

Network Working Group
Request for Comments: 101
NIC: 5762

Richard W. Watson
SRI-ARC
February 23, 1971

NOTES ON THE NETWORK WORKING GROUP MEETING

Wednesday Evening, February 17

Mike Sher opened by welcoming the group to Urbana and briefly indicated that ILLIAC IV was expected to be running this summer. The ILLIAC IV Project has been split into two projects; one on basic system hardware and software, and the other on applications. Their IMP is not yet connected to their PDP-11.

Steve Crocker asked for topics to be discussed at this meeting; these are indicated below.

Peggy Karp of Mitre has been summarizing the old RFC's. She has a list of about 30 topics and is summarizing their present status. She expects to finish around the end of February. See RFC #100, NIC(5761). It was suggested that someone write an RFC indicating which ones are obsolete. It was also suggested that the Network Information Center (NIC) help sites in organizing their hardcopy material.

There then followed brief discussions of experiences in using the Network. John Melvin (SRI-ARC) summarized SRI's experience in using the Utah PDP-10 to help in SRI's transfer from an XDS 940 to a PDP-10. In April-May 1970 it was clear that SRI was headed toward a PDP-10 in order to have the capacity and reliability to fulfill their role as the Network Information Center. They had had some previous experience in connecting with Utah, and so it seemed logical to try to use the Utah 10 to aid the transfer.

In June use of the Network began. SRI uses higher level languages extensively, so the first task was to transfer the compiler-compiler Tree Meta. Source code was generated on the 940 to run on the PDP-10. Binaries were then transmitted to Utah and run and debugged. Patches were performed where possible, and source changes accumulated. A new source and binaries would then be made periodically.

Once Tree Meta was running, a new high level language (called L-10) for programming the On-Line System (NLS) was implemented in the same way. When L-10 was running the core device independent parts of NLS were rewritten and debugged. NLS was completely reorganized during the transfer.

At the SRI and Utah ends a control program allowing three users to connect to Utah was written, which ran as a user process and allowed character interaction and files to be transmitted. The scheme worked well and much useful work was accomplished in the July--December period with some people on 4-5 hours per day. The voice link was used when something would go wrong in trying to determine where the problem existed and to reset. At times they would go 2 weeks with no problems. SRI has an IMP interface diagnostic which ran as a T/S process.

Generally, echoing was handled at the SRI end. DDT was used at Utah end. Round trip character delays of 4 seconds were not uncommon, and at certain points delays of 8 or 10 minutes were experienced. These delays were the result of the implementation used which involved multiple processes at each end, each to be scheduled. Utah was heavily loaded at 2:00 PM and the SRI people took to running at night and on weekends.

When the SRI PDP-10 came in in December, use of the Network slowed.

Users would have liked a more constant response time instead of the widely varying one so that their work habits could adapt to it even if it was slow.

Gerry Cole reported on some results of measurements made during the SRI-Utah work. Measurements were also made at SRI to help in interpreting the data obtained by UCLA. Gerry wrote a paper summarizing these statistics which is available from him care of SDC.

Gerry requested that when people are set up to use the Network, they inform him so that he can gather statistics. UCLA will eventually have a program to scan the Network for utilization, but if people could tell him when they were going to use the Network, it would be easier to measure meaningful things and interpret the data from a knowledge of type of usage.

Bob Kahn indicated that BBN is interested in the Statistics on overall flow to see if the Network is configured properly. Gerry said that UCLA is interested in the statistics for Network modeling studies. Measurements are taken by remote control by use of a feature designed into the IMP's by BBN for such a function.

Jim White of UCSB said that UCSB and RAND had begun to experiment in use of the Network for the climate study at RAND. The UCSB NCP has been up the last 3 or 4 weeks during the day. A document, NIC (5480) is available in the NIC collection describing it. UCSB is also using their NCP for local interprocess communication experiments. RAND is using the Remote Job Entry facility of the UCSB 360-75. They are using UCSB to check out their NCP. Now that UCSB is running their NCP during normal usage hours, they have uncovered some bugs in their hardware interface to their IMP. The software at both UCSB and RAND seems to be working. Typical jobs being sent back and forth are just test jobs of a few source statements. The UCSB NCP is about 39K bytes and runs in a 60K byte partition. Users access it through assembly language, Fortran or PL/I calls.

Steve Crocker now returned to the discussion of the agenda for the meeting and longer range organization of the NWG. Steve felt that Working committees on various topics were required as the open meeting was good for bringing up problems, general discussion and education, but was too large to prepare detailed specifications on various topics.

The following topics requiring work were listed:

1. Graphics
2. Data Transformation Languages
3. Host-Host Protocol -- long range study
4. Host-Host Protocol -- Short term maintenance and modifications
5. Accounting
6. Logger Protocol
7. Typewriter connection protocol
8. Documentation
9. Data Management

In #1 Al Vezza of MIT is organizing an NWG meeting in graphics April 25-27 which can accommodate 31 people. People desiring to come should prepare for their institution a working paper. Al sees three classes of problems:

- i) two hosts, each with computing and graphics facilities, wanting to use special facilities at the other

ii) one host with graphic facilities but no number crunching facilities wanting to use computing capabilities of a second host

iii) a node with a graphic terminal not having picture processing or computing capability desiring to obtain these from other nodes.

With respect to #2 John Heafner of RAND indicated RAND wants to provide data rearrangement services of the type indicated in RFC #83, NIC (5621). More on this topic below.

With respect to #3 a group under A. N. Habermann of CMU has been formed to look at the Host-Host protocol. Toward the end of March they are planning a paper discussing their ideas. The group consists of:

A. N. Habermann, CMU

G. B. Hansen, CMU

W. Wulf, CMU

R. Chen, CMU

R. Kalin, Lincoln Lab

The group welcomes suggestions of topics.

With respect to #4 a group is to be set up to evaluate present protocol and produce needed changes to the protocol. The group is to be conservative and produce only changes needed to solve known problems and leave esthetic changes until later.

With respect to the other problems discussion was put off until later (see below).

Two people interested in the Network who were observers at the meeting spoke briefly.

C. D. (Terry) Shepard of the Computer Communication Task Force, Canadian Government, outlined the goals of his group. These goals are:

- 1) establish a plan to link up various Canadian computers and establish a network
- 2) develop what the needs of Canada are for such a network

- 3) see that the benefits of such a network are distributed throughout Canada
- 4) prevent control of computing in Canada from being totally dependent on foreign sources.
- 5) see that critical computer facilities exist in Canada.

Doug McKay of IBM then described briefly a network project started in IBM about 2 years ago. Basic network is completed. Users are coming on. The network is to be used heavily to send files back and forth for program updating. IBM is trying to look at the network as a multiprocessor machine. They are trying to handle all IBM system possibly heterogeneous such as 360's, 370's, CP '67, the 91, a 44, and a NYU CDC 6600.

There is another project linking TSS systems using a 91 for remote job entry. IBM has taken a centralized control point of view using one central machine for control and flow distribution. They are not entirely happy with this approach and are moving toward a more decentralized approach like the ARPA Network. IBM presently has about 14 people involved in the project.

Thursday morning, February 18

Thursday morning started with the various sites reporting their status. Alex McKenzie of BBN prepared a status form later in the day which was filled out by the representatives of the sites Thursday evening. BBN and NIC will prepare a procedure for keeping this information at the sites and up to date.

STATUS

BBN, TENEX PROJECT: Final stage of incorporating NCP in TENEX. A connection was attempted to Utah, but some bugs were found. The NCP treats the network as a file in a way integral with other types of files. The NCP includes a teletype interface. They hope to incorporate the NCP in SRI'S TENEX system by the end of the month.

BBN, NETWORK GROUP: reported that they were working on three areas

1. Improving the current network
2. Working on a 316 version of the IMP and as a Terminal Interface Processor (TIMP)
3. Accounting

There are currently 15 IMPs connected to the network. A new software system with minor changes is expected by March.

The TIMP uses the 316. A hardware design exists, but they are still defining the software. A TIMP can handle up to 64 variable speed terminals both sync and async. The first machine is to go to MITRE in September.

BBN emphasized that there are 3 products: a 516 IMP, a 316 IMP, and a 316 TIMP. The 316 IMP is less expensive than the 516 IMP and can connect to one host. BBN is not planning at the moment to exchange 316 IMPs for 516 IMPs. The two are plug-plug compatible.

SDC: In the debug phase for their NCP and expect it up in 4 to 6 weeks. Maybe by 8 weeks their T/S available for network use. Their T/S is a 360/65 running the ADEPT system.

CASE WESTERN RESERVE UNIVERSITY: The IMP has been connected for about one month, but as yet have no NCP. They are planning to use the NCP implemented at Harvard. Case has a PDP-10/50 system. They expect to be up in two to three months.

HARVARD: Harvard has a PDP-1 and PDP-10 connected to the IMP. The NCP for the 10 is in final debugging. The PDP-1 is for refreshing displays. The PDP-10 is for linguistic research and students. Expect to be up in one to two months.

SRI-ARC: SRI has been in the final stage of conversion from an XDS 940 to a PDP-10. They plan to use the BBN TENEX NCP and should be up in three or four weeks.

MIT DYNAMIC MODELING - PDP-10: They expect an NCP to be working by March.

MIT MULTICS: They are connected to the IMP and expect their NCP to be in the final debug stage in four weeks. As Multics is a service machine, they don't have unlimited access and must perform checkout at off hours. They expect to offer regular service to the network in three or four months.

UTAH: PDP-10/50 probably going to be running TENEX in the fall. Their NCP is being written in a higher level language (Euler run interpretively) and are debugging in conjunction with BBN. They have connected to and logged into themselves and expect a debugged version within a month.

LINCOLN LABORATORY, TX-2: They are testing the IMP interface, found errors in Lincoln hardware. Currently, no data errors, but have errors with message IDs. They expect their NCP with logger to be up by April 15. They indicated that for testing purposes, they would like to bring up their IMP without being open to network traffic. BBN says that there is a way to echo to yourself without being open to the network (contact BBN for details).

LINCOLN LABORATORY, 360/97: Running CP/CMS. The IMP interface was completed last month. The NCP and Logger are working. They are planning to put up the NCP as a regular service in April. On request experiments with them can be run sooner.

UCSB: Has had their NCP since October. The NCP runs as independent batch job. They plan to provide service to their On-Line System (a manual is in the NIC collection at each site NIC (5480)). They plan to be on the air morning to evening on a regular basis. There are some interface bugs as indicated earlier.

RAND: 360/65. Their NCP is a user process and can be resident. It requires 8K bytes and does not have a logger.

UCLA, Sigma-7: Their NCP is in final debugging. They expect to be up by March 1 with NCP, logger, and typewriter connection program.

COMPUTER CORPORATION OF AMERICA (CCA): Has just started a project to create a node for the 10-12 bit laser store. They are going to use a PDP-10 as a front end. They are developing a language for data manipulation. The store will also be connected to the B-6500-ILLIAC IV. They are planning data compression as part of their language to ease the problems in use of the network's 50 kilobit line. They are concentrating on security and privacy measures. Initial emphasis will be on shared files. Installation is planned during 1972.

The following projects did not have representatives at the meeting. Steve Crocker reported on their status.

CMU: PDP-10/50: Their IMP is connected, and they are planning to use the Harvard NCP.

SRI-AI PROJECT: PDP10. They are planning to run TENEX in the fall.

STANFORD-AI: They are not connected yet, but expect to be on by summer.

The above completed the review of status. Steve Crocker then indicated that the old NWG mailing list was no longer to be used and that the list maintained by the NIC (5731,) was the one to be used or that the NIC would handle distribution by sending things through your station agent to them. If your station agent or liaison person should change, please let the NIC know immediately.

HOST-HOST LOGGER PROTOCOL DISCUSSION: Tom Skinner of Multics opened the discussion of the logger by indicating that they wanted at least an interim protocol so that use of the network could get started. They had handed out RFC #98 NIC(5744), containing their ideas Wednesday night. SRI-ARC had a similar document, RFC #97, NIC (5740), handed out Wednesday night also. Multics recommended the revised logger protocol of RFC #80, NIC (5608).

Some discussion on the relative merits of the logger protocol of RFC #66, NIC (5409), versus RFC #80 was given. The protocol of #80 had some potential problems due to assumptions which must be made after the initial contact was established.

The result of the discussion was that the logger protocol of RFC #66 was adopted with the correction that the allocate commands were to be issued after the connection was established.

There seemed to be a need for an official document to be issued with the correct logger specification given.

Tom also recommended that initial communication to the logger be in 7-bit ASCII in a 8-bit field. There was some discussion as to whether the eighth bit should be a 0 or a 1. It was finally decided that it should be a 0.

Steve then listed some known problems or questions about the host-host protocol.

- 1) Echo
- 2) Message Type
- 3) Interrupts
- 4) Marking-Padding
- 5) Half Duplex vs. Full Duplex communication during the establishing of a socket.

With regard to marking the following choices existed

- a) leave alone
- b) separate the heading and data into two messages
- c) have message by multiples of 72 bits

With regard to interrupts (INS, INR), there was a synchronization problem with regard to message transmission. That is, a message could be sent and then an interrupt issued. The interrupt could arrive before the message, in the middle of the message. Some way of marking the point in the data stream where an interrupt was sent is needed.

A subgroup was appointed to consider the above Host-Host problems. Shortly, they would issue an RFC with modifications to the Host-Host protocol, then collect comments and then issue an official revision. People with suggestions should contact the committee. The committee would also be contacting the sites. The committee is:

- S. Crocker, UCLA (Chairman)
- R. Tomlinson, BBN
- T. Barkalow, Lincoln Lab
- G. Grossman, University of Illinois
- J. White, UCSB
- R. Bressler, MIT, Project MAC

The discussion then returned to problems of typewriter access to the network. The problems are presented in RFC #97, NIC (5740). Some are:

- a) Character set
- b) End of Line
- c) Interrupts
- d) Message Format
- e) Half Duplex, Full Duplex

These problems were given to a committee on typewriter connection protocol for solution:

Tom O'Sullivan, Raytheon (Chairman)

Ed Meyer, MIT-MAC

John Melvin, SRI-ARC

Bob Long, SDC

Bob Metcalfe, Harvard

Wil Crowther, BBN

This committee will come up quickly (within a week) with an interim protocol and within several weeks a longer term protocol.

Thursday afternoon, February 18

Thursday afternoon was open to a presentation by the University of Illinois on the ILLIAC IV and a demonstration of the Plato project. The initial test in November of the transmission lines to the ILLIAC IV processors indicated no timing problems. The ILLIAC IV hardware is to be up the fall as is the software. The system will be located in California at NASA Ames Research Center. The connection to the network from the University of Illinois will be a PDP-11 with storage CRTs, 2400 baud character CRTs, typewriters attached. It will have a Gould Clevite printer, DECTapes and small disc. The B6500 at the University will also be connected to the Network.

Thursday evening, February 18

The initial topic was a discussion of status and plans for the Network Information Center. Dick Watson of SRI reviewed the present off-line system consisting of a Station Agent and Network Liaison person. The function of the Station Agent is to aid in the use of the NIC services. The function of the Network Liaison person is to be a point of contact for technical questions about his site which may be asked by people at other sites, and to see that the appropriate people see relevant documents and information received by the site. If the network is really going to develop the feeling of a community, people need to be aware of what people are doing and thinking at the various sites. Therefore, people were encouraged to send reports, memos, notes, records of conversations of general interest through the NIC. Any kind of information can be sent through the NIC from formal reports to informal handwritten notes. In order to encourage people to send out initial thoughts and ideas

as well as those having had much thought, the question was raised as to whether or not there should be titles for different classes of documents which would help to make clear the level of informality or formality of the communication.

There did not seem to be a need for such an arrangement. The question of privacy and security was then raised. There was some feeling among a few people that if letters or records of conversations were entered in the NIC collection that there might be compromise of some privacy. The NIC was asked if it would check all parties involved in such a communication before entering it in the collection. Dick felt that given NIC's resources, it would be better if the parties involved gave their approval before giving the letter or other communication to the NIC.

The initial online services to be provided by the NIC are access to a typewriter version of the SRI-ARC On-Line system (NLS), provision of a message service, access to the NIC catalog and probably files of site status, network personnel, etc. Services will be provided later to aid station agents in communities at their sites. At the principal investigators meeting there seemed to be considerable interest in having NIC obtain a collection of ARPA project reports and working papers. To handle storage from such an expanded collection, use of microfilm seemed important. There are number of problems with use of microfilm, such as a single or limited number of readers and need for hardcopy facilities. The NIC will be looking into these problems and begin experimenting with use of microfilm material.

The NIC is experimenting with remote access to NLS using an IMLAC terminal. Considerable interest in graphic access to NIC was indicated. The NIC feels graphic access is not an immediate high priority requirement, but will as soon as possible provide specifications to those sites with programming resources waiting to experiment with graphic access.

Steve Crocker brought up the problem of how people are to gain access and learn to use service facilities at various sites. The question of what additional information needed to be included with or appended to user documentation to use service facilities over the network was discussed. The question of what material should be in hardcopy, and what online was raised. The NIC will study these problems and produce a set of recommended procedures for handling user manuals, and a list of information needed to enable network access.

Dick Watson indicated that users of the NIC would feel most comfortable using typewriter terminals running at 30 char/sec and having upper and lower case graphics, although service would be available for slower terminals and terminals with single-case graphics. RFC #97, NIC (5740), described an initial protocol for connection to the NIC. As a result of the formation of a committee to produce a standard typewriter connection protocol, the protocol of RFC #97 will be modified to conform to an interim protocol suggested by that committee. A new RFC will be issued shortly with the interim protocol. Since the meeting the typewriter connection protocol committee has decided not to issue an interim protocol.

The discussion turned to file transfer between sites by name and without users being required to log into each site involved in the transfer. Gary Grossman of the University of Illinois will produce an initial RFC on this subject.

Friday morning, February 19

There are several aspects of Data Management associated with the network. The following aspects and the people responsible for them were indicated:

Data Machine 10^{12} bit store

Data Management Language

The Form Machine

ILLIAC IV Information Management System

Interim File System

File Transfer Protocol

The Data Machine is Computer Corporation of America's responsibility, but close coordination with the ILLIAC IV Information Management System and network efforts toward a Data Management Language is required.

The work on a Data Management Language is to be coordinated by J. Madden of University of Illinois, Bob Metcalfe of Harvard, J. Heafner of RAND, Jim White of UCSB, and Doug McKay of IBM.

John Heafner indicated that he plans to implement his plans for the Form Machine, RFC #83, NIC (5609) UCSB, Multics, and Lincoln Lab also indicated that they are interested in getting a version running.

A number of sites, UCLA, SRI, RAND, University of Illinois, Raytheon, MITRE, indicated interest in the range 1-3 months in storing files on UCSB 360/75 disc packs. Jim White said he would produce a system within the next 4-6 weeks to allow network users to store files at UCSB.

The problems of file transfer by name between host systems was again raised and G. Grossman of University of Illinois indicated he would start a dialog on the subject by producing an RFC.

The question of user names and the meaning of user IDs in socket numbers was raised. At present socket numbers have no structure, but several people felt that for accounting, file transfer, and interprocess communication some structure was probably valuable. A committee consisting of:

J. Heafner, RAND (chairman)

E. Meyer, MIT-Multics

G. Grossman, University of Illinois

will produce an RFC stating the issues behind alternate proposals for socket number structures.

UCLA indicated it wanted a link number in the experimental range of link numbers for use in measurements experiments with the network. Link number 223 was assigned to this function. (Link 223 was later discovered to be assigned. Link 191 was chosen instead. See RFC #104, NIC (5768,)).

The problem of accounting was raised as a number of machine or systems on the network will provide service functions. The present service facilities being the 360/91 at UCLA, the 360/75 at UCSB, the NIC at SRI, Multics at MIT, the ILLIAC IV, the 360/67 at Lincoln Lab, and the Data Machine. The advanced Host-Host protocol study committee is looking at the accounting problem. There was brief mention made of a network banking system. Bob Kahn of BBN indicated that he would start a dialog on the subject of accounting by producing a paper putting down the issues as he sees them.

The question was then raised about handling of administrative procedures such as obtaining accounting numbers on foreign systems. Dick Watson said he would look into this problem and see how the NIC can help in its solution.

The final question to be considered was the frequency and utility of these NWG meetings. The general consensus was that this had been a useful meeting, but that more preparation on specific topics to be discussed at the meeting should be done ahead of time. People who want to bring up topics at the meeting were asked to distribute position or introductory papers about a month ahead of the next meeting, if possible. Peggy Karp will handle trying to obtain a block of rooms for the NWG during the Spring Joint. She will send out a request for reservations to the sites soon.

[This RFC was put into machine readable form for entry]
[into the online RFC archives by Kelly Tardif, Viagénie 10/99]