Danny

IEN 63
Section 1.4.6

J. Postel
ISI
14 November 1978

Internet Meeting Notes - 30 & 31 October 1978

EDITORS REMARKS - Jon Postel

In this report, remarks are attributed to various people; in almost no case will this be their actual words, but rather the editor's reconstruction of the intent of their remarks. Throughout this report the term "IN" is used for "internet protocol," and the term "TCP" is used for "transmission control protocol."

OPENING REMARKS - Vint Cerf

Vint welcomed us to the meeting, and Jim Mathis told us about the local facilities.

Vint: A number of feasibility demos have been done. We need to show an operational capability. In June 1979, eighty users will be online via PRNET at Ft. Bragg. In April 1979, there will be a PRNET demo at Ft. Sill. In May-June 1979, UCL will be disconnected from the rest of the ARPANET and will depend on the Internet system. Thus, stress testing is needed now. TCP-4 and IN-4 are being discussed as standards for DOD wide use. Preliminary specifications are to be ready in December 1978, and final in April 1979. Thus, all this is extremely visible in DOD and its contractors. Today we should focus on the technical details that are not completely worked out and the opportunities for stress testing.

Col. Russell: Confirmed the visibility of this effort, and the realization in DOD that protocols and internetting are now very important. The DOD is now about to make a commitment in the computer-communication area. There is a window (in time) in which input is considered and decisions are taken.

Copies of the Agenda were distributed. A brief discussion of it resulted in no significant changes.

STATUS REPORTS

BBN - Ginny Strazisar

The SATNET-ARPANET gateway and the PRNET-ARPANET gateway are now operating. New gateway software release in November 1978. Vint wants to get a list of all equipment needs - e.g., LSI-11's, etc.

for gateways and port expanders. BBN will maintain all gateways. Bill Plummer's note IEN 57 reports TCP status at BBN. Also, some testing programs are now available: (1) traffic generator, (2) testing gateway. Is IN available directly without TCP? Some programs can be specially run, but random user programs can't because the "underlying mechanisms" need to be changed. The initial IN probably will not support fragmentation and type of service. There could be a Unix TCP-4 (J. Haverty) two weeks after the Jim Mathis version is available.

SRI - Jim Mathis

TCP-4 for the 11 is coming along slowly and will be ready approximately mid-November. DMA-1822 is coming along, ready in mid-December to early January for the first two copies. The interrupt version can be made in three weeks per copy. Scheduling impacts debugging since XNET uses IN headers, and the header changeover means changes in XNET. PRNET runs occasionally. Yet to do mobile tests with point-to-point labeler. Vint would like to have a big demo in April or May.

MIT - Dave Reed

Three LCSNET interface units are now operational and are being installed on pdp11's. Interface design is being repackaged. Some local net software in Unix is being tested. Planning to use port expander and gateway code developed by BBN and/or SRI. The demo port expander code is now available from Mathis. The production version will begin to be developed next week. Multics TCP-4 is being tested by Dave Clark. Will try the BBN Unix "C" version as a base. MIT has a "C" compiler for the ITS system. Gateway is waiting till the other stuff comes together and hope to use gateway code from BBN. Ginny said the BBN minigateway software package will be available in January.

UCL - Andrew Hinchley

Trying to get code from other people and have ordered equipment to use it. Currently testing the port expander, but having some problems. This has held up getting the TIU up. A new 11 system will be available that could run a copy of someone's Unix TCP. The RSRE network is to be interconnected at UCL to the internet system. Now have an X.25 hardware interrupt driver interface. Plan to use the DMA interface similar to the Mathis 1822-DMA interface. There is a X.25 interface from DEC for PDP11 for $6,000. This is a HDLC (level 2) only interface. As a side comment, Jim Mathis mentioned that he talked to a guy from Honeywell who claimed to have built a 1822-X.25 interface.

NDRE - Yngvar

NDRE supports the SATNET-ARPANET gateway and participates in SATNET speech testing. There have been line problems. TCP-4 development goes on, have been able to do some testing on NORD-10. Have had some problems with staffing.

LL - Jim Forgie

Internet Voice Conferencing design thinking. Approach is to interconnect an ARPANET conference and a SATNET conference rather than an internet conference. There has to be some memory in the gateway to convert a group address to a list of addresses. Hope to have an internet speech capability up by the end of the year.

CCA - Dave Low

Will have TCP on CCA's 11 a few months after the Mathis version becomes available.

Record & Relay Facility was described. Designed for use by speech people, but it is more general. The basic idea is to prearrange a transform to be on file at the record and relay facility (RRF). Then to send a message there to be recorded, have the transform performed, and the resulting message sent on.

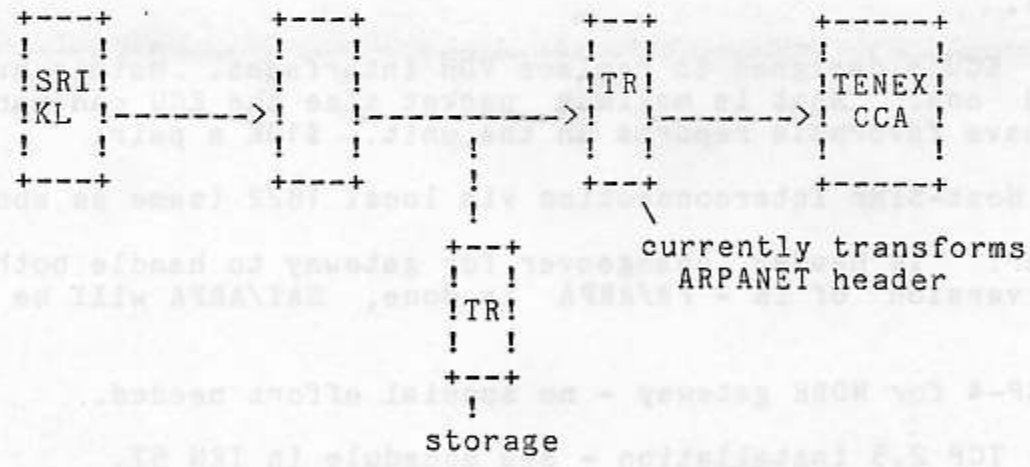

CCA has a "C" compiler under RSX running. A memo on the PSMF was distributed.

MITRE - Anita Skelton

MITRE is installing a local network called MITRE-BUS which is a cable bus that uses 5M-300M hertz frequency division (333KB digital portion), and can use some of the RF for video or etc.

The bus interface is microprocessor based. A connection between the MITRE-BUS and the ARPANET via PDP-11 Unix gateway is planned. The cable bus may start with just terminals and may later add hosts on the local net.

FORD - Ken Biba

FORD is building KSOS, a secure operating system for PDP11s based on Unix, and will include a TCP-4, now in detailed design phase. Multilevel security is important. The TCP will be written in C, the operating system will be written in either modula or pascal. The system will be running (single level security) in APRIL.

XEROX - Dave Boggs

Two PR units are interfaced to Alto-gateways to use PRNET as a bridge between two Ethernets. Network specific fragmentation is used on PARC internet packets to cut them into three pieces to get them thru the PRNET. About 25KB is best regular thruput one way thru PRNET (one hop) point to point. The bottleneck seems to be the packet processing time in the PRUs.

REVIEW OF ACTION ITEMS FROM LAST MEETING:

1. Postel: IN specification revision - done, see September version (IEN-54).

2. Mathis: ECU's designed to replace VDH interfaces. Mathis hasn't received one. What is maximum packet size the ECU can handle? Others have favorable reports on the unit. $10K a pair.

3. Binder: Host-SIMP Interconnection via local 1822 (same as above).

4. Strazisar: IN Header changeover for gateway to handle both old and new version of IN - PR/ARPA is done, SAT/ARPA will be done soon.

5. Cerf: TCP-4 for NDRE gateway - no special effort needed.

6. Plummer: TCP 2.5 installation - see schedule in IEN 57.

7. Plummer: TCP 4 installation - see schedule in IEN 57.

8. Cain: Provide EDN Unix to MIT - moot, since MIT has a Unix TCP now.

9. Cohen: TOS mapping - done, note distributed. Authorization for IN hosts to use type 3 messages in ARPANET may be necessary (ACTION: CERF). Forgie expresses concern that performance

problems will occur if so much use of type 3 is made such that IMP's actually throw away packets. It is suggested that we get information from Randy Cole of the ISI NSC project on type 3 data traffic performance statistics. (ACTION: COHEN).

10. Cerf: NCP to TCP translation strategies - a few service hosts would accept TCP connections then turn around and call on some other host via NCP, at least for terminal services (Telnet), on an interim basis. To avoid this, a special hack to get ARPANET stuff across Atlantic via SATNET can be instituted. Dick Binder explained how this could be done. The SIMP has a special interface for the "host" on the SIMP which is actually the IMP-IMP line of the UCL TIP. Then encapsulates the IMP-IMP traffic and delivers it to a similar interface at another SIMP. The idea is that the SATNET provides an IMP-IMP path between UCL-TIP and SDAC. This may be up by the end of the year.

11. Hoversten: SATNET stream setup - Stream setup, based on group names, a set of commands to create, add used to, delete users from, and destroy streams. A stream is really an allocation of satellite resources. A stream has parameters of interval, packet length, and a holding time. (The holding time is a measure of allowable variation in the interval.) The stream is designed to be used for structured conferences. Hosts in a stream conference are notified if hosts join or leave the stream, the stream is deleted, or the stream is suspended or resumed. Data may be in SATNET or IN format. Addresses may be Stream ID (conferencing), SATNET address (point-to-point), or SATNET group name (broadcast). Flow control is by SIMP accepting or refusing packets, and host can do the same thing. Some discussion about timing of response to commands (e.g., join) to be synchronized with interval availability times. Memo PSPWN 104 was distributed. A presentation on conferencing and high level protocol/network interactions is to be made at the next meeting by Jim Forgie (ACTION: FORGIE).

12. Cerf: Public Demo feasibility - it will not be attempted in 1979. Vint would still like to do a full scale demo in 1980.

13. Reed: Document checksum algorithms - done, see IEN 56.

INTERNET MAIL - Jon Postel

Jon described the the internet computer mail system being designed at ISI. The primary attributes are the use of a structured typed-data format for the exchange of mail items which allows for limited sharing of information between mail items. The mail items may be grouped into larger structures called mail bags. Mail bags are processed by programs called post offices which run as background or

daemon process on internet hosts. Post offices process mail bags and route mail items to other post offices or to local network mail systems. The specification is to be available in about three months.

WINDOW CONTROL/RETRANSMISSION - Steve Treadwell

Steve described the results from a set of window size vs retransmission time out measurements. When the retransmission time out is fixed at 30 seconds, thruput goes up as window goes up to until a cutoff where the window is larger than the capacity of the pipe. A rule of thumb from these experiments is that when delay goes up while thruput does not then reduce window.

INTERACTIONS WITH GATEWAYS - John Davidson

John described the results of a meeting held at BBN and documented in IEN 60 which was distributed.

1. Access Control - This is covered later on agenda.

2. Monitoring and Control - This is covered later on agenda. No general mechanism that had to be implemented, and monitoring center could fail without adverse effect on catenet.

3. Debugging - Gateway owners' responsibility but should be supported by internet conventions but need not allow random people to access this part of the gateway.

4. Statistics Gathering - up to gateway owner if more than needed in monitoring and control.

5. Stream Setup in a CATENET. What does it mean if the local network does not have any means to reserve resources.

6. Congestion Control, Flow Control, Minimal Monitoring

Some functions needed are source routing, echoing in gateway, echoing in a host on the "other side" of a gateway to bounce traffic thru a gateway.

MINIMAL GATEWAY MONITORING -Mike Brescia

Mike distributed "A Report Format for Gateways," then discussed some monitoring functions and the report formats. Vint suggested that measures be in terms of IN datagrams rather than "packets" and a measure of fragmentation should be included. The question "What can be inferred from the measures?" should be answered in the documentation of the measurement package.

Dave Boggs mentioned that PARC counts packets per from-to pair of networks in a matrix (including a "discard" column). Number of packets sent out each interface and number of bytes sent out interface. Gateways also support several service processes such as echo, daytime, name service, and statistics. The statistics service returns a message carrying the accumulated measurement data with prefix identifying gateway and up time and copies of the matrix, etc.

A small group is to consider the report format one last time, determine if reports should be sent based on a time out or based on polling. Also, consider increasing field sizes if statistics are to be sent at low frequency.

VDH & GATEWAY PERFORMANCE - Ginny Strazisar

Experiments were conducted and documented in two messages sent to SATNET group, which Ginny will forward to anyone who request them. Mainly changes to buffering in the VDH implementation to eliminate the retransmission of each message, but the results suggest that ELF is processor bound or could not respond to interrupts fast enough. This confirms the desirability of using the ECU's rather than VDH for interfaces. The following data on gateway performance was presented. Program is 12K leaving 14K for buffers.

	PACKET/SEC	SEC/PACKET	BYTES/MSG	DATA RATE
MOS on 11/40	345	2.9 ms	220	600 KB
ELF 11/45	8-12	?	?	?

INTERNET HEADER FORMAT - Jon Postel

Jon presented the current IN and TCP header formats. This led into a discussion of segment size.

Sizes of segments in various networks:

ARPANET	8000 bits or	1000 octets
PRNET	2000 bits or	250 octets
LCSNET	500000 bits or	62500 octets
SATNET	4000 bits or	500 octets

It was pointed out that the size should probably be 2**N+E and the +E is important. This would allow a 2**N block to be sent as a message if the header could fit into the E.

The two questions are "What number will not get fragmented?" and "What number will be reassembled?" Some calculation followed:

```
   INH     20        20
   TCPH    20        20
   Slop    24        24
           --        --
           64        64
   DATA  1024       512
         ----   -------
         1088   ! 576 ! Octets
                -------
```

A packet of 576 octets is to be handled by all IN users (hosts or gateways); that is, any host can send datagrams of up to that size with the assurance that any and all destinations are prepared to accept datagrams of up to that size even if reassembly is required. The 576 octets is the value in the IN header total length field.

MONITORING - Vint Cerf

(notes taken by Mike Brescia)

After looking at a proposed gateway reporting format distributed by Brescia, there was discussion on various points:

Vint suggested that the kinds of things we want to know are: changes in catenet connectivity, traffic from net to net, gateways up and down. We need tools for finding out why things don't work well. We should use an internet metric so measures are comparable.

Whether reports should be spontaneous (sent automatically and regularly) or on demand (polled by a monitoring center). Automatic reporting can handle new gateways as soon as they come up. Routing algorithm handles new gateways.

Perhaps the monitoring center can use the routing messages. Routing tables store information about reaching nets, not gateways. Monitoring should be expandable to the case when there are 100's of nets and 1000's of gateways, with many different monitoring points, and each may want to use its own reporting format. Polling gives the option of getting traffic information from a few gateways on a finer time scale during an experiment. Some nets and gateways may not want to give out traffic or error information, e.g., PTT's.

Cerf then spoke to the question "What is it we are trying to solve?" This is an internet experiment, not an attempt to set global policy on network interconnection for all time. Monitoring of gateways is not an essential component for passing internet traffic, but a tool - for recording connectivity and topology of the catenet and to warn when fault isolation procedures need to be started. Should not impose restriction or much more traffic to accomplish monitoring job.

In experimenting, traffic summaries are helpful, and performance problems should not be hidden. So, since there will be (at least one) monitoring center, gateways should provide - status information, traffic statistics between attached nets and between source and destination nets, tools for fault isolation, and a way for users to get this information.

Specific format issues - count fields should be large (32 bits) fixed format independent of type of reporting machine (e.g., pdp11's will have to swap bytes to report 16 bit integers). Traffic reports will be gotten by polling infrequently, status by frequent, short messages sent automatically.

Statistics should be accumulated by Source-DEST pair of NETS and by Source-DEST pair of Interfaces. Cumulative statistics need longer fields, 2's complement overflow. It is preferable to have a common format. A polling mechanism should be used for asking for data. Routing table should be retrievable via monitoring probe. There should be a echoing monitoring probe. There should be a discard monitoring probe. Yet another draft to be prepared by Mike Brescia and John Davidson (ACTION: BRESCIA).

Brescia to propose new formats for polling and reporting, and tools for fault detection and isolation (ACTION: BRESCIA).

ACCESS CONTROL - Radia Perlman

Radia distributed copies of IEN 58 and discussed access control in the catenet. Access control impacts more than security; it has routing effects. A list of basis for choosing a route (restrictive routing) might be: source and destination nets, type of service, message length, amount of traffic. Routing is a global problem. The no shared information approach is random routing, i.e., trial by error, is very poor. Thus, a global plan is needed. One suggestion is to break into routing into a per "routing group" basis, where a routing group is a set of restrictions. There are potentially a large number of routing groups. There are two types of algorithms: ARPANET style, complete link table.

Col. Russell brought up a point about protecting against use of ARPA sponsored nets as transit net between commercial nets.

This in turn raised a question about the granularity of access control? That is, is it net, host, or user? One suggestion for access checking was to provide a "seal of approval" on a packet that is a permission that gateways interpret, like a capability. It was noted that restrictions are likely to be not very dynamic. Two strategies can be used: 1) run into blocks, 2) know by routing algorithm ahead time that it will go a legal path all the way.

CONGESTION CONTROL - Ginny Strazisar

Ginny suggested that there are two approaches to congestion control: 1) use alternate routing (to avoid congestion), 2) source quenching. Alternate routing may alleviate the problems for awhile; but, ultimately, source quenching will be necessary. One way to do it is to send a message to source from gateway when packets are dropped, saying "please slow down." A question about the relationship between the packets discarded and the offending sources was raised.

Dave Boggs pointed out that in the Ethernet gateways send a message to the source saying, "I discarded a packet of yours," with the implication that if the source slows down it is less likely to loose.

Other questions: Can a gateway detect congestion before packets are lost? Can it detect queue build up before loss occurs and send a quench message?

A congested gateway can report the problem to neighbor gateways (or to the gateway the particular gateway the traffic came from) which then will repartition the division of traffic routing.

SOURCE ROUTING - Danny Cohen

Source Routing put off to next meeting.

TRANSNET FILE TRANSFER - Chris Bennett

Chris discussed an experiment to transmit files from a host on ARPANET in NCP world to a host on EPSS via the catenet, and distributed IEN 59. The experiment is based on the Network Independent File Transfer Protocol (NI-FTP) from EPSS (INWG Protocol Note 86). FTP support modules would reside in a gateway between host-to-host protocol worlds. The FTP support module converts connection establishment procedures, and relies on reliability and ordering of underlying host-to-host protocols.

If not everyone uses NI-FTP, then one can stage the transfers by having a local net FTP to a host that does NI-FTP to remote network then that NI-FTP can turn it over the destination local net FTP. Mapping between local and NI-FTP's is about like mapping between an operating systems file system and a network FTP.

Transnet FTP will have to do two things:

(1) open a connection to a FTP service at host x in a local address (which may be a FTP support module in a gateway).

(2) send the first message to say, "I hope I am talking to host

x." The response by host x is "ok"; the response by a gateway FTP support module is to then open a connection to that host (i.e. host x) possibly via another gateway and forward the message... recurse.

Implementation is in progress on ISIE.

NAME SERVER - Jon Postel

Jon distributed IEN 61 and discussed it. The name server can operate on the IN datagram level. Most of the discussion was spent on the possible elaborate kinds of data a name server could keep and hand out.

MULTIPLEXING - Danny Cohen

Danny's message is that messages cost more than bits. Combining messages routed the same direction at the time they are routed is like factoring.

(HD)+(HD)+(HD)
(H(D)(D)(D))

The problem with this seems to be the fragmentation id. In fact, fragmentation is based on the fields: ID, PROTOCOL, and pair of ADDRESSES.

The selection of the ID should be unique across all uses of a protocol, but this could be done by the IN module. But in TCP if we want to let identical retransmission add to the probability of successful reassembly, then the TCP module must set the ID field value.

It is suggested that multiplexing may be of interest in the context of the wide band communication experiment. The payoff of multiplexing teriminal traffic (i.e., telnet data on TCP connections) is suggested and questioned. The idea of having TOS rather than (or in addition to) Protocol as a basis for limiting combination is suggested. It is pointed out that this type of combining of traffic is similar to "blocking" in IBMs SNA. It is suggested that there be a value of Protocol that means multiplexing, then there could be a multiplexing header applied to each chunk in the datagram.

ENRICHED ARPANET ADDRESSING - John Davidson

John distributed IEN 62 and discussed some of the issues. The ARPANET currently provides for one address per physical interface (i.e., host). The internet calls for the possibility of multiple internet addresses per internet host. Multiple addresses provide for

the following cases: multiple TCP's on some host, port expander, subnets as "hosts."

If 24 bit REST of IN address is used directly for ARPANET host/IMP address, then we can't do these things. One idea is to use a few of the bits to name pseudo hosts or logical hosts.

Various proposals were discussed and the following one selected:

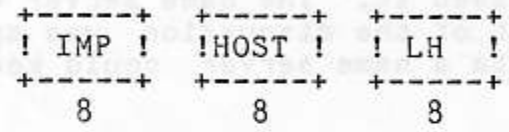

We need to document the mapping of internet protocol to each real network protocol. The following assignments are made to carry out this task:

D. Reed	LCSNET
J. Davidson	ARPANET
V. Strazisar	BBN RCCNET
J. Mathis	PRNET
D. Binder	SATNET
E. Cain	EDN

DATAGRAM PROTOCOL - Danny Cohen

There was discussion of the need for a Datagram protocol on top of IN to multiplex IN datagrams to various datagram application processes, such as the name server process.

Dave Reed suggested the following and was then asked to prepare a memo on it (ACTION: REED):

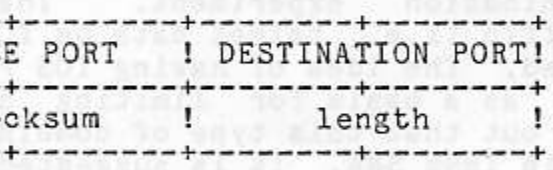

NCP-TCP TRANSITION - Vint Cerf

The immediate problem is the hosts in Norway and London which will be come isolated when the Atlantic line goes away in a few months. There are two things to be done to mitigate this: one is to use the SATNET kludge, and the other is to provide one or more "relay hosts" which implement both NCP and TCP.

TENTATIVE AGENDA - Vint Cerf

Status Reports:

Each of BBN, SRI, MIT, UCL, NDRE, LL, UCLA, CCA, PARC, MITRE, FACC, COMSAT should report on the status of any Network, Gateway, IN or TCP implementation they are responsible for.

Action Item reports:

1. Cerf & Postel - TCP and IN specification status.

2. Internet to real net mapping notes:

   Davidson - ARPANET
   Strazisar - RCCNET
   Mathis - PRNET
   Binder - SATNET
   Reed - LSCNET
   Cain - EDN

3. Access Control small group meeting - Nov. 21 - ARPA

   Perlman, McFarland, Cerf, Cain, Dlugos, Davidson, Postel

4. Datagram Protocol - D. Reed

5. Gateway Minimal Reporting - Davidson, Brescia

6. Internet Name Server - Postel

Other Agenda Items:

-Internet/SATNET Final Operational Configuration - Cerf

-Internet Mail Server - Postel

-Fault Isolation Facilities in the Catenet - Davidson

-Stream Setup/Conferencing - Hoversten/Cohen/Binder

-LL Method of Conferencing and Alternatives - Forgie

-DOD operational needs for standards - Cerf

-AUTODIN FTP - Wingfield

NEXT MEETING - Vint Cerf

Next TCP meeting is scheduled at ARPA on December 4 and 5.

Next IN meeting is tentatively scheduled at ISI on January 25 and 26.

MEMOS DISTRIBUTED

Agenda - Postel

IEN 58 "Access Control - An Informal Discussion" - Perlman

IEN 59 "The UCL Transnet File Transfer Implementation" - Bennett

IEN 60 "Boston Area Meeting of the Internet Working Group to Discuss Interactions with Gateways" - Davidson

IEN 61 "Internet Name Server" - Postel

IEN 62 "Enriched Internet Addressing of ARPANET Resources" - Davidson

RFC 750 "Assigned Numbers" - Postel

PSPWN 104 "Host/SATNET Stream Access Protocol" - Binder

------ "PSMF Record and Relay Facility" - Low

------ "A Report Format for Gateways" - Brescia

------ "TOS Mapping for ARPANET" - Postel

ANSI X3533 "Functions" - Cerf

IEN Index - Postel

Notebook Table of Contents - Postel

ATTENDEES

Vint Cerf	ARPA	CERF@ISIA
David E. Russell	ARPA	RUSSELL@ISIA
Richard Binder	BBN	BINDER@BBNE
Mike Brescia	BBN	BRESCIA@BBNE
John Davidson	BBN	DAVIDSON@BBNE
Radia Perlman	BBN	PERLMAN@BBN
Virginia Strazisar	BBN	STRAZISAR@BBN
David Low	CCA	LOW@CCA
David L. Mills	COMSAT	MILLS@ISIE
Ed Cain	DCA	DCEC-R850@BBNB
Ray McFarland	DOD	MCFARLAND@ISIA
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 Reed	MIT	DPR@MIT-ML
Anita Skelton	MITRE	MITRE@BBN
Yngvar Lundh	NDRE	Yngvar@SRI-KA
Paal Spilling	NDRE	PAAL@SRI-KA
Ken Biba	FORD-AERO	BIBA@SRI-KL
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
Chris Bennett	UCL	UKSAT@ISIE
Andrew Hinchley	UCL	UKSAT@ISIE
Steve Treadwell	UCL	UKSAT@ISIE
David Boggs	XEROX	BOGGS@PARC-MAXC
Yogen Dalal	XEROX	DALAL@PARC-MAXC
John Shoch	XEROX	SHOCH@PARC-MAXC