

NETWORK MAIL PATH SERVICE

STATUS OF THIS MEMO

This RFC proposes a new service for the ARPA-Internet community and requests discussion and suggestions for improvements. Distribution of this memo is unlimited.

INTRODUCTION

The network mail path service fills the current need of people to determine mailbox addresses for hosts that are not part of the ARPA-Internet but can be reached by one or more relay hosts that have Unix To Unix Copy (UUCP) mail, CSNET mail, MAILNET mail, BITNET mail, etc.

Anyone can use the service if they have TCP/TELNET to one of the hosts with a mail path server.

DISCUSSION

Currently many hosts that are not connected to the ARPA-Internet network can send mail to and receive mail from the ARPA-Internet community. The ARPA-Internet community sends mail using mailbox addresses of the form "user@host" or "local-part@domain" [1, 5]. In an effort to provide service to hosts not connected directly to the ARPA-Internet, mail maintainers have used the feature that the "local-part" of the mailbox address is locally interpreted to imbed specially encoded mail routing or relaying information. These encoded mailbox addresses have a variety of forms and have become common practice. For example:

demco%ucb-ean.cdn%ubc.csnet@CSNET-RELAY.ARPA

"Rudy.Nedved%CMCCTE@CARNEGIE.MAILNET"@MIT-MULTICS.ARPA

ihnp4!cmucsg!ern@UT-SALLY.ARPA

mss.dartmouth@CSNET-RELAY.ARPA

nedved%CMCCTF.BITNET@WISCVM.ARPA

It is important that people be able to communicate, but it is clear from the rampant confusion and frustration that something must be

provided to make it easier for people to address mail to non-ARPA-Internet hosts. The result, for a variety of reasons, has been the work and development of the Domain Name system and facilities [2, 3, 7, 9], and it is expected to make mailbox addresses be as simple as the current ARPA-Internet mailbox format (e.g., "user@domain").

How do people discover the special encoded addresses for non-ARPA-Internet host mailboxes until the domain name system is working and covering the majority of hosts in the mail world? The proposed solution to this problem is to provide a network service for the ARPA-Internet and a mail service for the non-ARPA-Internet hosts that, given a host and an optional addressing system or communication protocol or some other piece of information, supplies the mailbox address format for sending mail to that host. For example, "nedved@Carnegie.MAILNET" would be translated by the server to "nedved%Carnegie.MAILNET@MIT-MULTICS.ARPA". This memo covers the proposed network service.

DOCUMENT CONVENTIONS

Unless otherwise noted, all numbers are in decimal.

The term "host", as used in this document, describes one computer system which may have more than one name associated with it. It may have a name for each network or mail connection it supports and may have several nicknames or aliases for the computer and/or for each set of network names that the computer has acquired.

OVERVIEW

The network service is a connection based application on TCP [4]. A server listens for TCP connections on the assigned port of 117 [8]. It responds to the connection with a coded greeting message and waits for a command line. For each command line sent to the server, the server will respond with a coded message. The special command QUIT causes the server to respond with a coded closing message and closes the connection.

DESIGN GOALS

One of the goals is to provide the service to as many ARPA-Internet hosts as possible. In the current ARPA-Internet, experience has shown that software people first implement TELNET/TCP [6] before any other network application or protocol. Therefore, it is a sub-goal that people be able to access the service using available programs (with minimal modifications) that implement TELNET/TCP. Therefore, TELNET/TCP on port 117 will work correctly. The server understands TELNET options but refuses all option negotiations that disagree with the NVT characteristics defined by the TELNET protocol (see [6]), does not echo, and expects command lines to end with <CRLF> (ASCII code 13 (octal 15) followed by code 10 (octal 12)). Character echoing and line editing is expected to be handled by the user host for the benefit of the user.

Mail systems and other programs are also expected to be able to access and understand the service. Each command reply can have multiple line responses with text understandable by the novice user. Each command is encoded so as to make it easy for a program to parse the lines and extract interesting information, such as whether the operation was successful.

THE PROTOCOL

Given the developing nature of the protocol and its intent, the command lines are composed of a command (case ignored) followed by white space, the argument(s) and a <CRLF>. The white space is required if any arguments are supplied but the arguments are optional. White space following the command and any optional arguments are ignored.

<cmdline> := <cmd> [<WS> <args>] [<WS>] <CRLF>

<WS> := [<WS>] <WS> | <TAB> | <SPACE>

Coded response lines have the rigid format of a 3-digit decimal code followed by a space or a dash followed by text composed of characters within the ASCII range 32 to 126 (octal 40 to 176) with <CRLF> at the end of the line. The dash after the 3-digit code indicates at least one more response line will be supplied while the space indicates the current response line is the last one.

<rspline> := <digit><digit><digit><cont><rtext><CRLF>

<cont> := <SPACE> | "-"

<rtext> := ASCII characters in the range 32 to 126.

Some of the successful response text to certain commands have rigid formats so programs can extract path information. The commands that have format restrictions are clearly noted and the response format is documented with the command.

The response codes are in the range from 200 to 599 inclusively. The following paragraphs provide the break down for each digit.

The first, most significant, digit is the success indicator. It breaks down into the simple success and total failure responses but includes the ability to communicate a temporary failure condition and the need for further information that has worked so well for SMTP [5] and other similiar protocols. The codes are:

2xx Positive reply.

3xx Intermediate reply. Positive acknowledgement but more information is neccessary.

4xx Temporary error. Try again later.

5xx Permanent error. Do not retry.

The second digit is used to classify the response to provide a flavor for certain types of success. The flavor is apparent in providing the response on whether a host name is known by a domain name server or not. The codes are:

x0x Command related response.

x1x Connection related response.

x2x Database related response.

x3x Domain transition related response.

x4x Data added response.

x5x Data deleted response.

x6x Data modified response.

BASIC IMPLEMENTATION

The minimum implementation is the support of three commands: HELP, PATH and QUIT. The HELP command provides some level of documentation and possibly lists the known addressing or communication protocols. The PATH command takes as a required argument a user name or id followed by a "@", followed by a domain style host name whose domain components may be an addressing protocol, a communication environment, or an unofficial or colloquial domain.

```
S: (server listens on port 117)
U: (user connects to port 117)
S: 210-Welcome to the CMU network mail path service.
S: 210 Type 'HELP' for help.
U: help
S: 200-The server currently knows about the following mail worlds:
S: 200-      BITNET,UUCP,CSNET,.AC.UK,EARNET,JANET,CDNNET
S: 200-Use the PATH command with "user@host.world" to get the
S: 200 ARPA-Internet mail address.
U: path root@inria.uucp
S: 220 philabs!mcvax!inria!root@SEISMO.ARPA
U: quit
S: 211 Bye bye.
S: (server closes connection)
```

DETAILED PROTOCOL DESCRIPTION

The protocol is designed to provide a flexible but conservative mechanism for providing responses and adding experimental or extended commands.

<user connects to server>

The server responds with a message indicating the status of the server and optional information.

210 Greeting message indicating the server is ready.

410 The server is down for some unknown reason for a short time.

510 The server is unavailable.

HELP [<arg>]

The server can respond with general help information about the server, about the specific topic described by "arg", or it can

indicate that something is temporarily wrong with the HELP facility. It is strongly recommended that the general HELP command documentation be implemented and expanded.

200 General or specific documentation given.

220 Documentation given from a database.

421 Service temporarily unavailable.

501 Command not implemented or topic not known.

PATH <user>@<host>

The server normally responds with either the mail path that will work for the given mailbox address or indicates the domain style host name is unknown. If the database is in transition or inconsistent, a temporary or permanent error can be supplied.

220 Rigid format route given.

230 Rigid format route given. Domain servers should be used.

420 Database problems. Try again later.

501 Invalid argument form or null argument given.

520 No such host found in database.

521 Host name is ambiguous.

When a route is supplied with the 2xx success responses. It has a fixed format with a one-line response. The format is as follows:

<3-digit-code><SP><local-part>@<domain><CRLF>

The "local-part" and "domain" components are defined under the SMTP protocol [5] and are intended to be used over SMTP connections.

QUIT

Respond and close the server down.

211 Close the connection down.

One special code is reserved and is used for a special case. The code is 412 and is sent when the server has been waiting for a response for more than 2 minutes and has decided to timeout the connection. After the "412 <timeout msg>" is sent, the server may close or possibly abort the connection.

Because of the somewhat experimental nature of the server, additional commands are expected to be added as they become needed. No restrictions are placed on the names of these experimental commands other than they must not conflict with the basic commands and are not allowed to be abbreviated (i.e., "SEAR" can not be used for "SEARCH").

PATH COMMAND ARGUMENTS

It is important to understand that the server is an aid to users that may have minimal amount of information about the host. Therefore the PATH command takes domain style host names that may be complete or incomplete specifications for the host and may be common or colloquial domain names. The servers look through the entire database for anything that matches and try to find the best answer disregarding any local domain information. If several hosts have the same nickname or alias and lack distinguishing domain components, the server returns an error response containing all of the hosts found. Some implementation may even break down the host name and indicate in error messages that even though it did not find the host, it found something else that might be what the user wanted.

MAIL PATH SERVICE AND DOMAINS

As mentioned previously, the mail path service is not intended to be a replacement or a parallel service to the domain name system. It is a stop gap measure and, when most of the domain name system is in place, will probably be disabled on some or most of the hosts with the service.

Mail systems should check the domain name servers for the specified host before trying a mail path server. The mail path servers should be modified when one or more domain servers are in place to check if a host is part of the domain system and to generate an error or an indication (but still include the path information) if a host is found to be a part of the domain system.

The names used by the mail path servers have no official standing in the ARPA-Internet community and have colloquial origins. The domain name components are based on the administrative entities involved whereas many of the current unofficial common domain style names for

non-ARPA-Internet hosts are based on the protocol used, the relay host used, or some acronym that someone dreamed up. Only a few of the current domain style names that are privately in use are expected to be used by the ARPA-Internet community when the domain name service is in use by the majority of the ARPA-Internet community.

CAVEATS

The greatest problem with the new service, as implemented, is that it reports paths from the service host rather than from the user's host. This is due to the nature of software. It would be more convenient if it reported a correct path from the caller's host, but this would require a different method of database management (a method which could quickly compute the path from the caller's machine or a machine which would be willing to keep updated databases for each host (which is impractical)).

Two minor problems exist with the database used by the software. Many relay hosts exist in several different protocol or addressing name spaces but under different names. The current software cross referencing for the multiple protocol relay hosts is done by hand, but, given the seeming reliability of these relay hosts, the problem does not appear to be significant. The second problem is that the data should be collected from the actual relay hosts to ensure correctness, but in many cases this is impossible.

EXAMPLES

Find a route to CMU-CC-TE in the CARNEGIE part of MAILNET for user id ENOC:

```
S: (server listens on port 117)
U: (user connects to port 117)
S: 210-Welcome to the CMU network mail path service
S: 210 Type 'HELP' for help.
U: path ENOC@CMU-CC-TE.CARNEGIE.MAILNET
S: 220 ENOC%CMU-CC-TE%CARNEGIE.MAILNET@MIT-MULTICS.ARPA
U: quit
S: 211 Bye bye.
S: (server closes connection)
```

Find a route to a host which has an unknown addressing system or communication protocol and for which the name may be an alias:

```
S: (server listens on port 117)
U: (user connects to port 117)
S: 210-Welcome to the CMU network mail path service
S: 210 Type 'HELP' for help.
U: path mss@dartvax
S: 220 mss%dartmouth@CSNET-RELAY.ARPA
U: quit
S: 211 Bye bye.
S: (server closes connection)
```

Find a route to a host that is known by a very long domain style name but is not in the current ARPA-Internet host tables:

```
S: (server listens on port 117)
U: (user connects to port 117)
S: 210-Welcome to the CMU network mail path service
S: 210 Type 'HELP' for help.
U: path rob@vax1.cent.lanc.ac.uk
S: 220 rob%vax1.cent.lanc@UCL-CS.ARPA
U: quit
S: 211 Bye bye.
S: (server closes connection)
```

Find a route to a host without any additional information and the name is discovered to be ambiguous:

```
S: (server listens on port 117)
U: (user connects to port 117)
S: 210-Welcome to the CMU network mail path service
S: 210 Type 'HELP' for help.
U: path brad@pitt
S: 521-Several hosts found under the name of 'pitt', try one of:
S: 521-brad@pitt.UUCP
S: 521-brad@pitt.CSNET
U: path brad@pitt.CSNET
S: 220 brad%pitt@CSNET-RELAY.ARPA
U: quit
S: 211 Bye bye.
S: (server closes connection)
```

ACKNOWLEDGEMENTS

The original protocol was documented by Marc Elvy for a server that he and Alan Langerman built. The server used the pathalias software created by Steve Bellovin, as modified by Peter Honeyman and Robert T. Morris, to maintain the host to host connection database. The software provided a way for people to make sense out of the jungle of UUCP hosts. The Info-Nets@MIT-MC mailing list, created and maintained by Robert Krawitz, made the CMU and Harvard mail path projects aware of each other and the people on the list provided many of the mail relay databases that are in use by the mail path servers. The original server may be accessed through TCP port 117 on harvard.arpa -- the "pathto" program that runs under 4.2BSD UNIX may be obtained as a front end to the server from RFC915@HARVARD.ARPA.

The current protocol scope was changed by Rudy Nedved to cover BITNET, CSNET, MAILNET and other "mail networks" and further refined by Marc Elvy, Alan Langerman and others.

Comments should be sent to RFC-915@HARVARD.ARPA or mailed (via the U.S. Postal Service) to:

Marc A. Elvy
108 Aiken Computation Laboratory
33 Oxford Street
Harvard University
Cambridge, MA 02138

(617) 495-5849

Rudy Nedved
Department of Computer Science
Carnegie-Mellon University
Schenley Park
Pittsburgh, PA 15213

(412) 578-7685

REFERENCES

- [1] Crocker, D. "Standard for the Format of ARPA Internet Text Messages". RFC 822, Department of Electrical Engineering, University of Delaware, August, 1982.
- [2] Mockapetris, P. "Domain Names - Concepts and Facilities". RFC 882, USC/Information Sciences Institute, November, 1983.
- [3] Mockapetris, P. "Domain Names - Implementation Specification". RFC 883, USC/Information Sciences Institute, November, 1983.
- [4] Postel, J. "Transmission Control Protocol- DARPA Internet Program Protocol Specification". RFC 793, USC/Information Sciences Institute, September, 1981.
- [5] Postel, J. "Simple Mail Transfer Protocol". RFC 821, USC/Information Sciences Institute, August, 1982.
- [6] Postel, J., and J. Reynolds. "Telnet Protocol Specification". RFC 854, USC/Information Sciences Institute, May, 1983.
- [7] Postel, J. "Domain Name System Implementation Schedule". RFC 897, USC/Information Sciences Institute, February, 1984.
- [8] Reynolds, J., and J. Postel. "Assigned Numbers". RFC 923, USC/Information Sciences Institute, October, 1984.
- [9] Su, Z., and Postel, J. "The Domain Naming Convention for Internet User Applications". RFC 819, SRI International, August, 1982.

