

Post Office Protocol - Version 3
Extended Service Offerings

Status of This Memo

This memo suggests a simple method for workstations to dynamically access mail from a discussion group server, as an extension to an earlier memo which dealt with dynamically accessing mail from a mailbox server using the Post Office Protocol - Version 3 (POP3). This RFC specifies a proposed protocol for the Internet community, and requests discussion and suggestions for improvements. All of the extensions described in this memo to the POP3 are OPTIONAL. Distribution of this memo is unlimited.

Introduction and Motivation

It is assumed that the reader is familiar with RFC 1081 that discusses the Post Office Protocol - Version 3 (POP3) [RFC1081]. This memo describes extensions to the POP3 which enhance the service it offers to clients. This additional service permits a client host to access discussion group mail, which is often kept in a separate spool area, using the general POP3 facilities.

The next section describes the evolution of discussion groups and the technologies currently used to implement them. To summarize:

- o An exploder is used to map from a single address to a list of addresses which subscribe to the list, and redirects any subsequent error reports associated with the delivery of each message. This has two primary advantages:
 - Subscribers need know only a single address
 - Responsible parties get the error reports and not the subscribers

o Typically, each subscription address is not a person's private maildrop, but a system-wide maildrop, which can be accessed by more than one user. This has several advantages:

- Only a single copy of each message need traverse the net for a given site (which may contain several local hosts). This conserves bandwidth and cycles.
- Only a single copy of each message need reside on each subscribing host. This conserves disk space.
- The private maildrop for each user is not cluttered with discussion group mail.

Despite this optimization of resources, further economy can be achieved at sites with more than one host. Typically, sites with more than one host either:

1. Replicate discussion group mail on each host. This results in literally gigabytes of disk space committed to unnecessarily store redundant information.
2. Keep discussion group mail on one host and give all users a login on that host (in addition to any other logins they may have). This is usually a gross inconvenience for users who work on other hosts, or a burden to users who are forced to work on that host.

As discussed in [RFC1081], the problem of giving workstations dynamic access to mail from a mailbox server has been explored in great detail (originally there was [RFC918], this prompted the author to write [RFC1081], independently of this [RFC918] was upgraded to [RFC937]). A natural solution to the problem outlined above is to keep discussion group mail on a mailbox server at each site and permit different hosts at that site to employ the POP3 to access discussion group mail. If implemented properly, this avoids the problems of both strategies outlined above.

ASIDE: It might be noted that a good distributed filesystem could also solve this problem. Sadly, "good" distributed filesystems, which do not suffer unacceptable response time for interactive use, are few and far between these days!

Given this motivation, now let's consider discussion groups, both in general and from the point of view of a user agent. Following this, extensions to the POP3 defined in [RFC1081] are presented. Finally, some additional policy details are discussed along with some initial experiences.

What's in a Discussion Group

Since mailers and user agents first crawled out of the primordial ARPAnet, the value of discussion groups have been appreciated, (though their implementation has not always been well-understood).

Described simply, a discussion group is composed of a number of subscribers with a common interest. These subscribers post mail to a single address, known as a distribution address. From this distribution address, a copy of the message is sent to each subscriber. Each group has a moderator, which is the person that administrates the group. The moderator can usually be reached at a special address, known as a request address. Usually, the responsibilities of the moderator are quite simple, since the mail system handles the distribution to subscribers automatically. In some cases, the interest group, instead of being distributed directly to its subscribers, is put into a digest format by the moderator and then sent to the subscribers. Although this requires more work on the part of the moderator, such groups tend to be better organized.

Unfortunately, there are a few problems with the scheme outlined above. First, if two users on the same host subscribe to the same interest group, two copies of the message get delivered. This is wasteful of both processor and disk resources.

Second, some of these groups carry a lot of traffic. Although subscription to an group does indicate interest on the part of a subscriber, it is usually not interesting to get 50 messages or so delivered to the user's private maildrop each day, interspersed with personal mail, that is likely to be of a much more important and timely nature.

Third, if a subscriber on the distribution list for a group becomes "bad" somehow, the originator of the message and not the moderator of the group is notified. It is not uncommon for a large list to have 10 or so bogus addresses present. This results in the originator being flooded with "error messages" from mailers across the Internet stating that a given address on the list was bad. Needless to say, the originator usually could not care less if the bogus addresses got a copy of the message or not. The originator is merely interested in posting a message to the group at large. Furthermore, the moderator of the group does care if there are bogus addresses on the list, but ironically does not receive notification.

There are various approaches which can be used to solve some or all of these problems. Usually these involve placing an exploder agent at the distribution source of the discussion group, which expands the name of the group into the list of subscription addresses for the

group. In the process, the exploder will also change the address that receives error notifications to be the request address or other responsible party.

A complementary approach, used in order to cut down on resource utilization of all kinds, replaces all the subscribers at a single host (or group of hosts under a single administration) with a single address at that host. This address maps to a file on the host, usually in a spool area, which all users can access. (Advanced implementations can also implement private discussion groups this way, in which a single copy of each message is kept, but is accessible to only a select number of users on the host.)

The two approaches can be combined to avoid all of the problems described above.

Finally, a third approach can be taken, which can be used to aid user agents processing mail for the discussion group: In order to speed querying of the maildrop which contains the local host's copy of the discussion group, two other items are usually associated with the discussion group, on a local basis. These are the maxima and the last-date. Each time a message is received for the group on the local host, the maxima is increased by at least one. Furthermore, when a new maxima is generated, the current date is determined. This is called the last date. As the message is entered into the local maildrop, it is given the current maxima and last-date. This permits the user agent to quickly determine if new messages are present in the maildrop.

NOTE: The maxima may be characterized as a monotonically increasing quantity. Although successive values of the maxima need not be consecutive, any maxima assigned is always greater than any previously assigned value.

Definition of Terms

To formalize these notions somewhat, consider the following 7 parameters which describe a given discussion group from the perspective of the user agent (the syntax given is from [RFC822]):

NAME	Meaning: the name of the discussion group Syntax: TOKEN (ALPHA *[ALPHA / DIGIT / "-"]) (case-insensitive recognition) Example: unix-wizards
ALIASES	Meaning: alternates names for the group, which are locally meaningful; these are typically used to shorten user typein Syntax: TOKEN (case-insensitive recognition) Example: uwiz
ADDRESS	Meaning: the primary source of the group Syntax: 822 address Example: Unix-Wizards@BRL.MIL
REQUEST	Meaning: the primary moderator of the group Syntax: 822 address Example: Unix-Wizards-Request@BRL.MIL
FLAGS	Meaning: locally meaningful flags associated with the discussion group; this memo leaves interpretation of this parameter to each POP3 implementation Syntax: octal number Example: 01
MAXIMA	Meaning: the magic cookie associated with the last message locally received for the group; it is the property of the magic cookie that it's value NEVER decreases, and increases by at least one each time a message is locally received Syntax: decimal number Example: 1004
LASTDATE	Meaning: the date that the last message was locally received Syntax: 822 date Example: Thu, 19 Dec 85 10:26:48 -0800

Note that the last two values are locally determined for the maildrop associated with the discussion group and with each message in that maildrop. Note however that the last message in the maildrop have a different MAXIMA and LASTDATE than the discussion group. This often occurs when the maildrop has been archived.

Finally, some local systems provide mechanisms for automatically archiving discussion group mail. In some cases, a two-level archive scheme is used: current mail is kept in the standard maildrop, recent mail is kept in an archive maildrop, and older mail is kept off-line. With this scheme, in addition to having a "standard" maildrop for each discussion group, an "archive" maildrop may also be available. This permits a user agent to examine the most recent archive using the same mechanisms as those used on the current mail.

The XTND Command

The following commands are valid only in the TRANSACTION state of the POP3. This implies that the POP3 server has already opened the user's maildrop (which may be empty). This maildrop is called the "default maildrop". The phrase "closes the current maildrop" has two meanings, depending on whether the current maildrop is the default maildrop or is a maildrop associated with a discussion group.

In the former context, when the current maildrop is closed any messages marked as deleted are removed from the maildrop currently in use. The exclusive-access lock on the maildrop is then released along with any implementation-specific resources (e.g., file-descriptors).

In the latter context, a maildrop associated with a discussion group is considered to be read-only to the POP3 client. In this case, the phrase "closes the current maildrop" merely means that any implementation-specific resources are released. (Hence, the POP3 command DELE is a no-op.)

All the new facilities are introduced via a single POP3 command, XTND. All positive responses to the XTND command are multi-line.

The most common multi-line response to the commands contains a "discussion group listing" which presents the name of the discussion group along with its maxima. In order to simplify parsing all POP3 servers are required to use a certain format for discussion group listings:

NAME SP MAXIMA

This memo makes no requirement on what follows the maxima in the listing. Minimal implementations should just end that line of the response with a CRLF pair. More advanced implementations may include other information, as parsed from the message.

NOTE: This memo STRONGLY discourages implementations from supplying additional information in the listing.

XTND BBOARDS [name]

Arguments: the name of a discussion group (optionally)

Restrictions: may only be given in the TRANSACTION state.

Discussion:

If an argument was given, the POP3 server closes the current maildrop. The POP3 server then validates the argument as the name of a discussion group. If this is successful, it opens the maildrop associated with the group, and returns a multi-line response containing the discussion group listing. If the discussion group named is not valid, or the associated archive maildrop is not readable by the user, then an error response is returned.

If no argument was given, the POP3 server issues a multi-line response. After the initial +OK, for each discussion group known, the POP3 server responds with a line containing the listing for that discussion group. Note that only world-readable discussion groups are included in the multi-line response.

In order to aid user agents, this memo requires an extension to the scan listing when an "XTND BBOARDS" command has been given.

Normally, a scan listing, as generated by the LIST, takes the form:

MSGNO SIZE

where MSGNO is the number of the message being listed and SIZE is the size of the message in octets. When reading a maildrop accessed via "XTND BBOARDS", the scan listing takes the form

MSGNO SIZE MAXIMA

where MAXIMA is the maxima that was assigned to the message when it was placed in the BBoard.

Possible Responses:

+OK XTND

-ERR no such bboard

Examples:

```
C: XTND BBOARDS
S: +OK XTND
S: system 10
S: mh-users 100
S: .
C: XTND BBOARDS system
S: + OK XTND
S: system 10
S: .
```

XTND ARCHIVE name

Arguments: the name of a discussion group (required)

Restrictions: may only be given in the TRANSACTION state.

Discussion:

The POP3 server closes the current maildrop. The POP3 server then validates the argument as the name of a discussion group. If this is successful, it opens the archive maildrop associated with the group, and returns a multi-line response containing the discussion group listing. If the discussion group named is not valid, or the associated archive maildrop is not readable by the user, then an error response is returned.

In addition, the scan listing generated by the LIST command is augmented (as described above).

Possible Responses:

+OK XTND

-ERR no such bboard Examples:

C: XTND ARCHIVE system

S: + OK XTND

S: system 3

S: .

XTND X-BBOARDS name

Arguments: the name of a discussion group (required)

Restrictions: may only be given in the TRANSACTION state.

Discussion:

The POP3 server validates the argument as the name of a discussion group. If this is unsuccessful, then an error response is returned. Otherwise a multi-line response is returned. The first 14 lines of this response (after the initial +OK) are defined in this memo. Minimal implementations need not include other information (and may omit certain information, outputting a bare CRLF pair). More advanced implementations may include other information.

Line	Information (refer to "Definition of Terms")
----	-----
1	NAME
2	ALIASES, separated by SP
3	system-specific: maildrop
4	system-specific: archive maildrop
5	system-specific: information
6	system-specific: maildrop map
7	system-specific: encrypted password
8	system-specific: local leaders, separated by SP


```
9      ADDRESS
10     REQUEST
11     system-specific: incoming feed
12     system-specific: outgoing feeds
13     FLAGS SP MAXIMA
14     LASTDATE
```

Most of this information is entirely too specific to the UCI Version of the Rand MH Message Handling System [MRose85]. Nevertheless, lines 1, 2, 9, 10, 13, and 14 are of general interest, regardless of the implementation.

Possible Responses:

```
+OK XTND
-ERR no such bboard
```

Examples:

```
C:      XTND X-BBOARDS system
S:      + OK XTND
S:      system
S:      local general
S:      /usr/bboards/system.mbox
S:      /usr/bboards/archive/system.mbox
S:      /usr/bboards/.system.cnt
S:      /usr/bboards/.system.map
S:      *
S:      mother
S:      system@nrtc.northrop.com
S:      system-request@nrtc.northrop.com
S:
S:      dist-system@nrtc-gremlin.northrop.com
S:      01 10
S:      Thu, 19 Dec 85 00:08:49 -0800
S:      .
```

Policy Notes

Depending on the particular entity administrating the POP3 service host, two additional policies might be implemented:

1. Private Discussion Groups

In the general case, discussion groups are world-readable, any user, once logged in (via a terminal, terminal server, or POP3, etc.), is able to read the maildrop for each discussion group known to the POP3 service host. Nevertheless, it is desirable, usually for privacy reasons, to implement private discussion groups as well.

Support of this is consistent with the extensions outlined in this

memo. Once the AUTHORIZATION state has successfully concluded, the POP3 server grants the user access to exactly those discussion groups the POP3 service host permits the authenticated user to access. As a "security" feature, discussion groups associated with unreadable maildrops should not be listed in a positive response to the XTND BBOARDS command.

2. Anonymous POP3 Users

In order to minimize the authentication problem, a policy permitting "anonymous" access to the world-readable maildrops for discussion groups on the POP3 server may be implemented.

Support of this is consistent with the extensions outlined in this memo. The POP3 server can be modified to accept a USER command for a well-known pseudonym (i.e., "anonymous") which is valid with any PASS command. As a "security" feature, it is advisable to limit this kind of access to only hosts at the local site, or to hosts named in an access list.

Experiences and Conclusions

All of the facilities described in this memo and in [RFC1081] have been implemented in MH #6.1. Initial experiences have been, on the whole, very positive.

After the first implementation, some performance tuning was required. This consisted primarily of caching the datastructures which describe discussion groups in the POP3 server. A second optimization pertained to the client: the program most commonly used to read BBoards in MH was modified to retrieve messages only when needed. Two schemes are used:

- o If only the headers (and the first few lines of the body) of the message are required (e.g., for a scan listing), then only these are retrieved. The resulting output is then cached, on a per-message basis.
- o If the entire message is required, then it is retrieved intact, and cached locally.

With these optimizations, response time is quite adequate when the POP3 server and client are connected via a high-speed local area network. In fact, the author uses this mechanism to access certain private discussion groups over the Internet. In this case, response is still good. When a 9.6Kbps modem is inserted in the path, response went from good to almost tolerable (fortunately the author only reads a few discussion groups in this fashion).

To conclude: the POP3 is a good thing, not only for personal mail but for discussion group mail as well.

References

- [RFC1081] Rose, M., "Post Office Protocol - Version 3 (POP3)", RFC 1081, TWG, November 1988.
- [MRose85] Rose, M., and J. Romine, "The Rand MH Message Handling System: User's Manual", University of California, Irvine, November 1985.
- [RFC822] Crocker, D., "Standard for the Format of ARPA-Internet Text Messages", RFC 822, University of Delaware, August 1982.
- [RFC918] Reynolds, J., "Post Office Protocol", RFC 918, USC/Information Sciences Institute, October 1984.
- [RFC937] Butler, M., J. Postel, D. Chase, J. Goldberger, and J. Reynolds, "Post Office Protocol - Version 2", RFC 937, USC/Information Sciences Institute, February 1985.

Author's Address:

Marshall Rose
The Wollongong Group
1129 San Antonio Rd.
Palo Alto, California 94303

Phone: (415) 962-7100

Email: MRose@TWG.COM

