

EMAILCORE WG IETF 113, Vienna March 2022

Chairs:

Alexey Melnikov <alexey.melnikov@isode.com>

Todd Herr <todd.herr@valimail.com>

Note Well

- This is a reminder of IETF policies in effect on various topics such as patents or code of conduct. It is only meant to point you in the right direction. Exceptions may apply. The IETF's patent policy and the definition of an IETF "contribution" and "participation" are set forth in BCP 79; please read it carefully.
- As a reminder:
 - By participating in the IETF, you agree to follow IETF processes and policies.
 - If you are aware that any IETF contribution is covered by patents or patent applications that are owned or controlled by you or your sponsor, you must disclose that fact, or not participate in the discussion.
 - As a participant in or attendee to any IETF activity you acknowledge that written, audio, video, and photographic records of meetings may be made public.
 - Personal information that you provide to IETF will be handled in accordance with the IETF Privacy Statement.
 - As a participant or attendee, you agree to work respectfully with other participants; please contact the ombudsteam (<https://www.ietf.org/contact/ombudsteam/>) if you have questions or concerns about this.

Note Well

(continued)

- Definitive information is in the documents listed below and other IETF BCPs. For advice, please talk to WG chairs or ADs:
 - BCP 9 (Internet Standards Process)
 - BCP 25 (Working Group processes)
 - BCP 25 (Anti-Harassment Procedures)
 - BCP 54 (Code of Conduct)
 - BCP 78 (Copyright)
 - BCP 79 (Patents, Participation)
 - <https://www.ietf.org/privacy-policy/> (Privacy Policy)

IETF Code Of Conduct Guidelines RFC 7154

- Treat colleagues with respect
- Speak slowly and limit the use of slang
- Dispute ideas by using reasoned argument
- Use best engineering judgment
- Find the best solution for the whole Internet
- Contribute to the ongoing work of the group and the IETF

Administrivia

- This session is being recorded
- Meetecho:
 - <https://www.conf.meetecho.com/conference/?group=emailcore>
- Jabber room (discussions/back channel):
 - emailcore@jabber.ietf.org
- Shared note taking:
 - <https://notes.ietf.org/notes-ietf-113-emailcore>
- ***Note taker?***

Agenda (1 of 2)

- Agenda bashing, administrivia, note well (chairs) - 5 mins
- Tickets for the SMTP draft:
- #11 (G.7.4. Possible clarification about mail transactions and transaction state) [**https://github.com/ietf-wg-emailcore/emailcore/issues/11**](https://github.com/ietf-wg-emailcore/emailcore/issues/11)
- #56 (Relax IANA registration policy for SMTP extensions) [**https://github.com/ietf-wg-emailcore/emailcore/issues/56**](https://github.com/ietf-wg-emailcore/emailcore/issues/56)
- #63 (VRFY in required commands in 4.5.1) [**https://github.com/ietf-wg-emailcore/emailcore/issues/63**](https://github.com/ietf-wg-emailcore/emailcore/issues/63)
- #60 (Restricted-capability clients?) [**https://github.com/ietf-wg-emailcore/emailcore/issues/60**](https://github.com/ietf-wg-emailcore/emailcore/issues/60)
- #62 (null mx vs server domain in 4.2.4.2) [**https://github.com/ietf-wg-emailcore/emailcore/issues/62**](https://github.com/ietf-wg-emailcore/emailcore/issues/62)

Agenda (2 of 2)

- #15 (G.7.9. Discussion of 'blind' copies and RCPT) [**https://github.com/ietf-wg-emailcore/emailcore/issues/15**](https://github.com/ietf-wg-emailcore/emailcore/issues/15)
- #55 (G.14. The FOR Clause in Received header field: Semantics, Security Considerations, and Other Issues) [**https://github.com/ietf-wg-emailcore/emailcore/issues/55**](https://github.com/ietf-wg-emailcore/emailcore/issues/55)
- #3 (G.3. Meaning of "MTA" and Related Terminology) [**https://github.com/ietf-wg-emailcore/emailcore/issues/3**](https://github.com/ietf-wg-emailcore/emailcore/issues/3)
- Tickets for the A/S draft:
- Text about ENHANCEDSTATUSCODES requirements in A/S
- #54 (G.7.17 Hop-by-hop Authentication and/or Encryption) [**https://github.com/ietf-wg-emailcore/emailcore/issues/54**](https://github.com/ietf-wg-emailcore/emailcore/issues/54)

RFC 5321

G.7.4. Possible clarification about mail transactions and transaction state)

<https://github.com/ietf-wg-emailcore/emailcore/issues/11>

John Klensin wrote:

Section 3.3 would be improved by being more specific about where mail transactions begin and end and then talking about "transaction state" here, rather than specific prior commands.

rfc5321bis currently says (1st para of Section 3.3):

There are three steps to SMTP mail transactions. The transaction starts with a MAIL command that gives the sender identification. (In general, the MAIL command may be sent only when no mail transaction is in progress; see Section 4.1.4.) A series of one or more RCPT commands follows, giving the receiver information. Then, a DATA command initiates transfer of the mail data and is terminated by the "end of mail" data indicator, which also confirms (and terminates) the transaction.

Proposal: no change, already correct.

RFC 5321

Relax IANA registration policy for SMTP extensions

<https://github.com/ietf-wg-emailcore/emailcore/issues/56>

SMTP extension registry registration procedure currently says:

Standards-Track or **Experimental** RFCs specifically approved by the IESG for this purpose.

With the removal of restriction on extensions starting with "X", at least **IETF Informational** RFCs should be allowed. What about **non IETF stream Informational/Experimental**?

Or even **Specification Required** registration procedure?

Or even FCFS + registration template requiring a specification?

RFC 5321

Relax IANA registration policy for SMTP extensions

OLD

Entries may be made only for service extensions (and associated keywords, parameters, or verbs) that are defined in Standards-Track or Experimental RFCs specifically approved by the IESG for this purpose.

NEW-1

New service extensions (and associated keywords, parameters, or verbs) may be registered using the **Specification Required** policy as specified in BCP 26 [12.1].

NEW-2

New service extensions (and associated keywords, parameters, or verbs) may be registered using the **First Come First Served** policy as specified in BCP 26 [12.1], with the provision that **a reference must also be provided** that describes the extension's purpose, format, and usage.

RFC 5321

VRFY in required commands in 4.5.1

<https://github.com/ietf-wg-emailcore/emailcore/issues/63>

John Levine wrote:

Nobody implements VRFY any more other than by always responding 252. It seems pointless still to require it.

Thoughts (Alexey): It is already advertised as a capability in EHLO, so not requiring it is fine. MTAs that still implement it don't need to change.

ACTION: register VRFY capability with IANA. It is mentioned in RFC 821, but was not registered.

Alternative: keep text as-is and recommend in A/S to do stub VRFY implementation that always returns 252.

Are requirements on Submission different? E.g. VRFY after a successful AUTH might be of use and safer to use? What about enterprise network?

RFC 5321

Restricted-capability clients?

<https://github.com/ietf-wg-emailcore/emailcore/issues/60>

Section 3.3 says:

Restricted-capability clients MUST NOT assume that any SMTP server on the Internet can be used as their mail processing (relaying) site.

Should clarify what such a client is, e.g., a printer that sends notifications to a single address which could be configured to send to a submission server, or to the MX for that address.

Is this text about **open relays**? If yes, is it in a wrong section?

Opinions?

RFC 5321

null mx vs server domain in 4.2.4.2

<https://github.com/ietf-wg-emailcore/emailcore/issues/62>

Mail server handles mail for another domain but not itself:

example.com. MX 10 mail.example.com.

mail.example.com. A 192.0.2.1

mail.example.com. MX 0 .

RFC 5321

null mx vs server domain in 4.2.4.2

<https://github.com/ietf-wg-emailcore/emailcore/issues/62>

```
220 mail.example.com ESMTP
EHLO client.com
250-mail.example.com
250 ENHANCEDSTATUSCODES
MAIL FROM:<bob@client.com>
250 2.1.0 sender OK
RCPT TO:<mary@mail.example.com>
??? 550 5.1.0 Not our mailbox
??? 556 5.1.10 Not anyone's mailbox
```

- Mail hosts know what domains they handle
 - But not why they don't handle other domains
- 556 useful in submission when relay sees null MX

RFC 5321

G.7.9. Discussion of 'blind' copies and RCPT

<https://github.com/ietf-wg-emailcore/emailcore/issues/15>

In Section 7.2. "Blind" Copies

Addresses that do not appear in the message header section may appear in the RCPT commands to an SMTP server for a number of reasons. The two most common involve the use of a mailing address as a "list exploder" (a single address that resolves into multiple addresses) and the appearance of "blind copies".

OLD (remainder of this paragraph):

Especially when more than one

RCPT command is present, and in order to avoid defeating some of the purpose of these mechanisms, SMTP clients and servers ***SHOULD NOT*** copy ***the full set*** of RCPT command arguments into the header section, either as part of trace header fields or as informational or private-extension header fields.

SHOULD NOT appears to be too weak, if no extra context is provided.

RFC 5321

G.7.9. Discussion of 'blind' copies and RCPT

<https://github.com/ietf-wg-emailcore/emailcore/issues/15>

Possible way forward: discuss this in terms of privacy versus debuggability tradeoff:

- * Favour privacy, debugging and tracing be damned
- * Favour debugging and tracing, privacy be damned
- * Alter SMTP protocol(?) to signal whether privacy or debugging is preferred

Any volunteers to propose some text?

RFC 5321

G.14. The FOR Clause in Received header field: Semantics, Security Considerations, and Other Issues

<https://github.com/ietf-wg-emailcore/emailcore/issues/55>

4.4.1. Received Header Field

4th paragraph:

- * If the FOR clause appears, it MUST contain exactly one <path> entry, even when multiple RCPT commands have been given. Multiple <path>s raise some security issues and have been deprecated, see Section 7.2.

Proposal to add:

- * If the FOR clause appears, it MUST contain exactly one <path> entry, even when multiple RCPT commands have been given, **and that <path> entry MUST contain one of the addresses that caused the message to be routed to the recipient of this message copy.** Multiple <path>s raise some security issues and have been deprecated, see Section **7.2**.

Should this also point to **Section 7.2 ("Blind" Copies)**, **Section 7.6 (Information Disclosure in Trace Fields)**, or even both?

RFC 5321

G.14. The FOR Clause in Received header field: Semantics, Security Considerations, and Other Issues

<https://github.com/ietf-wg-emailcore/emailcore/issues/55>

7.6. Information Disclosure in Trace Fields

In some circumstances, such as when mail originates from within a LAN whose hosts are not directly on the public Internet, trace (e.g., "Received") header fields produced in conformance with this specification may disclose host names and similar information that would not normally be available. This ordinarily does not pose a problem, but sites with special concerns about name disclosure should be aware of it. ***Also, the optional FOR clause should be supplied with caution or not at all when multiple recipients are involved lest it inadvertently disclose the identities of "blind copy" recipients to others.***

Remove the last sentence or fix it? What does "supplied with caution" means?

RFC 5321

Exploders seem to be prohibited from adding List-* header fields

<https://trac.ietf.org/trac/emailcore/ticket/4>

OPTIONS to choose from:

(1) Leave the RFC 5321 text (in Section 3.9 there and Section 3.4.2 of 5321bis-10 unchanged. In the latter, that sentence reads:

However, in this case, the message header section (RFC 5322 [12]) MUST be left unchanged; in particular, the "From" field of the header section is unaffected.

(2) Remove that sentence and replace it with nothing.

(3) Replace that sentence with

This change to MAIL FROM doesn't affect the header section of the message.

(4) Replace the sentence with something like

This change to MAIL FROM does not affect any header field that is already present in the message.

RFC 5321

G.3. Meaning of "MTA" and Related Terminology

<https://trac.ietf.org/trac/emailcore/ticket/3>

G.3. Meaning of "MTA" and Related Terminology

A terminology issue has come up about what the term "**MTA**" actually refers to, a question that became at least slightly more complicated when we formalized RFC 6409 Submission Servers. Does the document need to be adjusted to be more clear about this topic? Note that the answer may interact with the question asked in Section 2 above.

Possibly along the same lines, RFC 2821 changed the RFC 821 terminology from "**sender-SMTP**" and "**receiver-SMTP**" to "**SMTP client**" and "**SMTP server**" respectively. As things have evolved, it is possible that newer terminology is a source of confusion and that the terminology should be changed back, something that also needs discussion.

Question 1: "**sender-SMTP**" and "**receiver-SMTP**" versa "**SMTP client**" and "**SMTP server**". Proposal: no change.

Question 2: definition of MTA (next slide)

RFC 5321

G.3. Meaning of "MTA" and Related Terminology

<https://trac.ietf.org/trac/emailcore/ticket/3>

2.3.3. Mail Agents and Message Stores

Additional mail system terminology became common after RFC 821 was published and, where convenient, is used in this specification. In particular, SMTP servers and clients provide a mail transport service and therefore act as "**Mail Transfer Agents**" (MTAs). "**Mail User Agents**" (MUAs or UAs) are normally thought of as the sources and targets of mail. At the source, an MUA might collect mail to be transmitted from a user and hand it off to an MTA or, more commonly in recent years, a specialized variation on an MTA called a "**Submission Server**" (MSA) [42]. . At the other end of the process, the final ("delivery") MTA would be thought of as handing the mail off to an MUA (or at least transferring responsibility to it, e.g., by depositing the message in a "message store"). However, while these terms are used with at least the appearance of great precision in other environments, the implied boundaries between MUAs and MTAs often do not accurately match common, and conforming, practices with Internet mail. Hence, the reader should be cautious about inferring the strong relationships and responsibilities that might be implied if these terms were used elsewhere

Proposal: no change, unless the above text is broken.

A/S

What does it mean for servers to "support" ENHANCEDSTATUSCODES?

Summary of the mailing list discussion:

- Not universally supported
- Useful for senders in describing errors to users
- Want to encourage support without making existing implementations non-compliant
- Disconnect between the enhanced status codes and the SMTP service extension

Suggested text:

While Enhanced Mail System Status Codes [RFC 3463, RFC 5248] are widely supported, they are not ubiquitous. Nevertheless, they have been found to be useful to SMTP senders in determining the exact reason for a transmission failure in a machine-readable, language-independent manner, thus allowing them to present more detailed and language-specific error messages to users.

Given the usefulness of these enhanced codes, SMTP receivers are **RECOMMENDED** to implement the SMTP Service Extension for Returning Enhanced Error Codes [RFC 2034] utilizing the codes registered in RFC 5248.

A/S

G.7.17 Hop-by-hop Authentication and/or Encryption

<https://github.com/ietf-wg-emailcore/emailcore/issues/54>

Todd has posted updated text to the mailing list. Please check it if it is correct.