attributes cannot be predictably and correctly rendered.

Recommendation: Instead of the current specification, permit either that the <date> element may have text content, or an alternative attribute to be used for rendering if year, month, or day cannot be specified exactly.

4.1.2. In Section 2.47, <seriesInfo>

The possible and forbidden combinations of attributes for this element has now become so convoluted that it's really hard to understand how to use it correctly. This needs a serious reconsideration.

The 'name' attribute is mandatory, and only 3 values are permitted: "RFC", "Interned-Draft", and "DOI". But it is also mandatory to set the name to "" for a <seriesInfo> with a status attribute. Hmm...

So there are 4, not 3 permitted values: "RFC", "Internet-Draft", "DOI", and "".

This means that all reference files which has things like name="ISO", name="W3C Recommendation", etc., etc., have become illegal.

This limitation on <seriesInfo> "name" attributes has not been enforced in the current implementation.

4.1.3. In Appendix A.1.1: TLP switch-over date discrepancies

There are discrepancies between the specified switch-over dates in the specification, and those given by the Trust statements:

* TLP3.0: The specification says 2009-11-01 but the TLP statement says effective date 2009-09-12.

* TLP4.0: The specification says 2010-04-01 but the TLP statement says effective date 2009-12-28. The dates on which TLP 4 started to be use in published RFCs seems to match the stated effective date of 2009-12-28, based on a scan of some RFCs around that date.

The current implementation uses the official dates in the preptool, not the dates in RFC 7991.

RFC 7991 also states this about the pre5378 text: this text appears under "Copyright Notice", unless the document was published before November 2009, in which case it appears under "Status of This Memo".

This does not agree at all with what actual RFCs contain; they seem to consistently have this text under Copyright Notice.

4.1.4. Index

There is no guidance on the structure of an index, if one is to be generated by the preptool.

4.1.5. Anchors

Section 5.1 of RFC 7992 says in part:

> "The prep tool produces XML with anchor attributes in all elements that need them."

This is rather vital information regarding the content of the prepped xml when building a formatter, unfortunately it is not mentioned in RFC 7991.

4.2. RFC 7992

4.2.1. In Section 8.1.1, Index Contents

The index has an extra enclosing the contents, starting directly after <h2>, while sections explicitly does not have a div here. This irregularity seems quite unnecessary, but makes the formatter code more complex than need be. Could we please align the two?

4.3. RFC 7994

4.3.1. Additional Guidance

* <aside>: Guidance requested on the rendering. Now rendered with an indentation of 9 relative to surrounding text

* <blockquote>: Guidance requested on the rendering. Now rendered with an indentation of 3 spaces, pipe(|), two spaces relative to surrounding text.

* <sub>: Guidance requested. Now rendered as _(text)

* <sup>: Guidance requested. Now rendered as ^(text)

* <tt>: Guidance requested. Now rendered as "text"

* Guidance for <eref> rendering. In the html formatter, handling of <eref> is straightforward and is specified; it simply translates

to an external link. In the legacy text formatter, <eref> was handled by inserting an extra <references> subsection called "URLs", and adding reference entries for the URLs there, while the <eref> citation point got a trailing numeric reference number. With the preptool output becoming the authoritative published document, this difference won't be reflected in the xml. The two formats would be more aligned if the text formatter renders <eref> URLs inline.

Recommendation: Change the rendering of <eref> in text to render the URL inline within parentheses instead of adding the 'URLs' reference subsection.

4.4. RFC 7998

4.4.1. In Section 5.2.3, <date> Insertion

Error if any of year, month, day is missing:

It is an unnecessary and unwanted restriction when not in RFC processing mode to given an error for missing date elements. Missing date elements is permitted because they make it easier for draft authors to rev drafts without having to pay attention to the date values every time they generate new output. This requirement should apply only to RFC prepping mode.

Additionally, in RFC processing mode, this implicitly changes the RFC-Editor policy regarding publication dates, which earlier have specified only year and month (except for April 1st RFCs). Is this intentional?

4.4.2. In Section 5.2.4, "prepTime" Insertion

This is under-specified, given the detailed requirements on the <date> attributes. Should probably be RFC3339.

4.4.3. In Section 5.2.6, Attribute Default Value Insertion

All the default values in 7991 are also expressed in the v3.rnc schema. Remove text indicating otherwise. And by the way, it was very helpful to extract these from the schema programmatically; having them specified otherwise would make it much harder to follow a changing schema.

A number of attributes which are deprecated have default values. The current specification will cause those to be inserted, even if they have been removed earlier by the v2v3 converter because they are deprecated. This seems inconsistent.

Recommendation: Omit deprecated attributes from the default-setting.

4.4.4. In Section 5.2.7, "toc" Attribute

It's specified that sections with <boilerplate> ancestors should have toc="exclude", but this won't then affect <boilerplate> sections which are inserted as part of the processing in 5.4.2. It would make more sense to move this processing to after 5.4.2.

The logic in the second bullet is flawed. First it says to set elements with children with toc="include" to "include", but then it says that it is an error if they are set to "exclude". Either there should be a warning, and the toc= attribute should be updated, or there should be an error and termination. Not both.

4.4.5. In Section 5.2.8, "removeInRFC" Warning Paragraph

This potentially inserts a new <t> element, but after the default setting in 5.2.6. Maybe place default setting after all potential element insertions have taken place.

4.4.6. In Section 5.3.1, "month" Attribute

> "Normalise the values of "month" attributes in all <date> elements in <front> elements in <rfc> elements to numeric values."

Is that 'in' a direct descendant relationship, or any descendant? I.e., does this affect <date> elements in included <reference> elements? Unclear. (RFC7991 is much clearer on this point, but that's not an excuse for being unclear here).

4.4.7. In Section 5.3.2, ASCII Attribute Processing

The uppercasing of 'ascii' in the section <name> is incorrect in this case; the attribute name is explicitly 'ascii', not 'ASCII'. The section name should be '"ascii" Attribute Processing'.

> "In every <author> element"...

After the earlier XInclude processing, this will include all the author elements in the included references, which the document author cannot normally do an anything about. Is this the intention?

Recommendation: Limit it to /rfc/front/author' elements.

<title> and <postalLine> also has an ascii attribute -- is it a mistake that they are not mentioned here? Assuming so, for the preptool implementation.

What about the ascii* attributes on author? Assuming they should be processed the same way.

4.4.8. New Section: "keepWithNext" Normalisation

This should specify normalisation of keepWithNext/keepWithPrevious such as to replace all keepWithNext with an equivalent keepWithPrevious on the following <t>

4.4.9. In Section 5.4.2, <boilerplate> Insertion

> "Create a <boilerplate> element if it does not exist. If there are any children of the <boilerplate> element, produce a warning that says "Existing boilerplate being removed. Other tools, specifically the draft submission tool, will treat this condition as an error" and remove the existing children."

Should this be done in both I-D mode and RFC mode? The trouble is that the following subsections only describes the boilerplate relevant to an RFC; there's additional boilerplate that is needed for drafts. I don't think it's reasonable to have a draft with only parts of the boilerplate contained in a boilerplate section.

Recommendation: The boilerplate-element insertion parts of 5.4.2 be done in both RFC and draft mode, with the appropriate boilerplate for each case. Add text to describe the appropriate boilerplate for drafts, or remove the sections specific to RFC boilerplate.

This section also specifies an error message to be used verbatim; the troublesome thing is that it's not clear what it means. The message is: "Existing boilerplate being removed. Other tools, specifically the draft submission tool, will treat this condition as an error". What is it that the draft submission tool is going to treat as an error? The presence of boilerplate? Why? The removal of boilerplate? How is that related to draft submission? This is very jumbled.

4.4.10. In Section 5.4.2.1, Compare <rfc> submissionType and <seriesInfo> "stream".

This comes too late. It is specified that if either is missing, it should be added. But the default attribute setting earlier has set stream="IETF" on all <seriesInfo> elements that didn't have it. If a document is read without submissionType, and stream set correctly to something else than "IETF" on one of the <seriesInfo> elements, then the default-setting will have created a conflict which cannot be resolved purely from the document at this point.

It doesn't seem like a good fit to have tag attributes that all have to be set to the same value. This is not DRY, and unnecessarily introduces the possibility of conflict, as a result of multiple <seriesInfo> elements being permitted (Relevant to the v3 schema, not the preptool).

Recommendation: Remove the default value for stream, and make it subordinate to submissionType.

4.4.11. In Section 5.4.2.2, "Status of this Memo" Insertion

It specifies that one should consider both submissionType and <seriesInfo> stream value; but those have just been set equal in 5.4.2.1. The text should be adjusted to not sound as if these two should be both be considered as if they could be different.

4.4.12. In Section 5.4.3, <reference> "target" Insertion

> "Insert "target" attributes for RFC, DOI, and Internet-Draft references that lack them."

It is indicated that the rfc-editor will provide the URL patterns. What are they?

The order of <seriesInfo> determines the rendering order. These should be sorted in the desired rendering order (currently 'BCP', 'RFC', 'DOI'. The current implementation does so.

4.4.13. In Section 5.4.4, <name> Slugification

The 'n-' prefix for slugs is unnecessarily opaque.

Recommendation: Use slugs with prefix "name-" rather than "n-", to be more self-documenting.

Should the slugs be unique? Assuming yes, but guidance would be good. The current implementation enforces unique slugs, with the following algorithm:

* remove non-ascii letters

* replace-non-letters with dash, compacting multiple dashes to one

* reduce length to 32, but insure uniqueness by increasing length or adding numerical suffixes, up to length 40 with suffixes numbered 2 to 99.

4.4.14. In Section 5.4.6, "pn" Numbering.

What does 'pn' mean? Cryptic is never good when humans have to deal with it. At least explain as "part number" in text. Possibly even change pn="" to part="".

<back><section> is not mentioned. Assuming numbering as section-appendix.1.2

<iref> elements are not mentioned (but covered in 7991). Should be listed in 7998.

The numbering scheme is inconsistent between notes/boilerplate and other sections, in that attempting to split a pn on dashes (which external tools might want to do) the boilerplate/note sections contain an additional dash.

Recommendation: Change that to a dot, for better consistency with other sections. This also makes the <t> part numbers less confusing: "section-boilerplate.1-1" instead of "section-boilerplate-1-1"

4.4.14.1. RFC format anchors / fragment identifiers

The anchor prefixes described unnecessarily break with existing links to document sections. Wikipedia has (2018-02-19) about 84 000 pages that link to RFCs; with most pages having multiple links. A small manual sampling indicates that about 1 link in 10 has a #section-fragment identifier. All of these will break if the new tools are used to generated content linked from these pages.

How much larger than Wikipedia is the whole of the internet, in terms of links to RFCs? Hard to tell (though searching for 'rfc' on Google indicates 'about 10 000 000 results). In any case, we are talking about breaking a substantial number of links using fragment identifiers of the format #section- and #appendix- if the new tools are used to replace the old html content that sites currently point to.

Recommendation: update rfc7998 preptool to use these prefixes, instead:

- "section-xxx"
- "figure-xxx"
- "table-xxx"

- "appendix-xxx"
- "index-xxx"
- "para-xxx"
- "name-xxx"

4.4.15. In Section 5.4.7, <iref> Numbering

Numbering of <iref> talks about setting the 'pn' attribute. Mixed into this is a mention of 'irefid', which isn't a valid attribute. The current implementation assumes that 'pn' is meant.

The item and sub-item text is not constrained to slug format; in order to deliver useful pn values, slugification should be done. On the other hand, the explicit prescription of how to ensure uniqueness clashes with the total lack of uniqueness attention under 5.4.4.

Recommendation: Remove the details of how to ensure uniqueness.

4.4.16. In Section 5.4.8.2, "derivedContent" Insertion (without Content)

There's a formatting mistake:

The last sentence of the last bullet ("Issue a warning...") should not be part of the bullet, but a separate final paragraph for the Section.

4.4.17. In Section 5.5.1, <artwork> Processing

RFC791 specifies that the <artwork> content is a fallback if there is external <svn> content, but 7998 says to drop the fallback and insert the external <svn>. This deletes information, and makes the fallback unavailable. This needs a better handling.

For now, if there is fallback content, the external URL content is converted to a data: URL for the src, which pulls it in and makes it immutable, but retains the fallback.

4.4.18. In Section 5.5.2, <sourcecode> Processing

List item 4 says:

> "fill the content of the <sourcecode> element with the resolved XML from the URI in the "src" attribute"

However, the URI should not be assumed to resolve to xml, but instead treated like CDATA.

4.4.19. In Section 5.4.8.2, "derivedContent" Insertion.

It is not clear from the description if the derived content text should contain square brackets when an <xref> would be rendered with square brackets in current output formats.

It is not clear if the derived content should include the 'Figure', or 'Table' label when pointing to such objects. When rendering such a reference in the current output formats, the generated text would include the label, but the current text seems to lean towards not making this part of the derived content, which would cause incompatibility with the output of v2 formatters.

The purpose of this is insufficiently explained. If the intention is to use this when generating derived formats, there are problems: If, for instance, the derived format with a <reference> target is set to 'RFC1234', the text inserted in a derived format should have surrounding square brackets; but if the target is a section, it should not. If on the other hand the derived format includes the square brackets when appropriate, the link in a derived format with internal link capability will use the whole of the bracketed string, rather than the more appropriate text within the brackets.

The current implementation works around this by using different formatter code for different cases, which is not good from the viewpoint of using the prepped XML as the archival format. The whole "derivedContent" handling and specification needs a thorough rework, with specification of the intended use of the attribute by formatters.

4.4.20. In Section 5.4.9, <relref> Processing

Why doesn't <relref> have the same format options as <xref>? Surely they must be just as relevant here. But more importantly, <relref> overlaps <xref> so much that it would be better to just add section, relative, and displayFormat to <xref>. Maybe change displayFormat to the earlier proposed 'sectionFormat'.

Recommendation: Deprecate <relref>, and fold the functionality into <xref>.

4.4.21. In Section 5.6.3, <link> Processing

Bullet 4.: Bad grammar s/RFC the form/RFC, in the form/

Bullet 4.: Hmm. The <link rel="convertedFrom" href="draft-...."> should ideally be created automatically, but there is no clear path of how to do that.

Recommendation: Require docName to be set to the draft name, and use that to create this link.

4.4.22. New Section for Index

RFC7998 does not say a word about index, but it seems counter-intuitive not to produce one, given all other prepping being done. What's more, in Section 2.27 of RFC 7991 there's this text:

> "When the prep tool is creating index content, it collects the items in a case-sensitive fashion for both the item and sub-item level."

5. Informative References

[RFC4228] Rousskov, A., "Requirements for an IETF Draft Submission Toolset", RFC 4228, DOI 10.17487/RFC4228, December 2005, <https://www.rfc-editor.org/info/rfc4228>.

[RFC6087] Bierman, A., "Guidelines for Authors and Reviewers of YANG Data Model Documents", RFC 6087, DOI 10.17487/RFC6087, January 2011, <https://www.rfc-editor.org/info/rfc6087>.

[RFC7749] Reschke, J., "The "xml2rfc" Version 2 Vocabulary", RFC 7749, DOI 10.17487/RFC7749, February 2016, <https://www.rfc-editor.org/info/rfc7749>.

[RFC7991] Hoffman, P., "The "xml2rfc" Version 3 Vocabulary", RFC 7991, DOI 10.17487/RFC7991, December 2016, <https://www.rfc-editor.org/info/rfc7991>.

[RFC7992] Hildebrand, J., Ed. and P. Hoffman, "HTML Format for RFCs", RFC 7992, DOI 10.17487/RFC7992, December 2016, <https://www.rfc-editor.org/info/rfc7992>.

[RFC7993] Flanagan, H., "Cascading Style Sheets (CSS) Requirements for RFCs", RFC 7993, DOI 10.17487/RFC7993, December 2016, <https://www.rfc-editor.org/info/rfc7993>.

[RFC7994] Flanagan, H., "Requirements for Plain-Text RFCs", RFC 7994, DOI 10.17487/RFC7994, December 2016, <https://www.rfc-editor.org/info/rfc7994>.

[RFC7995] Hansen, T., Ed., Masinter, L., and M. Hardy, "PDF Format for RFCs", RFC 7995, DOI 10.17487/RFC7995, December 2016, <https://www.rfc-editor.org/info/rfc7995>.

[RFC7996] Brownlee, N., "SVG Drawings for RFCs: SVG 1.2 RFC", RFC 7996, DOI 10.17487/RFC7996, December 2016, <https://www.rfc-editor.org/info/rfc7996>.

[RFC7997] Flanagan, H., Ed., "The Use of Non-ASCII Characters in RFCs", RFC 7997, DOI 10.17487/RFC7997, December 2016, <https://www.rfc-editor.org/info/rfc7997>.

[RFC7998] Hoffman, P. and J. Hildebrand, ""xml2rfc" Version 3 Preparation Tool Description", RFC 7998, DOI 10.17487/RFC7998, December 2016, <https://www.rfc-editor.org/info/rfc7998>.

[XML2RFC] Levkowetz, H., "xml2rfc", 2018, <https://pypi.org/pypi/xml2rfc>.

Author's Address

Henrik Levkowetz
Elf Tools AB
Ollonstigen 8
Sweden

Email: henrik@levkowetz.com