

Internet Meeting Notes - 14 & 15 May 1980

I. INTRODUCTION

The meeting was held at the Laboratory for Computer Science (LCS) at MIT. Dave Clark was the host.

II. OVERVIEW AND OBJECTIVES

Vint Cerf put forward the problem of internet computer mail as a topic that must be addressed in the next few months. This would be an interim system not expected to support multimedia mail but to operate on TCP and to provide some mechanism for interworking NCP supported mail with TCP supported mail.

Another issue is gateway-host interaction, especially for congestion control.

A third topic is routing in a large packet radio environment with optional use of backbone point-to-point links.

III. STATUS REPORTS

A. ARPA

Vint Cerf reported that the IP and TCP specifications (IENs 128 and 129) are now stamped as DOD standards. DCEC will serve as the focus for coordinating this within DOD. A seminar for DOD people is planned for July at NBS to explain TCP.

The monthly reports are very useful, but still needed are some milestone schedules, for planning and cross project coordination.

B. BBN

Dale McNeill reviewed the work on the VAN gateway, noted the arrival of the FAX equipment, described the installation of the PSP terminal at TANUM, and noted that improvements have been made to the CMCC software.

Mike Brescia described the VAN gateway in more detail. The operation of the gateway is still some way off. Current activity

focuses on the X.25 interfaces. At this point there was an extended discussion of low level interface options.

Jack Haverty described the work on the Unix IP and TCP. Most recently a version using shared memory for user/TCP communication has been tested and is 2 to 3 times faster than the older "pipes" version. Work is just beginning on an IP and TCP for VAX Unix. This version will be in the kernel and will include the fragmentation and reassembly features. Also in progress are a TCP for a HP 3000, and a TCP-TIP in a 316. There was also some discussion of which versions of Unix are involved: The VAX Unix is version 7, the current 11/70 Unix is version 6. Jack also mentioned some performance measurements using a traffic generator. This will be documented in a forthcoming IEN.

Ruth Nelson briefly discussed a local net project at BBN which is building a net based on the CHAOS net design. There is now a system with two interfaces which is tested by linking two terminals.

Bill Plummer reported on the TENEX/TOPS20 IP and TCP status. A number of improvements have been made to the debugging version at BBNF and these should soon be distributed and installed at other sites. Among the improvements are: a resolution of the "data stream capture problem," an update to the internet user queues, an improvement to the gateway functions - especially updating the gateway table. Bill has also prepared an "Installation Guide for Wizards."

Ginny Strazisar reported that the last ELF based gateway has been converted to MOS, and that the first gateway running in an LSI-11 has been delivered to SRI. Also a three port gateway is working.

C. COMSAT

Hoi Chong reported that there are currently a few problems with the COMSAT gateway which seem to be related to a power supply. Also it was noted that the line to NSRDC will be removed at the end of May. The IP and TCP used in the COMSAT hosts has been improved and these programs now support the timestamp option. The FAX machine and protocols are in place and ready to be used, what is needed is a common FTP to move the bits to other sites. COMSAT was particularly active in the "Bakeoff" as reported in several messages. Within COMSAT there is an investigation of a local net to link some PDP-11 and IBM computers. Hoi also distributed a memo on COMSATs internet activities and milestones.

D. DCA

Ed Cain reported that the DCEC gateway had been considerably improved with help from Bill Plummer, Dave Mills and others. Ed also reported that DCEC role in the DOD protocol standards will have three aspects: (1) Executive Agent, (2) Review Panel, (3) Laboratory. This Laboratory will be based on the EDN.

E. DOD

Ray McFarland reported that a group in his office building a TCP 2.5 for a PDP 11/34 with Unix (v.6). This group needs software for Telnet and Echo processes. Ray also discussed his role as an information coordinator in DOD. He receives many requests for information from within DOD. He needs help in fielding these questions. In particular more data on performance would be very helpful.

F. ISI

Jon Postel reported that multi-media mail system is being redesigned, and that a new draft of the MPM specification is available. Also good progress has been made in the investigation of program verification tools for protocol analysis - especially AFFIRM. ISI has ordered 4 computer readable clocks which maintain their time by listening to WWVB. Jon also described the recent movement of users between the various machines at ISI.

G. Linkabit

Estil Hoversten reported that Linkabit is not much involved in the internet activities directly, but rather in the details of SATNET and the WBC project. Estil and Danny Cohen are working on an overview paper.

H. LL

Jim Forgie noted that one goal of his project is to demonstrate point-to-point speech via the WBCNET in FY 80. This may be hard to do due to the slippage in the schedule for some of the key pieces. Lincoln is going ahead with the development of the voice terminals and their access path -- the LEXNET. In cooperation with ISI programs are being written to support NVP2, ST, and IP. Lincoln is using the ISI developed EPOS system for this project.

I. MIT

Dave Clark reported on the various systems at MIT and their interconnections. Things are getting quite complex. (1) Unix: MIT has a copy of the IP/TCP from BBN, which MIT modified to have a user accessible IP interface. A TFTP was installed in this host. (2) Multics has had only small changes since last time, there is now a Name Server. (3) The ALTOs now have IP, UDP and TFTP in BCPL running; IP, TCP and Telnet (with SUPDUP) in MESA are being programmed.

Dave noted that the lack of implementation of fragmentation and reassembly is causing problems.

Dave gave some performance measures of the local network at MIT (LCSNET v1) which indicate very high reliability and very low average load. However there is some trouble interfacing MIT-XX (a TOPS20) to this local net.

Dave noted that version 2 of this network is in hardware testing and that the Nu Machine (a personal computer being developed at MIT) will interface to the version 2 LCSNET.

J. MITRE

Anita Skelton reported that starting with a half completed Z8000 C cross-compiler obtained from MIT, the compiler was finished, an assembler and loader was written. Then starting with BBN's C version of SRI's MOS, the interprocess communication was modified, the terminal handler was rewritten, and device drivers were added; MOS is running in the Z8000. MITRE is now converting the MOS-TCP to C, rewriting large portions and following the DOD spec.

There are two Z8000 development boards interfaced to the MITRE cable, with the cable contention algorithms coded, and packets have been exchanged on the cable. MITRE hopes to have TCP running in the Z8000s soon.

The next step is to implement Telnet and interface the 11/70 to the Z8000. The intention is to have a high speed DMA interface to the 11/70 via the UMC-280 and also a parallel interface via the DR11-C, so that comparable measurements with TCP running over the cable can be made. (These measurements for the older interface unit have been made).

In addition, a cable bus test bed will be installed at DCA in Reston. The interface units will be Z8000 based and built by Reaction Instruments; the interface units will sell for about

Internet Meeting Notes

\$4000 with a parallel host port and two terminal ports (asynchronous and synchronous). The cable will sit between the IMP and the NFE, with terminals attached to the cable, and one Z8000 unit will act as a gateway.

K. NDRE

Yngvar Lund reported that NDRE is investigating a local network with HDLC interfaces to the connected computers. The connected computers will be built up of modules based on Z80 processors and programmed to support voice protocol, HDLC, TCP, etc.

L. RSRE

Andy Bates noted the work on Extended Memory MOS was reported in IEN 136. IP was recoded in CORAL66 for EMOS and a recoding of TCP is planned. Modems to upgrade the RSRE-UCL line are on order. Work is in progress on a gateway between PPSN and PSS.

M. SRI

Ron Kunzelman gave an overview of SRI's activities and a brief rundown on hardware available. Ron discussed some instrumentation efforts underway at SRI. A PDP 11/44 is on order to aid in this work.

Ed Perry discussed some of the problems with the Ft. Bragg installation of the Packet Radio net. The gateway (or attached nets) seem to have a low packet/second throughput limit. This is causing the higher level protocols to be modified to use fewer packets (e.g., a "half duplex" user mode).

There was some discussion at this point of XNET and again the problem of "big" datagrams. It really is important to have fragmentation and reassembly implemented everywhere.

Holly Nelson reported on the port expander. New versions were delivered to Ft. Bragg and UCL.

A new directory has been set up for the current version of TIU software. It is <TIU-SOFTWARE> on SRI-KL. News of interest to TIU programmer or others concerned with MOS software should be sent to the MOS-Users-List run by Noel Chiappa.

N. UCL

Rob Cole related the recent activities at University College. There are many pieces of equipment at UCL connected in complicated

ways. The current focus is a local "Cambridge Ring" network. The NIFTP now works on ISIE and an LSI-11 host. There is a revised specification of NIFTP forthcoming. UCL has a grant to provide a computer mail service between UK and US university users. There is a major concern about cost of transmission via IPSS, so current focus is on multi-destination mail in single transmission.

The PDP-9s which serve as a gateway between the ARPANET and EPSS will die if either EPSS changes or the 32-bit leaders in the ARPANET go away. The PIXIE device is no longer in use.

O. UCLA

Bob Braden reported on the IP and TCP in the IBM 3033 at UCLA. The operating system is VM with OS-MVT. The IP and TCP are running and the 96-bit leader (24 bit address) support is installed. Further work is needed in the area of "gateway" functions in the IP. In TCP some work is needed in the ACK policy area. One can poke this system via telnet. Try "NETSTAT" and "HELP-TCP".

P. XEROX

Vint Cerf read the following report sent in by John Shoch:

1. We have contributed our efforts to measure the effectiveness of the Packet Radio Network when used as part of the Pup environment. IEN 138 describes some of the (disappointing) results obtained with the new IPRs; we are in the process of moving one of the radios, however, and plan to replicate the tests in the near future.
2. IEN 140 is a paper prepared with Danny Cohen and Ed Taft, outlining the use of "mutual encapsulation" as a means to support the coexistence of the Pup protocols and the IP protocols.
3. On Tuesday, May 13, Xerox joined with DEC and Intel to announce the results of an effort to establish a common specification for a local network, based upon the Ethernet technology.

Internet Meeting Notes

IV. MEASUREMENTS

Andy Bates reported that RSRE conducted further measurements (since the last meeting). Their findings again indicate a large spread in the delay through the concatenated ARPANET and SATNET. Some speculated that the ARPANET introduces this spread. The RSRE group suggest that it would be appropriate to introduce a provision for a negotiated retransmission option. Also a larger window would help.

Rob Cole presented the results of some recent measurements made by UCL of the gateway and SATNET. The gateway at UCL was able to echo packets at 34 packets/second. When the path was extended to the BBN gateway, the UCL gateway accepted 14 packets/second, but very few echo replies were received from the BBN gateway. A port expander in the path did not pose a limit on either measurement.

V. GATEWAY PROTOCOLS AND HOSTS

Jim Mathis presented his procedure for routing. The main points are first pick any gateway, second refine the chance to the best gateway and third detect the failure of that gateway should it occur.

- o Pick a Prime gateway
- o Poll it at a slow rate
- o Send to the Prime gateway
- o Accept and act on a Redirect message
- o Ping gateway in use if higher level protocol complains
- o Periodically change the Prime gateway

Does this procedure get unstable in high load?

IENs 109 and 131 should be reviewed by host IP implementors.

VI. TELNET AND FTP FOR TCP

Jon Postel reported on draft specifications for Telnet and FTP for TCP. For Telnet the key changes are that ICP is eliminated and the single full duplex connection is between ports U and L. Another change is that the Telnet SYNCH becomes DM + Urgent, but one can't count Urgents in the way one used to count Interrupts.

For FTP, again ICP is eliminated for the control connection, which is now the full duplex connection between ports U and L. The major change here is the elimination of the BYTE command and all that implies. Other changes affect the defaults for the data connection and the "third party" transfers. The MAIL and MLFL commands are included in the specification, and the new reply codes are used.

Internet Meeting Notes

VII. NEW IMPS

Vint Cerf reported that BBN has formed a separate BBN Computer Company. This company sells a C-30 machine (formerly known as an MBB). The current C-30 emulates an H316.

Some things on the queue for IMP improvements are:

1. extend the memory capacity of the C-30 IMP program
2. make software changes to support more than 4 hosts
3. provide HDLC interfaces for hosts
4. investigate the non blocking interface
5. provide logical addressing for hosts

Note also that a TCP-TIP is being developed with the TCP/Telnet code in a H316 and the IMP part in a C-30.

VIII. THE CMCC

David Flood Page gave a brief review and a demonstration of the CMCC functions. Basically the CMCC programs collect data from cooperating gateways and display the results on a terminal. The programs reside at ISIE in directory <CMCC>. The key program is CMCC-DISPLAY. Some files in this directory give helpful information about the program. These are: HELP.TXT, NEWS.TXT, and SESSION.TXT. Further information can be obtained from David via sndmsg to DFloodPage@BBNE. Please see also IEN 131 and 132. The demo went very well and these programs should be useful to any gateway builder.

IX. INTERNET PROTOCOL SPECIFICATIONS

Ken Shottling presented some results of studying the IP specification for a formal description. The key result is the identification of ambiguous areas in the specification. In particular, the interaction of the fragmentation procedure and the return route option is one cause for concern. Another issue is the use of the identification field.

X. EGG BREAKING

Danny Cohen led a discussion of the problems arising from assumptions about which end of a word/page/..., bytes are transmitted from. This is a holy war between the big-endians and the little-endians (see IEN 137). Noel Chiappa and Danny were appointed as a small group to argue about it.

XI. FLYING PACKET RADIOS

Radia Perlman presented some ideas on how to handle mobile host (e.g., flying packet radios) based on having the gateways do most of the work. This approach is based on the gateway using a "link-state" routing procedure and a method of handling partitioned nets as if each partition were a separate net. The method presented is related to but not identical with that presented in IEN 120. A revision of IEN 120 will be forthcoming.

Carl Sunshine presented an alternate strategy for dealing with mobile hosts which makes the hosts do most of the work. This scheme uses the existing IP source routing option and a new GCP message. It calls for a new host (or special functions in an existing host or gateway) to support (1) forwarding and (2) a global name server. New messages are needed in the gateway protocol (IEN 109) to relay information about a mobile host's current location to the name server and the "connected" hosts. This procedure is described in IEN 135. This procedure may also be workable for multihomed hosts and as a least effort solution to the partition problem.

Clearly a decision must be made as to whether these problems are to be solved within the gateways or not.

Vint Cerf described another routing problem. This problem arises when a destination can be reached either within a network or via another network which is connected to the first network in two (or more) places. For example, the ARPANET and the WBCNET will be connected in four places. It may be better for messages from Boston to Los Angeles to go one hop in the ARPANET then via WBCNET then one hop in the ARPANET, than cross country via many hops in the ARPANET. Another example, is a large Packet Radio environment surrounding a small ARPANET style network. Messages from one packet radio to another packet radio on the opposite side of the environment might best be forwarded through the ARPANET style net rather than via many PRNET repeaters. This out-of-net crossnet routing is a difficult decision to make by current procedures. (The latter example was described as a "cloud of packet radios." First we had a flying PR, now we have a whole cloud! I hope it doesn't rain!!!)

XII. SOURCE ROUTING IN A CAMPUS ENVIRONMENT

Jerry Saltzer presented his ideas on routing in a collection of local network in a campus environment. Much of the strategy is based on the lack of central control over the environment.

A Name Server will be provided that will supply route information to be used in the source route. Source routing permits gateways to be

very simple, and allows the end user some control over the route. The latter is good for trouble shooting, and performance control.

These ideas are described in IENs 143 and 144 which were distributed at the meeting.

XIII. CONGESTION CONTROL

Dave Clark discussed some studies about what to do when a Source Quench message arrives. A simulation program was used to try various methods. All were bad. The delay in the feedback loop is quite large and some information needs to be supplied about how far away (in time) the bottleneck is.

XIV. LOGICAL HOST SUPPORT

Jon Postel and Bill Plummer discussed some of the issues in logical host support. The problem came up with a desire to run the testing gateway at the same time as the regular IP. This is not possible but points out the need for the IP layer to support the multiple logical host concept.

XV. BAKEOFF

The "distributed bakeoff" was too distributed in space and time to be as effective as one would like. Some good things did come from it. (E.g. see Dave Mills messages). It was felt that the next time it should be scheduled for one day and with the participants in two or three locations.

Internet Meeting Notes

XVI. NEXT MEETING

The next meeting will be held at RSRE in Malvern, England, on 7,8, and 9 October 1980. Attendance will be restricted so if you plan to attend please clear that with Vint Cerf and notify Linda. John Laws will be the host and information about the local arrangements will be distributed at an early date. Malvern is about two and one half hours by train from London.

AGENDA ITEMS

1. Resolution of the Partitioned Net Problem - Cerf
2. Proposal for Controlled Routing - Cohen
3. Experience with VAN Gateways - Brescia, Kirstein
4. Demonstration of Interim Internet Mail - Postel
5. Performance Evaluation Parameters - SRI
6. Name Server Demonstration - SRI
7. New ST Document and Explanation - Cohen, Forgie, Hoversten
8. Tiny Pipe Nets vs. the Catenet - Mills
9. Congestion Control - Clark

ACTION ITEMS:

1. XNET Specification - Haverty, Strazisar, Mathis, Tomlinson
2. Tenex running TCP4 - Plummer

Internet Meeting Notes

APPENDICES: Small Group Discussions

A. FTP and Mail

Need to extend current style of computer mail to work in the internet. The extension of "Mailbox@Host" to "Mailbox@Host@Net" may not be acceptable because too many programs must be changed. Rather something like "Mailbox@Place" is suggested. This gives all hosts at least one global name for mail purposes. (Will there be a name czar?) "Place" is going to have to map to a 32-bit internet address. The mailer will have to try to send the mail via TCP if both source and destination know that or via NCP if both know that or if one is TCP only host and other is a NCP only host via a forwarding host that knows both and provides a special forwarding function.

Postel will produce an IEN on this topic.

B. FTP

Brief discussion of FTP and default data ports. The data ports should be U and L-1 so that port pairs are not required in the users portspace. Also it seems OK to use the new reply codes.

C. Measurements

Tools are needed for measurements, for example, echo servers, and time stamps. Echo servers can be TCP based, UDP based, or use the GGP; but more are needed. Timestamp needs and desirable properties will be described in a memo by Rob Cole and Andy Bates. Jack Haverty's forthcoming IEN on a traffic generator should help too. Ed Cain has a report on a TCP Tester. Also note the CMCC facilities.

D. Fault Isolation

This discussion was subtitled "What to do when things go wrong." The conclusion was "try another gateway."

E. Transport Layer Specification

Richard Tenney is working on a task at BBN (under contract from NBS) to specify a Transport Protocol. He is developing a specification methodology along the way. Richard had some questions about the TCP specification and the actual behavior of real TCPs. The main questions had to do with out-of-order data, resets, and closing cleanly.

Richard noted that we all should look at some NBS reports:

"Draft Features Analysis of Transport Protocols"

"Draft Service Specifications of Transport Protocols"

obtainable from John Heafner of NBS (Heafner@NBS-10).

F. Fragmentation and Reassembly

Fragmentation and Reassembly must be implemented so that datagrams of up to 576 octets (including header) may traverse the catenet. It is proposed to hold a special bakeoff to test this capability in early September.

G. Gateway-Gateway Protocol

Discussion of how much of this must be known by the hosts. The messages: Source Quench, Destination (Host/Net) Unreachable, Redirect, Echo/Echo Reply. Much discussion of error reporting should it be in GGP or IP? (The current IP error option seems useless). Postel will write a memo on error reporting (including who processes which errors).

H. IP Option Overflow

The problem is that some IP header option fields (e.g., return route, timestamp) may expand, causing the maximum header size to be exceeded. Ways to handle this might be: (1) discard the packet, (2) expandable options should have an "overflowed" flag which gets set when they can't grow any more and the packet is forwarded without expanding the option, (3) expandable option should be sent with filler so header is "right" or max length to start with. On header overflow both source and destination hosts should be told out about the error. The issue was not resolved. But the problem should be discussed in an "Implementers Guide."

An additional problem concerns making full length fragments when variable length options are present making it likely that the first fragment will overflow the next maximum packet size and have to be further fragmented later. One suggestion is to leave the first fragment not quite full in such a case.

I. Name Server

SRI will implement a name server. It will be an extension of the one specified in IEN 116.

Internet Meeting Notes

J. Acknowledgement Algorithms - Adaptive Timeouts

There seems to be a lot of ACK traffic so maybe ACKs should be sent on a periodic basis rather than on an event basis. The retransmission strategy also needs to be smarter. Could it be negotiated? Type of service should be a consideration in these algorithms. Bill Plummer and Andy Bates will conduct an experiment. Idea: set retransmission time according to the Network at the other end of the connection.

K. ARPANET Problems

The "8 messages outstanding at a time" problem was discussed. It turns out that the maximum messages in transit can often be less than 8 due to heavy loads at the destination IMP. It seems that the only solution to this is to provide more buffers and message blocks in the busy IMPs.

RECENT DOCUMENTS

IEN	Author	Title
---	-----	-----
131	Flood Page	Gateway Monitoring Protocol
132	Flood Page	The CMCC Terminal Process
133	Sollins	The TFTP Protocol
134	Postel	Internet Meeting Notes-4,5, & 6 February 1980
135	Sunshine	Addressing Mobile Hosts in the ARPA Internet Environment
136	Wiseman	Memory Management Extensions to the Micro Operation System for PDP-11/23/34/35/40
137	Cohen	On Holy Wars and a Plea for Peace
138	Shoch	Initial Comparison of EPRs and IPR in the Pup Internet Environment
139	Haverty	HOSTs as IMPs
140	Shoch	Mutual Encapsulation of Internetwork Protocols
141	Bennett	Message System Issues
142	Postel	Time Server

DOCUMENTS DISTRIBUTED

IEN	Author	Title
---	-----	-----
143	Saltzer	Environment Considerations for Campus-Wide Networks
144	Saltzer	Source Routing for Campus-Wide Internet Transport

Internet Meeting Notes

ATTENDEES

Vint Cerf	ARPA	Cerf@ISIA
Mike Brescia	BBN	Brescia@BBNE
Ross Callon	BBN	RCALLON@BBND
Len Evenchik	BBN	EVENCHIK@BBNE
Gil Falk	BBN	
Jack Haverty	BBN	JHAVERTY@BBND
Dale McNeill	BBN	DMCNEILL@BBNE
David Flood Page	BBN	DFLOODPAGE@BBNE
Ruth Nelson	BBN	RNelson@BBND
Radia Perlman	BBN	PERLMAN@BBN
William Plummer	BBN	Plummer@BBNA
Virginia Strazisar	BBN	STRAZISAR@BBNA
Hoi Y. Chong	COMSAT	Chong@ISIE
Chris Elliott	CTEC	CTEC@BBNC
Ed Cain	DCEC	Cain@EDN-UNIX
Jim Showalter	DCEC	gamma@EDN-UNIX
Michael Begun	DEC	BEGUN@DEC-MARLBORO
Ray McFarland	DoD	McFARLAND@ISIA
Ken Shottling	DoD	Shottling@SRI-KL
Danny Cohen	ISI	Cohen@ISIB
Jon Postel	ISI	Postel@ISIF
Carl Sunshine	ISI	Sunshine@ISIF
Estil Hoversten	Linkabit	Hoversten@ISIA
Jim Forgie	Lincoln Lab	FORGIE@BBN
Noel Chiappa	MIT	JNC@MIT-XX
David Clark	MIT	Clark@MIT-Multics
Steve Kent	MIT	STK@MIT-XX
Jerry Saltzer	MIT	Saltzer@MIT-Multics
Karen Sollins	MIT	Sollins@MIT-XX
Anita Skelton	MITRE	Anita@MITRE
Frank Deckelman	NAVELEX	DECKELMAN@ISIA
Yngvar Lundh	NDRE	Yngvar@SRI-KA
Oyvind Hvinden	NDRE	Oyvind@SRI-KA
Glen Allgaier	NOSC	ALLGAIER@ISIC
Merle Neer	NOSC	Neer@ISIA
Andrew Bates	RSRE	RSRE-T4@ISIE
John Laws	RSRE	RSRE-T4@ISIE
Ron Kunzelman	SRI	Kunzelman@SRI-KL
Jim Mathis	SRI	Mathis@SRI-KL
Holly Nelson	SRI	HNelson@SRI-KL
Ed Perry	SRI	Perry@SRI-KL
Robert Cole	UCL	UKSAT@ISIE
Peter Kirstein	UCL	PKIRSTEIN@ISIA
Bob Braden	UCLA OAC	Braden@UCLA-CCN