

EMAILCORE WG

IETF 120 – July 2024

A/S Open Issues

Editors:

John Klensin <john-ietf@jck.com>

Ken Murchison <murch@fastmail.com>

#40 - Recommended SMTP Extensions

<https://github.com/ietf-wg-emailcore/emailcore/issues/40>

Current text:

- MUST support RFC 6152 - 8-bit MIME
- MUST support RFC 3461 - Delivery Status Notifications
- SHOULD support RFC 2920 - Command Pipelining
- SHOULD support RFC 6531 - Internationalized Email

John Levine wrote in part:

RFC 3461 is about the EHLO keyword DSN which enables MAIL FROM keywords RET and ENVID and RCPT TO keywords NOTIFY and ORCPT. The only large mail system I know that announces it is Microsoft. Please just take it out.

Lyndon Nerenberg wrote in part:

And requiring it for embedded devices is nonsensical.

#78 - Advice against using URL %-encoding on non ASCII email addresses to create ASCII version of them

<https://github.com/ietf-wg-emailcore/emailcore/issues/78>

New text per John Levine and discussion thread:

[4.2. Use of non-ASCII Characters](#)

Proper generation and transmission of email addresses containing non-ASCII characters is discussed in [RFC6530]. [Section 9](#) of [RFC6530] says: "a downgrade mechanism that transforms the local part of an email address cannot be utilized in transit." Hence SMTP clients and servers MUST NOT try to encode non-ASCII email addresses as ASCII addresses. In particular, they MUST NOT use web URI percent encoding (see [Section 2.1](#) of [RFC3986]) nor Internationalized Domain Names for Applications (IDNA) (see [Section 4.4](#) of [RFC5891]) [Punycode](#) [RFC3492] in the local-part of an address, nor the former in the domain-part, since neither will produce a valid address.

In some cases, servers or clients may be able to use local knowledge to substitute ASCII addresses for specific non-ASCII addresses, but that is beyond the scope of this memo. See [Section 8](#) of [RFC6530] for further discussion.

#79 - Add Internationalization Considerations section

<https://github.com/ietf-wg-emailcore/emailcore/issues/79>

Pete Resnick wrote:

Are references to 653[01] that are already in A/S, along with pointers to MIME RFCs for internationalizing the contents of messages in the A/S sufficient?

John Levine wrote:

Unlikely we have anything useful and new to say at this point.

#85 - Add text to A/S about what mail agents should do/not do with Received header fields

<https://github.com/ietf-wg-emailcore/emailcore/issues/85>

Current text:

3.2 Use of Received Header Fields

3.2.2. Consumption

Received header fields are primarily for use when there are concerns about a message, such as to analyze handling or delivery problems, or to aid evaluation of a message with suspicious content or attributes. Received header fields are easily created and have no direct security or privacy protections.

Therefore, the fields do not warrant automatic trust. They should be used with care, for whatever information is deemed valuable, and especially when syntax or values occur that are not defined by the specifications

[\[I-D.ietf-emailcore-rfc5321bis\]](#) [\[I-D.ietf-emailcore-rfc5322bis\]](#).

#85 - Add text to A/S about what mail agents should do/not do with Received header fields

<https://github.com/ietf-wg-emailcore/emailcore/issues/85>

Dave Crocker's suggested edits:

Received header fields ~~are primarily for use when there are concerns about a message, such as to analyze~~ support analysis of handling ~~or~~ and delivery problems, ~~or to~~ as well as aiding evaluation of a message with suspicious content or attributes. ~~Received header~~ The fields are easily created and have no direct security or privacy protections, and the fields can contain personally identifiable information.

Therefore, the fields do not warrant automatic trust and do warrant thoughtful disclosure to others. They should be used with care, for whatever information is deemed valuable, and especially when syntax or values occur that are not defined by the specifications

[\[I-D.ietf-emailcore-rfc5321bis\]](#) [\[I-D.ietf-emailcore-rfc5322bis\]](#).

#84 - Add text about handling of Trace Header Fields by MUAs

<https://github.com/ietf-wg-emailcore/emailcore/issues/84>

New text per Alexey Melnikov, with support from list:

3.3. Handling of Trace Header Fields

A mail user agent (MUA) that uses an existing email message as a template for editing with the intention of sending it to a new set of recipients (this is sometimes implemented as "edit as a new message" feature) SHOULD strip trace header fields (see [Section 3.6.7](#) of [[I-D.ietf-emailcore-rfc5322bis](#)]).

#86 - Expand on operational meaning of being a trace header field

<https://github.com/ietf-wg-emailcore/emailcore/issues/86>

Pete Resnick wrote:

Right now, being a trace field is syntactically to be part of the Return-Path/Received block, and semantically means to be information about this messages movement through the system from submission to delivery. But what is the operational meaning? Is the only operational implication of being a trace field that you delete it in certain circumstances (like edit-as-new, mailing list distribution, etc.), or does it mean something more than that?

If it only means "delete on MUA resend" (or some variant of that), then the registry should probably be labeled that instead of "trace". If it means something more than that, then the A/S should probably have a much richer discussion than text suggested for ticket #84.

#95 – Received header needs definition

<https://github.com/ietf-wg-emailcore/emailcore/issues/95>

Doug Foster wrote in part:

Received headers are difficult to parse because they do not have clear boundaries between clauses, but they can be parsed because the use and sequence of keywords is very consistent. Additionally, many Received headers use standardized comments, and these are sometime in standardized places. Most notable of these is the convention of following the HELO name with a comment containing the IP address and Reverse DNS name, but there are others. **At minimum, the Applicability Statement should document these defacto standards about the structure of the Received header, and should provide guidance to encourage more consistency between implementations.**

One area with significant inconsistency is how cipher information is documented. The content and placement varies widely, which makes it much harder for an evaluator to assess the quality of encryption across the Received chain. If the purpose of the A-S is something other than reducing variability of implementations, what is that purpose?

#92 – CNAME handling in “5.1. Locating the Target Host”

<https://github.com/ietf-wg-emailcore/emailcore/issues/92>

Alexey Melnikov wrote:

In 5.1:

The lookup first attempts to locate an MX record associated with the name. If a CNAME record is found, the resulting name is processed as if it were the initial name.

Is CNAME expansion only done once? Or multiple times?
If multiple times, what the minimal expected limit?

#92 – CNAME handling in “5.1. Locating the Target Host”

<https://github.com/ietf-wg-emailcore/emailcore/issues/92>

John Klensin wrote:

RFC 1123 Section 5.2.2 prohibits CNAMEs in MAIL or RCPT commands. At some point we changed that and 2821 contains the text above, with no hints about iteration or recursion. I vaguely recall the change being made in some document between 1123 and 2821, but that is a dozen or so very eventful years.

I think being more specific requires involving the WG and recognizing that whatever we say is likely to render some existing implementation non-conforming.

How would you feel about a statement in the A/S pointing out that long CNAME chains are a source of trouble, that, when possible, one should stick to 1123-style canonical names, and when that is not possible, to try to keep the chains short.

#93 – “7.3. VRFY, EXPN, and Security” should point to SMTP AUTH RFC

<https://github.com/ietf-wg-emailcore/emailcore/issues/93>

RFC5321bis section 7.3 says:

On the public Internet, the contents of mailing lists have become popular as an address information source for so-called "spammers." The use of EXPN to "harvest" addresses has increased as list administrators have installed protections against inappropriate uses of the lists themselves. However, VRFY and EXPN are still useful for authenticated users and within an administrative domain. For example, VRFY and EXPN are useful for performing internal audits of how email gets routed to check and to make sure no one is automatically forwarding sensitive mail outside the organization. Sites implementing **SMTP authentication** may choose to make VRFY and EXPN available only to authenticated requestors. Implementations SHOULD still provide support for EXPN, but sites SHOULD carefully evaluate the tradeoffs.

#93 – “7.3. VRFY, EXPN, and Security” should point to SMTP AUTH RFC

<https://github.com/ietf-wg-emailcore/emailcore/issues/93>

Alexey Melnikov wrote:

I think it would be useful to add an Informative Reference to SMTP AUTH RFC here, “SMTP authentication (such as RFC 4954)”

John Klensin wrote in part:

If we want to address this, I think a section of the A/S on the subject of making VRFY and EXPN more useful, starting where 7.3 leaves off, would be entirely reasonable.

#94 – Use of Quoted Strings

<https://github.com/ietf-wg-emailcore/emailcore/issues/94>

Current text:

[3.1. Use of Empty Quoted Strings](#)

The quoted-string ABNF non-terminal is used in various places in [[I-D.ietf-emailcore-rfc5322bis](#)] grammar. While it allows for empty quoted string, such construct is going to cause interoperability issues when used in certain header fields. In particular, use of empty quoted strings is NOT RECOMMENDED in "received-token" (a component of a Received header field), "keywords" (a component of a Keywords header field) and "local-part" (left hand side of email addresses). Use of empty quoted strings is in particular problematic in the "local-part". For example, all of the following email addresses are non-interoperable:

`"".bar@example.com`

`foo.""@example.net"`

`"@example.com"`

Use of empty quoted strings is fine in "display-name".

#94 – Use of Quoted Strings

<https://github.com/ietf-wg-emailcore/emailcore/issues/94>

John Levine suggests replacing the current text with the following:

3.1 Use of Quoted Strings

The quoted-string ABNF non-terminal is used in various places in rfc5321bis and rfc5322bis grammar, as a way to allow strings that are otherwise forbidden in mailbox local parts. In practice, quoted strings do not work reliably in SMTP. For example, one large mail provider rejects mail from any address with two consecutive dots even if they are quoted, e.g.

"a..b"@example.com.

Local parts that require quoted strings are NOT RECOMMENDED.