

Network Working Group
Internet Draft
Obsoletes: [2197](#), [1854](#)

Ned Freed
Innosoft
<[draft-freed-smtp-pipe-01.txt](#)>

SMTP Service Extension for Command Pipelining

March 2000

Status of this Memo

This document is an Internet-Draft and is in full conformance with all provisions of [Section 10 of RFC 2026](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

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."

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/1id-abstracts.txt>

The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>.

Copyright Notice

Copyright (C) The Internet Society (2000). All Rights Reserved.

1. Abstract

This memo defines an extension to the SMTP service whereby a server can indicate the extent of its ability to accept multiple commands in a single TCP send operation. Using a single TCP send operation for multiple commands can improve SMTP performance significantly.

The present document is an updated version of [[RFC-2197](#)], which in turn was a revision of [[RFC-1854](#)]. Only textual and editorial changes have been made; the protocol has not changed in any way.

2. Introduction

Although SMTP is widely and robustly deployed, certain extensions may nevertheless prove useful. In particular, many parts of the Internet make use of high latency network links. SMTP's intrinsic one command-one response structure is significantly penalized by high latency links, often to the point where the factors contributing to overall connection time are dominated by the time spent waiting for responses to individual commands (turnaround time).

In the best of all worlds it would be possible to simply deploy SMTP client software that makes use of command pipelining: batching up multiple commands into single TCP send operations. Unfortunately, the original SMTP specification [[RFC-821](#)] did not explicitly state that SMTP servers must support this. As a result a non-trivial number of Internet SMTP servers cannot adequately handle command pipelining. Flaws known to exist in deployed servers include:

- (1) Connection handoff and buffer flushes in the middle of the SMTP dialogue. Creation of server processes for incoming SMTP connections is a useful, obvious, and harmless implementation technique. However, some SMTP servers defer process forking and connection handoff until some intermediate point in the SMTP dialogue. When this is done material read from the TCP connection and kept in process buffers can be lost.
- (2) Flushing the TCP input buffer when an SMTP command fails. SMTP commands often fail but there is no reason to flush the TCP input buffer when this happens. Nevertheless, some SMTP servers do this.
- (3) Improper processing and promulgation of SMTP command failures. For example, some SMTP servers will refuse to accept a DATA command if the last RCPT TO command fails, paying no attention to the success or failure of prior RCPT TO command results. Other servers will

Expires September 2000

[Page 2]

accept a DATA command even when all previous RCPT TO commands have failed. Although it is possible to accommodate this sort of behavior in a client that employs command pipelining, it does complicate the construction of the client unnecessarily.

This memo uses the mechanism described in [[RFC-1869](#)] to define an extension to the SMTP service whereby an SMTP server can declare that it is capable of handling pipelined commands. The SMTP client can then check for this declaration and use pipelining only when the server declares itself capable of handling it.

2.1. Requirements Notation

This document occasionally uses terms that appear in capital letters. When the terms "MUST", "MUST NOT", "SHOULD", "SHOULD NOT", and "MAY" appear capitalized, they are being used to indicate particular requirements of this specification. A discussion of the meanings of the terms "MUST", "SHOULD", and "MAY" appears in [[RFC-1123](#)]; the terms "MUST NOT" and "SHOULD NOT" are logical extensions of this usage.

3. Framework for the Command Pipelining Extension

The Command Pipelining extension is defined as follows:

- (1) the name of the SMTP service extension is Pipelining;
- (2) the EHLO keyword value associated with the extension is PIPELINING;
- (3) no parameter is used with the PIPELINING EHLO keyword;
- (4) no additional parameters are added to either the MAIL FROM or RCPT TO commands.
- (5) no additional SMTP verbs are defined by this extension; and,
- (6) the next section specifies how support for the extension affects the behavior of a server and client SMTP.

Expires September 2000

[Page 3]

4. The Pipelining Service Extension

When a client SMTP wishes to employ command pipelining, it first issues the EHLO command to the server SMTP. If the server SMTP responds with code 250 to the EHLO command, and the response includes the EHLO keyword value PIPELINING, then the server SMTP has indicated that it can accommodate SMTP command pipelining.

4.1. Client use of pipelining

Once the client SMTP has confirmed that support exists for the pipelining extension, the client SMTP may then elect to transmit groups of SMTP commands in batches without waiting for a response to each individual command. In particular, the commands RSET, MAIL FROM, SEND FROM, SOML FROM, SAML FROM, and RCPT TO can all appear anywhere in a pipelined command group. The EHLO, DATA, VRFY, EXPN, TURN, QUIT, and NOOP commands can only appear as the last command in a group since their success or failure produces a change of state which the client SMTP must accommodate. (NOOP is included in this group so it can be used as a synchronization point.)

Additional commands added by other SMTP extensions may only appear as the last command in a group unless otherwise specified by the extensions that define the commands.

The actual transfer of message content is explicitly allowed to be the first "command" in a group. That is, a RSET/MAIL FROM sequence used to initiate a new message transaction can be placed in the same group as the final transfer of the headers and body of the previous message.

Client SMTP implementations that employ pipelining MUST check ALL statuses associated with each command in a group. For example, if none of the RCPT TO recipient addresses were accepted the client must then check the response to the DATA command -- the client cannot assume that the DATA command will be rejected just because none of the RCPT TO commands worked. If the DATA command was properly rejected the client SMTP can just issue RSET, but if the DATA command was accepted the client SMTP should send a single dot.

Expires September 2000

[Page 4]

Command statuses MUST be coordinated with responses by counting each separate response and correlating that count with the number of commands known to have been issued. Multiline responses MUST be supported. Matching on the basis of either the error code value or associated text is expressly forbidden.

Client SMTP implementations MAY elect to operate in a nonblocking fashion, processing server responses immediately upon receipt, even if there is still data pending transmission from the client's previous TCP send operation. If nonblocking operation is not supported, however, client SMTP implementations MUST also check the TCP window size and make sure that each group of commands fits entirely within the window. The window size is usually, but not always, 4K octets. Failure to perform this check can lead to deadlock conditions.

Clients MUST NOT confuse responses to multiple commands with multiline responses. Each command requires one or more lines of response, the last line not containing a dash between the response code and the response string.

4.2. Server support of pipelining

A server SMTP implementation that offers the pipelining extension:

- (1) MUST respond to commands in the order they are received from the client.
- (2) SHOULD elect to store responses to grouped RSET, MAIL FROM, SEND FROM, SOML FROM, SAML FROM, and RCPT TO commands in an internal buffer so they can sent as a unit.
- (3) SHOULD issue a positive response to the DATA command if and only if one or more valid RCPT TO addresses have been previously received.
- (4) MUST NOT, after issuing a positive response to a DATA command with no valid recipients and subsequently receiving an empty message, send any message whatsoever to anybody.

Expires September 2000

[Page 5]

- (5) MUST NOT buffer responses to EHLO, DATA, VRFY, EXPN, TURN, QUIT, and NOOP.
- (6) MUST NOT buffer responses to unrecognized commands.
- (7) MUST send all pending responses immediately whenever the local TCP input buffer is emptied.
- (8) MUST NOT make assumptions about commands that are yet to be received.
- (9) MUST NOT flush or otherwise lose the contents of the TCP input buffer under any circumstances whatsoever.
- (10) SHOULD issue response text that indicates, either implicitly or explicitly, what command the response matches.

The overriding intent of these server requirements is to make it as easy as possible for servers to conform to these pipelining extensions.

5. Examples

Consider the following SMTP dialogue that does not use pipelining:

```
S: <wait for open connection>
C: <open connection to server>
S: 220 innosoft.com SMTP service ready
C: HELO dbc.mtview.ca.us
S: 250 innosoft.com
C: MAIL FROM:<mrose@dbc.mtview.ca.us>
S: 250 sender <mrose@dbc.mtview.ca.us> OK
C: RCPT TO:<ned@innosoft.com>
S: 250 recipient <ned@innosoft.com> OK
C: RCPT TO:<dan@innosoft.com>
S: 250 recipient <dan@innosoft.com> OK
C: RCPT TO:<kvc@innosoft.com>
S: 250 recipient <kvc@innosoft.com> OK
C: DATA
S: 354 enter mail, end with line containing only "."
...
C: .
```

Expires September 2000

[Page 6]

```
S: 250 message sent
C: QUIT
S: 221 goodbye
```

The client waits for a server response a total of 9 times in this simple example. But if pipelining is employed the following dialogue is possible:

```
S: <wait for open connection>
C: <open connection to server>
S: 220 innosoft.com SMTP service ready
C: EHLO dbc.mtview.ca.us
S: 250-innosoft.com
S: 250 PIPELINING
C: MAIL FROM:<mrose@dbc.mtview.ca.us>
C: RCPT TO:<ned@innosoft.com>
C: RCPT TO:<dan@innosoft.com>
C: RCPT TO:<kvc@innosoft.com>
C: DATA
S: 250 sender <mrose@dbc.mtview.ca.us> OK
S: 250 recipient <ned@innosoft.com> OK
S: 250 recipient <dan@innosoft.com> OK
S: 250 recipient <kvc@innosoft.com> OK
S: 354 enter mail, end with line containing only "."
...
C: .
C: QUIT
S: 250 message sent
S: 221 goodbye
```

The total number of turnarounds has been reduced from 9 to 4.

The next example illustrates one possible form of behavior when pipelining is used and all recipients are rejected:

```
S: <wait for open connection>
C: <open connection to server>
S: 220 innosoft.com SMTP service ready
C: EHLO dbc.mtview.ca.us
S: 250-innosoft.com
S: 250 PIPELINING
C: MAIL FROM:<mrose@dbc.mtview.ca.us>
C: RCPT TO:<n timer>
C: RCPT TO:<nsb@thumper.bellcore.com>
C: RCPT TO:<galvin@tis.com>
C: DATA
```

Expires September 2000

[Page 7]

```
S: 250 sender <mrose@dbc.mtview.ca.us> OK
S: 550 remote mail to <nsb@thumper.bellcore.com> not allowed
S: 550 remote mail to <galvin@tis.com> not allowed
S: 554 no valid recipients given
C: QUIT
S: 221 goodbye
```

The client SMTP waits for the server 4 times here as well. If the server SMTP does not check for at least one valid recipient prior to accepting the DATA command, the following dialogue would result:

```
S: <wait for open connection>
C: <open connection to server>
S: 220 innosoft.com SMTP service ready
C: EHLO dbc.mtview.ca.us
S: 250-innosoft.com
S: 250 PIPELINING
C: MAIL FROM:<mrose@dbc.mtview.ca.us>
C: RCPT TO:<nsb@thumper.bellcore.com>
C: RCPT TO:<galvin@tis.com>
C: DATA
S: 250 sender <mrose@dbc.mtview.ca.us> OK
S: 550 remote mail to <nsb@thumper.bellcore.com> not allowed
S: 550 remote mail to <galvin@tis.com> not allowed
S: 354 enter mail, end with line containing only "."
C: .
C: QUIT
S: 554 no valid recipients
S: 221 goodbye
```

6. Security Considerations

This RFC does not discuss security issues and is not believed to raise any security issues not endemic in electronic mail and present in fully conforming implementations of [\[RFC-821\]](#).

7. Acknowledgements

This document is based on the SMTP service extension model presented in [RFC 1425](#). Marshall Rose's description of SMTP command pipelining in his book "The Internet Message" also served as a source of inspiration for this extension.

Expires September 2000

[Page 8]

8. References

[RFC-821]

Postel, J., "Simple Mail Transfer Protocol", STD 10, [RFC 821](#), August, 1982.

[RFC-1123]

Braden, R., "Requirements for Internet Hosts -- Application and Support", STD 3, [RFC 1123](#), October, 1989.

[RFC-1854]

Freed, N., "SMTP Service Extension for Command Pipelining", [RFC 1854](#), October, 1995.

[RFC-1869]

"Klensin, J., Freed, N., Rose, M., Stefferud, E., Crocker, D., "SMTP Service Extensions", STD 10, [RFC 1869](#), November, 1995.

[RFC-2197]

Freed, N., "SMTP Service Extension for Command Pipelining", [RFC 2197](#), September, 1997.

9. Author's Address

Ned Freed

Innosoft International, Inc.

1050 Lakes Drive

West Covina, CA 91790

USA

tel: +1 626 919 3600

fax: +1 626 919 3614

email: ned.freed@innosoft.com

This document is a product of work done by the Internet Engineering Task Force Working Group on Messaging Extensions, Alan Cargille, chair.

Expires September 2000

[Page 9]

10. Full Copyright Statement

Copyright (C) The Internet Society (2000). All Rights Reserved.

This document and translations of it may be copied and furnished to others, and derivative works that comment on or otherwise explain it or assist in its implementation may be prepared, copied, published and distributed, in whole or in part, without restriction of any kind, provided that the above copyright notice and this paragraph are included on all such copies and derivative works. However, this document itself may not be modified in any way, such as by removing the copyright notice or references to the Internet Society or other Internet organizations, except as needed for the purpose of developing Internet standards in which case the procedures for copyrights defined in the Internet Standards process must be followed, or as required to translate it into languages other than English.

The limited permissions granted above are perpetual and will not be revoked by the Internet Society or its successors or assigns.

This document and the information contained herein is provided on an "AS IS" basis and THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIMS ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Expires September 2000

[Page 10]