

Mark Crispin
SU-AI
2 January 1979

[illegible]

system at Stanford, where it eventually replaced the former host table. Recently, support for multiple networks has been added, including allowing hosts to be on more than one network, and a more flexible compiler than assembler macros was written (the program which generates the binary file now does the compiling).

THE HOST TABLE:

In the descriptions below, angle brackets and lower case are used as a meta-linguistic device. It should be noted that spaces and tabs are often ignored and may be used freely in the source format, while commas are always explicit delimiters. In addition, semicolon always begins a commentary and everything after the semicolon on the line is ignored; however, any text before the semicolon on the line is processed as usual. The syntax rules should be obvious by examining the text of the host table in the appendix. Names are alphanumeric strings, consisting of the set (A-Z, 0-9, and - (i.e., dash)). Quoting is used to separate examples from the text and is not part of the example.

The host table consists of commentary and two types of text lines. The commentary lines begin with a semicolon and are ignored by the compiler. They are intended to provide information for a human reader or editor of the host table. The commentary lines may be in mixed case, however the text lines are by tradition entirely in upper case. There are two kinds of text lines: host and network.

Network text lines begin with the word "NET" followed by a space or tab. These specify a network name and the network number (as assigned by Postel) for that network. As there currently are no officially assigned network names, suitable names were assigned more or less based on the English names in Postel's "Assigned Numbers", RFC 750. These names may be changed in the future (however, some software has come to depend on the names ARPA, CHAOS, and DIAL for the ARPANET, Chaos net, and Dialnet).

The format of a network text line is:

NET <name>,<decimal-number>

For example, the ARPANET's entry would look something like:

NET ARPA, 10

Host text lines begin with the word "HOST" followed by a space or tab. These specify a host name, a host address list, whether this host is a "user" or a "server", the name of the host's operating system, the name of the host's machine type, and a nick name list. The operating system, machine type, and/or nick name list may be omitted, in which case they default to unknown or null.

The host name is a unique name string for that host. For ARPANET sites, it is the official name for that host as assigned by the NIC. For other networks, it is whatever name is in common use on that network. So far there haven't been any naming conflicts on multiple-network hosts.

The host address list is either a single host address, or a list of host addresses in square brackets and delimited by commas. A host address consists of a network name, a space, and the host's address on that network. If the network name is not specified, it defaults to ARPA (i.e., "ARPA 0/11" and "0/11" are equivalent). Different networks parse host addresses in different ways:

ARPANET addresses are in BBN-style host number slash IMP number notation, with both numbers being decimal. Hence host 2 on IMP 6 is represented as "2/6". Of course, this format is backwards, but it has become enough of a network standard to force its use. Old-style octal addresses are allowed (e.g. "206" for "2/6") but are no longer used or supported.

CHAOS net addresses are a single octal number, e.g. "CHAOS 2026", and specify the host's address on the CHAOS net.

Dialnet addresses are a ten-digit decimal number, and specify the TelCo (phone) number of the host's Dialnet port.

The definition of user vs. server is generally taken to mean "according to the NIC" for ARPANET hosts. A server is considered to be a host for which making a connection to a remote service is a meaningful operation. For some hosts with limited servers, the definition often is changed from the official one, depending upon the individual circumstances. For example, "users" who have an FTP server and occasionally a TELNET server may be called "servers". On the other hand a "server" which does not accept MAIL and rejects MAIL in a pathological way (e.g. by hanging) might be labelled a "user".

The name of the host's operating system is a string much as the host name is, such as "ITS", "TOPS-20", or "MULTICS". Some subsystems use this information to predict certain behavior of the remote server. For example, a MAIL user subsystem knows that for operating system "MULTICS" it has to log in as user NETML before attempting to deliver the mail.

The name of the host's machine type is a string as well. For the convenience of several subsystems, all DEC "PDP-n" machines are entered without the dash, and all PDP-10 like machines (e.g., KL-20, MAXC, etc.) are considered to be PDP-10's, which by the way gets entered as "PDP10" since that is a single 36-bit word in 7-bit ASCII. Like the operating system name, several subsystems use this information as well. For example, a PDP-10 FTP user process will try to negotiate 36-bit image

The nick name list is in square brackets and consists of a series of names delimited by commas. There may be any number of nick names.

HOST <name>,<address-list>,<status>,<system>,<machine>,<nickname-list>

HOST MIT-AI, [ARPA 2/6, CHAOS 2026], SERVER, ITS, PDP10, [AI, MITAI]

THE HOST TABLE BINARY FILE:

FILE HEADER:

```
word 0      The name of this file in SIXBIT.  Currently HOSTS2.
word 1      The name of the source file in SIXBIT.  Always HOSTS.
word 2      The version of the source file in SIXBIT  if compiled on
an ITS site, otherwise the name of the site in SIXBIT.
word 3      The directory name of the source, usually in SIXBIT.
word 4      The name of the site in SIXBIT.
word 5      The user name who compiled the file, usually in SIXBIT.
word 6      Date of compilation as SIXBIT YYMMDD.
word 7      Time of compilation as SIXBIT HHMMSS.
word 8      Address in file of NAME table.
word 9      Address in file of SITE table.
word 10     Address in file of NETWORK table.
            <words after this are reserved for future use>
```

```
word 0      Number of entries in table.
word 1      Number of words per entry, currently 2.
  entry word 0  Network number assigned by Postel.
  entry word 1  Left half: Address in file of name of network in ASCIZ.
                Right half: Address in file of network's ADDRESS table
                (zero means no ADDRESS table, i.e. no hosts).
```



```

word 0      Number of entries in table.
word 1      Number of words per entry, currently 2.
entry word 0 Network address of this entry, including network number.
              For ARPANET addresses this is in the format:
              xxx000,,000000  Network number
              000xxx,,xxx000  IMP number
              000000,,000xxx  Host number
              Each number is right justified.
              For CHAOS net addresses it is an octal number.
              For Dialnet addresses it is the address in the file of
              the TelCo number in ASCIZ.
entry word 1 Left half: Address in file of SITE table entry.
              Right half: Address in file of next ADDRESS table entry
              for this site (zero means end of list).
```

word 0	Number of entries in table.
word 1	Number of words per entry, currently 3.
entry word 0	Left half: Address in file of official name in ASCIZ. Right half: Address in file of first ADDRESS table entry for this site.
entry word 1	Left half: Address in file of operating system name in ASCIZ (zero means unknown). Right half: Address in file of machine type in ASCIZ (zero means unknown).
entry word 2	Left half: Flags. The 400000 bit means a server site. Right half: reserved

```
word 0      Number of entries in table.
word 1      Number of words per entry, currently 1.
  entry word 0 Left half: Address in file of SITE table entry for  this
               host.
               Right half: Address in file of host name in ASCIIZ.
```

A host table capable of supporting the full host addressing of the ARPANET and additional networks has been presented, along with a binary file format for efficient manipulation of this host table data.

We are documenting this format in order to present it to the outside world as a suggested replacement for the current host table. The advantage of our host table is that it has already been implemented and is in use at MIT and Stanford. We have established some conventions

