                      Computer Mail Meeting Notes

Introduction

    A meeting was held on the 11th of January 1982 at USC Information
    Sciences Institute to discuss addressing issues in computer mail.
    The attendees are listed at the end of this memo. The major
    conclusion reached at the meeting is to extend the
    "username@hostname" mailbox format to "username@host.domain", where
    the domain itself can be further structured.

Overview

    The meeting opened with a brief discussion of the objectives of the
    meeting and a review of the agenda.

        The meeting was called to discuss a few specific issues in text
        mail systems for the ARPA Internet. In particular, issues of
        addressing are of major concern as we develop an internet in which
        mail relaying is a common occurance. We need to discuss
        alternatives in the design of the mail system to provide high
        utility at reasonable cost. One scheme suggested is to create
        "mail domains" which are another level of addressing. The ad hoc
        scheme of source routing, while effective for some cases, is seen
        to lead to some problems. A key test of addressing schemes is the
        procedure for sending copies of a reply to a message to the people
        who received copies of the original message. The key reference
        documents for the meeting were RFCs 788, 799, and 801.

    Jon Postel gave a brief review of the NCP-to-TCP transition plan (RFC
    801). The emphasis was on mail, the internet host table, and the
    role of a Host Name Server.

    The major part of the meeting was devoted to a wide ranging
    discussion of the general mailbox identification problem. In
    particular, the notion of a hierarchial structure of name domains was
    discussed, and the issues associated with name servers were discussed
    including the types of information name servers should provide.

Name Domains

    One of the interesting ideas that emerged from this discussion was
    that the "user@host" model of a mailbox identifier should, in

principle, be replaced by a "unique-id@location-id" model, where the
unique-id would be a globally unique id for this mailbox (independent
of location) and the location-id would be advice about where to find
the mailbox. However, it was recognized that the "user@host" model
was well established and that so many different elaborations of the
"user" field were already in use that there was no point in persuing
this "unique-id" idea at this time.

Several alternatives for the structuring and ordering of the
extensions to the "host" field to make it into a general
"location-id" were discussed.

   These basically involved adding more hierarchical name information
   either to the right or the left of the @, with the "higher order"
   portion rightmost or leftmost. It was clear that the information
   content of all these syntactic alternatives was the same, so that
   the one causing least difficulty for existing systems should be
   chosen. Hence it was decided to add all new information on the
   right of the @ sign, leaving the "user" field to the left
   completely to each system to determine (in particular to avoid the
   problem that some systems already use dot (.) internally as part
   of user names).

The conclusion in this area was that the current "user@host" mailbox
identifier should be extended to "user@host.domain" where "domain"
could be a hierarchy of domains.

   In particular, the "host" field would become a "location" field
   and the structure would read (left to right) from the most
   specific to the most general.

      For example: "Postel@F.ISI.IN" might be the mailbox of Jon
      Postel on host F in the ISI complex of the Internet domain.

   Formally, in RFC733, the host-indicator definition rule would
   become:

   host indicator = ("at" / "@") domains

   domains = node / node "." domains

      Note only one "at" or "@" is allowed, and that the domains
      form a hierarchy with the most general in scope last.

      And note that the choice of domain names must be
      administratively controlled and the highest level domain
      names must be globally unique.

The hierarchial domain type naming differs from source routing in that the former gives absolute addressing while the latter gives relative adressing.

Name Servers

The discussion of name servers identified three separate name server functions: "white pages", "unique-id to location-id", and "location-id to address".

The "white pages" service is a way of looking up a user by name and other properties using pattern matching and may return several data base "hits". Each hit must have an associated unique-id.

The "unique-id to location-id" service returns the character string location-id where the unique-id is currently found.

The "location-id to address" service returns a network address (numeric) corresponding to the location-id.

If the location-id is the name of a host in the current domain it is clear that the address returned will be the address to send the mail to, but if the location-id is that of some other domain then the address returned may be either the address to send the mail to, or the address of a name server for that domain, and these two cases must be distinguished.

The conclusion of this discussion was that a location-id to address name service must be defined soon. The other types of name servers were not further discussed, and are not required in the implemenation.

Another aspect of the name server is returning additional information besides the address. In particular, for mail it is important to know which mail procedures the destination implements (NCP/FTP, TCP/SMTP, etc.). Two approaches were discussed: one is coding the information as service names (e.g., NCP/SMTP), and the other is by reference to protocol and port numbers (e.g., PROTOCOL=6, PORT=25). Another suggestion was that the request ought to be "location-id,service" (e.g., "ISIF.IN,MAIL") and the response ought to be the location-id, address, protocol, and port. A different way of getting this information was suggested that instead of (or in addition to) having this information in the name server, one should get this data from the host itself via some sort of query or "who are you" protocol.

Also discussed was the initial provision for name service. It seems useful to start with a text file that can be accessed via FTP, and to have both "Telnet-Like" (i.e., based on TCP) and "Datagram" (i.e.,

based on UDP) access to a query server. This might be possible as an
extension of the IEN-116 name server.

Another issue was the central vs. distributed implementation of the
name look up service. It is recognized that separate servers for
each domain has administrative and maintenance advantages, but that a
central server may be a useful first step. It is also recognized
that each distinct database should be replicated a few times and be
avialiable from distinct servers for robust and reliable service.

An Example:

   Suppose that the new mailbox specification is of the form
   USER@HOST.ORG.DOMAIN.

      e.g., Postel@F.ISI.IN

   A source host sending mail to this address first queries a name
   server for the domain IN (giving the whole location "F.ISI.IN").

   The result of the query is either (1) the final address of the
   destination host (F.ISI), or (2) the address of a name server for
   ISI, or (3) the address of a forwarder for ISI. In cases 1 and 3,
   the source host sends the mail to the address returned. In case
   2, the source host queries the ISI name server and ... (recursive
   call to this paragraph).

Action Items:

   RFC 733 Revision

      To include the hierarchial host and domain naming procedure, and
      to delete the features decommitted at the Computer Mail meeting on
      10-JAN-79.

      By: Dave Crocker

      Due: 15-Feb-82

   Host Name Server Description

      To specify a way to get name to address conversions and to find
      out about services offered. Also how to get info on domain names.

      By: Jon Postel

      Due: 15-Feb-82

Transition Plan Revision

   To include new host and domain names.

   By: Jon Postel

   Due: 15-Feb-82

SMTP Revision

   To include new host and domain names.

   By: Jon Postel

   Due: Unspecified

Mail System Description Revision

   How to do mail systems, including use of SMTP and Host Name
   Server.

   By: Jon Postel

   Due: Unspecified

Conversion of User Programs and Mailer Programs.

   Programs have to handle dots in the "host" field. Many programs
   on many hosts will have to be modified to a greater or lesser
   extent. In many cases the modifications should be quite simple.

   By: A Cast of Thousands

   Due: Unspecified (See the Following Item)

Set a date when it ok to send messages with dots in "host" field.

   The must be a date after which it is ok to send host fields with
   dots throughout the ARPANET and Internet world without the
   recipients complaining.

   By: DARPA (Duane Adams)

   Due: 1-Mar-82

Attendees:

```
Duane A. Adams      DARPA/IPTO      Adams@ISI           (202) 694-8096
Vint Cerf           DARPA/IPTO      Cerf@ISI            (202) 694-3049
Harry Forsdick      BBN             Forsdick@BBN        (617) 497-3638
Eric Schienbrood    BBN             shienbrood@bbn-unix (617) 497-3756
Bob Thomas          BBN             BThomas@BBND        (617) 497-3483
Bob Fabry           Berkeley        Fabry@Berkeley      (415) 642-2714
Bill Joy            Berkeley        unj@berkeley        (415) 642-7780
Gene Ball           CMU             Ball@CMUA           (412) 578-2569
Anil Agarwal        COMSAT          Agarwal@ISID        (301) 863-6103
David L. Mills      COMSAT          Mills@ISID          (202) 863-6092
Dave Crocker        Univ. Del       DCrocker@Udel       (302) 738-8913
Ray McFarland       DoD             McFarland@ISIA      (301) 796-6290
Dave Lebling        MIT             PDL@MIT-XX          (617) 253-1440
Paul Mockapetris    ISI             Mockapetris@ISIF    (213) 822-1511
Jon Postel          ISI             Postel@ISIF         (213) 822-1511
Carl Sunshine       ISI             Sunshine@ISIF       (213) 822-1511
Mark Crispin        Stanford U.     Admin.MRC@SCORE     (415) 497-1407
Bob Braden          UCL[A]          braden@ISIA    (uk) (01)387-7050
Steve Kille         UCL             UCL-Netwiz@ISIE (uk) (01)387-7050
Bill Tuck           UCL             UKSAT@ISIE     (uk) (01)387-7050
Marv Solomon        Univ. Wisc      Solomon@UWisc
Ed Taft             Xerox Parc      Taft@Parc-Maxc      (415) 494-4419
```