

NCP/TCP TRANSITION PLAN

Introduction

ARPA sponsored research on computer networks led to the development of the ARPANET. The installation of the ARPANET began in September 1969, and regular operational use was underway by 1971. The ARPANET has been an operational service for at least 10 years. Even while it has provided a reliable service in support of a variety of computer research activities, it has itself been a subject of continuing research, and has evolved significantly during that time.

In the past several years ARPA has sponsored additional research on computer networks, principally networks based on different underlying communication techniques, in particular, digital packet broadcast radio and satellite networks. Also, in the ARPA community there has been significant work on local networks.

It was clear from the start of this research on other networks that the base host-to-host protocol used in the ARPANET was inadequate for use in these networks. In 1973 work was initiated on a host-to-host protocol for use across all these networks. The result of this long effort is the Internet Protocol (IP) and the Transmission Control Protocol (TCP).

These protocols allow all hosts in the interconnected set of these networks to share a common interprocess communication environment. The collection of interconnected networks is called the ARPA Internet (sometimes called the "Catenet").

The Department of Defense has recently adopted the internet concept and the IP and TCP protocols in particular as DoD wide standards for all DoD packet networks, and will be transitioning to this architecture over the next several years. All new DoD packet networks will be using these protocols exclusively.

The time has come to put these protocols into use in the operational ARPANET, and extend the logical connectivity of the ARPANET hosts to include hosts in other networks participating in the ARPA Internet.

As with all new systems, there will be some aspects which are not as robust and efficient as we would like (just as with the initial ARPANET). But with your help, these problems can be solved and we

can move into an environment with significantly broader communication services.

Discussion

The implementation of IP/TCP on several hosts has already been completed, and the use of some services is underway. It is urgent that the implementation of IP/TCP be begun on all other ARPANET hosts as soon as possible and no later than 1 January 1982 in any case. Any new host connected to the ARPANET should only implement IP/TCP and TCP-based services. Several important implementation issues are discussed in the last section of this memo.

Because all hosts can not be converted to TCP simultaneously, and some will implement only IP/TCP, it will be necessary to provide temporarily for communication between NCP-only hosts and TCP-only hosts. To do this certain hosts which implement both NCP and IP/TCP will be designated as relay hosts. These relay hosts will support Telnet, FTP, and Mail services on both NCP and TCP. These relay services will be provided beginning in November 1981, and will be fully in place in January 1982.

Initially there will be many NCP-only hosts and a few TCP-only hosts, and the load on the relay hosts will be relatively light. As time goes by, and the conversion progresses, there will be more TCP capable hosts, and fewer NCP-only hosts, plus new TCP-only hosts. But, presumably most hosts that are now NCP-only will implement IP/TCP in addition to their NCP and become "dual protocol" hosts. So, while the load on the relay hosts will rise, it will not be a substantial portion of the total traffic.

The next section expands on this plan, and the following section gives some milestones in the transition process. The last section lists the key documents describing the new protocols and services. Appendices present scenarios for use of the relay services.

The General Plan

The goal is to make a complete switch over from the NCP to IP/TCP by 1 January 1983.

It is the task of each host organization to implement IP/TCP for its own hosts. This implementation task must begin by 1 January 1982.

IP:

This is specified in RFCs 791 and 792. Implementations exist for several machines and operating systems. (See Appendix D.)

TCP:

This is specified in RFC793. Implementations exist for several machines and operating systems. (See Appendix D.)

It is not enough to implement the IP/TCP protocols, the principal services must be available on this IP/TCP base as well. The principal services are: Telnet, File Transfer, and Mail.

It is the task of each host organization to implement the principal services for its own hosts. These implementation tasks must begin by 1 January 1982.

Telnet:

This is specified in RFC 764. It is very similar to the Telnet used with the NCP. The primary differences are that the ICP is eliminated, and the NCP Interrupt is replaced with the TCP Urgent.

FTP:

This is specified in RFC 765. It is very similar to the FTP used with the NCP. The primary differences are that in addition to the changes for Telnet, that the data channel is limited to 8-bit bytes so FTP features to use other transmission byte sizes are eliminated.

Mail:

This is specified in RFC 788. Mail is separated completely from FTP and handled by a distinct server. The procedure is similar in concept to the old FTP/NCP mail procedure, but is very different in detail, and supports additional functions -- especially mail relaying, and multi-recipient delivery.

Beyond providing the principal services in the new environment, there must be provision for interworking between the new environment and the old environment between now and January 1983.

For Telnet, there will be provided one or more relay hosts. A Telnet relay host will implement both the NCP and TCP environments and both user and server Telnet in both environments. Users requiring Telnet service between hosts in different environments

will first connect to a Telnet relay host and then connect to the destination host. (See Appendix A.)

For FTP, there will be provided one or more relay hosts. An FTP relay host will implement both the NCP and TCP environments, both user and server Telnet, and both user and server FTP in both environments. Users requiring FTP service between hosts in different environments will first connect via Telnet to an FTP relay host, then use FTP to move the file from the file donor host to the FTP relay host, and finally use FTP to move the file from the FTP relay host to the file acceptor host. (See Appendix B.)

For Mail, hosts will implement the new Simple Mail Transfer Protocol (SMTP) described in RFC 788. The SMTP procedure provides for relaying mail among several protocol environments. For TCP-only hosts, using SMTP will be sufficient. For NCP-only hosts that have not been modified to use SMTP, the special syntax "user.host@forwarder" may be used to relay mail via one or more special forwarding host. Several mail relay hosts will relay mail via SMTP procedures between the NCP and TCP environments, and at least one special forwarding host will be provided. (See Appendix C.)

Milestones

First Internet Service	already
------------------------	---------

A few hosts are TCP-capable and use TCP-based services.

First TCP-only Host	already
---------------------	---------

The first TCP-only host begins use of TCP-based services.

Telnet and FTP Relay Service	already
------------------------------	---------

Special relay accounts are available to qualified users with a demonstrated need for the Telnet or FTP relay service.

Ad Hoc Mail Relay Service	already
---------------------------	---------

An ad hoc mail relay service using the prototype MTP (RFC 780) is implemented and mail is relayed from the TCP-only hosts to NCP-only hosts, but not vice versa. This service will be replaced by the SMTP service.

Last NCP Conversion Begins	Jan 82
----------------------------	--------

The last NCP-only host begins conversion to TCP.

Mail Relay Service Jan 82

The SMTP (RFC 788) mail service begins to operate and at least one mail relay host is operational, and at least one special forwarder is operational to provide NCP-only host to TCP-only host mail connectivity.

Normal Internet Service Jul 82

Most hosts are TCP-capable and use TCP-based services.

Last NCP Conversion Completed Nov 82

The last NCP-only host completes conversion to TCP.

Full Internet Service Jan 83

All hosts are TCP-capable and use TCP-based services. NCP is removed from service, relay services end, all services are TCP-based.

Documents

The following RFCs document the protocols to be implemented in the new IP/TCP environment:

IP	RFC 791
ICMP	RFC 792
TCP	RFC 793
Telnet	RFC 764
FTP	RFC 765
SMTP	RFC 788
Name Server	IEN 116
Assigned Numbers	RFC 790

These and associated documents are to be published in a notebook, and other information useful to implementers is to be gathered. These documents will be made available on the following schedule:

Internet Protocol Handbook	Jan 82
Implementers Hints	Jan 82
SDC IP/TCP Specifications	Jan 82
Expanded Host Table	Jan 82

Implementation Issues

There are several implementation issues that need attention, and there are some associated facilities with these protocols that are not necessarily obvious. Some of these may need to be upgraded or redesigned to work with the new protocols.

Name Tables

Most hosts have a table for converting character string names of hosts to numeric addresses. There are two effects of this transition that may impact a host's table of host names: (1) there will be many more names, and (2) there may be a need to note the protocol capability of each host (SMTP/TCP, SMTP/NCP, FTP/NCP, etc.).

Some hosts have kept this table in the operating system address space to provide for fast translation using a system call. This may not be practical in the future.

There may be applications that could take alternate actions if they could easily determine if a remote host supported a particular protocol. It might be useful to extend host name tables to note which protocols are supported.

It might be necessary for the host name table to contain names of hosts reachable only via relays if this name table is used to verify the spelling of host names in application programs such as mail composition programs.

It might be advantageous to do away with the host name table and use a Name Server instead, or to keep a relatively small table as a cache of recently used host names.

A format, distribution, and update procedure for the expanded host table will be published soon.

Mail Programs

It may be possible to move to the new SMTP mail procedures by changing only the mailer-daemon and implementing the SMTP-server, but in some hosts there may be a need to make some small changes to some or all of the mail composition programs.

There may be a need to allow users to identify relay hosts for messages they send. This may require a new command or address syntax not now currently allowed.

IP/TCP

Continuing use of IP and TCP will lead to a better understanding of the performance characteristics and parameters. Implementers should expect to make small changes from time to time to improve performance.

Shortcuts

There are some very tempting shortcuts in the implementation of IP and TCP. DO NOT BE TEMPTED! Others have and they have been caught! Some deficiencies with past implementations that must be remedied and are not allowed in the future are the following:

IP problems:

Some IP implementations did not verify the IP header checksum.

Some IP implementations did not implement fragment reassembly.

Some IP implementations used static and limited routing information, and did not make use of the ICMP redirect message information.

Some IP implementations did not process options.

Some IP implementations did not report errors they detected in a useful way.

TCP problems:

Some TCP implementations did not verify the TCP checksum.

Some TCP implementations did not reorder segments.

Some TCP implementations did not protect against silly window syndrome.

Some TCP implementations did not report errors they detected in a useful way.

Some TCP implementations did not process options.

Host problems:

Some hosts had limited or static name tables.

Relay Service

The provision of relay services has started. There are two concerns about the relay service: (1) reliability, and (2) load.

The reliability is a concern because relaying puts another host in the chain of things that have to all work at the same time to get the job done. It is desirable to provide alternate relay hosts if possible. This seems quite feasible for mail, but it may be a bit sticky for Telnet and FTP due to the need for access control of the login accounts.

The load is a potential problem, since an overloaded relay host will lead to unhappy users. This is another reason to provide a number of relay hosts, to divide the load and provide better service.

A Digression on the Numbers

How bad could it be, this relay load? Essentially any "dual protocol" host takes itself out of the game (i.e., does not need relay services). Let us postulate that the number of NCP-only hosts times the number of TCP-only hosts is a measure of the relay load.

Total Hosts	Dual Hosts	NCP Hosts	TCP Hosts	"Load"	Date
200	20	178	2	356	Jan-82
210	40	158	12	1896	Mar-82
220	60	135	25	3375	May-82
225	95	90	40	3600	Jul-82
230	100	85	45	3825	Sep-82
240	125	55	60	3300	Nov-82
245	155	20	70	1400	Dec-82
250	170	0	80	0	31-Dec-82
250	0	0	250	0	1-Jan-83

This assumes that most NCP-only hosts (but not all) will become to dual protocol hosts, and that 50 new host will show up over the course of the year, and all the new hosts are TCP-only.

If the initial 200 hosts immediately split into 100 NCP-only and 100 TCP-only then the "load" would be 10,000, so the fact that most of the hosts will be dual protocol hosts helps considerably.

This load measure (NCP hosts times TCP hosts) may over state the load significantly.

Please note that this digression is rather speculative!

Gateways

There must be continuing development of the internet gateways.
The following items need attention:

- Congestion Control via ICMP

- Gateways use connected networks intelligently

- Gateways have adequate buffers


- Gateways have fault isolation instrumentation

Note that the work in progress on the existing gateways will provide the capability to deal with many of these issues early in 1982. Work is also underway to provide improved capability gateways based on new hardware late in 1982.

APPENDIX A. Telnet Relay Scenario


Suppose a user at a TCP-only host wishes to use the interactive services of an NCP-only service host.

- 1) Use the local user Telnet program to connect via Telnet/TCP to the RELAY host.
- 2) Login on the RELAY host using a special account for the relay service.
- 3) Use the user Telnet on the RELAY host to connect via Telnet/NCP to the service host. Since both Telnet/TCP and Telnet/NCP are available on the RELAY host the user must select which is to be used in this step.
- 4) Login on the service host using the regular account.



Suppose a user at a NCP-only host wishes to use the interactive services of an TCP-only service host.

- 1) Use the local user Telnet program to connect via Telnet/NCP to the RELAY host.
- 2) Login on the RELAY host using a special account for the relay service.
- 3) Use the user Telnet on the RELAY host to connect via Telnet/TCP to the service host. Since both Telnet/TCP and Telnet/NCP are available on the RELAY host the user must select which is to be used in this step.
- 4) Login on the service host using the regular account.

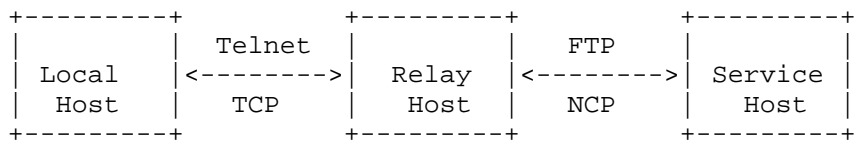


APPENDIX B. FTP Relay Scenario

Suppose a user at a TCP-only host wishes copy a file from a NCP-only donor host.

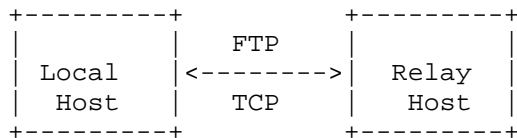
Phase 1:

- 1) Use the local user Telnet program to connect via Telnet/TCP to the RELAY host.
- 2) Login on the RELAY host using a special account for the relay service.
- 3) Use the user FTP on the RELAY host to connect via FTP/NCP to the donor host.
- 4) FTP login on the donor host using the regular account.
- 5) Copy the file from the donor host to the RELAY host.
- 6) End the FTP session, and disconnect from the donor host.
- 7) Logout of the RELAY host, close the Telnet/TCP connection, and quit Telnet on the local host.



Phase 2:

- 1) Use the local user FTP to connect via FTP/TCP to the RELAY host.
- 2) FTP login on the RELAY host using the special account for the relay service.
- 3) Copy the file from the RELAY host to the local host, and delete the file from the RELAY host.
- 4) End the FTP session, and disconnect from the RELAY host.




Note that the relay host may have a policy of deleting files more than a few hours or days old.

APPENDIX C. Mail Relay Scenario


Suppose a user on a TCP-only host wishes to send a message to a user on an NCP-only host which has implemented SMTP.

- 1) Use the local mail composition program to prepare the message. Address the message to the recipient at his or her host. Tell the composition program to queue the message.
- 2) The background mailer-daemon finds the queued message. It checks the destination host name in a table to find the internet address. Instead it finds that the destination host is a NCP-only host. The mailer-daemon then checks a list of mail RELAY hosts and selects one. It send the message to the selected mail RELAY host using the SMTP procedure.
- 3) The mail RELAY host accepts the message for relaying. It checks the destination host name and discovers that it is a NCP-only host which has implemented SMTP. The mail RELAY host then sends the message to the destination using the SMTP/NCP procedure.



Suppose a user on a TCP-only host wishes to send a message to a user on an NCP-only non-SMTP host.

- 1) Use the local mail composition program to prepare the message. Address the message to the recipient at his or her host. Tell the composition program to queue the message.
- 2) The background mailer-daemon finds the queued message. It checks the destination host name in a table to find the internet address. Instead it finds that the destination host is a NCP-only host. The mailer-daemon then checks a list of mail RELAY hosts and selects one. It send the message to the selected mail RELAY host using the SMTP procedure.
- 3) The mail RELAY host accepts the message for relaying. It checks the destination host name and discovers that it is a NCP-only non-SMTP host. The mail RELAY host then sends the message to the destination using the old FTP/NCP mail procedure.



Suppose a user on a NCP-only non-SMTP host wishes to send a message to a user on an TCP-only host. Suppose the destination user is "Smith" and the host is "ABC-X".

- 1) Use the local mail composition program to prepare the message. Address the message to "Smith.ABC-X@FORWARDER". Tell the composition program to queue the message.
- 2) The background mailer-daemon finds my queued message. It sends the message to host FORWARDER using the old FTP/NCP mail procedure.
- 3) The special forwarder host converts the "user name" supplied by the FTP/NCP mail procedure (in the MAIL or MLFL command) to "Smith@ABC-X" (in the SMTP RCTP command) and queues the message to be processed by the SMTP mailer-daemon program on this same host. No conversion of the mailbox addresses in made in thr message header or body.
- 4) The SMTP mailer-daemon program on the forwarder host finds this queued message and checks the destination host name in a table to find the internet address. It finds the destination address and send the mail using the SMTP procedure.



APPENDIX D. IP/TCP Implementation Status

Please note that the information in this section may become quickly dated. Current information on the status of IP and TCP implementations can be obtained from the file <INTERNET-NOTEBOOK>TCP-IP-STATUS.TXT on ISIF.

BBN C70 UNIX

Date: 18 Nov 1981
From: Rob Gurwitz <gurwitz at BBN-RSM>

The C/70 processor is a BBN-designed system with a native instruction set oriented toward executing the C language. It supports UNIX Version 7 and provides for user processes with a 20-bit address space. The TCP/IP implementation for the C/70 was ported from the BBN VAX TCP/IP, and shares all of its features.

This version of TCP/IP is running experimentally at BBN, but is still under development. Performance tuning is underway, to make it more compatible with the C/70's memory management system.

BBN GATEWAYS

Date: 19 Nov 1981
From: Alan Sheltzer <sheltzer at BBN-UNIX>

In an effort to provide improved service in the gateways controlled by BBN, a new gateway implementation written in macro-11 instead of BCPL is being developed. The macro-11 gateway will provide users with internet service that is functionally equivalent to that provided by the current BCPL gateways with some performance improvements.

ARPANET/SATNET gateway at BBN (10.3.0.40),
ARPANET/SATNET gateway at NDRE (10.3.0.41),
Comsat DCN Net/SATNET gateway at COMSAT (4.0.0.39),
SATNET/UCL Net/RSRE Net gateway at UCL (4.0.0.60),
PR Net/RCC Net gateway at BBN (3.0.0.62),
PR Net/ARPANET gateways at SRI (10.3.0.51, 10.1.0.51),
PR Net/ARPANET gateway at Ft. Bragg (10.0.0.38).

BBN H316 and C/30 TAC

Date: 18 November 1981
From: Bob Hinden <Hinden@BBN-UNIX>

The Terminal Access Controller (TAC) is user Telnet host that supports TCP/IP and NCP host to host protocols. It runs in 32K H-316 and 64K C/30 computers. It supports up to 63 terminal ports. It connects to a network via an 1822 host interface.

For more information on the TAC's design, see IEN-166.

BBN HP-3000

Date: 14 May 1981
From: Jack Sax <sax@BBN-UNIX>

The HP3000 TCP code is in its final testing stages. The code includes under the MPE IV operating system as a special high priority process. It is not a part of the operating system kernel because MPE IV has no kernel. The protocol process includes TCP, IP, 1822 and a new protocol called HDH which allows 1822 messages to be sent over HDLC links. The protocol process has about 8k bytes of code and at least 20k bytes of data depending on the number of buffers allocated.

In addition to the TCP the HP3000 has user and server TELNET as well as user FTP. A server FTP may be added later.

A complete description of the implementation software can be found in IEN-167.

BBN PDP-11 UNIX

Date: 14 May 1981
From: Jack Haverty <haverty@BBN-UNIX>

This TCP implementation was written in C. It runs as a user process in version 6 UNIX, with modifications added by BBN for network access. It supports user and server Telnet.

This implementation was done under contract to DCEC. It is installed currently on several PDP-11/70s and PDP-11/44s. Contact Ed Cain at DCEC <cain@EDN-UNIX> for details of further development.

BBN TENEX & TOPS20

Date: 23 Nov 1981
From: Charles Lynn <CLynn@BBNA>

TCP4 and IP4 are available for use with the TENEX operating system running on a Digital KA10 processor with BBN pager. TCP4 and IP4 are also available as part of TOPS20 Release 3A and Release 4 for the Digital KL10 and KL20 processors.

Above the IP layer, there are two Internet protocols within the monitor itself (TCP4 and GGP). In addition up to eight (actually a monitor assembly parameter) protocols may be implemented by user-mode programs via the "Internet User Queue" interface. The GGP or Gateway-Gateway Protocol is used to receive advice from Internet Gateways in order to control message flow. The GGP code is in the process of being changed and the ICMP protocol is being added.

TCP4 is the other monitor-supplied protocol and it has two types of connections -- normal data connections and "TCP Virtual Terminal" (TVT) connections. The former are used for bulk data transfers while the latter provide terminal access for remote terminals.

Note that TVTs use the standard ("New") TELNET protocol. This is identical to that used on the ARPANET with NCP and in fact, is largely implemented by the same code.

Performance improvements, support for the new address formats, and User and Server FTP processes above the TCP layer are under development.

BBN VAX UNIX

Date: 18 Nov 1981
From: Rob Gurwitz <gurwitz at BBN-RSM>

The VAX TCP/IP implementation is written in C for Berkeley 4.1BSD UNIX, and runs in the UNIX kernel. It has been run on VAX 11/780s and 750s at several sites, and is due to be generally available in early 1982.

The implementation conforms to the TCP and IP specifications (RFC 791, 793). The implementation supports the new extended internet address formats, and both GGP and ICMP. It also supports multiple network access protocols and device drivers. Aside from ARPANET 1822 and the ACC LH/DH-11 driver, experimental drivers have also been developed for ETHERNET. There are user interfaces for

accessing the IP and local network access layers independent of the TCP.

Higher level protocol services include user and server TELNET, MTP, and FTP, implemented as user level programs. There are also tools available for monitoring and recording network traffic for debugging purposes.

Continuing development includes performance enhancements. The implementation is described in IEN-168.

COMSAT

Date: 30 Apr 1980
From: Dave Mills <Mills@ISIE>

The TCP/IP implementation here runs in an LSI-11 with a homegrown operating system compatible in most respects to RT-11. Besides the TCP/IP levels the system includes many of the common high-level protocols used in the ARPANET community, such as TELNET, FTP and XNET.

DCEC PDP-11 UNIX

Date: 23 Nov 1981
From: Ed Cain <cain@EDN-UNIX>

This TCP/IP/ICMP implementation runs as a user process in version 6 UNIX, with modifications obtained from BBN for network access. IP reassembles fragments into datagrams, but has no separate IP user interface. TCP supports user and server Telnet, echo, discard, internet mail, and a file transfer service. ICMP generates replies to Echo Requests, and sends Source-Quench when reassembly buffers are full.

Hardware - PDP-11/70 and PDP-11/45 running UNIX version 6, with BBN IPC additions. Software - written in C, requiring 25K instruction space, 20K data space. Supports 10 connections.

DTI VAX

Date: 15 May 1981
From: Gary Grossman <grg@DTI>

Digital Technology Incorporated (DTI) IP/TCP for VAX/VMS

The following describes the IP and TCP implementation that DTI plans to begin marketing in 4th Quarter 1981 as part of its VAX/VMS network software package.

Hardware: VAX-11/780 or /750. Operating System: DEC standard VAX/VMS Release 2.0 and above. Implementation Language: Mostly C, with some MACRO. Connections supported: Maximum of 64.

User level protocols available: TELNET, FTP, and MTP will be available. (The NFE version uses AUTODIN II protocols.)

MIT MULTICS

Date: 13 May 1981
From: Dave Clark <Clark@MIT-Multics>

Multics TCP/IP is implemented in PL/1 for the HISI 68/80. It has been in experimental operation for about 18 months; it can be distributed informally as soon as certain modifications to the system are released by Honeywell. The TCP and IP package are currently being tuned for performance, especially high throughput data transfer.

Higher level services include user and server telnet, and a full function MTP mail forwarding package.

The TCP and IP contain good logging and debugging facilities, which have proved useful in the checkout of other implementations. Please contact us for further information.

SRI LSI-11

Date: 15 May 1981
From: Jim Mathis <mathis.tscb@Sri-Unix>

The IP/TCP implementation for the Packet Radio terminal interface unit is intended to run on an LSI-11 under the MOS real-time operating system. The TCP is written in MACRO-11 assembler language. The IP is currently written in assembler language; but is being converted into C. There are no plans to convert the TCP from assembler into C.

The TCP implements the full specification. The TCP appears to be functionally compatible with all other major implementations. In particular, it is used on a daily basis to provide communications between users on the Ft. Bragg PRNET and ISID on the ARPANET.

The IP implementation is reasonably complete, providing fragmentation and reassembly; routing to the first gateway; and a complete host-side GGP process.

A measurement collection mechanism is currently under development to collect TCP and IP statistics and deliver them to a measurement host for data reduction.

UCLA IBM

Date: 13 May 1981
From: Bob Braden <Braden@ISIA>

Hardware: IBM 360 or 370, with a "Santa Barbara" interface to the IMP.

Operating System: OS/MVS with ACF/VTAM. An OS/MVT version is also available. The UCLA NCP operates as a user job, with its own internal multiprogramming and resource management mechanisms.

Implementation Language: BAL (IBM's macro assembly language)

User-Level Protocols Available: User and Server Telnet

