Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: 24 December 2025

M. Richardson
Sandelman Software Works

22 June 2025

A policy on third-party links in IETF emails and archives
draft-richardson-no-trackers-in-archives-00

Abstract

This document established a hyper-linking policy for IETF mailing lists and IETF mailing list archives.

The IETF will not tolerate posts that contain third-party links that can track users and contributors.

About This Document

This note is to be removed before publishing as an RFC.

Status information for this document may be found at https://datatracker.ietf.org/doc/draft-richardson-no-trackers-in-archives/.

Discussion of this document takes place on the ietf Working Group mailing list (mailto:ietf@ietf.org), which is archived at https://mailarchive.ietf.org/arch/browse/ietf/. Subscribe at https://www.ietf.org/mailman/listinfo/ietf/.

Source for this draft and an issue tracker can be found at https://github.com/mcr/no-trackers.

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 24 December 2025.

Table of Contents

1. Introduction

The IETF does most of it's work via email lists.

Indeed unlike many organizations, face to face meetings are not considered final, and process documents [RFC2026] emphasis that it is on mailing lists that rough consensus is sought: [RFC7282].

To that end, an accurate, searchable, safe and public archive is maintained by the IETF. Viewing the archives should not subject participants to monitoring by third parties.

An increasing occurance is (enterprise) email systems that attempt to make all links "safe" for their employees, screening their employees from the phising attack of the day. This is often accomplishing by filtering all email for HTTP links that appear in email, and then rewriting them to go through a filtering service.

Such services attempt to then block links that the enterprise considers harmful. When contributors from

So, instead of a link to https://example.com/, a link to something like https://filtering-service.example/https%3A/example.com/ is left, often with HTML adjusted so that users that use the text/html part never even know they are going through a filtering service.

This is fine for the IETF contributor themselves, but then the contributor might reply to the email, quoting the original email, including the rewritten link. This is then sent back to the IETF mailing list.

1.1. Risks to other Participants

Other readers, particularily those whose mail user agent prefers an HTML rendering, will then get links which go through the filtering service.

Should the reader decide to follow that link: it might point some important discussion on an IETF mailarchive, or reference data, or other standards organization, then the reader will disclose their activity to this third party filtering service. (The author of this document suspects that this may well violate privacy proviions of the GDPR, but the author has no expertise to know)

The above is the _best case_ situation. The user’s privacy is violated.

In some other cases, the filtering service in question will refuse to process the request, and the link is broken.

This could because the filtering service has no contract with the reader. Filtering services would be well advised to not process link requests from random strangers on the Internet. Answering such things is essentially creating an open proxy by which attackers can mount destributed denial of service attacks on others. These services should be encouraged not to allow this, but closing this loophole would render all the links broken to others.

Some filtering services are not stateless: that is, they do not store the entire method to access the target URL in the URL they provide. Instead, there is some amount of state that is created in a database, and that state is used to find the target URL. Such systems have some advantages to the enterprises that use them, but at some point the database will fill up, and the state will be removed. When that occurs the link will also break, and end users will not be able to recover the original link.

1.2. Risks to the IETF

The IETF maintains an extensive archive for all emails accessible by web browser, and by IMAP.

While the archive stores the entire email, including all renderings (text/plain, text/html, and any other attachments), the web interface to the archive does not display the text/html part. Users who are reading the archives will therefore see that the link will redirect them. Knowing not to follow that link depends upon the sophistication of the user, and their knowledge of current filtering systems. Some users will know how to edit the link to remove the filtering part, but this is at best an arcane activity.

Users that access the archives by IMAP have access to all parts of the email, and if they rely on the text/html part, will also never see that they are being redirected.

2. Solutions

2.1. Filter URLs in mail archive

Recognizing URLs in email, even text/plain parts is already done by the mailarchive system. They are turned into HTML links in the web presentation.

A denylist system could filter links that simply not turn links that include trackers/filters into HTML on the web interface.

This still leaves direct readers and IMAP users vulnerable.

2.2. Filter URLs when processing emails

The same URL recognizing system could operate as a mailman filter. It could either render the filtering URLs unuseable.

Alternatively, it could reject the email. While the 451 result code used in HTTP [RFC7725] might seem appropriate, 4xx code in SMTP are retriable, some 5xx code would need to be used. There would need to be a FAQ page cited back to the user so that they could understand what the problem was.

2.3. No HTML email

Since the trackers are hardest to see for users when the email is HTML encoded, one answer is to just filter out the text/html part of emails. This is already a feature of mailman3, and there are other large volunteer organizations, such as the Linux Kernel Community where this is routinely done: [tsohtmlkernel]

Some contributors send only the text/html part, and once it is removed, there is nothing left. Such emails would then need to be rejected, again with a clear message, and a reference to a FAQ.

3. Objections

The list of all the filtering systems is unbounced, which ones are a problem is a problem. Like all security questions, there are no promises of absolute security, but rather incremental improvements that make this better. One advantage of the increasingly centralization of email is that the set of filtering systems that these centralized solutions use is not that large. A great deal of improvement can be made with only a few rules.

The filtering options that would reject email will be confusing to many users. In particular, if the filter rewriting is not visible to the user in their Sent box, then it might be very difficult for them to understand what has gone one.

In many cases, the tracking URL was not placed in the user themselves, but was in some of the text that was quoted. Lack of editing of quoted parts in email, particularly by those in top-quote and do not read the email itself, is unfortunately endemnic. Learning to quote properly is a key part of contributing well to the IETF. In many cases, the presence of the tracking URL is part of quoted that text that does not even need to be present.

It is not a new problem that many enterprises have email systems that are incompatible with the IETF. Contributors who can not get their local systems fixed have wound up going to email providers which do not have these problems. This has often had a positive effect on stability of email references in documents, however, this represents a significant barrier to new contributors.

4. Privacy Considerations

This entire document is about a privacy violation that the lack of a policy is perpetuating.

5. Security Considerations

This document establishes a policy. This policy has significant possibility to inappropriately filter email, which may appear to some as censorship.

However, as explained by [RFC8890], when there is a conflict, the needs of the end users are considered more important than the needs of expert IETF contributors.

One purpose of this document is to allow IETF contributors to point to a policy document in order to get their local enterprise entities to make appropriate arrangements.

6. IANA Considerations

This document establishes a policy and requires no IANA activities. However, the list of links that would be filtered by the IETF mail processing system needs to be placed in a public place. An IANA Registry was considered, but was deemed in appropriate. Instead the IETF Tools team is asked to make the filter list public in some form.

7. Acknowledgements

Hello.

8. Changelog

9. References

9.1. Normative References

[RFC7282] Resnick, P., "On Consensus and Humming in the IETF", RFC 7282, DOI 10.17487/RFC7282, June 2014, <https://www.rfc-editor.org/info/rfc7282>.

[RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <https://www.rfc-editor.org/info/rfc8174>.

9.2. Informative References

[RFC2026] Bradner, S., "The Internet Standards Process -- Revision 3", BCP 9, RFC 2026, DOI 10.17487/RFC2026, October 1996, <https://www.rfc-editor.org/info/rfc2026>.

[RFC7258] Farrell, S. and H. Tschofenig, "Pervasive Monitoring Is an Attack", BCP 188, RFC 7258, DOI 10.17487/RFC7258, May 2014, <https://www.rfc-editor.org/info/rfc7258>.

[RFC7725] Bray, T., "An HTTP Status Code to Report Legal Obstacles", RFC 7725, DOI 10.17487/RFC7725, February 2016, <https://www.rfc-editor.org/info/rfc7725>.

[RFC8890] Nottingham, M., "The Internet is for End Users", RFC 8890, DOI 10.17487/RFC8890, August 2020, <https://www.rfc-editor.org/info/rfc8890>.

[tsohtmlkernel] T’so, T., "IETF@IETF.org email archive, posting from Theodore T’so", 19 August 2023, <https://mailarchive.ietf.org/arch/msg/ietf/8HJKm4hqhhMDN_5D-c1ZvDkn5Ws/>.

Author’s Address

Michael Richardson
Sandelman Software Works
Email: mcr+ietf@sandelman.ca