

Out-of-Net Host Addresses for Mail

There is now interest in substantially extending the scope of the computer mail system used in the ARPANET to allow communication of voice, fax, graphics, as well as text information between users in different networks as well as within the ARPANET.

The discussion of a transition from the current ARPANET sndmsg environment and mechanisms to a more general internet environment and richer mechanisms must consider techniques for continued activity during the transition. In addition, there is a current need for a mechanism to support the interaction of the several already existing NSW-like message environments with the ARPANET message environment.

This memo discusses some possible alternatives for computer mail addressing for hosts outside the ARPANET in the short term. This memo is hopelessly Tenex oriented in its descriptions and examples.

It helps to keep a few goals in mind while considering the alternative solutions:

Goals:

- 1) Minimum Change to Existing Software.
- 2) Maximum User Acceptance.
- 3) Maximum Compatibility with the future Internet Message Environment.
- 4) Minimum Special Transition Software.

These goals are to some degree incompatible, so the evaluation should be expected to involve a trade off.

At this point, it would be good to have a model of the current situation and mechanisms of the ARPANET message environment. It is assumed the reader understands it well enough to dispense with a long description of how a message gets from A to B. The important thing is to note the types of players in the picture. There are:

message composition (or sending) programs (e.g., Hermes, SNDMSG), in general there are several message composition programs for each type of operating system or host in the network,

mailers,

mail servers (i.e., FTP servers) that receive the mail coming into at host and deposit it in mailboxes,

message processing (or reading) programs (e.g., Hermes, MSG, RD), in general there are several message processing programs for each type of operating system or host in the network, and note that the more developed mail are both reading and sending programs.

Messages are transmitted as a character string to an address which is specified "outside" the message. The destination host ("YYY") is specified to the sending (or user) FTP as the argument of the "open connection" command, and the destination user ("XXX") is specified to the receiving (or server) FTP as the argument of the "MAIL" (or "MLFL") command. In Tenex, when mail is queued this outside information is saved in the file name ("[--].XXX@YYY").

The proposed solutions are briefly characterized.

Proposed Solutions:

This first pass at describing the solutions is rather brief and intended to set the scene for a subsequent discussion based on examples.

A) SINGLE MAILBOX

This solution suggests that all mail for another network be routed to a single mailbox on a forwarding host on the ARPANET. The FTP server would naturally put all the mail for this mailbox into a single file to be examined by a routing daemon process. The routing daemon process would use information in new header lines to determine the actual destination.

Format:

Outside: [--].NSW-MAIL@FWDR

Inside: To: NSW-MAIL@FWDR
 From: Sam@ISIB
 NSW-User: Joe

B) GLOBAL NAMES INSIDE

This proposal suggests that all mail for users in another network be sent to a single mailbox on a forwarding host. The FTP server would naturally put all the mail for this mailbox into a single file to be examined by a routing daemon process. The routing daemon process would use information in existing header lines to determine the actual destination.

Format:

Outside: [---].NSW-MAIL@FWDR

Inside: To: Joe@NSW
From: Sam@ISIB

C) GLOBAL NAMES OUTSIDE

This proposal suggests that mail for users in another network be sent to distinct per user mailbox names on a forwarding host. The FTP server would somehow put all the mail for these mailboxes into a single file to be examined by a routing daemon process. The routing daemon process would use information in existing header lines to determine the actual destination.

Format:

Outside: [---].Joe@FWDR or [---].Joe@NSW

Inside: To: Joe@NSW
From: Sam@ISIB

D) STRUCTURED NAMES

This proposal suggests that mail for users in another network be sent to distinct per user mailbox names on a forwarding host, however, these mailbox names would have a common "network" part and a unique "user" part. By recognizing the common part the FTP server would put the mail and the mailbox name into a single file to be examined by a routing daemon process. The routing daemon process would use mailbox name information to determine the actual destination.

6 April 1979

Format:

Outside: [---].NSW-Joe@FWDR

Inside: To: NSW-Joe@FWDR
From: Sam@ISIB

Before further examination of the advantages and disadvantages of these proposals, it would be well to have some more detailed criteria in mind to help expose the degree to which the goals are met.

Criteria:

- 1) What changes are needed?
- 2) How many instances of the change need to be implemented?
- 3) What information does the routing deamon use?
- 4) How does the "answer" command work?
- 5) How is the name space used?

It is particularly instructive to work through examples with a mixture of mailbox destinations in the ARPANET and other networks in each of the "To:" and "CC:" fields and to see what happens when one wants to send an answer to all, just the "To:", or just the "CC:", or just the "From:" or "Sender:" mailboxes.

Solutions Reconsidered:

It is easier to talk about these things in terms of examples. In the following "NSW" is an example of a network name. "FWDR" is a host name, or nickname for the forwarding host. Also note that for all of these solutions it is assumed that host tables can have alternate or nicknames for hosts, e.g., FWDR could map to 86 while ISI also maps to 86, although this is not essential.

In addition, all these solutions provide a single forwarding point from the ARPANET into the destination net.

All forwarded messages are handled by a routing deamon which lives in the FWDR host.

Also note that the information shown as the "outside" information is the Tenex representation. The key thing is the mailbox argument value that is passed to the FTP server is the one in the string

"[---].XXX@YYY", not anything from the header. Only the string "XXX" is passed to the FTP server.

A) SINGLE MAILBOX

Example:

Outside: [---].NSW-MAIL@FWDR

Inside: To: NSW-MAIL@FWDR,Bill@ISIA
CC: Jeff@ISIB
From: Joe@ISIB
NSW-User-To: SAM,Fred
NSW-User-CC: Bob,Mike

or

Outside: [---].NSW-MAIL@FWDR

Inside: To: NSW-MAIL@FWDR,Bill@ISIA
CC: Jeff@ISIB
From: NSW-MAIL@FWDR
NSW-User-To: SAM,Fred
NSW-User-CC: Bob,Mike
NSW-User-From: Paul

Every mail composition program has to change to make it easy for users to put the "NSW-User:" line in the header. Every mail reading program has to change to notice and make use of this line. In an "answer" command the mail processing program has to know to copy this line into the answer message. The daemon has to examine the inside message header to find the "NSW-User:" line and forward the message to the users listed there. If there is a message that has both NSW and ARPANET mailboxes in both the "To:" and "CC:" lines, then it seems there must be both a "NSW-Users-To:" and a "NSW-Users-CC:" lines if it is to be possible to send an answer to just the users in the "To:" lines. If there is another network, e.g. PRNET, then another set of header lines must be introduced, e.g. PRNET-USER-To: etc., that is up to four new lines per network (To, CC, From, Sender).

This solution has the advantage of saving some transmissions: when several of the destination mailboxes are in NSW, the sending program sends just one copy to the FWDR and routing daemon, the routing daemon sends copies to all NSW users it finds. If this is not done, the daemon would have difficulty avoiding sending multiple copies to each destination user.

A problem arises about acknowledgements of mail receipt. First the normal ARPANET message delivery mechanisms will say the mail is delivered when the FTP server puts the mail in the file for the routing deamon to examine. Second if the routing deamon discovers an message is to be forwarded to a nonexistent user, care must be used to notify the original sender unambiguously.

Changes:

all composition programs

B) GLOBAL NAMES INSIDE

Example:

Outside: [---].NSW-MAIL@FWDR

Inside: To: Joe@NSW, Bill@ISIA, Fred@NSW
CC: Mike@NSW, Paul@NSW, John@ISIB
From: Sam@ISIB

Every mail composition program has to know that NSW is a very special host name, for which it uses a different mailbox argument and sends to the FWDR host. The FTP server naturally puts all the NSW mail into a single mailbox file which the routing deamon examines. The "answer" command works fine. The routing deamon has to look at the inside header to determine where to forward the messages. It has to check the "To:" and "CC:" lines.

The sending programs must also send just one copy to the FWDR and routing deamon, the routing deamon will send copies to all NSW users it finds. If this is not done, the deamon would have difficulty avoiding sending multiple copies to each destination user. This is an advantage in terms of number of transmissions.

A problem arises about acknowledgements of mail receipt. First the normal ARPANET message delivery mechanisms will say the mail is delivered when the FTP server puts the mail in the file for the routing deamon to examine. Second if the routing deamon discovers an message is to be forwarded to a nonexistent user, care must be used to notify the original sender unambiguously.

Changes:

all sending programs

C) GLOBAL NAMES OUTSIDE

Example:

Outside: [---].Joe@NSW

Inside: To: Joe@NSW, Bill@ISIA, Fred@NSW
CC: Mike@NSW, Paul@NSW, John@ISIB
From: Sam@ISIB

No changes to mail composition or processing programs are needed. The FTP server has to put all the NSW users mail into a single mailbox file which the routing daemon examines. The cheapest way to do this is to put all the names of the NSW users in the ARPANET user forwarding file with the same destination ARPANET mailbox. This means the local users of the FWDR host and the users in the destination networks share the name space for user names. The routing daemon has to look at the inside header to determine where to forward the messages. It has to check the "To:" and "CC:" lines.

This appears to be the solution with the minimum change to existing software. The "answer" command works fine.

There is a problem with the name space, for example, if ISIA serves as FWDR host, then Fred@ISI and Fred@NSW cannot co-exist. Further, there is the database update problem. Every time a new user is added to NSW or any of the hosts in any of the nets that the FWDR host serves the forwarding file at the FWDR host has to be updated. The names added have to be unique so all user names assigned in NSW and all the hosts on all the networks served by the same FWDR host have to be oked by the "forwarding file data base administrator" before they can actually be used. Also note that Fred@NSW and Fred@PRNET cannot be routed through the same FWDR host.

This doesn't work too well, if the sending programs are not changed they will send one copy of this message for each NSW user and all these copies will end up in the file to be examined by the routing daemon. If the FTP server code is not changed the outside information will be lost and the routing daemon will have no idea which NSW user this copy is for. To do the job right with the information available the routing daemon would have to keep a substantial record about each message it handled checking to see if it received for, and send a copy to, each intended destination user.

A problem arises about acknowledgements of mail receipt. First the normal ARPANET message delivery mechanisms will say the mail is delivered when the FTP server puts the mail in the file for the routing deamon to examine. Second if the routing deamon discovers an message is to be forwarded to a nonexistent user, care must be used to notify the original sender unambiguously.

Changes:

ARPANET user forwarding file at FWDR host

D) STRUCTURED NAMES

Example:

Outside: [---].NSW-Joe@NSW

Inside: To: NSW-Joe@NSW, Bill@ISIA, NSW-Fred@NSW
CC: NSW-Mike@NSW, NSW-Paul@NSW, John@ISIB
From: Sam@ISIB

No changes to mail composition or processing programs are needed. The FTP server has to put all the NSW-x users mail into a single file which the routing deamon examines. The FTP server can do this on the recognition of the "NSW-" prefix without knowing all the legal individual users. In addition the FTP server puts the mailbox argument into the file with the message. This is necessary to avoid the loss of the "outside" information. The routing deamon can then look at the mailbox argument to determine where to forward the messages. It need not look at the inside of the message at all. The "answer" command works fine.

A problem arises about acknowledgements of mail receipt. First the normal ARPANET message delivery mechanisms will say the mail is delivered when the FTP server puts the mail in the file for the routing deamon to examine. However, if the routing deamon discovers an message is to be forwarded to a nonexistent user, the deamon can easily tell the original sender the exact destination user that is unreachable.

Changes:

FTP server at FWDR host

Summary:

A	B	C	D
Single Mailbox	Global Names Inside	Global Names Outside	Structured Names

Criteria:

1) What changes?	Composer	Composer	None	FTP server
2) How many?	100	100	0	1
3) Routing information?	New Inside	Old Inside	Old Inside	Old Outside
4) "Answer" command?	Changes	Same	Same	Same
5) ARPANET name space use?	1 per FWDR	1 per FWDR	1 per user	1 per user

Goals:

1) Software Change	Bad	Bad	Good	Good
2) User Acceptance	Bad	Good	Good	Poor
3) Future Compatibility	Bad	Poor	Poor	Fair
4) Transition Software	Fair	Good	Bad	Good

Conclusions:

Solution D is recommended.

Only solution D is based on the use of strictly "outside" information. Please note that the existing ARPANET message DELIVERY system is based strictly on the use of "outside" information only. Also note that the problems that keep coming up in ARPANET message processing & composition programs have to do with the different possibilities for syntax (and semantics) of the "inside" information. This is a major advantage of solution D.

Please note that the syntax NET-USER@FWDR in the examples is not the only form that could be used. Any of the following (or even others) would be fine:

Net-User@FWDR	User-Net@FWDR
Net/User@FWDR	User/Net@FWDR
Net.User@FWDR	User.Net@FWDR
Net.and.User@FWDR	User.on.Net@FWDR

