

A BCP for MTAs

Murray S. Kucherawy
<msk@cloudmark.com>

draft-kucherawy-mta-malformed

- SMTP, IMF, MIME, etc., all specify how to format and transport a message
- Lots of support in the ABNF for obsolete formats
- The old Postelism, “Be liberal in what you accept, conservative in what you send”, has been and still is extensively applied in mail software development

draft-kucherawy-mta-malformed

- What the RFCs don't do is provide solid guidance on what to do with a message that is non-conformant in terms of format or transport
 - And maybe explain what pain will be suffered if one doesn't do so
- The default is generally to accept the malformed message or transport anyway and make a best effort at delivery

draft-kucherawy-mta-malformed

- This leads to a bunch of bad things
 - Inconsistent user experience, where one service provider or software package delivers or renders “properly” a piece of email while others reject or get it wrong
 - Exploitable code, where a spam or virus filter fails to catch something because it can’t anticipate malformations that other agents will forgive in an effort to be “helpful”

draft-kucherawy-mta-malformed

- And some other bad things
 - We specify what things like IDNs are supposed to look like, but people still send UTF-8 URLs and the browsers happily do the intended thing even though they're not "legal" on numerous levels
 - "The great thing about standards is that there are so many to choose from."

draft-kucherawy-mta-malformed

- There is by now lots of distributed wisdom about how to handle common malformations for security and performance reasons
- We should collect it, write it down, and make it available to implementers so they don't have to re-learn it all on their own
 - That's what this effort is about
- Several implementers have been polled to ask what they do so that a consensus can be published, but more should participate

draft-kucherawy-mta-malfomed

- Scenario #1: Invalid header line
 - Consensus is to act like the body starts at line 4.

```
From: user@example.com  
To: userpal@example.net  
Subject: This is your reminder  
about the football game tonight  
Date: Wed, 20 Oct 2010 20:53:35 -0400
```

Don't forget to meet us for the tailgate party!

draft-kucherawy-mta-malfomed

- Scenario #2: Unusual extra spaces
 - Consensus is to modify the message outbound to correct unusual spacing

```
From: user@example.com
To: userpal@example.net
Subject: This is your reminder
MIME-Version : 1.0
Content-Type: text/plain
Date: Wed, 20 Oct 2010 20:53:35 -0400
```

Don't forget to meet us for the tailgate party!

draft-kucherawy-mta-malformed

- Scenario #3: Weakly-enforced rules
 - Section 3.6 of RFC5322 says, among other things, that a message must have exactly one From: field, but several implementations forgive this
 - Leads to one part of a system (e.g., a filter) keying off the first while others (e.g., an MUA) key off some other one
 - Consensus appears to be take some corrective action, whether that's removing all but one or rejecting the message outright

draft-kucherawy-mta-malformed

- Scenario #4: More weakly-enforced rules
 - RFC2045 says a message isn't MIME unless it has a MIME-Version field, but there's lots of mail intended to contain MIME parts that lacks a MIME-Version field
 - Consensus appears to be to act as if MIME-Version is there if the message otherwise appears to have MIME structure

draft-kucherawy-mta-malformed

- Probably many other cases that should be discussed with preferred solutions published
- You're all invited to read the draft and propose other scenarios and recommended practices
- When we have enough we can seek publication

draft-kucherawy-mta-malformed

- Discussion?