

Internet Engineering Task Force (IETF)
Request for Comments: 7735
Category: Informational
ISSN: 2070-1721

R. Sparks
Oracle
T. Kivinen
INSIDE Secure
January 2016

Tracking Reviews of Documents

Abstract

Several review teams ensure specific types of review are performed on Internet-Drafts as they progress towards becoming RFCs. The tools used by these teams to assign and track reviews would benefit from tighter integration to the Datatracker. This document discusses requirements for improving those tools without disrupting current work flows.

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 Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Not all documents approved by the IESG are a candidate for any level of Internet Standard; see Section 2 of RFC 5741.

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/rfc7735>.

Copyright Notice

Copyright (c) 2016 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
- 2. Overview of Current Workflows 3
- 3. Requirements 5
 - 3.1. Secretariat Focused 5
 - 3.2. Review-Team Secretary Focused 5
 - 3.3. Reviewer Focused 8
 - 3.4. Review Requester and Consumer Focused 10
 - 3.5. Statistics Focused 11
- 4. Security Considerations 12
- 5. Informative References 12
- Appendix A. A Starting Point for Django Models Supporting the
Review Tool 14
- Appendix B. Suggested Features Deferred for Future Work 15
- Acknowledgements 16
- Authors' Addresses 16

1. Introduction

As Internet-Drafts are processed, reviews are requested from several review teams. For example, the General Area Review Team (Gen-ART) and the Security Directorate (SecDir) perform reviews of documents that are in IETF Last Call. Gen-ART always performs a follow-up review when the document is scheduled for an IESG Telechat. SecDir usually performs a follow-up review, but the SecDir secretary may choose not to request that follow-up if any issues identified at Last Call are addressed and there are otherwise no major changes to the document. These teams also perform earlier reviews of documents on demand. There are several other teams that perform similar services, often focusing on specific areas of expertise.

The secretaries of these teams manage a pool of volunteer reviewers. Documents are assigned to reviewers, taking various factors into account. For instance, a reviewer will not be assigned a document for which he is an author or shepherd. Reviewers are given a deadline, usually driven by the end of Last Call or an IESG Telechat date. The reviewer sends each completed review to the team's mailing list and to any other lists that are relevant for the document being reviewed. Often, a thread ensues on one or more of those lists to resolve any issues found in the review.

The secretaries and reviewers from several teams are using a tool developed and maintained by Tero Kivinen. Much of its design predates the modern Datatracker. The application currently keeps its own data store and learns about documents needing review by inspecting Datatracker and tools.ietf.org pages. Most of those pages are easy to parse, but the Last Call pages, in particular, require

some effort. Tighter integration with the Datatracker would simplify the logic used to identify documents that are ready for review, make it simpler for the Datatracker to associate reviews with documents, and allow users to reuse their Datatracker credentials. It would also make it easier to detect other potential review-triggering events, such as a document entering Working Group (WG) Last Call or when an RFC's standard level is being changed without revising the RFC. Tero currently believes this integration is best achieved by a new implementation of the tool. This document captures requirements for that reimplementation, with a focus on the workflows that the new implementation must take care not to disrupt. It also discusses new features, including changes suggested for the existing tool at its issue tracker [art-trac].

For more information about the various review teams, see the following references.

General Area Review Team	[Gen-ART] [RFC6385]
Security Directorate	[SecDir]
Applications Area Directorate	[AppsDir]
Operations Area Directorate	[OPS-dir]
Routing Area Directorate	[RTG-dir]
MIB Doctors	[MIBdoctors]
YANG Doctors	[YANGdoctors]

2. Overview of Current Workflows

This section gives a high-level overview of how the review team secretaries and reviewers use the existing tool. It is not intended to be comprehensive documentation of how review teams operate. Please see the references for those details.

For many teams, the team's secretary periodically (typically once a week) checks the tool for documents it has identified as ready for review. The tool compiles a list from Last Call announcements and IESG Telechat agendas. The secretary creates a set of assignments from this list and enters them into the reviewer pool, choosing the reviewers in roughly a round-robin order. That order can be perturbed by several factors. Reviewers have different levels of availability. Some are willing to review multiple documents a month. Others may only be willing to review a document every other month. The assignment process takes exceptional conditions such as reviewer vacations into account. Furthermore, secretaries are careful not to assign a document to a reviewer that is an author, shepherd, responsible WG chair, or has some other already existing association with the document. The preference is to get a reviewer with a fresh

perspective. The secretary may discover reasons to change assignments while going through the list of documents. In order to not cause a reviewer to make a false start on a review, the secretaries complete the full list of assignments before sending notifications to anyone. This assignment process can take several minutes and it is possible for new Last Calls to be issued while the secretary is making assignments. The secretary typically checks to see if new documents are ready for review just before issuing the assignments and updates the assignments if necessary.

Some teams operate in more of a review-on-demand model. The Routing Area Directorate (RTG-dir), for example, primarily initiates reviews at the request of a Routing AD. They may also start an early review at the request of a WG chair. In either case, the reviewers are chosen manually from the pool of available reviewers driven by context rather than using a round-robin ordering.

The issued assignments are either sent to the review team's email list or are emailed directly to the assigned reviewer. The assignments are reflected in the tool. For those teams handling different types of reviews (Last Call vs. Telechat, for example), the secretary typically processes the documents for each type of review separately, and potentially with different assignment criteria. In Gen-ART, for example, the Last Call reviewer for a document will almost always get the follow-up Telechat review assignment. Similarly, SecDir assigns any re-reviews of a document to the same reviewer. Other teams may choose to assign a different reviewer.

Reviewers discover their assignments through email or by looking at their queue in the tool. The secretaries for some teams (such as the OPS-dir and RTG-dir) insulate their team members from using the tool directly. These reviewers only work through the review team's email list or through direct email. On teams that have the reviewers use the tool directly, most reviewers only check the tool when they see they have an assignment via the team's email list. A reviewer has the opportunity to reject the assignment for any reason. While the tool provides a way to reject assignments, reviewers typically use email to coordinate rejections with the team secretary. The secretary will find another volunteer for any rejected assignments. The reviewer can indicate that the assignment is accepted in the tool before starting the review, but this feature is rarely used.

The reviewer sends a completed review to the team's email list or secretary, as well as any other lists relevant to the review, and usually the draft's primary email alias. For instance, many Last Call reviews are also sent to the IETF general list. The teams typically have a template format for the review. Those templates usually start with a summary of the conclusion of the review.

Typical summaries are "ready for publication" or "on the right track but has open issues". The reviewer (or in the case of teams that insulate their reviewers, the secretary) uses the tool to indicate that the review is complete, provides the summary, and has an opportunity to provide a link to the review in the archives. Note, however, that having to wait for the document to appear in the archive to know the link to paste into the tool is a significant enough impedence that this link is often not provided by the reviewer. The SecDir secretary manually collects these links from the team's email list and adds them to the tool.

Occasionally, a document is revised between when a review assignment is made and when the reviewer starts the review. Different teams can have different policies about whether the reviewer should review the assigned version or the current version.

3. Requirements

3.1. Secretariat Focused

- o The Secretariat must be able to configure secretaries and reviewers for review teams (by managing Role records).
- o The Secretariat must be able to perform any secretary action on behalf of a review team secretary (and thus, must be able to perform any reviewer action on behalf of the reviewer).

3.2. Review-Team Secretary Focused

- o A secretary must be able to see what documents are ready for a review of a given type (such as a Last Call review).
- o A secretary must be able to assign reviews for documents that may not have been automatically identified as ready for a review of a given type. (In addition to being the primary assignment method for teams that only initiate reviews on demand, this allows the secretary to work around errors and handle special cases, including early review requests.)
- o A secretary must be able to work on and issue a set of assignments as an atomic unit. No assignment should be issued until the secretary declares the set of assignments complete.
- o The tool must support teams that have multiple secretaries. The tool should warn secretaries that are simultaneously working on assignments and protect against conflicting assignments being made.

- o It must be easy for the secretary to discover that more documents have become ready for review while working on an assignment set.
- o The tool should make preparing the assignment email to the team's email list easy. For instance, the tool could prepare the message, give the secretary an opportunity to edit it, and handle sending it to the team's email list.
- o It must be possible for a secretary to indicate that the review team will not provide a review for a document (or a given version of a document). This indication should be taken into account when presenting the documents that are ready for a review of a given type. This will also make it possible to show on a document's page that no review is expected from this team.
- o A secretary must be able to easily see who the next available reviewers are, in order.
- o A secretary must be able to edit a reviewer's availability, both in terms of frequency, not-available-until-date, and skip-next-n-assignments. (See the description of these settings in Section 3.3.)
- o The tool should make it easy for the secretary to see any team members that have requested to review a given document when it becomes available for review.
- o The tool should make it easy for the secretary to identify that a reviewer is already involved with a document. The current tool allows the secretary to provide a regular expression to match against the document name. If the expression matches, the document is not available for assignment to this reviewer. For example, Tero will not be assigned documents matching '^draft-(kivinen|ietf-tcpinc)-.*\$'. The tool should also take any roles, such as document shepherd, that the Datatracker knows about into consideration.
- o The tool should make it easy for the secretary to see key features of a document ready for assignment, such as its length, its authors, the group and area it is associated with, its title and abstract, its states (such as IESG or WG states), and any other personnel (such as the shepherd and reviewers already assigned from other teams) involved in the draft.
- o The tool must make it easy for the secretary to detect and process re-review requests on the same version of a document (such as when a document has an additional Last Call only to deal with new IPR information).

- o Common operations to groups of documents should be easy for the secretary to process as a group with a minimum amount of interaction with the tool. For instance, it should be possible to process all of the documents described by the immediately preceding bullet with one action. Similarly, for teams that assign re-reviews to the same reviewer, issuing all re-review requests should be a simple action.
- o A secretary must be able to see which reviewers have outstanding assignments.
- o The tool must make it easy for the secretary to see the result of previous reviews from this team for a given document. In SecDir, for example, if the request is for a revision that has only minor differences and the previous review result was "Ready", a new assignment will not be made. If the given document replaces one or more other prior documents, the tool must make it easy for the secretary to see the results of previous reviews of the replaced documents.
- o The tool must make it easy for the secretary to see the result of previous reviews from this team for all documents across configurable, recent periods of time (such as the last 12 months). A secretary of the RTG-dir, for example, would use this result to aid in the manual selection of the next reviewer.
- o The tools must make it easy for the secretary to see the recent performance of a reviewer while making an assignment (see Section 3.5). This allows the secretary to detect overburdened or unresponsive volunteers earlier in the process.
- o A secretary must be able to configure the tool to remind them to follow up when actions are due. (For instance, a secretary could receive email when a review is about to become overdue.)
- o A secretary must be able to assign multiple reviewers to a given draft at any time. In particular, a secretary must be able to assign an additional reviewer when an original reviewer indicates their review is likely to be only partially complete.
- o A secretary must be able to withdraw a review assignment.
- o A secretary must be able to perform any reviewer action on behalf of the reviewer.
- o A secretary must be able to configure the review team's set of reviewers (by managing Role records for the team).

- o Information about a reviewer must not be lost when a reviewer is removed from a team. (Frequently, reviewers come back to teams later.)
- o A secretary must be able to delegate secretary capabilities in the tool (similar to how a working group chair can assign a Delegate). This allows review teams to self-manage secretary vacations.

3.3. Reviewer Focused

- o A reviewer must be able to indicate availability, both in frequency of reviews and as "not available until this date". The current tool speaks of frequency in these terms:
 - Assign at maximum one new review per week
 - Assign at maximum one new review per fortnight
 - Assign at maximum one new review per month
 - Assign at maximum one new review per two months
 - Assign at maximum one new review per quarter
- o Reviewers must be able to indicate hiatus periods. Each period may be either "soft" or "hard".
 - A hiatus must have a start date. It may have an end date or it may be indefinite.
 - During a hiatus, the reviewer will not be included in the normal review rotation. When a provided end date is reached, the reviewer will automatically be included in the rotation in their usual order.
 - During a "soft" hiatus, the reviewer must not be assigned new reviews but is expected to complete existing assignments and do follow-up reviews.
 - During a "hard" hiatus, the reviewer must not be assigned any new reviews and the secretary must be prompted to reassign any outstanding or follow-up reviews.
- o Reviewers must be able to indicate that they should be skipped the next "n" times they would normally have received an assignment.

- o Reviewers must be able to indicate that they are transitioning to inactive and provide a date for the end of the transition period. During this transition time, the reviewer must not be assigned new reviews but is expected to complete outstanding assignments and follow-up reviews. At the end of the transition period, the secretary must be prompted to reassign any outstanding or follow-up reviews. (This allows review-team members that are taking on AD responsibility to transition gracefully to an inactive state for the team.)
- o Both the reviewer and the secretary will be notified by email of any modifications to a reviewer's availability.
- o A reviewer must be able to easily discover new review assignments. (The tool might send email directly to an assigned reviewer in addition to sending the set of assignments to the team's email list. The tool might also use the Django Message framework to let a reviewer that's logged into the Datatracker know a new review assignment has been made.)
- o Reviewers must be able to see their current set of outstanding assignments, completed assignments, and rejected assignments. The presentation of those sets should either be separate or, if combined, the sets should be visually distinct.
- o A reviewer should be able to request to review a particular document. The draft may be in any state: available and unassigned; already assigned to another reviewer; or not yet available.
- o A reviewer must be able to reject a review assignment, optionally providing the secretary with an explanation for the rejection. The tool will notify the secretary of the rejection by email.
- o A reviewer must be able to indicate that they have accepted and are working on an assignment.
- o A reviewer must be able to indicate that a review is only partially completed and ask the secretary to assign an additional reviewer. The tool will notify the secretary of this condition by email.
- o It should be possible for a reviewer to reject or accept a review either by using the tool's web interface or by replying to the review assignment email.
- o It must be easy for a reviewer to see when each assigned review is due.

- o A reviewer must be able to configure the tool to remind them when actions are due. (For instance, a reviewer could receive email when a review is about to become overdue).
- o A reviewer must be able to indicate that a review is complete, capturing where the review is in the archives and the high-level, review-result summary.
- o It must be possible for a reviewer to clearly indicate which version of a document was reviewed. Documents are sometimes revised between when a review was assigned and when it is due. The tool should note the current version of the document and highlight when the review is not for the current version.
- o It must be easy for a reviewer to submit a completed review.
 - The current workflow, where the reviewer sends email to the team's email list (possibly copying other lists) and then indicates where to find that review, must continue to be supported. The tool should make it easier to capture the link to review in the team's email list archives (perhaps by suggesting links based on a search into the archives).
 - The tool should allow the reviewer to enter the review into the tool via a web form (either as directly provided text or through a file-upload mechanism). The tool will ensure the review is posted to the appropriate lists and will construct the links to those posts in the archives.
 - The tool could also allow the reviewer to submit the review to the tool by email (perhaps by replying to the assignment). The tool would then ensure the review is posted to the appropriate lists.

3.4. Review Requester and Consumer Focused

- o It should be easy for an AD or group chair to request any type of review, but particularly an early review, from a review team.
- o It should be possible for that person to withdraw a review request.
- o It must be easy to find all reviews of a document when looking at the document's main page in the Datatracker. The reference to the review must make it easy to see any responses to the review on the email lists it was sent to. If a document "replaces" one or more other documents, reviews of the replaced documents should be included in the results.

- o It must be easy to find all reviews of a document when looking at search result pages and other lists of documents, such as the documents on an IESG Telechat agenda.

3.5. Statistics Focused

- o It must be easy to see the following across all teams, a given team, or a given reviewer, and independently across all time or across configurable recent periods of time:
 - how many reviews have been completed
 - how many reviews are in progress
 - how many in progress reviews are late
 - how many completed reviews were late
 - how many reviews were not completed at all
 - average time to complete reviews (from assignment to completion)
- o It must be easy to see the following for all teams, for a given team, or for a given reviewer, across all time or across configurable recent periods:
 - total counts of reviews in each review state (done, rejected, etc.)
 - total counts of completed reviews by result (ready, ready with nits, etc.)
- o The above statistics should also be calculated reflecting the size of the documents being reviewed (such as using the number of pages or words in the documents).
- o Where applicable, statistics should take reviewer hiatus periods into account.
- o Access to the above statistics must be easy to configure. Access will be initially limited as follows:
 - The Secretariat and ADs can see any statistic.
 - A team secretary can see any statistics for that team.

- A reviewer can see any team aggregate statistics or their own reviewer-specific statistics.
- o Where possible, the above statistics should be visible as a time-series graph.
- o The implementation should anticipate future enhancements that would allow ADs to indicate their position was informed by a given review. Such enhancements would allow reporting correlations between reviews and documents that receive one or more "discusses". However, implementing these enhancements is not part of the current project.

4. Security Considerations

This document discusses requirements for tools that assist review teams. These requirements do not affect the security of the Internet in any significant fashion. The tools themselves have authentication and authorization considerations (team secretaries will be able to do different things than reviewers).

5. Informative References

- [AppsDir] IETF, "Applications Area Directorate Review Process", <<http://trac.tools.ietf.org/area/app/trac/wiki/AppsDirReview>>.
- [art-trac] IETF, "Area Review Team Tool: {1} Active Tickets", <<https://wiki.tools.ietf.org/tools/art/trac/report/1>>.
- [Gen-ART] IETF, "General Area: General Area Review Team (GEN-ART) Guidelines", <<http://trac.tools.ietf.org/area/gen/trac/wiki>>.
- [MIBdoctors] IETF, "MIB Doctors", <<http://www.ietf.org/iesg/directorate/mib-doctors.html>>.
- [OPS-dir] IETF, "OPS Directorate (OPS-DIR)", <<https://svn.tools.ietf.org/area/ops/trac/wiki/Directorates>>.
- [RFC6385] Barnes, M., Doria, A., Alvestrand, H., and B. Carpenter, "General Area Review Team (Gen-ART) Experiences", RFC 6385, DOI 10.17487/RFC6385, October 2011, <<http://www.rfc-editor.org/info/rfc6385>>.

- [RTG-dir] IETF, "Routing Area Directorate Wiki Page",
<<http://trac.tools.ietf.org/area/rtg/trac/wiki/RtgDir>>.
- [SecDir] IETF, "Security Directorate",
<<http://trac.tools.ietf.org/area/sec/trac/wiki/SecurityDirectorate>>.
- [YANGdoctors]
IETF, "YANG Doctors",
<<http://www.ietf.org/iesg/directorate/yang-doctors.html>>.

Appendix A. A Starting Point for Django Models Supporting the Review Tool

```
from django.db import models
from ietf.doc.models import Document
from ietf.person.models import Email
from ietf.group.models import Group, Role

from ietf.name.models import NameModel

class ReviewRequestStateName(NameModel):
    """ Requested, Accepted, Rejected, Withdrawn, Overtaken By Events,
        No Response , Completed """

class ReviewTypeName(NameModel):
    """ Early Review, Last Call, Telechat """

class ReviewResultName(NameModel):
    """Almost ready, Has issues, Has nits, Not Ready,
        On the right track, Ready, Ready with issues,
        Ready with nits, Serious Issues"""

class Reviewer(models.Model):
    """
    These records associate reviewers with review teams and keep track
    of admin data associated with the reviewer in the particular team.
    There will be one record for each combination of reviewer and team.
    """
    role = models.ForeignKey(Role)
    frequency = models.IntegerField(help_text=
        "Can review every N days")
    available = models.DateTimeField(blank=True,null=True, help_text=
        "When will this reviewer be available again")
    filter_re = models.CharField(blank=True)
    skip_next = models.IntegerField(help_text=
        "Skip the next N review assignments")
```

```

class ReviewResultSet(models.Model):
    """
    This table provides a way to point out a set of ReviewResultName
    entries that are valid for a given team in order to be able to
    limit the result choices that can be set for a given review as a
    function of which team it is related to.
    """
    team          = models.ForeignKey(Group)
    valid         = models.ManyToManyField(ReviewResultName)

class ReviewRequest(models.Model):
    """
    There should be one ReviewRequest entered for each combination of
    document, rev, and reviewer.
    """
    # Fields filled in on the initial record creation:
    time          = models.DateTimeField(auto_now_add=True)
    type          = models.ReviewTypeName()
    doc           = models.ForeignKey(Document,
                                       related_name='review_request_set')
    team         = models.ForeignKey(Group)
    deadline     = models.DateTimeField()
    requested_rev = models.CharField(verbose_name="requested_revision",
                                    max_length=16, blank=True)
    state        = models.ForeignKey(ReviewRequestStateName)
    # Fields filled in as reviewer is assigned, and as the review
    # is uploaded
    reviewer     = models.ForeignKey(Reviewer, null=True, blank=True)
    review       = models.OneToOneField(Document, null=True,
                                       blank=True)
    reviewed_rev = models.CharField(verbose_name="reviewed_revision",
                                    max_length=16, blank=True)
    result       = models.ForeignKey(ReviewResultName)

```

Appendix B. Suggested Features Deferred for Future Work

Brian Carpenter suggested a set of author/editor-focused requirements that were deferred for another iteration of improvement. These include providing a way for the editors to acknowledge receipt of the review, potentially tracking the email conversation between the reviewer and document editor, and indicating which review topics the editor believes a new revision addresses.

Acknowledgements

Tero Kivinen and Henrik Levkowitz were instrumental in forming this set of requirements and in developing the initial Django models in the appendix.

The following people provided reviews of this document: David Black, Deborah Brungard, Brian Carpenter, Elwyn Davies, Stephen Farrell, Joel Halpern, Jonathan Hardwick, Russ Housley, Barry Leiba, Jean Mahoney, Randy Presuhn, Gunter Van De Velde, and Martin Vigoureux.

Authors' Addresses

Robert Sparks
Oracle
7460 Warren Parkway
Suite 300
Frisco, Texas 75034
United States

Email: rjsparks@nostrum.com

Tero Kivinen
INSIDE Secure
Eerikinkatu 28
Helsinki FI-00180
Finland

Email: kivinen@iki.fi