Internet Name Domains

1. Introduction

     In the long run, it will not be practicable for every internet
host to include all internet hosts in its name-address tables. Even
now, with over four hundred names and nicknames in the combined
ARPANET-DCNET tables, this has become awkward. Some sort of
hierarchical name-space partitioning can easily be devised to deal
with this problem; however, it has been wickedly difficult to find one
compatible with the known mail systems throughout the community. The
one proposed here is the product of several discussions and meetings
and is believed both compatible with existing systems and extensible
for future systems involving thousands of hosts.

2. General Topology

     We first observe that every internet host is uniquely identified
by one or more 32-bit internet addresses and that the entire system is
fully connected. For the moment, the issue of protocol compatibility
will be ignored, so that all hosts can be assumed MTP-competent. We
next impose a topological covering on the space of all internet
addresses with a set of so-called name domains. In the natural model,
name domains would correspond to institutions such as ARPA, UCL and
COMSAT, and would not be necessarily disjoint or complete. While in
principle name domains could be hierarchically structured, we will
assume in the following only a single-level structure.

     Every name domain is associated with one or more internet
processes called mail forwarders and the name of that domain is the
name for any of these processes. Each forwarder process for a
particular domain is expected to maintain duplicate name-address
tables containing the names of all hosts in its domain and, in
addition, the name of at least one forwarder process for every other
domain. Forwarder processes may be replicated in the interests of
robustness; however, the resulting complexities in addressing and
routing will not be discussed further here. A particular internet
host may support a number of forwarder processes and their collective
names represent nicknames for that host, in addition to any other
names that host may have. In the following an internet host
supporting one or more forwarder proceses will be called simply a
forwarder.

     Every host is expected to maintain name-address tables including
the names of at least one forwarder for every

domain together with additional hosts as convenient. A host may
belong to several domains, but it is not necessary that all hosts in
any domain, be included in its tables. Following current practice,
several nicknames may be associated with the principal name of a host
in any domain and these names need not be unique relative to any other
domain. Furthermore, hosts can be multi-homed, that is, respond to
more than one address. For the purpose of mail forwarding and
delivery, we will assume that any of these addresses can be used
without prejudice. The use of multi-homing to facilitate source
routing is a topic for future study.

3. Naming Conventions

     In its most general form, a standard internet mailbox name has
the syntax

                    <user>.<host>@<domain> ,

where <user> is the name of a user known at the host <host> in the
name domain <domain>. This syntax is intended to suggest a
three-level hierarchically structured name (reading from the right)
which is unique throughout the internet system. However, hosts within
a single domain may agree to adopt another structure, as long as it
does not conflict with the above syntax and as long as the forwarders
for that domain are prepared to make the requisite transformations.
For instance, let the name of a domain including DCNET be COMSAT and
the name of one of its hosts be COMSAT-DLM with Mills a user known to
that host. From within the COMSAT domain the name Mills@COMSAT-DLM
uniquely identifies that mailbox as could, for example, the name
Mills.COMSAT-DLM@COMSAT from anywhere in the internet system.
However, Mills@COMSAT-DLM is not necessarily meaningful anywhere
outside the COMSAT domain (but it could be).

     A typical set of name domains covering the current internet
system might include ARPA (ARPANET), COMSAT (DCNET), DCA (EDNET,
WBNET), UCL (UCLNET, RSRENET, SRCNET), MIT (CHAOSNET), INTELPOST
(INTELPOSTNET) and the various public data networks. The ARPA
forwarder would use a name-address table constructed from the latest
version of the HOSTS.TXT table in the NIC data base. The other
forwarders would construct their own, but be expected to deposit a
copy in the NIC data base.

4. Mail Transport Principles

     In the interests of economy and simplicity, it is expected that
the bulk of all mail transport in the internet system will take place
directly from originator to recipient

host and without intermediate relay. A technique of caching will
probably be necessary for many hosts in order to reduce the traffic
with forwarders merely to learn the internet address associated with a
correspondent host. This naturally encourages naming strategies
designed to minimize duplicate names in the various domains; however,
such duplicates are not forbidden.

     There are several reasons why some messages will have to be
staged at an intermediate relay, among them the following:

1. It may not be possible or convenient for the originator
    and recipient hosts to be up on the internet system at
    the same time for the duration of the transfer.

2. The originator host may not have the resources to
    perform all name-address translations required.

3. A direct-connection path may not be feasible due to
    regulatory economic or security constraints.

4. The originator and recipient hosts may not recognize the
    same lower-level transport protocol (e.g. TCP and NCP).

     A mail relay is an internet process equipped to store an MTP
message for subsequent transmission. A mail forwarder is a mail
relay, but not all relays are forwarders, since they might not include
the full name-address capability required of forwarders. In addition,
relays may not be competent in all domains. For instance, a MTP/TCP
relay may not understand NCP. In other words, the forwarders must be
fully connected, but the relays may not.

     The particular sequence of relays traversed by a message is
determined by the sender by means of the source route specification in
the MRCP command. There are several implications to this:

1. Advisory messages returned to the originator by a relay
    or recipient host are expected to traverse the route in
    reverse order.

2. Relay host names follow the same naming convention as
    all host names relative to their domain. Since it may
    not be possible (see below) to use internet addresses to
    dis-ambiguate the domain, the complete standard internet
    name .<host>@<domain> is required everywhere.

3. There is no current provision for strict/loose route
    specifications.   If, in fact, the "ordinary" host
    specification @<host> were used, each relay or forwarder

would use the rules outlined in the next section for
routing. This may result in additional relay hops.

5. Forwarder Operations

     This section describes a likely scenario involving hosts, relays
and forwarders and typical internet routes. When a forwarder receives
a message for <user>.<host>@<domain>, it transforms <host> if
necessary and forwards the message to its address found in the
name-address table for <domain>. Note that a single host can be a
forwarder for several independent domains in this model and that these
domains can intersect. Thus, the names Mills@USC-ISIE,
Mills.USC-ISIE@ARPA and Mills.USC-ISIE@COMSAT can all refer to the
same mailbox and the names USC-ISIE, ARPA and COMSAT can, conceivably,
all be known in the same domain. Such use would be permissable only
in case the name USC-ISIE did not conflict with other names in this
domain.

     In order for this scheme to work efficiently, it is desireable
that messages transiting forwarders always contain standard internet
mailbox names. When this is not feasible, as in the current ARPANET
mail system, the forwarder must be able to determine which domain the
message came from and edit the names accordingly. This would be
necessary in order to compose a reply to the message in any case.

     In the RFC-780 model a message arriving at a forwarder is
processed by the MTP server there. The server extracts the first
entry in the recipient-route field of an MRCP command. There are two
cases, depending on whether this entry specifies a domain name or a
host name. If a domain name, as determined by a search of a universal
table, it refers to one of the domains the server represents. If not,
it must a name or nickname of the server's host relative to ooe of the
domains to which the sender belongs. This allows a distinction to be
made between the domains COMSAT and INTELPOST on one hand and the
COMSAT host COMSAT-PLA on the other, all of which might be represented
by the same internet address, and implies that domain names must be
unique in all domains.

     The server next extracts the second entry in the recipient-route
field of the MRCP command and resolves its address relative to the
domain established by the first entry. If the second entry specifies
an explicit domain, then that overrides the first entry. If not and
the first entry specifies a domain, then that domain is effective.
However, if the first entry specifies the server's host, it may not be
apparent which domain is intended. For instance, consider the
following two MRCP commands:

MRCP to:<@COMSAT,Mills@HOST> and
MRCP to:<@INTELPOST,Mills@HOST> ,

where Mills.HOST@COMSAT and Mills.HOST@INTELPOST are distinct
mailboxes on different hosts. A receiving host supporting forwarders
for both COMSAT and INTELPOST can then preserve this distinction and
forward correctly using the above rules.

     Now let the forwarder host have the name FORWARDER in both the
COMSAT and INTELPOST domains and consider its options when receiving
the command

MRCP to:<@FORWARDER,Mills@HOST> .

The forwarder is being asked simply to relay within the domain of the
sender; however, it belongs to more than one domain! The obvious way
to resolve this issue would be to forbid the use of implicit domains,
as represented by Mills@HOST, and require the full internet mailbox
names Mills.HOST@COMSAT or Mills.HOST@INTELPOST. It is also possible
to dis-ambiguate the domain by inspecting the first entry of the
sender-route field of the MAIL command (see below).

6.   Source and Return Routing

     In the RFC-780 model, routes can be specified in the
recipient-route field of the MRCP command and in the sender-route
field of the MAIL command. In point of fact, neither the
recipient-route or sender-route is necessary if the originator
specifies standard internet mailbox names. So long as the routes,
when used, consist only of domain names, there is no conflict with the
current RFC-780 specification. If for some reason forwarding must be
done via other hosts, then the use of a complete and unambigous syntax
like .<host>@<domain> is required in order to avoid problems like that
described above.

     The present RFC-780 specification requires the receiver to
construct a name for the sender and insert this at the beginning of
the sender-route. Presumably, the only information it has to
construct this name is the internet address of the sender. Consider
the case, as in the example above, where multiple domains are
supported by a single server on a particular host. If hosts receiving
a message relayed via that server were to map its address into a name,
there would be no way to determine which domain was intended. We
conclude that the sending host must update the sender-route as well as
the recipient-route. It does this simply by copying the first entry
in the recipient-route as received as the new first entry in the
sender-route.

7.   Editing the RFC-733 Header

     Every effort should be made to avoid editing the RFC-733 header,
since this is an invasive procedure requiring extensive analysis. It
is expected that newly developed mail systems will be aware of the
standard internet mailbox syntax and ensure its use everywhere in the
RFC-733 and RFC-780 fields. On the occasions where this is not
possible, such as in many current ARPANET hosts, the necessary editing
should be performed upon first entry to the internet mail system from
the local mail system. This avoids the problems mentioned above and
simplifies reply functions.

     In the case of ARPANET hosts, the editing operations assume that
all names in the form <anything>@<domain>, where <domain> is the name
of a domain, are unchanged. Names in the form <anything>@<host>,
where <host> is the name of a host in the ARPA domain, are transformed
to the form <anything>.<host>@ARPA. Anything else is an error.
Before handing off to an ARPANET NCP mailer, an ARPA MTP forwarder
might optionally transform <anything>.<host>@ARPA to <anything>@<host>
in order to reduce the forwarder traffic when local mail systems are
available. Similar situations might exist elsewhere.

8. Concluding Remarks

     This memorandum is intended to stimulate discussion, not simulate
it.