sml
Internet-Draft
Intended status: Standards Track
Expires: 14 November 2026

H.-J. Happel
audriga
A. Gulbrandsen
ICANN

13 May 2026

Trust and security considerations for Structured Email
draft-ietf-sml-trust-01

Abstract

This document discusses trust and security considerations for structured email and provides recommendations for message user agents on how to deal with structured data in email messages.

Discussion Venues

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

Discussion of this document takes place on the Structured Email Working Group mailing list (sml@ietf.org), which is archived at https://mailarchive.ietf.org/arch/browse/sml/.

Source for this draft and an issue tracker can be found at https://github.com/arnt/sml-trust.

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 14 November 2026.

Copyright Notice

Table of Contents

1. Introduction

Structured email, as described in [I-D.sml-structured-email], makes the content of some email messages machine-readable, such that user agents can provide higher-level functions than displaying/replying, for example "add this to calendar".

Naturally, new functions bring new trust and security considerations, or bring new urgency to existing issues. This document recommends security and trust mechanisms that should be applied when processing machine-readable content in email messages, both by senders and receivers.

2. Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

3. Types of security concerns

This section gives an overview of the various types of security and privacy concerns that arise when email messages contain structured data. The same concerns often arise for other messages, of course.

This section is informative.

3.1. Spam/virus filters

Structured email increases the syntactical and semantic complexity of email messages. If a spam/virus filter parses structured email in order to block malevolent messages, the filter’s parser will necessarily differ from that of the MUA that will finally act on the structured data, creating a risk of misclassification.

These risks are elevated when a structured data format has complex syntax, syntax that offers several optional or alternative ways to express the same substance, and of course by parsers that deviate from the specification for better bug compatibiloty.

3.2. Formal display of data

A common example is displaying a received calendar invitation using dates/times in the recipient’s timezone, in a fixed format.

Formal display introduces additional possibilities of discrepancy between the different representations. For example, a single message might contains a multipart/alternative containing a text/plain description of a flight itinerary, a text/html description of the same itinerary, and a structured representation. All three may be different, leading to confusion (and in this example, perhaps to missing a flight).

Unintentional discrepancy is a risk for senders; some recipients may be misinformed.

If a message is sent to a group and there is a discrepancy, different members of the group may see it differently.

If a particular MUA displays the formal representation within the message, a malevolent sender could try to mimic the visual representation using HTML with CSS, but with misleading content.

3.3. Additional user interface options

Structured mail processing may provide the receiving user with additional commands. Returning to the calendar example, many MUAs provide the user with additional commands to add something to a calendar.

3.4. Automated processing

Automated processing covers actions that are taken as soon as the message arrives rather than when a human user reads the message. For example, a user might want flight reservations to be automatically added to a travel itinerary application and/or a calendar.

Such automation could be a custom MUA feature or a future extension of the Sieve email filtering language ([RFC5228]). A related example for abuse in automated processing is calendar spam ([CalSpam]).

3.5. External references

Email messages with a text/html body part ("HTML email messages") may contain image resources that link to web servers. Such links can be used for tracking user interaction with the message.

Similar concerns apply to structured data types which include image references, such as the cover image of a music album or the teaser image of a news article.

RDF structured data can be partial by design and include references to additional data. Using a "follow your nose" approach, tools can follow URL references to obtain further structured data concerning a resource. For example, a piece of structured data about an article could reference the article’s authors only by URL. For a meaningful processing of author information, one might try to obtain further data using that URL.

3.6. Social engineering

While the risks of social engineering are hardly new and the human-readable text in a message can in principle be used to persuade the human reader to do anything, structured data widens the variety of actions the human reader can easily perform. If there are more buttons to click, then there’s also a greater variety of attacks.

Put differently: A user who might not be able to follow the instructions in a long and involved text-based social engineering attack may be able to follow simple instructions such as "click this then that".

4. Trust mechanisms

Several implementations of structured email restrict processing to messages that are particularly trusted. That is to say, an incoming message is in one of these three categories:

1. Spam. Structured data is not processed.

2. Ordinary. Structured data is not processed.

3. Trusted. Structured data is processed.

This section gives an overview of the trust mechanisms used to differentiate between 2 and 3.

It does not attempt to describe whether a trust-based mechanism is appopriate in a particular case.

4.1. Processing structured data

MUAs SHOULD display structured data in incoming email messages only if any of these criteria hold:

* Processing the data offers no additional attack surface compared to displaying the HTML in which the structured data is embedded. This may often be the case for formal display.

* Only for MUAs that process calendar invitations/updates: The MUA would process a calendar invitation in the same message.

* The sender is trusted (e.g., part of the user’s address book) and the messsage contains a valid personal or domain signature.

* The message is part of an ongoing thread with a trusted sender.

* The message’s content indicate a connection with an older, trusted message. For example, if a calendar invitation was accepted, then updates or responses for the same event are connected with the original.

Structured data that requires or suggests automatic processing may benefit from additional precautions before acting on the message. Documents that specify such data types should discuss how recipients should decide whether to act.

Open issue: At some point this document needs to mention JSON Web Signatures and RFC 7519, ether to use or to ignore.

4.2. Inlining data

Structured data included in an email message SHOULD be self-contained in order to avoid privacy problems. This implies that if an MUA is able to provide meaningful user interaction (rather than mere display), then the data SHOULD be self-contained, such that the interaction will not need referenced resources from the web.

5. Security considerations

Security considerations are a core subject of this document.

6. Privacy considerations

Privacy considerations are a core subject of this document.

7. IANA Considerations

This document has no IANA actions at this time.

8. References

8.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <https://www.rfc-editor.org/rfc/rfc2119>.

[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/rfc/rfc8174>.

8.2. Informative References

[I-D.sml-structured-email] "*** BROKEN REFERENCE ***".

[RFC5228] Guenther, P., Ed. and T. Showalter, Ed., "Sieve: An Email Filtering Language", RFC 5228, DOI 10.17487/RFC5228, January 2008, <https://www.rfc-editor.org/rfc/rfc5228>.

[RFC6132] George, R. and B. Leiba, "Sieve Notification Using Presence Information", RFC 6132, DOI 10.17487/RFC6132, July 2011, <https://www.rfc-editor.org/rfc/rfc6132>.

[RFC6134] Melnikov, A. and B. Leiba, "Sieve Extension: Externally Stored Lists", RFC 6134, DOI 10.17487/RFC6134, July 2011, <https://www.rfc-editor.org/rfc/rfc6134>.

[RFC7942] Sheffer, Y. and A. Farrel, "Improving Awareness of Running Code: The Implementation Status Section", BCP 205, RFC 7942, DOI 10.17487/RFC7942, July 2016, <https://www.rfc-editor.org/rfc/rfc7942>.

[CalSpam] "Calendar operator practices—Guidelines to protect against calendar abuse", n.d., <https://standards.calconnect.org/csd/cc-18003.html>.

[MachineReadable] "NIST IR 7511 Rev 4", n.d., <https://csrc.nist.gov/glossary/term/Machine_Readable>.

Appendix A. Acknowledgements

The authors wish to thank Ben Bucksch, Alexey Melnikov, Phillip Tao, Lisa Dusseault, Orie Steele, Daniel Kahn Gillmor, and others whose suggestions were made before this paragraph was started.

Appendix B. Note to self

RFC Editor: Please remove this section.

The charter has this to say about what this document should contain: "Recommendations for security and trust mechanisms that should be applied when processing machine-readable content in email messages" and "security and trust recommendations to prevent abuse of structured email". No more, no less.

Authors' Addresses

Hans-Joerg Happel
audriga
Email: happel@audriga.com
URI: https://www.audriga.com

Arnt Gulbrandsen
ICANN
6 Rond Point Schumann, Bd. 1
1040 Brussels
Belgium
Email: arnt@gulbrandsen.priv.no
URI: https://icann.org/ua