

NWG/RFC#469
NIC 14798

Michael D. Kudlick MDK (SRI-ARC)
8-MAR-73

Network Mail Meeting Summary

Introduction

The purpose of this RFC is to briefly summarize, from the NIC's viewpoint, the principal conclusions reached at the Network Mail Meeting held Friday, February 23 1973, at SRI-ARC.

Please refer to RFC #475 (NIC 14919) for Abhay Bhushan's comprehensive summary of the issues discussed at the meeting.

There is no major disagreement between the present RFC and RFC #475.

RFC #453 (NIC 14317) contains background information on the meeting.

RFC #479 (NIC 14948) describes what the NIC would like to see included in the File Transfer Protocol for Network Mail purposes, and also describes briefly how the NIC would use the information.

The present RFC is organized as follows:

- Conclusions
- Discussion
- Attendees

Conclusions

Additional FTP mail requirements were decided upon. These would be implemented as a new mail command, with the following subcommands:

TO

This field is explicitly allowed to contain multiple addressees, with a standard syntax: user@host.

FROM

This field provides a return-address for notification of undeliverable mail, as well as a clearcut identification of the sender for the recipient's information..

AUTHOR

This field denotes the author of the mail. There may be multiple authors

TITLE

The "title" (i.e. subject) of the mail is to be terminated by period carriage return.

ACKNOWLEDEGMENT success / failure (time out) / normal

For use by the intermediate host, probably the NIC in most cases, to tell the sender what happened to his attempt to send mail. (Note: "normal" wasn't defined.)

RECORDED jnumber / null

Note: "jnumber" is the pre-assigned accession number (NIC number), to be used when known.

The "RECORDED" subcommand provides for the option of having the mail recorded. Information given with this subcommand would be recognized at the NIC. Options are:

- to be recorded (in NIC journal) only,
- to be recorded and distributed,
- to be distributed only.

This field would also be used to inform the recipient that the mail has been recorded.

(In retrospect, it may be preferable to have a separate command to inform the recipient of this fact, but no decision on this was made at the 23-Feb-73 meeting.)

TYPE long / urgent / ordinary

This allows the recipient site to take whatever action it thinks appropriate in storing the mail.

TEXT / FILE / CITATION

TEXT

This field is for the text of the mail message.

FILE

The purpose of the field is unclear to me. Does it contain a machine readable pointer to the file that the sender wishes the recipient to read?

CITATION

This field is a person-readable pointer to the file that the sender wishes the recipient to read. When the citation command is used, no mail is sent other than the citation.

Discussion

Introduction

The key aspects in the solution are:

- 1) It is based on FTP.
- 2) It uses the NIC without requiring direct use of NLS.
- 3) There is a mechanism for uniformity in the use of user identifications.
- 4) There is a mechanism for recording the mail for later reference.

These issues are covered in the discussion that follows.

New FTP Mail Subcommands

TO

Addressee Format

The standard form of the address is: user@host

"User" may be an individual's last name; or it may be whatever other identification the recipient has chosen AND has made known to the rest of the network.

If the intended host doesn't recognize the intended recipient's identification, then it sends back an "undeliverable" mail message to the sender's host. It is up to the individual to keep the NIC informed of his whereabouts [sic]; otherwise, he may not get his mail on time.

NIC Role

The NIC need have no role at all for mail sent from point A to point B, whenever that mail is not to be recorded at the NIC.

For mail that is to be recorded at the NIC, the RECORDED subcommand is to be used.

Also, when the sender does not know the standard address of the recipients, he may use the NIC to obtain this information.

Idents and Addresses

The NIC will modify its identification files to include the "user@host" standard address for each individual.

Sites may ask the NIC to translate from NIC Ident, or from a user's last name, to the standard address. A query facility will be made available at the NIC to do the translation on request. The translation service will also be available for "group idents".

This service would be FTP-like, in term of the prootocol [sic] it accepts, but would not be within FTP itself. A different server process would handle Ident translation requests.

Translation will also be done at the NIC when the NIC is used as an intermediate point on the delivery route.

The NIC could be an intermediate point for recording the mail as a NIC journal item, and for forwarding the mail to its ultimate destinations. During this process, the NIC would translate from NIC idents to standard addresses.

In the NIC ident files, provision already exists to specify hardcopy or on-line delivery of recorded (NIC Journal) mail.

This provision will be extended to include a "network" attribute, which means "deliver the mail to the host of this person".

The network attribute may be qualified by restricting all mail to be kept at the sender, with only a notification message actually mailed.

Notification would be in the form of a citation giving "to", "from", "title", "date of submission", and "location of mail".

TIP Users

To enable TIP users to have access to the mail system, both for sending and receiving mail, it was suggested that some hosts will have to be the "home" site for these users (but no more than one "home" site per user).

That is, an account that allows a TIP user to send and receive mail will have to be established at such a host.

For the present, any TIP user can use the SRI-ARC system for his mail requirements.

An alternate solution, that TIP's be equipped with a hardcopy device that is continuously available for printing mail, was discarded in favor of the above approach.

FROM

The "FROM" command in FTP, identifies the sender in "standard address" form.

This will allow "undeliverable" mail notices to be sent back to the originator.

The default condition is that the sender's host must retain the mail until it is "delivered" to the recipient's host.

"Delivered" means that the recipient's host has accepted the mail. It does NOT mean that the recipient has READ the mail.

Alternatively, the sender may designate that an intermediate host store the mail. Then the intermediate host has the responsibility of storing the mail until it is "delivered" to all intended recipients.

The "ACKNOWLEDGEMENT" command will allow an optional, positive acknowledgement to be given to the originator of the mail (the "FROM" addressee), stating that the mail was delivered.

AUTHOR

The AUTHOR may be several persons. For recorded documents the authors appear separately in the index of authors, to facilitate searching for mail when an author is known, but the title and location of the mail are unknown.

TITLE

The TITLE field is especially useful for recorded mail, since indexes on key words in the title can be produced relatively easily, and facilitate searching for mail.

For this reason, the title should be a succinct indicator of the contents.

ACKNOWLEDGEMENT

Acknowledgement of failure to deliver should be given to the sender.

An optional, positive acknowledgement of successful delivery to the recipient's sitename will be given on request of sender (like U.S. CERTIFIED mail).

No acknowledgement that the recipient actually saw the mail will be given (comparable to not having U.S. REGISTERED mail).

RECORDED

The concept of "recorded" mail is that a permanent record of the mail is kept centrally, to allow future references and re-readings of the mail to be made.

For example, in the NIC Journal system, a record is kept of all the items entered into the Journal. From this record, author, title-word, and NIC number indexes are produced to allow for references and re-readings.

The key to retrieval of recorded Journal items is the use of an accession number (the NIC number). This essentially removes the possibility of duplicate filenames being used.

The basic aspect of recorded mail which was discussed at the mail meeting is the assignment of an "accession" number.

It was decided to get the accession numbers from the NIC on an as-needed basis, without pre-assignment and without local assignment of numbers.

This subject may be reviewed in the future. Local assignment may be desirable to prevent the NIC from becoming a bottleneck in the mail process.

It was pointed out that local assignment of numbers would be un-ambiguous if the numbers included some information such as sitename, date, and time.

One other problem exists [sic], namely "where is the recorded document?".

Initially the document should be in the NIC, but ultimately it could be anywhere on the Network, provided only that there is a central mechanism for indexing and cataloging all the recorded documents.

The pathname to the recorded document would then include filename and sitename.

TYPE

The TYPE subcommand was a result of a discussion on the problems of large mail files, and the associated question of who would pay for the processing and storing of these files.

The main decisions made were:

- a) The processing, transmittal, and storage costs of sending mail should be borne at the sender's host.
- b) The processing and storage costs of receiving mail should be borne at the recipient's host initially, as a default.

Information to enable the recipient host to make an intelligent decision about where to store the incoming mail are passed along via the TYPE command.

The recipient host will have the local option of providing either of the following services:

- a) free use of system to send mail;
- b) free use of system to receive mail, i.e. login not required for delivery over the Network. (A possible alternative is use of a "mail" account, or use of the recipient's account, for processing and storage of the incoming mail.

TEXT / FILE / CITATION

TEXT

This field is for the text of the mail message.

FILE

The purpose of this field is unclear to me. Does it contain a machine readable pointer to the file that the sender wishes the recipient to read?

CITATION

The citation is a person-readable pointer to the file that the sender wishes the recipient to read.

An alternative to sending entire messages or files over the Network is to use the "CITATION" mechanism. With this, the sender sends a short message (the citation) saying, in effect, "please read file X at site Y".

This alternative would be especially useful for

- a) mail that is distributed with group idents (to all liaisons, for example), and
- b) "long" files (size not defined) that the recipient may not be immediately interested in.

However no method of enforcing use of this alternative was discussed. It will be up to the recipients to devise a scheme satisfactory to them.

Other General Discussion

Bob Kahn placed on the floor the following question (I paraphrase):

Can't the design of a mail system be made to include alternative sources of data and alternative modes of operation, unless exclusion of these alternatives can be quantitatively defended?

Particular aspects of this question are:

1) What is the desirability and difficulty of admitting different data sources into the mail system?

What are the "boundaries" that divide permitted from prohibited data sources?

What is the quantitative distinction between deferred and realtime mail?

Will the design we come up with allow such things as

a) handling a calendar that reflects the known and anticipated whereabouts of people so that meetings can be scheduled sensibly?

b) formatting the mail contents for later query and other information handling?

2) Whatever primitives we implement, can't they be designed so as not to preclude things like Tenex "linking"?

This requires two-way data communication paths.

How do we specify and get the attention of a "sink" for the data stream?

e.g., for interprocess communication, and for Tenex-type "linking".

The general reaction to this discussion was one of perspective:

In the scheme of things that could be considered "point-to-point communication", mailbox-type of communication is not the most general kind.

AKB listed several types of communication problems:

program-program communication
people-people real-time communication, e.g.
Tenex-type "links"
computer teleconferencing
mailbox communication: cataloging, storage
protocols: host-host, telnet, file transfer

A design for a mailbox-type system won't be required to encompass the problems of, say, a computer teleconferencing system, which has attributes (real-time, video, very large volume of data to be transferred, to name some) that are not attributes of a mail box system.

Attendees at the Network Mail Meeting 2/23/73 at SRI-ARC

	Nancy Mimno	BBN
ACB	Alan Bomberger	AMES-67
AKB	Abhay Bhushan	MIT-DMOG
AWH	Wayne Hathaway	AMES-67
CHI	Charles Irby	SRI-ARC
DHC	Dave Crocker	UCLA-NMC
JBP	Jon Postel	UCLA-NMC
JDH	Dave Hopper	SRI-ARC
JEW	Jim White	SRI-ARC
LPD	Peter Deutsch	PARC-MAXC
MCK	Mark Krilanovich	UCSB-MOD75
MDK	Mike Kudlick	SRI-ARC
REK2	Bob Kahn	ARPA
RKK	Rajendra Kanodia	MIT-MULTICS
RST	Ray Tomlinson	BBN-TENEX

[This RFC was put into machine readable form for entry]
[into the online RFC archives by Joseph Marshall 9/97]

