

Internet Engineering Task Force (IETF)
Request for Comments: 6293
Category: Informational
ISSN: 2070-1721

P. Hoffman
VPN Consortium
June 2011

Requirements for Internet-Draft Tracking by
the IETF Community in the Datatracker

Abstract

The document gives a set of requirements for extending the IETF Datatracker to give individual IETF community members, including the IETF leadership, easy methods for tracking the progress of the Internet-Drafts and RFCs of interest to them.

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/rfc6293>.

Copyright Notice

Copyright (c) 2011 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	3
1.1.	Usage Scenarios	3
1.2.	Context for This Document	4
1.3.	Definitions Used in This Document	5
1.4.	Expected User Interactions	6
2.	Requirements for Tools Features	6
2.1.	Lists	6
2.1.1.	Requirement: Lists of I-Ds and RFCs can be large	6
2.1.2.	Requirement: Every Datatracker user can create one list	7
2.1.3.	Requirement: Read-only views of private lists can be made visible to others	7
2.1.4.	Requirement: The Datatracker must support optional publicly-readable lists for WGs and Area Directors	7
2.1.5.	Requirement: Specifying the I-Ds and RFCs that are in a list must be simple	8
2.1.6.	Requirement: Adding groups of I-Ds to a list by attribute must be simple	8
2.1.7.	Requirement: Private information must not be exposed in lists	9
2.2.	Notifications	9
2.2.1.	Requirement: Users can be notified when an I-D changes status	9
2.2.2.	Requirement: Every list has Atom feeds associated with it	9
2.2.3.	Requirement: Every list has mail streams associated with it	10
2.2.4.	Requirement: Notifications need to specify which list caused the notification	10
2.3.	Display in the Datatracker	10
2.3.1.	Requirement: Users can define their Datatracker document view	10
2.3.2.	Requirement: Users can choose which attributes to display	11
2.3.3.	Requirement: Users can flag I-Ds with dates in the future	12
2.3.4.	Requirement: Users can specify highlighting of I-Ds and RFCs with recent changes	12
2.4.	File Output	12
2.4.1.	Requirement: Users can get their current list as a single file	12
3.	Security Considerations	13
4.	Acknowledgements	13
5.	Informative References	13

Appendix A. Possible Tracking of Other Data	15
A.1. Tracking WG Charter Changes	15
A.2. Tracking IANA Registry Changes	15
A.3. Tracking Changes in the Liason Statement Directory	15
A.4. Tracking Changes in Documents Outside the IETF Sphere	15
A.5. Tracking Additions to the IPR Statement Repository	16
Appendix B. Ideas that Might Be Implemented Later	16

1. Introduction

The IETF Datatracker is used by many IETF community members to find the status of Internet-Drafts (I-Ds) and RFCs, and view I-Ds and RFCs that meet particular criteria. The current Datatracker, found at <https://datatracker.ietf.org/>, allows anyone to search for active I-Ds and RFCs, and get a list matching the given criteria. (The Datatracker also allows for expired I-Ds, but those are not relevant to this discussion.)

Users can search in the Datatracker by the filename of the I-D, words in the I-D title, I-D author list, associated Working Group (WG), IETF area, the responsible Area Director (AD), or IESG status. They can search for RFCs by number or words in the title. The returned list of I-Ds and/or RFCs includes six columns: filename or RFC number, the document's title, the date it was published, its status in the IETF or RFC process, IPR statements, and the responsible AD (if any).

Instead of using the search capability of the Datatracker to manually find I-Ds and RFCs of interest, users might want to create a list of I-Ds that they normally follow. Some users will want to keep their list to themselves, but others will want to allow others to view their list.

Different users in the IETF community will have different ways that they want to get information on I-D and RFC updates and status. Many users will want to be notified immediately, such as through an Atom feed (see [RFC4287]) or automatically-generated email. Many users will want to only find out about updates when they go to a Web page. Many users might want to get the data for a list as input to other tools. And, of course, some users will want all three. All of these assist users in tracking I-Ds through their lifecycle.

1.1. Usage Scenarios

The main motivation for these proposed changes to the Datatracker is to allow a variety of potential users to be able to track I-Ds and RFCs, and thus be better able to see when important events happen. A few examples include:

- o A WG chair might want to keep a list of all the I-Ds from other WGs that relate to active I-Ds in his or her WG.
- o That same WG chair might want to help WG members be able to follow the same I-Ds that he or she is following.
- o Someone who cares about an established topic such as the DNS may want to follow the various I-Ds that might make changes to the DNS, as well as be aware if any of the DNS RFCs are later updated and/or have errata posted against them. This would include not only I-Ds that are in the many WGs that directly are changing the DNS (DNSEXT, DNSOP, BEHAVE, and so on), but also individual submissions, IAB I-Ds, IRTF I-Ds, and Independent submissions.
- o Developers who are not active in the IETF process might want to lightly follow I-Ds and RFCs on a particular topic to watch for things that might affect their implementations.
- o An IETF "regular" might want to follow parts of the process by focusing on all the I-Ds that are being shepherded by a particular Area Director.

1.2. Context for This Document

This document describes the requirements for extending the Datatracker for such capabilities. When complete, this document may be used to issue an RFP for the design and development of these enhancements to the Datatracker.

Some of the requirements in this document are listed as "later requirements". It is expected that items listed in this document would be part of the initial RFP because they provide the highest benefit to the community; the later requirements might be part of a later RFP.

The initial general requirements that led to the specific requirements this document described tools that include:

- o the ability to create one or more (possibly large) lists of I-Ds that community members want to follow
- o the ability to get notifications when particular I-Ds from a list change state
- o the ability to see all of the state changes that have occurred on all the I-Ds in a list over a specified range of dates

- o the ability to set the granularity of the changes (such as "every change", "just approvals and publication", and so on)
- o the ability to organize views of a list in many fashions that would be useful to different types of community members
- o the ability to share and merge lists with other community members

Note that [RFC2026] describes the process that I-Ds go through before they either become RFCs or are abandoned. The Datatracker does not control this process: instead, it simply reports on the current state of each I-D as it goes through the process.

1.3. Definitions Used in This Document

A "user" is an individual person who is a member of the IETF community.

A "list" is an unordered set of RFCs, I-Ds, and groups of I-Ds. Lists are specified by users. In some cases, the authors are role-based, such as a WG chair being the specifier of the list associated with that WG.

An "attribute" is a feature of an I-D or RFC, such as its filename or RFC number, its current state in the IETF or RFC process, and so on. Attributes are usually displayed as columns in the Datatracker.

A "row" is a set of attributes about a single I-D or RFC that is displayed in the Datatracker.

A "significant change in status" is all approvals and disposition of an I-D. Assuming that the changes to the Datatracker specified in [RFC6174], [RFC6175] and [ALTSTREAMS] are made, "all approvals" means the following:

- o IETF stream: the WG states "Adopted by a WG", "In WG Last Call", "WG Consensus: Waiting for Write-up", "Parked WG document", and "Dead WG document"; the IESG states "Publication Requested", "In Last Call", "IESG Evaluation", and "Sent to the RFC Editor"
- o IAB stream: "Active IAB Document", "Community Review", and "Sent to the RFC Editor"
- o IRTF stream: "Active RG Document", "In RG Last Call", "Awaiting IRSG Reviews", "In IESG Review", "Sent to the RFC Editor", and "Document on Hold Based On IESG Request"

- o ISE stream: "Submission Received", "In ISE Review", "In IESG Review", "Sent to the RFC Editor", and "Document on Hold Based On IESG Request"
- o All streams: in addition to the above, the disposition states "Approved", "RFC Published", and "Dead" are also included

An "update to an RFC" is the announcement of a newer RFC that updates or obsoletes the base RFC, an in-place change to the RFC's maturity level, the RFC's status being changed to historic, or an announcement of an errata posted for the base RFC.

1.4. Expected User Interactions

When a user wants to follow a group of I-Ds and/or RFCs, he or she goes to the Datatracker and creates a new list. The requirements for lists are given in Section 2.1. After a list is created, the user has three ways that he or she might see when I-Ds and/or RFCs in the list are updated:

- o By going to the Datatracker page for the list (see Section 2.3)
- o By subscribing to the Atom feed for the list (see Section 2.2.2) in a feed reader that automatically fetches updates
- o By subscribing to the mail stream for the list (see Section 2.2.3) and reading the mail stream in their mail reader

2. Requirements for Tools Features

This section defines the requirements for the tool described earlier in this document. The eventual tool, if implemented, may have more features than are listed here; however, before this document is finished, it should contain as many requirements as possible upon which the IETF community can agree.

2.1. Lists

2.1.1. Requirement: Lists of I-Ds and RFCs can be large

An active IETF participant might want to follow the status of hundreds of I-Ds and dozens of RFCs; for example, some ADs have 100 I-Ds in their area. Additionally, they may also want to follow I-Ds outside their area that affect documents in their area.

2.1.2. Requirement: Every Datatracker user can create one list

When a user gets a Datatracker account, that account comes with an empty list pre-defined. The list can normally be modified only by the owner of the account, although the Secretariat can also modify the list as part of its support role for the Datatracker. Each Datatracker user is restricted to having one list.

In order for this requirement to be met, it must be easy for any community member to get a Datatracker account. Account setup must not involve any direct action on the part of the Secretariat. However, the Secretariat will be responsible for support of Datatracker accounts (lost passwords, odd interactions, and so on), so this addition of more Datatracker accounts will potentially increase the amount of work the Secretariat must do.

The only person who can edit the contents of a private list is the person who knows the password to the account with which the list is associated.

2.1.3. Requirement: Read-only views of private lists can be made visible to others

Some users will want to make available a read-only view of their list. Each private list will have a URL that leads to the Datatracker view of the list; that URL must be able to be shared without giving others the ability to edit the list. Similarly, the Atom feed associated with a private list must be able to be shared without giving others the ability to edit the list.

2.1.4. Requirement: The Datatracker must support optional publicly-readable lists for WGs and Area Directors

It is common in the IETF for users to follow the work of an entire WG, not just single I-Ds and RFCs within a WG. It is also very common that some work that is related to a WG happens outside the WG, either in other WGs or as individual efforts. Many WG chairs monitor this outside-the-WG activity for various reasons.

A smaller number of community members follow an entire Area's worth of topics. Again, these topics often happen within the WGs of an area, but not always; for example, some topics related to the Security Area happen in WGs in the Applications Area.

Because of this, it would be useful for community members to be able to find a list that corresponds to the WGs or Areas in which they are interested. The WG lists could be maintained by the WG chairs; the Area lists would likely be maintained by the ADs. Note that such

lists are not mandatory; for example, a WG chair might not choose to maintain such a list for a WG whose topic is extremely broad.

Both Working Group chairs and Area Directors currently already have Datatracker accounts, so fulfilling this requirement only involves associating those accounts with the role that controls the list.

2.1.5. Requirement: Specifying the I-Ds and RFCs that are in a list must be simple

When a user creates a new list, it must be easy to add single I-Ds and RFCs to the list. This could be done using the Datatracker's current search facility, and simply adding an "add to list" option to the display of searched-for I-Ds. Further, when editing an existing list, it must be easy to add additional I-Ds and RFCs, and it must be easy to remove I-Ds and RFCs from a list.

2.1.6. Requirement: Adding groups of I-Ds to a list by attribute must be simple

I-Ds have many attributes, and some users might want to follow all of the I-Ds that have a particular attribute. Some, but not all, attributes have values that make sense in specifying lists. It should be easy to add each of the following attributes when adding to or editing a list:

- o All I-Ds associated with an particular WG
- o All I-Ds associated with all WGs in an particular Area
- o All I-Ds with a particular responsible AD
- o All I-Ds with a particular author
- o All I-Ds with a particular document shepherd
- o All I-Ds that have a reference to a particular RFC
- o All I-Ds that have a reference to a particular I-D
- o All I-Ds that are referenced by a particular RFC
- o All I-Ds that are referenced by a particular I-D
- o All I-Ds that contain a particular text string

These attributes are dynamic, and thus the list of I-Ds that have a particular attribute will change after the user adds that attribute to a list. The Datatracker should update lists with dynamic attributes as often as is sensible for the server environment, such as once an hour or more.

Note that some of these attributes are based on heuristics derived by programs that parse I-Ds, and are therefore inherently not completely reliable.

2.1.7. Requirement: Private information must not be exposed in lists

Any private information in the Datatracker must be excluded from any displays of the lists or mail streams. This private information includes private notes in the IESG balloting for an I-D, and probably other data that currently is restricted to being seen by certain members of the IETF leadership.

2.2. Notifications

2.2.1. Requirement: Users can be notified when an I-D changes status

Some users do not want to go to the Datatracker's display page to find out when an I-D or RFC has been updated. Instead, they want to be notified immediately after the change. The Datatracker needs to support this type of immediate notification, where "immediate" means within an hour of a change to any I-D or RFC in the list. This requirement can be met with Atom feeds and mail streams, as described in the next two sections.

The Datatracker might create a generic "notifications engine" that can be used to generate the Atom feeds and mail streams. This engine can then be used to later add other notification types, such as a Jabber feed.

2.2.2. Requirement: Every list has Atom feeds associated with it

The list will have two Atom feeds that are generated from the changes to the list: one for every change in status and another for significant change of status. Each Atom feed will have a stable URL that can be used by feed readers.

Many IETF users are already using Atom feeds created by the IETF Tools Team for single I-Ds. Using the new feeds for lists described here will allow them to have better selection capabilities to reduce the number of feeds they need to follow.

2.2.3. Requirement: Every list has mail streams associated with it

A user can subscribe to two mail streams that are generated from the changes to the list: one for every change in status, and another for significant change of status.

Note that the mail streams are for each change; they are not batched (such as one message per day). Users who want less frequent but batched notifications need to use the Atom feeds instead of the mail streams.

2.2.4. Requirement: Notifications need to specify which list caused the notification

Users might have feeds and/or subscriptions to multiple lists. In order to disambiguate duplicate notifications from multiple lists, the body of the message in the Atom feed or mail stream needs to say which list generated the notification. (Ideally, a user who wants notifications will make one list based on multiple lists, but if they subscribe to multiple lists, this requirement will at least suggest to them that they want to limit their overlapping subscriptions.)

2.3. Display in the Datatracker

2.3.1. Requirement: Users can define their Datatracker document view

There are many ways that a user might want to see the Datatracker's HTML view of a list. For example, a user might want the view displayed in alphabetical order by the I-Ds' filenames and RFC numbers, but after the user is off the net for a week, he or she might want the view displayed in order of changes of status so that those I-Ds and RFCs changed recently appear at the top.

The default is to list I-Ds in alphabetical order by I-D filename, with RFCs at the end. When displaying a list, the Datatracker should allow easy sorting of the I-Ds with the following collation orders:

- o Alphabetical by I-D filename and RFC number
- o Alphabetical by document title
- o Alphabetical by associated WG
- o Date of publication of current version of the document
- o Date of most recent change of status of any type
- o Date of most recent significant change of status

In displays, a particular I-D or RFC should only be included once; for example, if someone manually adds draft-ietf-cuteacronym-sometopic to his list and also specifies that all I-Ds from the "cuteacronym" WG are included in the list, that I-D should only appear once in the display. The column saying which included list(s) contain this I-D helps alleviate this loss of information.

The user might also want to group the I-Ds using the groupings in the list, such as "all I-Ds from this WG" and "all I-Ds that contain this word in the title".

The Datatracker should save the last-chosen sorting for display with the definition of the list.

2.3.2. Requirement: Users can choose which attributes to display

There are many attributes that might be displayed, and different users will have different information that they want to see. Also, users will have different display technologies: someone might normally use a Web browser on a large screen, but at other times use the browser on their phone.

Choosing which attributes should be displayed should be simple for the user. The Datatracker should save the last-chosen set of attributes for display with the definition of the list. The default is to display the I-D filename or RFC number, document title, date of current I-D or RFC publication date, status in the RFC queue or RFC process, the associated stream (IETF WG, IRTF RG, IAB, or ISE), whether it was changed within the last 7 days, and included list(s) that contain this I-D.

The Datatracker should support display of the following attributes:

- o I-D filename
- o I-D title
- o Date of current I-D
- o Status in the IETF process
- o Associated WG or RG
- o Associated AD, if any
- o Changed within the last 1 day

- o Changed within the last 2 days
- o Changed within the last 7 days

There is some leeway for how the Datatracker might display these attributes. For example, the "changed within" attributes might be shown with a check mark or a colored box.

2.3.3. Requirement: Users can flag I-Ds with dates in the future

When tracking I-Ds, some users want to be able to say "tell me if this I-D has not changed state by a particular date" such as when an I-D is starting a two-week last call or an I-D author has promised a new version by the end of the week. This feature gives the user a "dashboard" style capability.

For each I-D, the user should be able to set a marker date by which an update is expected. The Datatracker display will provide a visual indication if the marker date has passed but no change in status has occurred. It must be very easy for the user to remove these update-expected markers.

2.3.4. Requirement: Users can specify highlighting of I-Ds and RFCs with recent changes

The Datatracker cannot easily keep track of when a user last looked at the page for a particular list. Thus, it instead needs to let a user say which range of dates they are most interested in. To that end, the user needs to be able to easily specify the amount of time they consider recent, either as "the past nnn hours", "the past nnn days", or "since this particular date".

2.4. File Output

2.4.1. Requirement: Users can get their current list as a single file

Some users have their own tools for displaying and otherwise processing lists of I-Ds and RFCs. To make this easier, users should be able to get a machine-parsable file that has a well-known format and syntax that contains all the data that was used to create the current display. The order of the records in the file is not important because it is assumed that the user's program will sort the results themselves. All attributes will be included because it is assumed that the user's programs will only deal with the ones the user cares about.

When a list is marshaled into a data file, each record in the file format represents a single I-D or RFC. In a file, a particular I-D or RFC is only included once; for example, if someone manually adds draft-ietf-cuteacronym-sometopic to his list and also specifies that all I-Ds from the "cuteacronym" WG are included in the list, that I-D only appears once.

This feature will allow anyone to create mash-ups of their own and create their own Web sites based on the IETF data. This is significantly easier than adding features to the Datatracker, and is able to cater to narrow audiences. The format of this file has yet to be determined.

3. Security Considerations

A tool for tracking the status of I-Ds and RFCs can affect the privacy of its users. Someone could possibly determine relevant information about a user if they knew what that user was tracking.

Web applications, particularly those that store data on a Web server, are a common source of security issues such as cross-site scripting attacks. The tool described in this document might also use access control for lists, and access control and authentication also cause security issues if not implemented properly.

4. Acknowledgements

Ideas used in this document were contributed by Scott Bradner, Leslie Daigle, Spencer Dawkins, Aaron Falk, Russ Housley, Tero Kivinen, Barry Leiba, John Levine, Henrik Levkowetz, Kurtis Lindqvist, Andy Malis, Ray Pelletier, Blake Ramsdell, Julian Reschke, Jim Schaad, Yaron Sheffer, Robert Sparks, Andrew Sullivan, and Sean Turner.

5. Informative References

- [ALTSTREAMS] Hoffman, P., "Data Tracker States and Annotations for the IAB, IRTF, and Independent Submission Streams", Work in Progress, May 2011.
- [RFC2026] Bradner, S., "The Internet Standards Process -- Revision 3", BCP 9, RFC 2026, October 1996.
- [RFC4287] Nottingham, M., Ed. and R. Sayre, Ed., "The Atom Syndication Format", RFC 4287, December 2005.
- [RFC6174] Juskevicius, E., "Definition of IETF Working Group Document States", RFC 6174, March 2011.

- [RFC6175] Juskevicius, E., "Requirements to Extend the Datatracker for IETF Working Group Chairs and Authors", RFC 6175, March 2011.
- [RFC6292] Hoffman, P., "Requirements for a Working Group Charter Tool", RFC 6292, June 2011.

Appendix A. Possible Tracking of Other Data

It is not at all clear if any of these will be a requirement, a later requirement, or a non-requirement. Further, even if one or more of these non-I-D items is made a requirement, it is not clear whether they will be included in the same lists with I-Ds. That is, if tracking IANA registry changes are considered a requirement, it is not clear whether a user would include the registries in a list that also contains I-Ds, or whether they would need to create two lists, one for I-Ds and one for IANA registries.

A.1. Tracking WG Charter Changes

It will soon be easier to track changes in WG charters and milestones; see [RFC6292] for more information. Someone subscribing to the mail stream for a WG would be able to see each of these changes. With the expected changes, the Datatracker would be able to update WGs in a list without any polling.

A.2. Tracking IANA Registry Changes

Developers may need to get values from IANA registries for their software/hardware implementations. They might want to know when the registry changes, such as additional entries or updates to current entries. Thus, being able to be notified when a registry changes would be valuable to them.

Adding this functionality may be tricky for some registries. For example, if a developer cared about DKIM signature tags, they would have to subscribe to <http://www.iana.org/assignments/dkim-parameters/> which (currently) covers a handful of registries, all related to DKIM. Thus, a change to the DKIM hash algorithms would trigger a message showing that the registry had changed, even though the DKIM signature tags registry had not.

A.3. Tracking Changes in the Liason Statement Directory

Users might want to know when a new liaison statement is sent by the IETF or when one is received by the IETF.

A.4. Tracking Changes in Documents Outside the IETF Sphere

Users might want to track documents that relate to IETF activities but are produced by other standards development organizations (SDOs) such as the W3C, the IEEE, the Unicode Consortium, the ITU, and others. In order for the tracker to track these documents, it would need to poll occasionally and possibly scrape listings from HTML.

A.5. Tracking Additions to the IPR Statement Repository

Users might want to know when a new IPR statement is submitted.

Appendix B. Ideas that Might Be Implemented Later

The following are ideas for the new tool that are not currently being considered for the first round of development, but are being documented for possible future use. Items from this list may move to the list of requirements that are expected to be integrated during the first round of development.

- o The Datatracker could list all of the publicly-readable lists (or certainly at least the ones associated with IETF activities), and have links from WG pages in the Datatracker to the publicly-readable lists maintained by the WG chairs.
- o Draft versions of this RFC included a requirement to be able to include other lists. While this may still be desired, it was decided that implementing this in a safe and understandable way would be too difficult. In particular, there was a concern about detecting and handling loops. Later versions of the Datatracker might include this feature.
- o In public lists, it might be useful for someone to be able to understand why particular I-Ds and/or groups are added. Allowing the user who put together the list to add a comment field would help someone else understand the motivation.
- o The Datatracker might remove lists if it seems that storing them on the Datatracker is taking too many resources. The Datatracker can periodically send mail to the user reminding them to delete lists that are no longer needed.
- o The normal Datatracker display could have a button to add a particular I-D to the user's personal list.
- o Allow each user to determine what "significant change in status" is for the list they create. This could be done by a series of check boxes for every possible status change.
- o A list creator can add a list-level comment about who might be interested in following the list.
- o If the agendas for an upcoming meeting are scraped for I-D names, it would be possible to add an attribute to an I-D that lists that WG agenda(s) on which it appears.

- o In the section on "Adding groups of I-Ds to a list by attribute", add an attribute for "all I-Ds that are referenced by any I-D in a particular list".
- o Make it possible to add all I-Ds that have a certain section to a list (non-trivial IANA considerations, ASN.1 modules in appendices, MIBs, ABNF, XML modules, ...).
- o Even though Atom feeds have been around for years, they are new to many Internet users, and even experienced users only know how to use them in limited ways. The Datatracker should have at least a few paragraphs explaining how the Atom feeds that it provides can be used in different tools such as dedicated feed readers, online feed-display services, and so on.

Author's Address

Paul Hoffman
VPN Consortium

EMail: paul.hoffman@vpnc.org

