

COMMENTS ON NCP/TCP MAIL SERVICE TRANSITION STRATEGY

INTRODUCTION

This memo reviews and expands on the mail service transition plan [20].

The principal aim of the plan is to provide for the orderly support of the most commonly used network service (mail) during the period of transition from ARPANET to Internet Protocol-based operation.

The goal of the transition is, at the end, to provide in the internet environment service which is equivalent to or better than what has been available in the ARPANET environment. During the interim period, when both internet and the older ARPANET-based protocols are in use, the goal of the transition is to minimize user impact and, to the extent possible, to minimize software development or modification required to deal with transitional problems.

It is assumed that the reader is familiar with both the ARPANET and internet protocol hierarchies [1-17]. The internet hierarchy is designed to interface to many different packet networks (e.g., packet satellite, packet radio, Ethernet, LCS Ring net, X.25 public nets, ...), while the ARPANET hierarchy is limited to ARPANET IMPs (This is less true of the levels above NCP, but NCP itself is closely bound to ARPANET services).

The objective of the transition plan is to specify means by which the ARPANET electronic mail services may be supported across the boundary between the purely ARPANET environment and the more general internet environment during the period of transition by ARPANET hosts to the richer internet world.

ELECTRONIC MESSAGE SERVICES

DARPA is beginning a new phase of research into automatic electronic message handling systems. Ultimately, it is intended that electronic messages incorporate multiple media such as text, facsimile, compressed digitized voice, graphics and so on. Success in this new research will require substantial progress in developing multimode user interfaces to computer-based services (voice input/output, graphics, tablet/light pen, facsimile input/output, video/bit mapped displays, ...).

At the same time, progress must be made towards an environment based on internet protocols so as to avoid confining the results of the

multimedia effort to any one network. As a result, DARPA is planning to make several transitions over the next few years, from the existing, text-based ARPANET electronic message system to an internet-based, multimedia electronic message system.

This paper addresses only the first of the transitions from NCP-based text mail to TCP-based multimedia mail. The transition to the new multimedia mail system [7,19] lies ahead, but need not be planned in detail until we have some experience with the basic concepts. This first step only provides for the transition to TCP-based text mail.

The basic ground rules for transition from ARPANET-based electronic mail to internet electronic mail are the following:

1. ARPANET mailbox names must continue to work correctly.
2. No change required to mail editors which parse message headers to compose replies and the like.
3. Accommodation of non-ARPANET mailbox designators without change to the header parsing and checking mechanisms of mail composition programs.
4. Automatic forwarding of messages between NCP and TCP environments without user intervention.
5. During the transition, old style mail mechanisms must still work.

ELECTRONIC MESSAGE MECHANISMS

In order to make progress at all, it has been necessary to postulate fairly sophisticated changes to the "mailer" function which accepts as input an electronic text message and causes it to be delivered to the destination (or to an intermediate forwarder).

We also posit the existence of special, well-known mail forwarding hosts on the ARPANET which are responsible for accepting messages from NCP (TCP)-based message senders and forwarding them to TCP (NCP)-based message receivers.

In the ARPANET, electronic messages are transported via special procedures of the File Transfer Protocol: MAIL and MLFL. The former method sends electronic messages via the FTP Telnet command channel

while the latter achieves this by actual file transfer. In both cases, it is generally assumed that the receiving FTP server is colocated with the destination mailbox.

Thus, the sending procedure identifies to the receiver the destination mailbox identifier, but not the destination host (or network) identifier. For example, messages sent from Postel at USC-ISIF to Adams at USC-ISIA would arrive at ISIA with an indicator "Adams" but no indication of "ISIA". This creates some problems when messages must be staged at an intermediate host for further processing, as is the case when moving from an NCP-based sender to a TCP-based receiver, or vice-versa. Similar considerations arise when dealing with compatible, but different, message systems requiring re-formatting of messages at intermediate points.

In the following paragraphs, a mechanism is proposed for dealing with the naming, addressing and routing [18] of messages between systems.

At the source, it is assumed that the user has prepared the text of the message (including "To:" and "CC:" fields) in the conventional way [12]. The mailbox identifiers will continue to exhibit the format:

User@Host

but "host" may in fact be a compound name (which is not necessarily parsed), such as:

USC-ISIA
ARPANET-ISIA
SATNET-NDRE
PPSN-RSRE
HOST1.SRINET
LCSNET/MAILROOM

or even the name of an organization, such as:

BBN
ARPA
MIT
SRI

The only restriction is that the "@" not appear in either "user" or "host" strings in the mailbox identifier.

During message composition, the "user" or "host" portions of the

mailbox identifier may be verified for correctness (or at least for validity). The "user" string may incorporate parenthetical information such as

RAK(Richard A. Karp)@SU-AI

as is currently allowed.

After composition, messages are either sent immediately or left as "unsent mail" files to be sent later by mailer demons. The actual sending process uses the "host" string to determine where and how to send the message.

NEW MAIL MECHANISMS

At this point, we encounter the first critical new requirement to support the transition plan. A new table is needed within the mailer or in the host supporting the mailer or accessible to the mailer via the internet name server (for instance). This table must provide for mapping of the "host" string into an internet destination address (i.e., 32 bits: 8 bits of net, 24 bits of host), and must also indicate whether the destination is NCP or TCP capable.

In the event that the source and destination hosts do not have a compatible host level protocol (e.g. source is NCP only, destination is TCP only) then the message must be passed to a "forwarder" which can stage the transport by accepting via one protocol and forwarding by another.

This leads to a problem for the forwarding host since the basic FTP mail mechanism sends only the "user" portion of the mailbox identifier ("user@host") because the assumption is that the "host" is the destination. In the case of forwarding, the "host" is not the forwarder. Even if we cleverly arrange for "host" to translate into the internet address of a forwarder, we will have two problems. First, the forwarder may need the "host" information to figure where now to forward the message and second, depending on which network the source is in, "host" may need to translate into different forwarder addresses. The latter observation raises the spectre of many different mappings of a given "host" string which would require different tables for different mail sources. This would lead to considerable complexity in the maintenance and distribution of tables of forwarder addresses. Furthermore, a single-entry table mapping "host" to forwarder would limit reliability since only one forwarder would be bound to serve a given "host".

For the NCP/TCP transition, it may be sufficient to declare some set of well-known hosts to be NCP/TCP forwarders. Each mailer, when it discovers an incompatible destination, can send the message to any forwarder which is available. In addition, however, the mailer must provide full mailbox identifier information "user@host" to the forwarding host.

In the present mailers, only the "user" portion of the mailbox identifier is sent, so all mailers must change to send "user@host" when sending to a forwarder. The mailers all have to learn how to do table look-up a new way, also, to map "host" into internet addresses and to interpret the NCP or TCP capability information.

For purposes of this discussion, we postulate three different cases of electronic mail service implementation which must be made to interoperate during the transition:

1. Unchanged OLD NCP (RFC733) mail
2. NCP mail with new internet tables
3. TCP mail with new internet tables.

The second case assumes that the host has adopted a new host-string to address table (including NCP/TCP capability bits) and new mailer - mail server programs, but continues to use the old NCP host level protocol, modified to send "user@host" when sending to a forwarder. For such hosts, the only table entries which result in direct source-destination mail delivery are those showing NCP capability. If the destination is TCP capable only then the source host selects a forwarder address from another table and sends the message to it for further processing.

In the third case, the source host has fully transitioned to TCP, uses the new internet address tables to translate host-strings into internet addresses, and uses the new mailer - mail server. Destinations which are NCP-compatible only are reached via NCP/TCP forwarders.

Mail composition programs (e.g. SNDMSG, MSG, Hermes, MH,...) which today use ARPANET string-to-address tables to verify the legality of host names in mailbox entries can continue to use these "old" tables as long as these are updated to include internet host names as well as ARPANET host names.

Indeed, expanding the old tables is essential to handle the hardest

transition case: OLD NCP to new TCP mail. The three types of hosts lead to a 3 by 3 matrix of cases of mail transfer. In all but one case, mail is either handled directly or explicitly by forwarder. The only case needing further explanation is OLD NCP to NEW TCP which uses an "implicit forwarder."

IMPLICIT FORWARDING VS EXPLICIT FORWARDING

If the source host has adopted the new internet tables, it can tell whether the destination host has a compatible mail acceptance protocol. Incompatibility is explicitly resolved by selection of an intermediate forwarder.

If, however, the source host is still using pure NCP tables, it will not be able to tell that a particular destination host is only TCP-capable. To provide service for this case, it is proposed to expand the conventional NCP host table to include internet host names, but to map them into the addresses of implicit mail forwarders (i.e. Aliases).

Since we are postulating a case in which the NCP host has made no change (except for extending the host table). we also assume that the source host cannot send the "user@host" information via FTP to the intermediate forwarder.

This leaves the intermediate forwarder with the problem of figuring out where to forward a message identified by "user" only. In this case, we postulate that internet TCP-only mailboxes are registered at implicit forwarders so that incoming mail from conventional NCP sources can be forwarded successfully to the destination.

In the reverse direction, the source can use explicit forwarding because it is assumed that all TCP hosts use the new internet tables.

The use of registered names in the implicit forwarder raises two problems:

1. How can we deal with ambiguous mailbox names? (e.g. USERX@BBN and USERX@ISI look the same if only the string "USERX" is presented to the intermediate forwarder)
2. How can we collect, update and distribute changes to the registries at implicit forwarders?

In the first case, we propose to duck the problem by insisting on

unambiguous mailbox names everywhere. This may force some internet mail users to change their mailbox names, but we believe this will be rare.

The second problem can be solved by collecting information on a regular basis from all network mail users and cataloging this data in a database which can be accessed automatically (e.g. by mailer programs).

One possible mechanism is to make the data available through an internet mailbox name server analogous to the internet host name server [6]. This data might be collectible as a natural part of the TIP LOGIN database which is under development to permit expanded access to the ARPANET TIPs by legitimate ARPANET users.

In any case, internet mail users need supply their mailbox information to a single collection site which would disseminate it to all implicit forwarders on ARPANET. Note that such forwarders are only needed on ARPANET since all other systems are starting with the TCP-base. It is the internet mailbox users who must register, however, since they are the ones who cannot otherwise be reached via NCP.

FORWARDER CHARACTERISTICS

By their definition, NCP/TCP forwarders must be both NCP and TCP capable. Consequently, all NCP/TCP forwarders must be ARPANET hosts.

Implicit forwarders must accept conventional NCP/FTP mail [11] and be equipped with tables of valid internet user mailbox names which can be associated with the proper destination host. To allow implicit forwarders to also accept ordinary mail for users with mailboxes on the implicit forwarder, the forwarder should check first whether incoming mail is for a local user.

Explicit mail forwarders must be able to accept both conventional NCP-FTP mail commands (for local user mail) and both NCP-based and TCP-based mail server commands (whose arguments include the full destination mailbox strings "user@host").

To prevent potentially anomalous behavior, the NCP-based and TCP-based mail servers will offer service on socket/port 57 (71 octal). To summarize the communication patterns:

- (a) TCP sends/receives mail via well known port 57.

- (b) implicit forwarder receives conventional NCP/FTP mail on well-known socket 3, and sends TCP mail to port 57.
- c) explicit forwarder receives NCP mail on well-known socket 57, but sends NCP mail via NCP/FTP on socket 3. TCP mail is sent/received via port 57.

USER HOST CHARACTERISTICS

NCP hosts must at minimum, update host name tables to include aliases for internet hosts (i.e. map to NCP implicit forwarder host addresses).

The next most useful step is to update NCP hosts to include internet address tables and NCP/TCP capability bits so as to make use of explicit forwarders. This requires implementation of the mail server and modification of the mailer programs for sending mail to explicit forwarders. This also requires addition of explicit forwarder address tables.

Finally, a host can implement full TCP mail services, incorporating internet name tables and explicit forwarder address tables as well.

DANGLING PARTICIPLES

1. Error message handling needs to be worked out in detail to assure reasonable reporting of problems with the use of forwarders.
2. Designation of forwarding hosts.
3. Collection of internet mailbox names for implicit forwarders.
4. Format and distribution of internet name table and NCP/TCP capability information.
5. Dealing with mail systems not compatible with NCP, TCP or RFC733. (e.g. Telemail, On-Tyme, Phonenet, TWX, TELEX,...)

PLANS

To encourage this transition, the following schedule is proposed:

1. January 1, 1981 - implicit and explicit NCP/TCP forwarders made available on various service hosts (e.g. TOPS-20).
2. January 1, 1982 - implicit NCP/TCP forwarder service removed; explicit forwarding service continues.
3. January 1, 1983 - explicit NCP/TCP forwarding service terminated, transition to TCP complete.

ACKNOWLEDGEMENTS

A number of people have reviewed and commented on this contribution. Particular comments by J. Pickens, J. Postel, J. Haverty, D. Farber and D. Adams are gratefully acknowledged.

REFERENCES

1. DoD Standard Internet Protocol, IEN 128, RFC 760, NTIS ADA 079730, Jan 1980.
2. DoD Standard Transmission Control Protocol, IEN 129, RFC 761, NTIS ADA 082609, Jan 1980.
3. Postel, J., Telnet Protocol Specification, IEN 148, RFC 764, Jun 1980.
4. Postel, J., File Transfer Protocol, IEN 149, RFC 765, Jun 1980.
5. Postel, J., User Datagram Protocol, RFC 768, Aug 1980.
6. Postel, J., Internet Name Server, IEN 116, Aug 1979.
7. Postel, J., Internet Message Protocol, IEN 113, RFC 759, Aug 1980.
8. Postel, Sunshine, Cohen, The ARPA Internet Protocol, in preparation.
9. NCP: ARPANET Protocol Handbook, NIC 7104, Jan 1978.
10. Telnet: ARPANET Protocol Handbook, NIC 7104, Jan 1978.
11. FTP: ARPANET Protocol Handbook, NIC 7104, Jan 1978.
12. D. Crocker, J. Vittal, K. Pogran, A. Henderson, Standard for the Format of ARPA Network Text Messages, RFC 733, Nov 1977.
13. Crocker, et.al., Function-Oriented Protocols for the ARPA Computer Network, SJCC, May, 1972.
14. Carr, Crocker, Cerf, Host-Host Communication Protocol in the ARPA Network, SJCC, May, 1970.
15. Cerf, V., The Catenet Model for Internetworking, IEN 48, DARPA/IPTO, Jul 1978.
16. BBN 1822: Specifications for the Interconnection of a Host and an IMP, BBN Report No. 1822.
17. Heart, et.al., The Interface Message Processor for the ARPA Computer Network, SJCC, May, 1970.

18. Shoch, J., Inter-Network Naming, Addressing, and Routing, COMPCOM, Fall 1978.
19. Postel, J., A Structured Format for Transmission of Multi-Media Documents, RFC 767, Aug 1980.
20. Cerf, V. and, J. Postel, Mail Transition Plan, RFC 771, Sep 1980.
21. Sluizer, S. and, J. Postel, Mail Transfer Protocol, RFC 772, Sep 1980.

