

Network Working Group  
Internet-Draft  
Intended status: Informational  
Expires: April 24, 2011

P. Hoffman  
VPN Consortium  
October 21, 2010

Requirements for Draft Tracking by the IETF Community in the Datatracker  
draft-ietf-genarea-datatracker-community-01

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 of interest to them.

Status of this Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

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

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

This Internet-Draft will expire on April 24, 2011.

Copyright Notice

Copyright (c) 2010 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. Definitions Used in This Document . . . . .	4
1.2. Discussion of These Requirements . . . . .	4
2. Requirements for Tools Features . . . . .	5
2.1. Lists . . . . .	5
2.1.1. Requirement: Lists of drafts can be large . . . . .	5
2.1.2. Requirement: A user can create multiple lists . . . . .	5
2.1.3. Requirement: Some lists must be able to be private . . . . .	5
2.1.4. Requirement: Specifying the drafts that are in a list must be simple . . . . .	6
2.1.5. Requirement: Adding groups of drafts to a list by attribute must be simple . . . . .	6
2.1.6. Requirement: Lists can dynamically include other lists . . . . .	7
2.2. Notifications . . . . .	7
2.2.1. Requirement: Users can be notified when a draft changes status . . . . .	7
2.2.2. Requirement: Every list has Atom feeds associated with it . . . . .	7
2.2.3. Requirement: Every list has mail streams associated with it . . . . .	8
2.2.4. Requirement: Notifications need to specify which list caused the notification . . . . .	8
2.3. Display in the Datatracker . . . . .	8
2.3.1. Requirement: Users can define how the rows are sorted in a display . . . . .	8
2.3.2. Requirement: Users can choose which attributes to display . . . . .	9
2.4. File Output . . . . .	10
2.4.1. Requirement: Users can get their current list as a single file . . . . .	10
3. IANA Considerations . . . . .	10
4. Security Considerations . . . . .	10
5. Acknowledgements . . . . .	11
6. References . . . . .	11
6.1. Normative References . . . . .	11
6.2. Informative References . . . . .	11
Appendix A. Some Known Open Issues . . . . .	11
Appendix B. Differences between -00 and -01 . . . . .	12
Author's Address . . . . .	12

## 1. Introduction

IMPORTANT NOTE: This is a very early draft of a set of requirements. It has gone through no general community review, and thus probably is missing many things that should be included, and some of the things in this draft are wrong and will be changed in future drafts. Nothing in this draft should be considered solid.

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

Users can search in the Datatracker by the filename of the draft, words in the draft's title, author, associated Working Group (WG) or IETF area, the responsible Area Director (AD), or IESG status. The returned list of drafts includes five columns: draft filename (with an active link to an HTMLized version of the draft maintained by the IETF tools team), the draft's title, the date it was submitted, its status in the IETF process, and the responsible AD (if any). For example, the output of a search in the current Datatracker can be seen at <http://imgur.com/snfyl.png>.

Instead of using the search capability of the Datatracker to manually find I-Ds of interest, users might want to create lists of drafts that they normally follow. Different users in the IETF community will have different ways that they want to get information on draft 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 desires are related to the overall desire to track drafts through their lifecycle.

For example, a WG chair might want to keep a list of all the drafts from other WGs that relate to active drafts in his or her WG. Someone who cares about the DNS probably also wants to follow the various drafts in different areas that affect the DNS. Developers who are not active in the IETF process might want to follow drafts on a particular topic to watch for things that might affect their implementations.

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. This document was prepared at the request of the IAOC.

Note that [RFC2026] describes the process that Internet Drafts 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 individual drafts as they go through the process.

### 1.1. Definitions Used in This Document

- o A "user" is an individual person who is member of the IETF community. (Yes, that definition is purposely vague.)
- o A "list" is an unordered set of filenames of Internet Drafts. Lists are specified by users.
- o An "attribute" is a feature of a draft, such as its filename, its current state in the IETF process, and soon. Attributes are usually displayed as columns in the Datatracker.
- o A "row" is a set of attributes about a single draft that is displayed in the Datatracker.
- o A "significant change in status" is all approvals and disposition. In the current process for drafts in the IETF stream, "all approvals" means "publication requested" "in last call" (this is IETF last call, not WG last call), and "IESG evaluation"; disposition is "approved" (for publication as an RFC), "RFC published", and "dead".

### 1.2. Discussion of These Requirements

This document is being discussed on the [datatracker-rqmts@ietf.org](mailto:datatracker-rqmts@ietf.org) mailing list. See <https://www.ietf.org/mailman/listinfo/datatracker-rqmts> for more information.

There will be a BoF at IETF 79 in Beijing to discuss this draft. It is currently being called "iddtspec", which somehow stands for "Review of Datatracker Specifications to Follow Internet-Draft Activities".

There is a plan to have one or two virtual meetings after Beijing to discuss these requirements.

## 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 drafts can be large

An active IETF participant might want to follow the status of hundreds of drafts. For example, some ADs have 100 drafts in their area, and they may also want to follow drafts outside their area that affect documents in their area.

#### 2.1.2. Requirement: A user can create multiple lists

A user might have multiple areas of interest and would want to track each area on a different web page. Another example would be a WG chair who wants to track the drafts in his or her WG separately from the drafts in a different area of interest. An IETF participant might want to have a list of drafts that they are following closely, and another list of drafts written by work colleagues.

#### 2.1.3. Requirement: Some lists must be able to be private

Seeing a list of drafts that covers multiple areas of interest can tell you something about the person who created the list. For example, you might be able to guess that they might be looking for a job in a different field by looking at their list of drafts of interest. Of course, anyone can follow individual drafts today without having that be exposed; however, following a particular group of drafts can reveal information about a person.

Methods that might keep lists private include:

- o The lists might only be available using passwords or some other common authentication mechanism. This would require that the Datatracker have a subscription process for users that could assign passwords, and a per-user process for adding lists to a user account.
- o Lists might be assigned random URLs from a very large ( $2^{128}$ ) namespace, and the user who creates a list does not tell others the assigned URL. This method makes it impossible for someone to search the entire set of assigned lists. Given that the URLs for

lists are most likely going to be copy-and-pasted anyway, having long random strings in the list's URL is not an impediment.

Note that some lists will purposely be made public, so there will be two types of lists.

2.1.4. Requirement: Specifying the drafts that are in a list must be simple

When a user creates a new list, it must be easy to add individual drafts to the list. There needs to be a mechanism that searches for potential drafts by partial filename, by partial or full title, and by author. Further, when editing an existing list, it must be easy to add additional drafts, and it must be easy to remove drafts from a list.

2.1.5. Requirement: Adding groups of drafts to a list by attribute must be simple

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

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

These attributes are dynamic, and thus the list of drafts that have a particular attribute will change after the user adds that attribute to a list. The Datatracker should update lists with dynamic

attributes every hour.

Note that some of these attributes are derived by programs created by the IETF Tools Team that parse drafts and are therefore inherently not completely reliable.

#### 2.1.6. Requirement: Lists can dynamically include other lists

If a user is authorized to see the contents of a list, he or she can include that other list in a different list. When the referenced list changes, those changes are also reflected in the referring-to list; that is, if list A includes list B, and the set of drafts in list B changes, the set of drafts in list A is automatically changed.

This feature is expected to be useful for experts (particularly WG chairs) who create lists on topics that others might consider interesting. For example, if Alice creates a list that contains all the drafts that she thinks relate to TLS, and Bob has access to that list, Bob can add that list to his personal list of things for which he is interested. Bob might also create a list-of-lists about TLS that includes references to Alice's list as well as to a similar list that Eric put together.

### 2.2. Notifications

#### 2.2.1. Requirement: Users can be notified when a draft changes status

Some users do not want to go to the Datatracker's display page to find out when a draft has been updated. Instead, they want to be notified immediately after the draft is changed. The Datatracker needs to support this type of immediate notification, where "immediate" means "within an hour of a change to any draft in the list".

#### 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 individual drafts. 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 email streams that are generated from the changes to the list: one for every change in status, and another for significant change of status.

### 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

When a list is displayed to the user in the Datatracker's web interface, each row represents a single draft. In a display, a particular draft should only be included once; for example, if someone manually adds draft-ietf-cuteacronym-sometopic to his list and also specifies that all drafts from the "cuteacronym" WG are included in the list, that draft should only appear once in the display.

### 2.3.1. Requirement: Users can define how the rows are sorted in a display

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 to normally see it in alphabetical order by the drafts' filenames, but after the user is off the net for a week, he or she might want to see the list in order of changes of status so that those drafts changed recently appear at the top of the list.

When displaying a list, the Datatracker should allow easy sorting of the drafts with the following collation orders:

- o Alphabetical by draft filename
- o Alphabetical by draft title
- o Alphabetical by associated WG
- o Date of publication of current version of the draft



- o Date of most recent change of status of any type
- o Date of most recent significant change of status

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. Also, the user should also be able to specify the order in which the attributes are displayed. The Datatracker should save the last-chosen set of attributes for display with the definition of the list.

The Datatracker should support display of the following attributes:

- o Draft filename
- o Draft title
- o Date of current draft
- o Status in the IETF process
- o Associated WG
- o Associated AD
- o Changed within the last 1 day
- o Changed within the last 2 days
- o Changed within the last 7 days
- o Included list(s) which contain this draft

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. The "included lists" attribute might show a pop-up with the names of the lists, given that list names might be long.

## 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 drafts. 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 care about.

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

The format of the file will be XML or JSON or tab-separated fields in a text file. The decision on which format is supported will be based on the desires of the community while discussing this document. (Imagine how much fun that will be!) Regardless of the format chosen, a syntax will need to be specified.

## 3. IANA Considerations

None.

## 4. Security Considerations

A tool for tracking the status of Internet Drafts can affect the privacy of its users. The requirements for privacy of the Datatracker views are discussed earlier in the document.

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.

## 5. Acknowledgements

Early ideas used in this document were contributed by Russ Housley, John Levine, Ray Pelletier, Blake Ramsdell, Julian Reschke, Yaron Sheffer, and Andrew Sullivan.

## 6. References

### 6.1. Normative References

[RFC2026] Bradner, S., "The Internet Standards Process -- Revision 3", BCP 9, RFC 2026, October 1996.

### 6.2. Informative References

[RFC4287] Nottingham, M., Ed. and R. Sayre, Ed., "The Atom Syndication Format", RFC 4287, December 2005.

## Appendix A. Some Known Open Issues

Given the very early stage of this document, there are actually many more open issues than are listed here. This list is mostly meant to remind the author of topics that need to be updated in future versions of the document, and to spur readers to think of even more open issues. Many of these topics were offered before the -00 draft by early reviewers.

- o One big feature that is desired is a way to say "tell me if this draft does not change state in the next nnn days". This gives a "dashboard" style capability. Doing this will mean holding more state on the Datatracker.
- o People get confused about the states of non-IETF streams (IRTF, IAB, ISE). These should be covered explicitly. Also, need definitions of "significant change in status" for the three non-IETF streams.
- o There will be an interesting and difficult interplay between privacy and sharing lists. If someone shares a list, that person doesn't want anyone modifying the contents of the list. So, there might need to be "sharing a shadow list" or something similar.
- o There may be legal issues with keeping user data private if we use login accounts.

- o Is there a formal definition for "drafts associated with a particular WG"?
- o When an AD agrees to sponsor an individual submission, does the Datatracker consider that draft associated with the AD? If not, that needs to be dealt with here.
- o Should the file output be in all the interesting formats (XML and JSON and tab-separated text) or just one?
- o As people coalesce on requirements for display, maybe mock up some HTML examples and put them in the document.
- o Thought: add a button in the normal Datatracker output to add a particular draft to a particular list.
- o How prescriptive do we want to be? Should this say things like "JavaScript pop-up" and "CSS" and such?
- o Should paging be supported for long lists in the HTML display?

#### Appendix B. Differences between -00 and -01

Added info for the mailing list.

#### Author's Address

Paul Hoffman  
VPN Consortium

Email: paul.hoffman@vpnc.org

