

MINUTES OF NETWORK WORKING GROUP MEETING
5/16 through 5/19/71

Preface

These notes are for reference and recall by those in attendance of the NWG meetings. No attempt has been made toward completeness to make this an understandable document for those not in attendance.

The notes are ordered chronologically. You may notice discrepancies for particular schedules and tasks within the notes; the discrepancies represent a revision of those schedules and tasks, thus those dates given more recent in time are assumed to apply.

If you detect any gross errors in this report, please make corrections via the accepted NIC procedures.

CONTENTS

I. SUNDAY EVENING SESSION (5/10/71)	4
Introduction of Attendees	4
Site Status Reports	6
UCLA-Sigma 7	6
UCLA-CCN	6
UCSB	7
SRI-ARC/NIC	7
SRI-AI	7
Rand	8
SDC	8
Illinois-CAC	9
AMES	9
CCA	10
Case Western	10
Carnegie	10
Harvard	11
IBM Research	11
RADC	11
MIT-DM and MULTICS	11
Lincoln	12
BBN-NCC	12
BBN/TENEX	13
Mitre	13

NBS	14
ETAC	14
Air Force Sites	14
Other Reports	15
Dept. Comm., Canada	15
U. of Chicago	15
United Kingdom	15
Merit-Univ. Michigan	15
EDUCOM	16
Raytheon	16
Miscellaneous Topics	16
Graphics	16
NCP Protocols	16
IMLAC Users Group	17
Official Document Formats	17
II. MONDAY MORNING SESSION (5/17/71)	18
Network Information Center	18
Plans for NIC	18
Concepts & Recommendations for Documentation ...	18
TELNET	19
III. MONDAY AFTERNOON SESSION (5/17/71)	20
File Transfer Protocol (RFC #114)	20
File Protocol Status Report	20
Miscellaneous Topics	20
Sockets	20
Initial Connection Protocol	21
Testing and Validation	21
IV. MONDAY EVENING SESSION (5/17/71)	22
Operating Systems and Networks	22
V. TUESDAY MORNING SESSION (5/18/71)	24
DRS Working Group Meeting with Open Attendance ...	24
Data Management on Computer Networks	24
Open Discussion on Data Management	25

VI.	TUESDAY EVENING SESSION (5/18/71)	27
	Terminal IMP	27
	Comments by Dr. Roberts	27
VII.	WEDNESDAY MORNING SESSION (5/19/71)	30
VIII.	WEDNESDAY EVENING SESSION (5/19/71)	31
	Miscellaneous Issues	32
	NWG Organization	32

I. SUNDAY EVENING SESSION (5/10/71)

INTRODUCTION OF ATTENDEES

Attendees introduced themselves and stated their affiliation. The following list includes persons attending any of the sessions.

SITE	NAME
AMES-ILLIAC	John McConnell
AMES-67	Wayne Hathaway
ARPA	Bruce Dolan Cordell Green Larry Roberts
BBN-NCC	Will Crowther Frank Hart Robert Kahn Alex McKenzie
Carnegie	William Broadley H. Van Zoeren
Case	Patrick Foulk
CCA	Richard Winter
Dept. Comm., Canada	Terry Shepard
EDUCOM	Henry Chauncey John LeGates
Harvard	R. Metcalfe R. Sundberg
IBM Research	Douglas McKay
Illinois-CAC	Jack Bouknight G. R. Grossman Jim Madden
Lincoln Labs	Richard Kalin Joel Winett

Merit	Al Cocanower Brian S. Read
Merit-Univ. Mich.	W. Scott Gerstenberger
MIT-DM	Abhay Bhushan Robert Fleischer Albert Vezza
MIT-MULTICS	J. C. R. Licklider Mike Padlipsky
Mitre	P. Karp David Wood Gene Raichelson
NBS	G. Lindamood T. N. Pyke
RADC	Tom Lawrence Bob Walker
Rand	E. F. Harslem J. F. Heafner
Raytheon	T. O'Sullivan
SDC	Robert Long Arie Shoshani
SRI-ARC-NIC	Charles Irby John Melvin R. W. Watson
Stony Brook	Ralph Akkoyunlu Art Bernstein M. Inam Ul Haq Richard Schantz
Univ. of Chicago	R. Ashenhurst
UCLA-CCN	Robert Braden Steve Wolfe

UCLA-NMC	Vint Cerf Steve Crocker Ari Ollikainen John Postel Rollin Weeks
UCSB	Steve Lynch Jim White
U. K.	Eric Foxley
Univ. of London	Peter Kirstein
Univ. of Mo.	Dan Garigan
Univ. of Penna.	Don Bernard
Univ. of Waterloo	Don Cowan

SITE STATUS REPORTS

The following are summaries of reports given by affiliates of the indicated sites.

UCLA-Sigma 7

- o NCP (document 1) and user/server TELNET-like protocol have been operational for three months.
- o NCP (RFC #107) will be implemented by June 1.
- o TELNET will be implemented within one month of official specification.
- o Have been and will continue to gather Network measurements.
- o Will use UCSB file protocol in conjunction with above.

UCLA-CCN

- o Local hardware has been connected to IMP.
- o NCP (RFC #107) 70% complete.
- o NCP (RFC #107) and NETRJS will be in production on JULY 1.
- o Draft of NETCRT will be issued as RFC soon.

- o They can offer APL and are looking for interested users.

UCSB

- o UCSB is a service center; 360/75 + 2314 storage.
- o NCP (doc. 1) operational since mid-October 1970.
- o NCP (RFC #107) operational and verified remotely.
- o Services currently offered:
 - 1) UCSB on-line system
 - 2) RJE/RJ0
 - 3) file system
 - 4) local interfaces to net via FORTRAN, PL/1, etc.
- o Short range plans:
 - 1) DRS implementation and experimental service
 - 2) Interested in APL
 - 3) Distributed data base experiment with SRI

SRI-ARC/NIC

- o Now running NCP (doc. 1) in TENEX.
- o NCP (RFC #107) will be installed when BBN releases it.
- o They are currently examining and tuning TENEX parameters and resolving interface difficulties.
- o Stage 0 plans include initial work with West Coast sites to gain experience, in using NIC, with respect to response times, traffic loads. This will initially look like TENEX to a user. Later, this will appear as a virtual machine to remote users via a special NIC subsystem.

SRI-AI (reported by SRI-ARC personnel)

- o Configuration includes PDP-10, TENEX, high-speed drum, robot.
- o Probable Network connection in July.
- o Software includes various theorem-proving mechanisms.
- o Interested in working on the above over the Network.

RAND

- o Current configuration includes 360/65, 1800, IMP, video graphics system.
- o NCP (RFC #107) in 360/65 in daily production use in conjunction with UCSB.
- o Network Services Program (NSP) used in above features:
 - 1) dynamic access to local files
 - 2) access to video consoles
 - 3) access to NCP
 - 4) UCSB RJE/RJ0 protocol
 - 5) UCLA NETRJS protocol (not operational)
 - 6) Logger and TELNET-like protocols.
- o PDP-10 to be connected into Net (directly to IMP) in about two months -- will eventually run TENEX and be a service center.
- o Short range software development includes TELNET and DRS.
- o Will continue to provide production support for Climate Dynamics.
- o Hardware is 360/65 connected to Honeywell 516 connected to IMP.
- o NCP (doc. 1) verified with Rand.

SDC

- o 360/65 runs Adept timesharing, has 10 users, can do protocols from a user process.
- o NCP (RFC #107) will be coded by end of month.
- o Logger fairly close to being checked out.
- o Will have TELNET similar to RFC #137 by June 1; will offer it by 20 August.

- o SDC not meant to be a general service; will allow 1 to 4 users; can use Adept to run a job.
- o Experimentation plans include:
 - 1) Voice I/O will use Net to communicate with speech researchers (will respond to specific programs only)
 - 2) Eventually graphic I/O
 - 3) Man/machine synergism
 - 4) Network data management
 - 5) Network resources Notebook use and update.

ILLINOIS-CAC

- o IMP works.
- o Will use B5500 on campus since B6500 doesn't work.
- o Will link to PDP-11 at Paoli for I4 software development -- not a general link for everyone.

AMES

- o Probably will be two separate nodes ILLIAC and 360/67.
- o Plans are in rumor stage.
- o Plan TIP by end of summer to gain access to I4 simulator.
- o Duplex 360/67 will be regular host node.
- o Are looking for an NCP implementer.
- o Will use other services; laser store and UCSB.
- o Their general research includes an interest in Network accounting and management.
- o Will go onto Net as soon as possible to ILLINOIS.
- o Will go on via TIP if it can support two nodes.
- o Will define Net protocol for interactive graphics for I4.
- o Plan to get on Net before I4 comes up.

CCA

- o CCA is responsible for data computer project.
- o It is special purpose computer with large storage device.
- o Data computer (PDP-10 and laser memory) should be viewed as one black box.
- o Will have two ports -- high-speed to I4 and low-speed to IMP.
- o Data language and data will pass over ports. This will include requests of files and portions of files, update, manipulation, and transformation of data. It will not include number crunching. Only access to the laser store will be through the black box data computer.
- o Hardware and communications are specified.
- o Data computer services are beginning to be defined, (600 questionnaires were sent out, 10 responses were received). With regard to services please call Dick Winter on (617) 491-3670. Dick wishes to hear from any potential user.
- o Laser store is one trillion bits of on-line storage packaged in 40 packs containing 10 mylar strips each. Strips are kept in a carousel that can be rotated and mounted in 10 seconds. Access to any track is a maximum of 400 ms.
- o Laser will arrive at CCA in early 1972 and move in late '72. Services to laser while located at CCA will be offered.
- o CCA will send out plans (as feasible) as RFCs.

CASE WESTERN

- o Hardware has been built (PDP-10) and starting to test it.

CARNEGIE

- o Hardware includes PDP-10 240K 36-bit words, TTY, etc.
- o Currently using DEC monitor.
- o Will modify Harvard's version of the monitor.

HARVARD

- o Hardware includes PDP-10 with PDP-1 as graphics devices handler.
- o Hardware is up, talking back and forth.
- o NCP (RFC #107) for PDP-10 will be completed by the end of the month.
- o PDP-10 has 48K 36-bit words; expect more core after July; will then make NCP resident.
- o Interested in file transfer, graphics, extensible languages, experiments of distributed processes.

IBM RESEARCH

- o IBM will buy an IMP, now negotiating it.
- o They are designing new concept of Networking.
- o They will become an active node with either a 67 or a 91. The candidate 91 now has 300-400K.

RADC

- o TIP is scheduled for delivery in October.
- o They will provide local access to the SRI on-line system.
- o They are interested in the July graphics meeting sponsored by MIT. (Al Vezza)

MIT DM AND MULTICS

- o GE-645 runs MULTICS: PDP-10 is for dynamic modeling of graphics systems.
- o NCP (RFC #107) on PDP-10 by end of week.
- o Logger on PDP-10 available by 15 June.
- o NCP (RFC #107) available on MULTICS by 11 June.
- o Logger and TELNET available on MULTICS by first week of July.

- o They have been conducting file transfer experiments of simplified ASCII transfers (not RFC #114).
- o Interests include:
 - 1) File transfers
 - 2) E & S processing to IMLAC and ARDS formats
 - 3) DRS service.

LINCOLN

- o Three connections are planned.
- o 360/67 has NCP, logger, TELNET, and some file transfer protocol.
- o TX-2 NCP is being developed. An interim NCP is working with file transfer experiments, error detection and correction.
- o They are experiencing hardware difficulties with the third connection.
- o 360/67 has been running NCP (doc. 1) and a TELNET-like protocol since 1 March. UCLA and Rand have logged in. The TELNET-like protocol supports ASCII and EBCDIC.
- o NCP (RFC #107), TELNET, and logger will be operational on 67 by 1 July.

BBN-NCC

- o NCC is responsible for maintaining the subnet and interfacing with Honeywell and AT&T. They are planning a mechanism for repair of the Net from their homes. The subnet has experienced a 98% up time.
- o An operational 316 version of the IMP has been hooked into the Net since February. It is a production item at 1/2 cost of 516; the IMP programs are identical.
- o The 316 will be incorporated into the TIP. The multi-line controller has been fabricated and is being debugged. A common language for the TIP has been developed. TIP will be delivered to NASA Ames on 1 August. The TIP will handle 64 lines up to 19.2K bits. BBN is still shaking down the kinds of terminals that will be compatible (will be bit serial at least). They are working with sites to determine terminal requirements.

- o The resource notebook has been compiled and distributed. Twelve of 19 sites are included in the notebook. SDC has since reported. Stanford, SRI/AI, MIT/DM, UTAH, Carnegie, LL/67 have not provided an entry. BBN again made a call for responses and asked that each site stay up-to-date. A comment was made regarding the Notebook that a mechanism is needed for a) date of information, b) complaints of information, c) reporting that advertised procedures do not work.
- o TIP delivery for the rest of the year is to the following five sites: Ames, Mitre, Rome, ETAC in Washington, D.C., and NBS.
- o BBN has been studying performance of Network to learn ways to improve it. An earlier Rand RFC reported a very low rate for a total Network experiment -- Rand re-ran the experiment to examine just the subnet performance and reported the subnet rate to be in the 13-15K bit range for that test. MIT/DM has reported a 5KB rate that will be examined further. BBN made a general offer to the Net community to phone NCC in the event that such measurements are taken and low data rates are discovered.
- o The BBN-Honeywell relationship is a straightforward maintenance contract that includes P.M. and other problems which arise. BBN stated that a 24-hour delay in service would be unusually long. BBN would like to be kept informed of views and opinions of the Honeywell service. BBN asks that each site keep a maintenance log and copy it to Marty Thrope at BBN.

BBN/TENEX (reported by BBN/NCC)

- o BBN has a three-host IMP that includes NCC and two PDP-10s. One 10 is for operational use, the other for experimentation.
- o Software for the PDP-10 is TENEX.
- o They are interested in getting involved in Network experimentation.

MITRE

- o Expect TIP delivery in September.
- o Now using PDP-10 at BBN.

- o Will work on sample data management system using UCSB file storage.

NBS

- o Expect TIP delivery in December.
- o Have selected PDP-11 as host; ready by December.
- o Will build from the U. of Illinois operating system.
- o Contemplating attaching their UNIVAC 1108 to Network.
- o Will provide experimental access to Network for services:
 - 1) measurement (performance)
 - 2) graphics
 - 3) personal communication
 - 4) lab automation
 - 5) support NBS with services appropriate
 - 6) provide local hardcopy from PDP-11.

ETAC (Environmental Technical Application Center)

- o ETAC, located in the Washington area, is a branch of Global Weather Service.
- o Air Weather Service in Omaha, Nebraska, has seven 1108s providing weather data; one will be connected to Net to provide daily weather information to ETAC in Washington.
- o ETAC has 1401 and 7040 in Washington to produce summaries for longer range use; the 7040 may go on the Network.

AIR FORCE SITES (reported by ARPA)

- o To export technology to other regions, two Air Force sites (Oklahoma City, Oklahoma, and Sacramento, California) will temporarily join the ARPA Net to talk only to each other in an operation to parallel their AUTODIN connection. If it materializes, they will run for the first six months of '72 but probably won't continue. The motivation is to examine the Net ideas for developing a wholly autonomous network in three or four years.

OTHER REPORTS

DEPT. COMM., CANADA

The Canadian Government wishes to optimize the use of all computers in Canada.

They now have a banking network.

They are interested in a small net for universities.

Their largest problem is the size of the country in relation to the sparseness of population.

U. OF CHICAGO

The University of Chicago has no current time schedule but they have definite ideas about what they wish to accomplish and they are seeking funding. They are applying to NSF to support a local net on campus for lab automation. They have good people and good equipment; the idea is to make it coherent.

Their interest in the ARPA Network is to make shared software available to their people and to a limited extent, make local services available to other ARPA nodes. Their proposed host is a PDP-11 to the mini computers and a second host (PDP-10) as a big software engine to make data available to the mini computers. The PDP-10 and PDP-11 will perhaps be linked together. They also expect to get a TIP to provide remote number crunching for their people.

UNITED KINGDOM

They have proposed three main machines and three terminal nodes. They have in mind the 906A, approximately the 360/75 in power.

Their Post Office also has plans for a digital network in the distant future.

MERIT

MERIT-UNIV. MICH.

Most of the bugs are out of the hardware.

Most of the software is written.

Will have PDP-11s by the end of the summer that are capable of transmission from one to another.

They will need and are now studying a TELNET-like protocol.

They are concerned with orderly communication of the two processors; will later become concerned about process-process communication.

EDUCOM

EDUCOM involves 100 major universities including most of those now in the ARPA Network. For two years, they have been running a network without wires. They assume the ARPA Network can resolve the technical issues. They are looking into marketing, contracts, documenting, etc., for running the network. They have conducted a survey of 70 universities, polled about their interests in the ARPA network: 60 of 70 are interested, 14 have money and are ready to become sites.

RAYTHEON

They will access the Net through the four Boston nodes.

Their interests include:

- 1) experiments of file transfer conversions.
- 2) indexing behavioral data to allow one to search an index to see if the body of data of interest is within the Network.

MISCELLANEOUS TOPICS

Graphics

Al Vezza will host a July meeting of a small group interested in Network graphics. The price of admission is a sincere interest, working background, and a prepared talk.

NCP Protocols

A new official document will replace document 1 and RFC #107; implementation should not be held up because of the absence of this new document.

The long range protocol committee chaired by Carnegie has been disbanded.

IMLAC Users Group

A quick survey was taken to determine which sites had or planned to get an IMLAC. The plan is to form an IMLAC users group. The following sites have or plan to get one: UCLA-S7, AMES-67, BBN, SRI-ARC, Stanford, MIT-DM, MIT-MULTICS, Mitre, Case, Raytheon, U. of Illinois.

Official Document Formats

The notion of a functional document was suggested, one of which would be the document of official protocols with divisions of levels of protocols.

II. MONDAY MORNING SESSION (5/17/71)

NETWORK INFORMATION CENTER

Plans for NIC

Two activities are planned for this summer, off-line mail and on-line access. The off-line service will continue after the on-line service has come into being. Plans for getting on the Net via PDP-10 (replaced XDS-940) are almost complete. Response times for display use are marginal.

The activities will be developed in stages. Stage 0 (June 18) NIC will work with West Coast sites. This will involve providing NLS facilities to allow people to create messages with initial delivery as hardcopy, etc., with automatic generation of catalog entry and NIC #. This system has been used locally for about a month. Stage 1 (August 2) NIC will be open to the Net community as a whole. Remote users will come in directly to the on-line system and will have on-line access to the catalog. Users will be trained either at SRI or at their own sites before coming on. Four to eight concurrent terminals will be supported. Stage 2 will include file transfer protocol, on-line delivery of messages, remote editing of SRI-located text. Prior to stage 0, a course will be offered (on June 16, 17) for UCLA, Rand, SDC, UCSB, Ames, and RADC for the use of Stage 0. The second group of users (after stage 0) will use NIC to do their own site documentation.

Concepts & Recommendations for Documentation

The NIC # is a unique "name" for reference -- it has no other meaning. Other numbering schema such as RFC numbers will eventually go away. However, the subgroups, such as RFCs, will remain. Appropriate set manipulators will be provided for assisting in storage and retrieval.

The notion of functional documents was introduced (see RFC #115). This is to be a document whose purpose is reasonably stable over time. It can have subdocuments that change more frequently. A current list of functional documents includes the NIC Catalog, Directory of People, Resources Notebook, Protocols, and Site Facilities (one for each site).

The mechanism of documentation is the responsibility of NIC; the document contents are the responsibility of the author. There are two cases of document revision; replace part of the document and replace the entire document. In general, NIC would like the document

to be re-issued in its entirety with a new NIC # rather than issuing errata. The functional documents are in looseleaf form, new pages can be issued with the same number and a revision date.

Documents are reproduced and mailed to site liaisons 24-48 hours after receipt. They are mailed to station agents on a weekly basis. When mailing is handled directly by a site, a copy of the document and a distribution list should also be sent to NIC. In the past, NIC has supplied abstracts of documents for the catalog; NIC requests that the authors include an abstract.

TELNET

The purpose of TELNET is to provide an immediate mechanism for communication between keyboard terminals and serving processes, with sufficient platform for later expansion and sophistication.

Tom O'Sullivan described TELNET as delineated in RFC #137. (Later in these NWG meetings, Tom issued RFC #158, a new TELNET protocol.) After the description, many issues and questions were raised, viz., can TELNET expect "recovery" from NCP, 128 vs. 256 character set, DLE + 7-bit code vs. high-order bit on, should protocol extend service beyond what level consoles see, human factors, if information is available at second level should it be passed to TELNET, TWX-like service from NIC, mailbox protocol, etc.

In large part, these issues were raised but not resolved. It was agreed that an RFC would be forthcoming (RFC #158, published later at the meetings) followed by a functional document.

III. MONDAY AFTERNOON SESSION (5/17/71)

FILE TRANSFER PROTOCOL (RFC #114)

The file transfer protocol (RFC #114) was described. See also RFC #133 and RFC #141.

A simplified version of RFC #114 is being implemented by MIT/DM and MIT/MULTICS in order to: 1) allow Dynamic Modeling access to MULTICS file storage facility and 2) conduct a pilot project to gain understanding of such protocols.

It was noted that RFC #114 was not simple enough to implement for TIPS.

A group was formed to meet Wednesday morning for more discussion and to exactly define the problems. The group would include representatives from UCLA, UCSB, BBN, MIT, Rand, SRI, Harvard.

FILE PROTOCOL STATUS REPORT

UCSB described the status of RFC #122, A Simple Minded File System, as an operational program; not a proposal. The basic concepts of the file system were described; the design objective was to provide a simple service quickly.

Currently one 2314 drive and pack is available. At most four drives will be made available during the next year. It is also not clear how long space will remain available. The storage is currently free.

Sites that will use the file system are Mitre, via BBN, UCLA, SRI, and Raytheon via one of the Boston hosts.

MISCELLANEOUS TOPICS

Sockets

Socket name structure was briefly discussed. Relevant RFCs that were mentioned were 1) RFC #129 whose purpose was to describe socket structures enumerated at the February NWG meetings, and 2) RFC #147, a recently proposed structure.

It was pointed out that there was a definite need to reduce the socket length from 32 to 16 bits (a TIP storage problem) regardless of its structure.

A committee (Bob Metcalfe, Chairman with Abhay Bhushan and Joel Winett) was appointed to produce a report in two weeks. The committee is to address the following three issues:

- 1) is a socket structure needed
- 2) are more than 16 bits needed
- 3) what procedures are recommended for managing socket numbers.

Initial Connection Protocol

Race conditions and time out problems were elucidated. See RFC #123, 127, and 151.

A committee (chaired by Jon Postel and including Steve Wolfe, Eric Harslem, and Arie Shoshani) was appointed to clean up the ICP specification.

Testing and Validation

Sites wishing a remote partner to exchange NCP, TELNET, and logger protocols can contact Rand. Rand was to collect status information before and during these exercises. Information was to be forwarded to Alex McKenzie to maintain and update status reports. (NOTE: A later steering committee decision reflects on the way in which this information is gathered, however. Rand is still available for testing and validation.)

IV. MONDAY EVENING SESSION (5/17/71)

NOTE: Minutes of this session were kindly prepared by Bob Walker, RADC.

OPERATING SYSTEMS AND NETWORKS

An attempt was made to study the ARPA Networks from an academic point of view. An analogy was drawn on the basis that the ARPA Network with its hosts and protocols is in a sense an "operating system" and that a study of what makes a good operating system might help define what makes a good ARPA Network.

Professor Art J. Bernstein of Stony Brook gave a presentation abstracting what he considered to be the features of a flexible operating system, the techniques for obtaining such; and when appropriate, a discussion of those aspects where a difference in techniques is required between dealing with an internal operating system and dealing with a network.

The features of a flexible operating system were cited as: (1) a flexible file structure, (2) a process hierarchy, and (3) an interprocess communication facility (IPC). The terminology and techniques described to obtain these three features were essentially those developed for the MULTICS system.

A file structure capability was defined in terms of hierarchy of directories, tree names, active file table, hold count, known file table, and reference number.

A process hierarchy was discussed in terms of father-son relationship and a father-node spawning a son node, creating an entry in the known file table and assigning resources, all embodied in the SPAWN primitive. Implementation of primitives as time independent was stressed as being crucial to Network activity whereas not necessarily so for an internal operating system. This lead into the concept subcontracting process, where executive type functions are treated on the same basis as user processes and as such are swappable. The "link process" was then described as the interface mechanism between two cooperating machines.

Interprocess communication was discussed in terms of channels, status return and software interrupts. Appropriate primitives were defined in detail as well as control type problems.

The discussion then went to file handling and a specification of the required primitives and thence to directory handling, specification of related primitives, and the mechanics of directory handling, specifically the outstanding operation entry table in the executive.

After a short recess, Bob Metcalfe gave a presentation from the ARPA Network point of view with reference to various points of Professor Bernstein's presentation. He noted the all pervasive tree structure in Bernstein's presentation which appears to be most efficient to internal operating systems (i.e., file system, process hierarchy, etc.), but that the ARPA Network is not a tree structure but rather a directed graph and that we should be careful not to impose tree structure thinking on a directed graph type situation.

A number of questions and problem areas were elicited from the group and listed on the blackboard:

- 1) How much does the operating system need to know about the Network to get how much and vice versa?
- 2) Degree of transparency to the user?
- 3) "Optimal" resource allocation on the Network?
- 4) Autonomy versus centralization of control.
- 5) Resiliency.

The group discussed the need for a committee on Theory, how it should function, how often should they meet, requirements for attendance, etc. Dave Walden was mentioned as a possible organizer of a related effort. Bob Metcalfe agreed to chair such a committee.

V. TUESDAY MORNING SESSION (5/18/71)

DRS WORKING GROUP MEETING WITH OPEN ATTENDANCE

The purpose of the Data Reconfiguration Service meeting was to resolve several lingering syntax and semantics issues and also to receive comments and discuss the DRS with the entire Net community.

A brief overview of the DRS (see RFC #138) was given.

Remaining technical issues were resolved. An implementation specification (replacing RFC #138) will be issued soon.

Initial implementers and users were polled for schedules and initial experiments, results are shown below.

MIT	No dates currently provided
U. of Ill.	One or two months will be required to reformat from remote formats to GOULD printer; also conversion of ARDS to COMPUTEC strings.
UCSB	Implementation of service in two months; will provide system forms for remote TTY-like devices to access UCSB on-line system.
MITRE	Will compare performance of DRS to current software of UCSB file experiment.
Rand	implement service by September; initial use to specify UCSB RJE/RJ0 and UCLA NETRJS formats for local users.
UCLA	will have a compiler of forms within one month unless serious difficulties arise.

DATA MANAGEMENT ON COMPUTER NETWORKS

SDC presented RFC #144 (see also RFC #146). Arie Shoshani presented considerations and approaches that can be taken to achieve data sharing. The considerations were common language, sharing of existing data, evolutionary/revolutionary, future and use facility, further development, implementation, and speed.

Approaches given were:

- 1) centralized
 - a) new data only
 - b) existing data
- 2) standardized data
- 3) integrated - common languages + interfaces
 - a) interface on different nodes
 - b) interface on service node
 - c) Data Reconfiguration Service
- 4) Unified

Dick Winter described the CCA approach. With several data computers it becomes decentralized. All data computers have identical hardware and software. Their objective is to dispose and restructure data throughout the Net to optimize its use, i.e., relocate it close to where it is used most heavily. For small files of wide interest multiple copies can be maintained.

Dr. Roberts commented that with respect to the Network, no distance/cost relationship exists if data is retrieved more than one link away. The reason for putting files in several places is reliability. He views the CCA approach as a Net-level language, thus the unified approach. Also the natural language approach is suitable as a research project but not suitable for data management for real Net experiments.

CCA will present a proposal of data language at the next NWG meeting.

OPEN DISCUSSION ON DATA MANAGEMENT

This time period was initially allocated to the description of a particular data management system being constructed by Mitre. It became, in fact, an open discussion of general principles and requirements for data management in the Network. The following were among the most recurrent comments made.

1. DRS, file protocol, and data management should be examined in a comprehensive way.
2. Important considerations of data management are to allow users to define and restructure files logically, to move towards transparency of the Net, and to move toward natural language.
3. A data management system should include functions for define, access, manipulate, analyze, store of files. For example, the data computer doesn't do formatting for output (like an RPG), it can take a number of conditions and do conditional retrieval but not RPG.
4. A data management system could be developed in stages where a) the user explicitly moves data around the Net, b) the user specifies the location but the access is integrated automatically, c) location is maintained by the service.
5. An area should be defined between file handling and application specific manipulation, and the area should be treated in a system wide way.

6. The super file (too large for any one individual to economically own) never came up before but it is reasonable for the Net. However it is just one use and there still will be need of many 10^9 files.
7. Privacy and security criteria should be applied at output rather than input, which is an argument for having processing capability at the location of the file.
8. Dr. Roberts indicated that the things that are important are what things are on the Net, and what things are there to say. The structure depends on what there is to say. Thus, one should concentrate on the language and not the structure.
9. The data management system can be viewed as having two parts: 1) the request, 2) the response and format. On the response side (operand side) there is the taxonomy of data types and a template of data followed by the data. A template is a string in which data types or their descriptions are given with knowledge of iteration, recursion, and data types. On the request (operator) side, templates can be used to precisely specify the data to be retrieved, assuming the structure of the file is already specified.
10. The disposition and request are over structures to the response.

A small group was established to continue discussion on data management.

VI. TUESDAY EVENING SESSION (5/18/71)

TERMINAL IMP

The TIP can either be configured with 1) one host and two phone lines or 2) three phone lines. Interfaces will provide 19.2KB to lowest TTY speeds for each line. It can handle various terminals and devices.

Normally the user speaks through the TIP but a primitive language exists for talking to the TIP. Commands will exist to do the particular protocols such as logger. Other commands will be present for terminate on line feed, on character, now, on nth char., at end of message, i.e., class of things to determine when message is sent. There is another class to determine echoing. Device rates can be set up. The serving site can also set up command such as capturing a printer.

The TIP is currently trying to comply with all second and third level protocols such as TELNET, file transfer (when defined).

Current plans are that the TIP cannot be reloaded through the Network.

When new terminals are added, BBN will supply the TIP routines as part of the service.

The TIP is intended to be used for RJE, terminal to process, and later tape to tape. The TIP is intended to be a switch rather than an operating system, under the assumption that power will reside in terminals and service centers.

The program limits the bandwidth -- the sum of input and output is 100KB.

Potential for TIP delivery is about one every three weeks after August. An upper figure for the TIP is \$100K; the leasable terms are \$40K/yr. for three years plus a residual of \$5K to own it, with a two-year minimum. This was designed as an alternate method of purchase.

"COMMENTS BY DR. ROBERTS"

The major cost benefit in the near term to getting on the Network will be to use other physical systems to access new resources. It will be a number of years before people enter the Network in order to get rid of machines or to boost CPU usage.

Regarding future Network growth, the University of California has proposed to enter seven universities into the Network. We should have the data and program sharing protocols fixed by that time. ETAC will be working on the past 10 years weather in 10¹¹ store. NCAR will be trading time (a 6600 and a 7600) with them and with ILLIAC; use is restricted to weather work. January or February are probable dates. This will be a third cross country connection through UTAH perhaps (second is via Omaha weather). SC will be added in March or April '72 for picture processing. England will join about February '72. There are other plans to tie in Mexico, France, Israel, Australia, Japan, Hawaii, Canada, etc. that could possibly all happen in '72.

With regard to operating the Network, ARPA will not operate it indefinitely. One plan is to have AT&T operate it since they can legally sell the services; this will not come about soon. A commercial organization (not a common carrier) can only operate the Net under Government sponsorship. The current plan is to have BBN run the Net as a service for the Government; this will be settled within the coming year.

On the question of resources, setting up contracts with the service people at each site to get one agent to ship money for various subcontracts is a basic legal framework; for ARPA purposes it is sufficient to have only one connection with each site.

On software development, the NCP progress has been extremely poor and slow. The second iteration should have been defined by now from experiences with the first. Towards the end of the year a new protocol should be defined to last for a couple of years. Accounting and billing protocol should also be defined. The NCP protocol is getting to be a critical problem -- everyone should be complete and consistent with the current protocol by July 1. Without it, there will be serious problems of bringing new people onto the Net. For example, the I4 and the laser store will be on the Net by March or April of '72 with serious people wanting to use it (80% of its use will be remote). By early '72 the Net must be a solid working entity.

The question of profit making time-sharing companies on the Net depends on whether or not AT&T takes over Net operations.

The capital arrangement for non-ARPA users to be on the Net is as follows. A federal agency can donate \$76K and get a TIP. Non-federal agencies can pay \$36K per year for the TIP for three years plus the \$5K residual to own it. ARPA will not decide casually to allow non-federal agencies on.

Regarding software support services, documentation will be upgraded so all sites need not keep complete NIC documentation (except service sites). In service centers it makes sense to add one or two personnel to work on net service programs, work with users, etc., if needed. Research centers will now have to concern themselves with reliability, integrity, and problems of access.

Regarding the charging mechanism for the data computer, the 10^{12} store cost one million, plus the cost of the PDP-10; thus 10^{-4} cents/bit is reasonable for permanent storage. The rate for short term storage strips (like two weeks) will be about the same. If medium term storage is needed, a rate will be worked out. ARPA will pay for this storage as backup for the sites.

The on-lineness of NIC is very important for initial use, but we must have something better than TTY or CRT. The Net is cheaper than the mails. (Electrostatic hard-copy devices were briefly mentioned).

Regarding new developments for AI symbolic processing, a plan hatched by Alan Kay is to have lots of processor, lots of core and a big switch with the capability of serving users in the Net. It is to provide low cost core space (economics of processing are not known). This may become associated with some experimental hardware development facility since the desire is to be able to build new architecture in a reasonable amount of time. It should be 10 to 100 times faster than the PDP-10 with earliest delivery in '73.

The speech effort is on the order of three million per year. The concern now is to be able to tie together pieces at various sites for comparative evaluation. The cross-testing can have an impact on the researcher, but everybody must maintain compatible interfaces.

The climatology program is for predicting future long-range climate of the World that comes about by perturbations. Various sites are involved at various levels and it is hard to get these people to big computers, to the data bases, and with each other. The Network provides their total communication path with the I4. Direct and effective use of the Network can be made without much more of an investment; the Rand/UCSB work is a good example.

VII. WEDNESDAY MORNING SESSION (5/19/71)

This session began with discussion of file transfer protocol, led by Abhay Bhusan. It was decided that the current file transfer protocol should be parsed into two pieces -- a data transfer protocol front-end that could be used for file transfer and other protocols, and the file mechanism protocol. This problem was referred to the committee which met for the remainder of the day to specify the data transfer and file protocols. An RFC will be forth-coming, describing these protocols.

The data management group met in parallel Wednesday. An RFC will be forthcoming on their results.

VIII. WEDNESDAY EVENING SESSION (5/19/71)

The following information was summarized by Steve Crocker.

Committees	Publication Date	Approval Date
ICP - Postel	5/27	6/3
File Transfer - Bhusan	6/7	---
Data Mgmt. - McKay	(7/21)	---
Socket Struc. - Metcalfe	6/22	---
Telnet - O'Sullivan	5/19	6/10
Theory - Metcalfe	---	---
DRS - Heafner	6/1	---
Graphics - Vezza	(7/18)	---

The following inputs were provided to Steve Crocker on implementation dates of NCP (RFC #107) and TELNET (RFC #158).

Service Hosts	NCP + TELNET
CCN	7/1
LL/67	6/15
SRI/NIC	(6/18)
MIT/MULTICS	7/1
BBN/10X	?
UCSB/75	Up

__Host__	NCP (RFC #107)	TELNET (RFC #158)
UCLA/S7	6/1	6/15
Rand	Up	6/15
Utah	Up	6/15
U. of Ill.	7/1	7/1
Harvard	?	?
MIT/DM	5/25	6/25

The following inputs were provided to Steve Crocker on schedules for current and pending work.

Users	Tasks
Mitre	data management in progress
Raytheon	data sharing (August)
NBS	PDP-11 via low-speed phone line (July)

BBN	validation of resource notebook (July 15)
UCLA	data store, retrieval, reduction (July 1)
DM/MULTICS/Harvard	graphics, file transfer (July 1)
Ames/67	I4 simulator (July 15) climate with UCSB (now) climate with UCLA (July 1) DRS (September) SRI/NIC (August) LL LISP (?)
LL	TX2 speech data TX2 data transfer (now) TSP compiler (September)
U. of Ill.	remote use (July 1) link to Paoli (July 1)

Miscellaneous Issues

Alex McKenzie will generate the NCP functional document in one month as an experiment.

Service documents to be sent to NIC include normal user documentation you would use at the site plus special conventions (if any) for remote users. Read RFC #115 and RFC #118.

NWG Organization

There is some concern over the size of the NWG. Its functions and reorganization were discussed. Nothing definitive resulted immediately. It was suggested by Steve Crocker that another NWG meeting would be held in August.

Dr. Roberts and Steve Crocker created a steering committee to examine this and other problems. More will be said about the steering committee by Steve Crocker, at a later date.

[This RFC was put into machine readable form for entry]
[into the online RFC archives by Nicholas Barnes 08/99]

