

Danny

Internet Meeting Notes - 25 & 26 January 1979

INTRODUCTION TO ISI - Jon Postel

Jon Postel welcomed the group to ISI and pointed out the relevant facilities.

OVERVIEW AND OBJECTIVES - Vint Cerf

Vint Cerf reviewed the need to have an operational Internet capability very soon, noting the Fort Bragg PRNET experiment and the cooperative experiments with RSRE (Royal Signals and Radar Establishment).

STATUS REPORTS

BBN - Strazisar, Plummer, Wingfield, Haverty

Ginny noted that both the SATNET and PRNET gateways will run both IN-4 and old protocols by the end of January. These gateways will also implement the alternate routing described in IEN-30. The prototype minigateway program is now being tested in the PRNET and will be tested in SATNET before the end of February. The minigateway is just like the regular gateway except it runs on MOS instead of ELF and eventually will run on an LSI-11 instead of an 11/40. Currently, there is no LSI-11 available for testing.

Bill reported that TCP-2.5.2 is up at SRI-KA, BBNC, and BBNE and is to be installed at ISIC this week. TCP-4.0.0 is up for testing at BBNB. UCL users uncovered a bug which will be fixed within a week.

Mike said that the DCEC sponsored "C" version of TCP-4 for Unix is running on BBN-UNIX, EDN-UNIX, and NPS-UNIX. This version is compatible with AUTODIN-II requirements for security precedence and user groups (S/P/T). The system runs on a Unix on a PDP-11 45 or 70 with the addition of a BBN interprocess communication package. The TCP is finished and debugged (except for rubber EOL). The AUTODIN II requirement for preemption is yet to be added. A rough measure of throughput indicates the TCP capable of supporting 11 KB/S while the NCP is capable of supporting 23 KB/S. This throughput increases with an increase in the number of connections handled concurrently, up to a point. The tests were

Internet Meeting Notes

made looping through the local IMP. The size of the TCP is roughly 30K bytes.

Jack reported that the Macro-11 version of TCP for Unix is a descendant of the Macro-11 TCP-2.5 implementation which we have been using for over a year. The TCP-2.5 implementation was based on an encapsulation of the SRI TCP11 program for MOS by Jim Mathis. In converting to version 4, the structure was changed to remove one 'layer' of implementation, with the goal of improving performance by tailoring the implementation to the Unix environment rather than simulating the MOS environment within a Unix process. In addition, other changes for performance improvement were made, such as in the buffer management scheme. The TCP 4 implementation can thus be viewed as a second-generation TCP implementation, which conforms to the version 4 protocol.

The Macro-11 TCP-4 implementation is compatible with the "C" implementation being done by Mike Wingfield, both at the system and user interfaces. This will permit sharing of other software, such as THP/TELNET, as well as testing software. The Macro-11 version additionally has been structured to make addition of reassembly software relatively simple.

The implementation is essentially complete, except for extensive testing and debugging, which will begin in the TCP testing sessions. It does not yet implement 'rubber-EOL' or reassembly, although the design has been done with those features in mind, so they will be easy to add.

SRI - Kunzelman, Mathis

Ron discussed the installation of a number of PRNET terminals (TIUs) at Fort Bragg and the intended use of the ISID Tops20 system via the PRNET and the ARPANET using IN-4 and TCP-4 protocols. A key piece is the gateway which currently resides in the same PDP-11 as the station, but work is underway to provide the gateway and station functions in separate hardware.

Jim reported that the SF area PRNET is temporarily down while the radio units are being upgraded. When the radios are available again, the system will continue to use TCP-2.5 until the TCP-4 is available on Tenex/Tops20 hosts. The number of radio networks in the SF area will be increased from 2 to 4.

MIT - Chiappa, Clark

Noel noted that the LCS net now has four running nodes. For implementations of IN-4 and TCP-4, MIT will try to utilize other

Internet Meeting Notes

peoples' code. Two LNIs are to be delivered to UCLA this week. Noel has provided the following summary:

Gateway

We are in the process of constructing two gateways: one internal gateway between the LCSNET and the CHAOSNET, and another based on Jim Mathis' Port Expander between the LCSNET and the ARPANET. The LSI-11 and all hardware for the latter has been obtained, although the 1822 boards are still not working. Convincing them to work will be the first thing done on return. Several new part-time staffers have been added, and one has been at work familiarising himself with Jim's code and the Twenex support software, and he is now ready to start on actual gateway work.

Network Hardware

Our hardware is now completely debugged and in service. We are now in the position to "crank out" LNI's for local use with a minimum of effort. Turn around time is several weeks for chassis fabrication and wirewrapping at outside service, and a day or two to actually "stuff chips" and find any miswires, bad chips, etc.

Internet Protocol

There are plans for putting an internet layer into ITS, but no manpower is currently available, so the plans have been temporarily shelved. Multics now has a full internet layer, with all the necessary features to support multinetwork usage, as well as support for other protocols in addition to TCP. It currently uses only the ARPANET, but this could be changed fairly easily if there was any need.

TCP

As with internet, no people are currently at work on a TCP for ITS. Multics TCP is now more or less fully up. There are several very preliminary servers available - New and Old TELNET and a Time Server. Those interested should contact Dave Clark for more details.

UCL - Kirstein

Peter discussed many aspects of the UCL networking activities, including the interconnection to RSRE. There are many problems of protocol conversion, both low-level and high-level. Several special interfaces have been developed to convert between 1822 type lines and various degrees of X.25 lines. There are numerous

delaying hardware difficulties in this type of device. There are also issues of higher level protocol conversion when attempting to set up an end-to-end terminal-to-host or file transfer path across multiple networks. Peter subsequently supplied the following summary:

UCL PROGRESS NOVEMBER 1978 - JANUARY 1979

1. HDLC HARDWARE

The drawings for the PCBs for the Program Interrupt version of the HDLC hardware for the LSI-11 were delayed. They will now get to the manufacturer only about Feb. 6. We expect prototype boards only towards the end of March, with production quantities only early in May.

2. DMA HDLC HARDWARE

We have not received all the information needed to design DMA boards for HDLC compatible with the SRI interface. In view of the number of high performance boards now becoming available, we will review the SRI, RSRE, Dublin U, and other designs before we decide definitely to develop HDLC boards to the SRI DMA Controller.

3. X.25 SOFTWARE TESTS

We have tested the X.25 level 2 software with the UK Post Office Tester (with our software running under RSX-11S). The assembler version of level 3 is coded, but not completely tested yet. Transfer into running under MOS has started.

4. X.25-ARPANET PACKET CONVERSION

The first step of our work with the RSRE will be to put a simple datagram gateway at UCL. All data packets from ARPANET will be stripped of their 1822 headers, have an X.25 level 2 header added, and be passed over the HDLC interface to RSRE. RSRE will mount TCP/MOS above the level 2.

5. TCP - X.25 CONVERSION

The problems of mapping TCP to a X.25 Transport Service are being studied in detail. Coding is expected to start in March.

6. NETWORK-INDEPENDENT FILE TRANSFER

The NIFTP code on the TOPS-20 works well enough to Retrieve and Put files on the same TOPS-20. Some preliminary documentation has

Internet Meeting Notes

been written. Interfacing to TCP requires resolution of some TOPS-20 TCP problems by Plummer.

7. INTERNET FACSIMILE

The code for connecting the DACOM facsimile to TCP on the LSI-11 has been written (but does not yet work correctly). The other end, to enable us to Retrieve and Put files onto the TOPS-20 file system, has been halted pending the resolution of the TCP problems of (6).

NDRE - Stensby

Aage reported that the NDRE NORD-10 TCP-4 has been running for about three months, though it does not yet implement Rubber EOL or send URGENT. The IN level is not a separate module and does not handle fragmentation. The maximum acceptable segment is 1008 octets. This has been tested in self loops via the testing gateway ast BBNB.

UCLA - Cerf

Vint noted that UCLA's internet work is primarily theoretical research on throughput and delay analysis. This work is headed by L. Kleinrock. UCLA is also building a small local network based on the LNIs provided by MIT. This work is managed by G. Popek. UCLA is also implementing a TCP-4 for the 360/91. Bob Braden is doing the implementation.

CCA - Postel

Jon distributed the following written report from CCA since their representative was unable to attend:

TCP for RSX-11m

CCA is integrating the SRI TCP-2.5 into an "ancillary control process" (ACP) for RSX-11m. This ACP is addressed very much as a device driver and employs executive facilities originally designed for DECNET. A server TELNET conversion and user TELNET are in the planning stage. When SRI TCP-4 is released, it will be incorporated into the ACP.

Current trials are in the connect-to-self stage. The schedule for completion is contingent on funding details, and could be as soon as the first half of 79.

Internet Meeting Notes

Internet Protocol under RSX-11m

In order to facilitate Internet usage by protocols other than TCP, CCA will isolate Internet functions into a separate "ancillary control process", which, under RSX-11m, is addressed very much as an I/O driver.

This work will commence when SRI TCP-4 is released and will be coterminous with the CCA revision of this software.

PARC - Shoch

John briefly reviewed the status of networks at XEROX. There are 7 types of nets, 20 nets, 500 hosts. The particular experiment of interest to the IN group is the use of the PRNET as a link between two ETHERNETs located about a mile apart. Adding the Radionet to the existing Parc internetwork architecture required building an 1822 interface, and adding a network-specific driver for the Radionet in the gateway at each end. This driver performs network specific fragmentation, splitting up large internet packets into smaller fragments that can be sent through the Radionet.

Over 1000 hours of gateway operation have been accumulated, and the system seems to work well; the radio link can support one-way byte stream traffic at about 12 KB/S. John noted that measurement shows that they can get about 8.3 KB/S from a line rated at 9.6 KB/S; and out of the 100 KB/S radio channel, they can usually get 12.5 KB/S; and if all the parameters are tuned, they can on occasion get 25 to 30 KB/S.

John also discussed a flow control problem that can seriously effect throughput. Basically, the variability in delay may trigger undesirable retransmissions that result in every message being sent twice. These effects may be repeated resulting in every message being sent many times. The serious part of the problem is that such states may be stable. This topic is further discussed later in the meeting.

This particular failure mode has since been corrected, by slightly modifying the flow control heuristics, and system behaves properly. This experience again highlights the difficulty of designing good flow control and congestion control mechanisms, especially when inter-connecting networks whose basic performance differs by several orders of magnitudes.

John also requested contributions to a bibliography on local networks that he is preparing.

MITRE - Skelton

Anita reported that the MITRE-Unix is now up, and the local network is installed with 6 interface units in place. The local net is based on cable TV technology using a 300 MB base band with a 300 KB subchannel for packet data transmission. TTY-to-TTY connections have been tested on the cable, and the 11/70 is being connected via a single DH-11 line with software multiplexing. Eventually, the 11/70 will interface to the cable via a UMC-Z80. There are three LSI-11s soon to be connected to the cable. These could run the MOS-TCP from Jim Mathis. Anita will provide some material from a MITRE working paper for an IEN describing the network.

FORD - Ken Biba

Ken discussed the work at FORD on the KSOS (Kernelized Secure Operating System). This is to be a UNIX-like system with multi-level security. The Critical Design Review is at the end of February. The system will be written in MODULA (possible fall back is "C"). Ken pointed out the problem of handling IN protocol fragment reassembly in a secure system. This is further discussed in IEN-73. The key issue seems to be to provide some demultiplexing information in every fragment so as to not require "shared" storage in the kernel.

Ken also reported that FORD is building a local net based on an 1MB channel with DMA connections to PDP-11s. The network hosts will use NCPs. The network interface is based on a UMC-Z80 from ACC.

COMSAT - Mills

David discussed the Earth station connection at Clarksburg and Washington. The Clarksburg ground station is a 16 KB/S station (the others are 64 KB/S). David distributed a memo titled "Issues and Plans in Internetworking" which describes the Washington connection in more detail.

COMSAT is continuing effort to bring up the small earth station for participation in Internet activities. The principal experimental tool in these activities is a small LSI-11 based system called the Demo Terminal. This system, which is to support SATNET and other activities as well, is equipped with packet speech, facsimile and graphics terminals suitable for use in a multimedia workstation mode.

The software which runs in the LSI-11 is derived both from ARPANET and DEC sources and designed specifically for real-time local

Internet Meeting Notes

networks applications. COMSAT current activities in support of this software include:

1. Integration of TELNET and TCP-2.5 into the system. This is substantially complete.
2. Design and integration of device drivers for packet speech, facsimile and gateway (1822). This is in progress.
3. Construction of an X.25 (HDLC) interface to support cooperative UCL experiments with facsimile. This is in final debugging.

The near-term plans include:

1. Design of suitable interface software to couple the gateway to the Demo Terminal (simple datagram exchange).
2. Design and implementation (petty larceny presumptions) of suitable high-level protocol modules (FTP, mail, etc.) designed to thrive at the end of a long thin flakey line (UET).
3. Design of high-whizzbang content demonstrations both for COMSAT management and public forums such as NTC 79 (which is the SATNET coming-out party).

CHAOS - Greenblat

Richard provided us with a special report on the status of the CHAOS net at MIT. There are currently 8 stations connected by CATV cable running at 8 MB/S. There are Lisp machines, ITS 10's and a 2040 Tops20 system on the net. The protocol was developed by Dave Moon. A feature of this protocol is the CONTACT NAME which is a text string. This mechanism allows one to send a "line" of text to a process which that process may treat as an initial argument. Another feature is the "Receipt" which is a flow control notification that a message has been buffered, but not yet accepted (it will be acknowledged when it is accepted). Richard notes that his group is interested in Interneting and has provided a way for higher level programs in host in CHAOS net to interact with their counterparts in the ARPANET through a bridge program that runs in a host connected to both nets.

FLOW CONTROL DISCUSSION

A discussion of the problem raised by John Shoch took place. The problem is that due to variable delay retransmission and acknowledgment mechanisms may move into stable states where each message is transmitted two (or more) times, resulting in heavy use of the channel with very poor efficiency. Aspects of this problem involve (1) determining the retransmission delay from measured round trip times, and (2) triggering retransmission on the arrival of "old" acknowledgments. The problem develops when the delay in the system suddenly increases so that synchronization between acks and messages is lost. There was also some discussion of the interactions between the levels of protocol. John Shoch pointed out that a paper by Donnelley and Yeh at the 1978 Berkeley Workshop was a good discussion of this problem.

ACTION ITEMS FROM LAST MEETING

- a. IN-4 and TCP-4 Specification Status - Cerf, Postel


Vint stated that the IN-4 plus TCP-4 have been adopted as a basis for standards for network protocols in the DOD.

Jon said the state of the specification is "in editing."


- b. Internet Address to Local Net Mappings

The following transformations are used to convert Internet addresses to local net addresses and vice versa:

ARPANET




SATNET




Internet Meeting Notes


PRNET



LCS NET



EDN



c. Report from the Access Control Meeting - Perlman

Radia discussed some of the issues in access control. The main point is that access control impacts routing. We should really talk about selective routing or augmented routing procedures. One can imagine enforcing routing restrictions on the basis of source (or destination, or both) host, net, user, or the type of service requested. A key issue is how to prevent the forgery of the source identification. The actual degree to which selective routing can be enforced may be impacted by the amount of data required to make an authorization check. Any fine-grain control will require impractically large data bases. A small group meeting will be held on this topic.

d. User Datagram Protocol - Clark

Dave presented the essence of IEN-71. A key point of contention is the existence or not of "well-known" ports (or sockets). The non-existence of "well-known" ports depends on the existence of a name server.

e. Internet Name Server - Postel

Jon reviewed the name server presented in IEN-61 but adapted to use the User Datagram protocol. The material in IEN-61 will have to be edited and reissued. This presentation assumed the name server could be found on a "well-known" port. It was pointed out that there are two models here:

(1) Static Binding = "well-known" ports

(2) Dynamic Binding = "name server"

however, every dynamic binding scheme has a root of static binding.

There is also an issue about what the name server is expected to know, that is how much detail --Nets, Hosts, Server Processes, etc.

f. Internet Speech Capability - Forgie

Jim reports the system is not yet operational. There is still some development and testing to be done before the speech system can use the SATNET stream capability.

g. Internet Mail Specification Status - Postel

Jon said the status is "in editing" to be ready the end of February.

h. BBN Minigateway Status - Strazisar

Ginny says it will be ready by the end of February.

i. Equipment Need - Cerf

Vint again asks that needs for equipment in the category of LSI-11s, port expanders, etc. be reported to him by 5 p.m.

j. ECU Packet Size - Mathis

Jim says an ECU will handle a whole ARPA net message (8096 bits).

k. Formats for Gateway Reporting - Brescia

Mike distributed a memo on this topic. Discussion was deferred until a later scheduled session.

1. Multiplexing Protocol - Postel

Jon noted that a simple-minded multiplexing protocol is proposed in IEN-72. The protocol number is 18 (22 octal). There is an issue of checksum protection of the multiplexing header.

m. 576 Octets - Postel

Jon said the question of just what was the meaning of the 576 octet size chosen at the last meeting had come up. The clarification is: an internet datagram of 576 octets may be sent by any internet host with expectation that all internet hosts will be able to accept internet datagrams that large. It may be that such datagrams will be fragmented to transit certain networks and that destination hosts will have to reassemble the fragments. Every internet host should be prepared to assemble a set of fragments totaling 576 octets.

It was also noted that every internet module should be prepared to pass a 58 octet datagram with out fragmentation. The maximum internet header is 60 octets, and the minimum data fragment is 8 octets.

SRI GENERAL STORE - Mathis

Jim reported on some hardware/software modules that SRI creates and can make available to other ARPA contractors under various arrangements.

- a. Port Expanders
 - . LSI-11
 - . 32K Memory
 - . 4 line Terminal Interface (DL115)
 - . Robustness Card
 - . DMA Controller + 5 1822 channels
- b. DMA 1822 Interface
- c. Interrupt driven 1822 Interface
- d. Robustness Card
 - . XNET Software
 - . Auto Loader
 - . Switch settable address

Internet Meeting Notes

PERIODIC STATUS REPORTS - Mathis

Jim reported that SRI would coordinate the periodic status reports to be sent in on Networks, Gateways, IN modules, and TCPs. Jim suggests the following formats:

a. Internet Program Status

- . XNET servers
 - ELF debugger
 - MOS debugger
 - boot strap
 - loader
- . XNET user (Tenex/Tops 20)
- . Internet Name Server
- . Measurement Center
- . Printer Spooler
- . Other

b. Gateway Status Reports

- . What is it?
 - 11/40 ELF gateway
 - Mini gateway
- . Where is it?
 - NET 10, HOST 363<->NET 2, HOST 11
- . What does it do?
 - Header Version Number
 - Time-to-Live
 - Header Checksum
 - GMCC Reporting
 - Fragmentation
 - Alternate Routing
 - Error Reporting
 - Address Mapping

c. TCP/IN Status Reporting

- . Where is it?
 - Net and Host address
- . What is it?
 - Port numbers of servers
- . What is left out?
 - No Checksum
 - No Urgent
 - No Fragmentation
- . When is it there?
 - By Appointment Only
 - All the time

Internet Meeting Notes

DOD OPERATIONAL NEED FOR STANDARDS - Cerf

Vint repeated the need that DOD has for a standard for the growing activity in networks within the DOD. Vint again emphasized the potential model our internet protocols provide.

INTERNET MAIL - Postel

Jon reviewed the proposed new internet mail system. The system is based on communication of a structured data object (the mail item) between mail routing programs. The system will not be based on the existing ARPANET Mail or FTP. The mail items will be built up of a small set of basic data elements such as integer, character string, etc., and a list type. In general, a transaction between mail routing programs will be a list of three things: first, an identifier; second, some control information; and third, the mail itself. The routing information will come from the control information. The Specification is to be available by the end of February.

AUTODIN II FTP - Wingfield

Mike discussed the FTP being designed for use in AUTODIN II. The basic model has 3 parties: the controller, the donor, and the recipient. This allows a person at one host to initiate file transfers between two other hosts. A key design issue (which is not really unique to FTP) is access control. The proposed FTP provides for (does not require) a very general access control scheme. The access control scheme will allow use of alternative techniques, such as normal log in, trusted FTPs or third party access controllers. There will be a foreground and a background mode for file transfers. Current activity focuses on developing scenarios for file transfer and specifying the protocol operations which make them happen. The specification is to be available at the beginning of March. DCA is sponsoring this work and would like it to become a DOD standard.

WIDE BAND COMMUNICATIONS - Cohen

Danny described very briefly the configuration of the Wide Band experiment. This is a 3 MB/S satellite channel with ground stations at LL, DCEC, SRI, and ISI. Danny then described the local distribution issues and potential uses of the system. The main thrust of the Wide Band experiment is to experiment with many voice users in simultaneous conversations, but other experiments which can use the available bandwidth are anticipated.

Internet Meeting Notes

STREAM SET UP - Hoversten

Estil discussed stream set up and group addressing. Both features are provided by SATNET. Stream set up is used to provide small variance in delay when use is periodic. The cost is the reservation of some resource.

CONFERENCING - Forgie

Jim talked about voice conferencing. A voice conference may make use of both stream set up and group addressing. There are two models of conferencing: central and distributed. There are really two levels of the conference: control and data. The control of the conference uses a central model; the data of the conference uses a distributed model. The problem of moving between a line-network (e.g., ARPANET) and a broadcast-network (e.g., SATNET) was discussed. One suggestion is to have a surrogate or representative in the conference on one network act on behalf of all the participants on the other network. The issue was raised about the necessity (or non-necessity) of the representative to be in the same physical host as the gateway between the networks. Two hard problems to be discussed in the future are internet multidestination delivery and internet stream set up. There seems to be a problem if two line-networks are connected to a broadcast network and are part of the same group of stream users. Under some cases their messages will conflict in the stream broadcast. This topic is to be discussed at future meetings.

FAULT ISOLATION - Brescia

There was a discussion of how to identify faults in the internet. The focus was on how to forward error reports or responses through the internet. Potentially, information about congestion or unreachable destinations could be sent via internet options or via a distinct protocol. It could be sent to gateways or hosts. There could be special datagram types for probing gateways for connectivity. One might also desire to trace the route of a datagram. A small working group will meet to discuss this topic.

ADDRESSING ISSUES - Cohen

Danny indicated he was not prepared to say much more about addressing than was discussed earlier during the conferencing session. There was an outcry that source routing was an important topic and where were the details. Danny promised to have something definite at the next meeting. Vint did say he expected all gateways to process the source routing options. There is some issue about how a TCP deals with source routing.

Internet Meeting Notes

MULTI-MEDIA ISSUES - Saunders

Steve briefly described a small multi-media terminal being designed at ISI. The "reality" of the terminal forces some decisions about multi-media protocol interactions. The protocol selected has impact on the dynamics of terminal use and on the memory and processing power needed in the terminal. Looking into synchronization between media, one finds that at the expected interaction rates speech and maps take about the same data rate except that speech comes as a regular flow of small packet, where maps come as infrequent bursts of large packets. The possibilities for synchronization range from no system provided synchronization to complete synchronization. These extremes are probably too restrictive, and a middle course should be taken. One middle course could be synchronized only when a terminal user has both the push to talk and the stylus engaged at the same time. In voice conferences, we coordinate by passing the floor. Drawing conferences may coordinate by passing the chalk. If copies for the record are kept, the synchronization information must be recorded too, so a play back will match what actually happened.

There are two overall models that could be used: conversation and mail. There was some discussion of existing examples of multi-media terminals or usage; but for each example suggested, someone else thought the coordination was done outside the system.

GATEWAY REPORTING - Brescia

Mike reviewed the note the distributed earlier. There was no detailed discussion of the note.

PROTOCOL STATUS REVIEW - Postel

Jon presented the current internet and TCP header formats. The specifications are being edited and will be published soon.

NEXT MEETING

The next meeting will be held 8-11 May at BBN.

- May 8 - Plenary - Status and Presentations
- May 9 - Working Groups and Reading
- May 10 - Plenary - Wrap-up
- May 11 - Giant Multinet Demo

This demonstration is to include ARPANET, SATNET, BOSTON PRNET, MIT LCSNET, SRI PRNET, and possibly EDN and EPSS. A point of contact for coordination will be established at BBN.

Meeting after that is scheduled for 10-13 September at UCL (London).

AGENDA FOR NEXT MEETING

Please note that the notation "(IEN)" associated with an item below indicates that an Internet Experiment Note is to be prepared on the topic and circulated at least one month before the next meeting.

1. Operational SATNET TRANSITION PLAN - Binder (IEN)
inputs from Strazisar and Mathis
2. Expanded Routing Strategy - Perlman
Presentation (IEN)
3. Source Routing - Cohen
Presentation and Discussion (IEN)
4. Stream/Conferencing Control, Set-Up - Forgie, Hoversten, Cohen
Presentation and Discussion (IEN)
5. Internet message Forwarding Specification - Postel
Presentation (IEN)
6. Fault Isolation - Strazisar
Presentation and Discussion (IEN)
7. TCP Bake-off Results - Postel
3. State of the Internet Report - Mathis
Presentation (IEN)
9. Equipment Report - BBN, SRI, UCL
10. Internet Usage Experience - Kirstein co-ordinates
11. DOD Standardization Efforts - McFarland
12. AUTODIN II + NIFTP Discussion - Wingfield, Bennett
13. Discussion of version 4 Internet Protocol Spec - Postel
14. MULTI-NET Demonstration
15. Status Reports

MEMOS DISTRIBUTED

1. Meeting Agenda - Postel
2. Format for Gateway Polling and Reporting - Brescia
3. Issues and Plans in Internetworking - Mills
4. TCP for RSX-11m - Low
5. Mapping Internet Addresses to "Real Net" Addresses on the Exploratory Data Network (EDN) - Cain
6. Internet Experimental Note (IEN) Index - Postel
7. IEN-71 "User Datagram Protocol" - Reed & Postel
8. IEN-72 "Multiplexing Protocol" - Postel
9. IEN-73 "A TCP Implementation Issue" - Abranovitz, Padlipsky, and Biba
10. IEN-74 "Sequence Number Arithmetic" - Plummer
11. IEN-75 "The TCP Reset Mechanism" - Plummer
12. "Issues in Packet-Network Interconnection" - Cerf & Kirstein
13. List of Attendees

Internet Meeting Notes

ATTENDEES

Mike Brescia	BBN	BRESCIA@BBNE
Jack Haverty	BBN	JHAVERTY@BBND
Dale McNeill	BBN	DMCNEILL@BBNE
Radia Perlman	BBN	PERLMAN@BBNA
William W. Plummer	BBN	PLUMMER@BBN
Virginia Strazisar	BBN	STRAZISAR@BBN
Ray Tomlinson	BBN	TOMLINSON@BBN
Mike Wingfield	BBN	WINGFIELD@BBND
David L. Mills	COMSAT	MILLS@ISIE
Vint Cerf	DARPA	CERF@ISIA
Ed Cain	DCA	CAIN@EDN-UNIX
Robert Lyons	DCA	LYONS@ISIA
Ray McFarland	DOD	MCFARLAND@ISIA
Norman Abramovitz	FORD	ABRAMOVITZ@SRI-KL
Ken Biba	FORD	BIBA@SRI-KL
Danny Cohen	ISI	COHEN@ISIB
Jon Postel	ISI	POSTEL@ISIB
Jim Forgie	LINCOLN LAB	FORGIE@BBN
Nick Abel	LINKABIT	ABEL@ISIE
Estil Hoversten	LINKABIT	HOVERSTEN@ISIA
Noel Chiappa	MIT	JNC@MIT-AI
David Clark	MIT	DCClark@MIT-AI
Anita Skelton	MITRE	skelton@BBN
Aage Stensby	NDRE	AAGE@SRI-KA
Carl Sunshine	RAND	CAS@RAND-UNIX
Brain Davies	RSRE	JTAYLOR@ISIA
Paul Masterman	RSRE	JTAYLOR@ISIA
Geoffrey Goodfellow	SRI	GEOFF@SRI-KA
Ron Kunzelman	SRI	KUNZELMAN@ISIE
Jim Mathis	SRI	MATHIS@SRI-KL
John Pickens	SRI	PICKENS@SRI-KL
Andy Poggio	SRI	POGGIO@SRI-KL
Bob Cringle	UCL	UKSAT@ISIE
Peter Kirstein	UCL	KIRSTEIN@ISIA
Alen Meyne	UCL	---
David Boggs	XEROX	BOGGS@PARC-MAXC
Yogen Dalal	XEROX	DALAL@PARC-MAXC
John Shooh	XEROX	SHOCH@PARC-MAXC
Larry Stewart	XEROX	LSTEWART@PARC-MAXC