J. Postel - ISI
21 August 1978

Section 1.4.4. Meeting Notes - 2,3 & 4 August 1978

REVISED AGENDA

- Arrangements - Forgie
- Introduction and Objectives - Cerf
- Working Groups
  - A: Symbolic Addressing - Postel
  - B: Error Handling - Strazisar
  - C: Type of Service - Cohen
  - D: Access Controls - Cerf
  - E: Demonstrations - Kirstein
- Status Reports
  - (i) BBN
    - Gateway Status
    - Host/SIMP Protocol Status
    - TOPS-20 and TENEX Internet & TCP Status
    - Gateway Monitoring Status
    - UNIX Internet & TCP Status
  - (ii) MIT
    - Multics Internet & TCP Status
    - LCSNET Status
    - Gateway Status
  - (iii) PARC
    - ETHERNET/PRNET Gateway Status
  - (iv) SRI
    - PR Network Status
    - Internet & TCP Status
  - (v) UCL
    - Internet & TCP Status
  - (vi) NDRE
    - NORD-10 Internet & TCP Status
  - (vii) CCA
    - RSX-11M Internet & TCP Status
- Working Groups Reports
- Checksums - Hinchley
- Retransmissions - Hinchley
- ARPANET Changes Planned - Cerf
- Multiplexing & Multiaddressing - Cohen
- Review of the Internet Protocol Header - Postel
- Agenda and Date for Next Meeting - Cerf

ARRANGEMENTS - Jim Forgie

Jim welcomed the group to Lincoln and told where various meeting related facilities were located.

INTRODUCTION AND OBJECTIVES - Vint Cerf

Vint put the stress on the need for the Internet to be a working system very soon. To emphasize this the following schedule of events which rely on working TCP and Internet Protocol implementations was presented.

[Editors note: from here on out IN stands for Internet Protocol.]

Milestone Dates

Jan 79
TCP & IN version 4s operational with Telnet for Tenex, Tops-20, 360, and Multics.
Apr 79
FTP operational, NCP - Internet "type x hack" installed on ARPANET, (NORSAR-TIP and LONDON-TIP disconnected from ARPANET, reconnected to SATNET), SATNET operational [64KB Atlantic], INTERNET mail system operational.
Jun 79
20 terminals, 8 PRU's, and a station/gateway at Ft. Bragg, plus a 2060 at ISI.
Dec 79
80 terminals, 27 PRU's, and a station/gateway at Ft. Bragg
Sometime in 79
NSW will use internet protocol

Vint also presented the following issues for the group to consider:

1. Operational Demand for Internetting
2. Gateway Performance
   -thruput
      existing systems too slow
   -congestion control
      gateway-to-gateway
      host-to-gateway
   -alternate routing
3. Symbolic Internet Addressing
   -FTP
   -Telnet
   -Internet Mail

STATUS REPORTS

i) BBN - Strazisar, Plummer, Davidson

Ginny reported that 3 gateways are up between SATNET & ARPANET, and an additional gateway at COMSAT is in progress. The host side of host/SIMP is done (about 4k of code - including buffers 2k buffers, 2k code), waiting for SIMP side (to be ready 1 October).

Vint wonders why code is so big? Some discussion of code size of various modules follows.

Concern about poor performance of gateways focus on problems with VDH thruput. The use of 4 channel vs 2 channel is something to investigate.

Some discussion about version of software and version of TCP and IN to be used in what tests. In general the latest versions are to be available as soon as possible.

Concern expressed about the impact of maintaining a system capable of doing demonstrations on the implementation of new versions of the protocols. Vint promises no demo's of INTERNET/PRNET are scheduled thru the end of the year.

Things that now use old versions of IN or TCP protocols are:

Gateway Monitoring, XNET, LSI-11, Bootstrap, Speech, SIMP fake hosts eg. XPAC, SRI-Loader, BCR-Loader, PTIP-gateways, GNOME.

Vint asks if gateways could be programmed to handle both new and old headers? Goal is to have gateway handle new headers by 1 Sep.

Every one should tell Ginny what software will be affected by the changeover to new internet headers and then all information will be circulated to the INTERNET and SATNET groups.

This was done during a break resulting in the following:
Things:

1) UCL Gnome (Kirstein)
2) User Programs, SIMPs, EXPAK (Binder)
3) SATNET Gateway Monitoring (Cole)
4) XNET, Bootstraps (Tomlinson, Burchfiel, Kirstein)
5) BCR Software (Bressler)
6) BBN Line Printer (Burchfiel)
7) TCP on ELF 11/40 for PRNET (Tomlinson, Burchfiel)

8) LL Speech (Forgie)
9) LSI-11 (Mathis)
10) PTIP (Burchfiel)
11) TCP (Burchfiel)
12) Tenex & Tops20 assign special queue changes (Burchfiel)
13) Gateways

Relationships:
1 <- 2,4
2 <- 12
3 <- 12
4 <- 10,12
5 <- 9,10,11,12
6 <- 10
7 <- 9
8 <- 13
9 <- 13
10 <- 13
11 <- 13
12 <- 13
13 = the gateways

The names in parenthesis are people to contact about scheduling.

Schedule for changeover is that on 1 Sept the new stuff becomes available, and on 1 Nov the old stuff goes away.

Some discussion about an alternative to the VDH interface took place, with the main point being the existence of a device made by ACC that looks to the host and the IMP like a local 1822 interface, but uses HDLC internally for the link protocol.

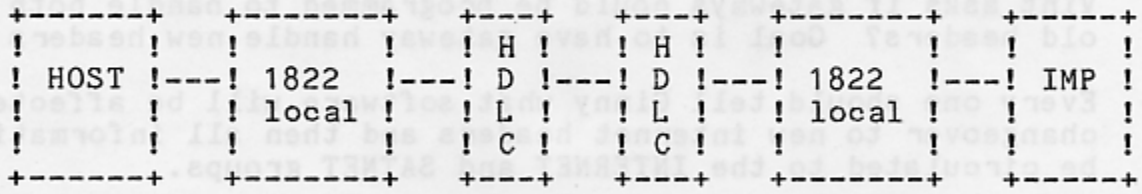

Bill Plummmer reported that TCP 2.5 is up at BBNC, BBND, and SRI-KA. The main task now is moving to TOPS-20 v3 and the model B processor.

At this point Mike Brescia distributed a Monitoring Information memo. It was decided to have a gateway monitoring information working group on Friday morning.

John Davidson reported on BBNs UNIX TCP. It is now running TCP 2.5, and should be running IN-4, and TCP-4 by 1 Nov, really aiming for 1 Oct.

BBN's EDN work may use the DTI version of TCP which is in C. This project, led by Wingfield, expects to have a C version for EDN by 1 Jan 79. This version will not be the same as the one used for the ARPA Internet project.

ii) MIT - Dave Reed

Dave reported that a Multics implementation of TCP version 3.1 was nearly completed, and now work is in progress on version 4. User side is straightforward, but there are policy problems to install server. MIT is also trying to get a UNIX TCP, an ITS TCP, and a TOPS 2040 TCP. On the LCS NET progress is being made. The 3rd interface has been ordered, and testing is now underway between two machines. A key problem at MIT is a shortage of IMP ports; in fact, they currently have two more hosts than ports.

iii) XEROX-PARC - John Shoch

John reported on PARCs experiment using the PRNET as transit net between two ETHERNETs. PARC now has about 22-25 net's (lost track of numbers). PARC is also doing a packet speech experiment using BYTE STREAM connection - up to 500 KB - so unencoded speech is sent.

iv) SRI - Jim Mathis

Jim reported that he has not really done much about converting to the new version of TCP and IN. He is waiting to see if version 4 turns into version 4.1! On the Port Expander idea things are progressing slowly also, the current thing works with the 32 bit 1822 leaders.

v) UCL - Andrew Hinchley

Andrew reported that the FTP standard from the EPSS group will be brought up on a Tenex so that expermients with end-to-end FTP between EPSS hosts and an ARPANET host can be performed, and possible expermients with hosts in other in X.25 networks.

vi) NDRE - Yngvar Lund

Yngvar discussed the state of development at NDRE. A TCP-3 is near completion for the NORD-10 and is nearly ready for testing. The need for a TCP on ELF was pointed out.

vii) CCA - Kou Mei Chuang

Kou Mei reported on the progress at CCA. Currently they are converting a TCP-11 supplied by Jim Mathis to run under RSX-11. This is version 2.5. When Jim can supply a version 4, they will cover that.

It is clear that Jim Mathis is the leader on TCP-4 & IN-4 for pdp-11s. Everyone that is waiting for his program to be available should contact him.

viii) LL - Jim Forgie

Jim Forgie remarked that speech conferencing had been demonstrated in SATNET with SIMP-1. SIMP-3 is up now, and speech is unusable due to new delay problems. There is a need to do internet speech tests. ISI is working on a new vocoder format on their FPS.

For some reason Danny Cohen asked "How do you talk to a gateway for stream setup, access control, X.25 setup?" There was a suggestion to handle the question in a small working group.

WORKING GROUP REPORTS

SYMBOLIC ADDRESSING - Jon Postel

A working group on symbolic addressing chaired by Jon Postel met on Wednesday morning. The issue seems to be "What syntax do we (the Internet Working Group) recommend for names of internet destinations as typed by users?".

An earlier suggestion was something like Net!Host/IMP/...

Some problems are:
(a) should protocol be identified in this, e.g. /TCP/
(b) how to deal with numbers?
(c) how to deal with process id?
(d) how to deal with port numbers?

Can't have a parser for every network!

User Syntax for Names

The working group recommendation is that names be character strings of the form:

! NET ! REST

Where max length for strings is 128 or 127 or 64 or 63 ...

The reason for being non-specific about the "REST" is that networks may have very different internal structures. For example, the ARPANET concepts of host and IMP may not apply. In general, the syntax of "REST" will follow the structure, and since we don't know the structure, we can't specify the syntax. The "! NET !" form was chosen as it restricts the "REST" less than the other proposed forms. The restriction on "REST" is that it not start with "!".

Name Server

Along with this, the working group recommends that for each network there be a name server process. The set of name server processes must be found at well known addresses. This name server is a simple minded process, not a general information service.

The basic operation is to send it a string, and get back the necessary address bits.

Each Host has a table of network/address pairs where the address points to a Name Server that does the name to address conversion for that network.

The name server would have to report errors for unmatchable names in a useful way. It seems reasonable to expect most hosts to maintain a local cache of recently used names to avoid repeated calls on the name server. Updating the name server could be a problem. Some felt automated update systems could be developed. Postel argued to have a person in the loop for data base reliability reasons. It seems quite possible to implement the name server at the IN datagram level (i.e. not using TCP).

No matter how well a name to address translation scheme is worked out, there must be a provision for a user to supply all the bits if the user wants to. The suggestion is that if a

user wants to enter a address directly a string of digits starting with the sharp sign be used, e.g. #1234...

Jon promised to prepare a memo specifying the Name Server and the syntax of names (action: Postel).

ERROR HANDLING - Virginia Strazisar

Ginny led a working group on error handling in the internet on Wednesday morning.

The kinds of errors that might have to be reported are:

- destination host dead
- destination host unreachable
- destination network unreachable
- can't meet type of service requested

Information to be reported in Error Messages:

- source in original msg triggering error
  INTERNET Address
  Protocol Identifier
  Error Code (type of error)
  Time stamp (when error detected)

Other information that would be nice to have reported:

- when service break reported
- when service will be restored
- what TOS is now available

Error reports should be generated for each datagram causing error. Use only information in IN header to compose error message. Must identify segment causing error uniquely.

Postel suggests simply reporting first N bits of packet in error, where N is large enough to cover all of the internet header.

TYPE OF SERVICE - Danny Cohen

Danny Cohen led a working group (of the whole) on Wednesday afternoon. We first characterized the service parameters in each of ARPANET, SATNET, and PRNET.

ARPANET

-Priority: (1 bit)
-Uncontrolled: (Subtype 3 vs Type 0)
(fast vs reliable)
-Packets: Single vs Multipacket (in Type 0)

1000 bits	8000 bits
low delay	high delay
low rate	high rate
Telnet	FTP

SATNET

-Type: block, stream (pre-allocated slots) 2 bits
-Priority: 2 bits
-Delay: 2 bits
-Holding time: 1 bit
-Reliability: 1 bit

PRNET

-Routing: Station vs. Point-to-Point
-Reliability: Ack vs. no Ack
(Use of Transparent vs. SPP protocol is a host-to-host protocol level issue.)

We then tried to come up with a characterization of service types for the Internet. It is apparent that each thing a user could ask for can be placed on a thing-cost trade off, e.g. speed-cost, or reliability-cost. We also see a speed-reliability trade off. We also talked about sequence or sorted delivery, but found no useful way of doing it. The set of parameters we came up with are:

INTERNET

Speed
-as fast as possible
-fast
-normal
-take your time
Reliability
-as reliable as possible
-reliable
-normal
-don't try too hard
Reliablility-Speed trade off
-which to try for if in conflict
Priority
-2 bits
Packaging
Stream vs datagram

The time to live field was discussed and it was decided to keep it as it is, a separate field. We also remarked on the trade off between low level protocol relibility and higher level (host-to-host) protocol reliability and the effects of that choice on effective user-to-user delay.

We then made up a chart of the type of service possiblities in the ARPANET, the SATNET, the PRNET, and the INTERNET for each of several applications.

Application	INTERNET	ARPANET	PRNET	SATNET
TELNET on TCP	P:stream S:fast R:normal P:speed	T: 3 S: S	R: ptp A: no	T: block D: min H: inf R: no
FTP on TCP	P:stream S:normal R:normal P:reliable	T: 0 S: M	R: ptp A: no	T: block D: normal H: inf R: no
interactive narrow band speech	P:stream* S:asap R:least P:speed	T: 3 S: S	R: ptp A: no	T: stream D: min H: short R: no
datagram	P:datagram S:fast R:normal P:speed	T: 3 or 0 S: S or M	R:station A: no	T: block D: min H: short R: no

key:	P=package S=speed R=relibility P=preference *=requires stream set up	T=type S=size	R=route A=ack	T=type D=delay H=holding time R=relibility

Danny then proposed the following layout for the Type of Service field.

PRIORITY	STRM	RELIABILITY	S/R	SPEED
11 highest	1-STREAM	11-highest	1-speed	11-highest
10 higher	0-DTGRM	10-higher	0-relib	10-higher
01 lower		01-lower		01-lower
00 lowest		00-lowest		00-lowest

One issue left open is the procedure for stream set up. SATNET

group is to consider this issue and to make recommendation at next meeting (action: Estil).

ACCESS CONTROL - Vint Cerf

Vint led a short working group (of the whole) on access control on Wednesday afternoon.

Access Control Buzzwords:
1. Policy
2. Implementation/Enforcement
3. control of entry/exit of traffic

Not all networks need to provide access control. Access controller host per network that referees access control requests from the gateway. A gateway can build up a local table. Some issues are:

-How long do tables stay set up?
-Can someone miss lead the gateway? Authentication?
-Efficiency of access control?
-Impact of use of access control in internet routing?

DEMONSTRATIONS - Peter Kirstein

Peter moderated a discussion of demonstrations on the Internet during which the following points were raised.

There was much discussion of putting together an Internet demo at the ICC conference to be held in spring 79 in Boston. Vint will determine if the necessary resources can be made available (action: Cerf).

Need TRIALS to stress the system in preparation for providing service, more than demonstrations.

TCP test September 18 & 19

1 September Gateway have NEW (+OLD) IN format
15 September LSI-11 TCP4 available from Mathis
30 September Ginny will have TCP4 in?

CHECKSUM - Andrew Hinchley

Andrew noted that there is current discussion of checksums as evidenced by IEN 45 by Plummer and IEN 49 by Higginson. The group then discussed the desirability of a strong checksum vs the cost of computing one. The two options considered were a one complement sum technique and a CRC technique. The group seemed to be pretty much set on using a CRC. A paper by Kirstein & Higginson describes a way of computing the CRC. The Internet checksum is only on the IN header (ten 16 bit words). Dave Reed is to supply listings of the 10 and 11 code to do the computation (action: Reed). Jon Postel is to put in the right words in the IN (and TCP) specifications (action: Postel).

[Editors note: Subsequent discussion via SNDMSG suggest that the computational costs of a CRC checksum are too high. So the version 4 IN and TCP will use the ones complement scheme.]

RETRANSMISSION RESULTS - Andrew Hinchley

Andrew discussed some of the findings of a study he and Steve Edge have done on the effects of various retransmission techniques as reported in IEN 50.

ARPANET CHANGES - Vint Cerf

Vint reviewed some plans for changes in the ARPANET under a project currently underway at BBN and supervised by John McQuillian. These changes might start to show up in Jan-Mar 79 time frame.

-Study of congestion type glitches:

"Waves of congestion that wash back and forth across the network."

-Routing:

Basic algorithm estimates delay (in an arbitrary way). Then makes routing changes; such decisions are made concurrently throughout the net.

Flaw is that the new routes are not at all related to the previous estimate of delay.

Shortest path first algorithm

-Flooding of routing information.
-Each node knows full net topology.

-Line Up/Down:

The current hard-nosed algorithms don't detect poor lines, only dead ones; and have a strong test for bringing a dead line up again.

Circulating token between IMPs. The token carries the count of the hellos so the two sides stay synchronized.

-Stability:

Control theory study in progress.

-Other Ideas

Broadcast
Multi Destination
Logical Host Addressing
-requires a leader change
-needs both source and destination addresses
-multihoming
Internet routing/delay/info.

-NCP internetting

Type X hack

-allow IMP to tell Gateway the destination as well as source across the 1822 interface

-also wanted a way to push RFNM's across the intermediate network

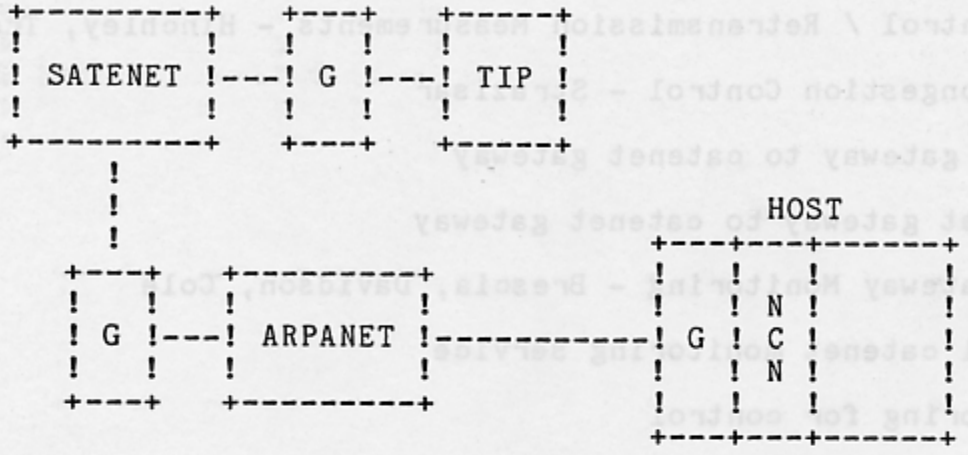

TIP is an isolated piece of the ARPANET.

MULTIPLEXING & MULTIADDRESSING - Danny Cohen

Danny gave a quick presentation of the role of multiplexing and multiaddressing in computer protocols, and discussed some transmission efficiencies that could be obtained by careful attention to these concepts. These remarks are presented in IEN 52.

INTERNET HEADER - Jon Postel

Jon gave a brief review of the layout of the internet header. It is as described in IEN 44 "Latest Header Formats", except that the TOS field has been defined as of this meeting and the DF bit has been moved to the center bit of the Flags field.

AGENDA AND DATE FOR NEXT MEETING - Vint Cerf

1. Access Control & Routing - Perlman

   by Network rather than Host, and by service class, classification, reliability, and cost.

2. Stream Set Up & Conferencing - Hoversten, Cohen, Binder

3. Internet Mail Forwarding - Postel

   What services are provided, what formats are used, how is compatibility with old systems maintained?

4. VDH & Gateway performance - Strazisar

5. Window Control / Retransmission Measurements - Hinchley, Treadwell

6. Gateway Congestion Control - Strazisar

   (a) local gateway to catenet gateway

   (b) catenet gateway to catenet gateway

7. Minimal Gateway Monitoring - Brescia, Davidson, Cole

   (a) global catenet monitoring service

   (b) monitoring for control

8. Source Routing - Cohen, Clark

   (a) are all pieces of source route internet addresses

   (b) how does source routing impact higher level protocols TCP, NVP, IN, and gateways?

   (c) impact on hidden gateways? on hosts without catenet ids.

   (d) impact on hidden networks?

9. Internet Name Translation Service - Postel

   What service is provided, how is it reached, and what formats are used?

10. Interactions with Gateways - Cohen, Ginny, Binder, Cole, Davidson, Brescia, Chiappa (plus UCL representation)

   small group meeting before next meeting

   (1) Access Control
   (2) Monitoring & Control
   (3) Debugging
   (4) Statistics Gathering
   (5) Stream Setup
   (6) Gateway-Gateway Congestion, Flow, & Routing Control

The date of the next meeting is 30 & 31 October and 1 November 1978.
The place of the next meeting is SRI.

ACTION ITEMS

1. Revised Internet Protocol Specification - Postel
checksum + TOS, DF to flags, error reporting

2. ACC ECU [error control unit] availability - Mathis

3. Host-SIMP interconnection via Local Host 1822 - Binder

4. Coordination of Changeover to use new Internet Header - Cerf, Strazisar

5. Provision of TCP version 4 to NDRE Gateway - Mathis, Strazisar, Tomlinson

6. TCP V2.5 installation ASAP @SRI,BBN,ISI - Plummer

7. Install TCP v4 at BBNC for testing - Plummer

8. EDN UNIX to MIT - Cerf & Cain

9. Document TOS Mappings - Cohen

10. NCP to TCP translation stratagies - Cerf & Kirstein

11. Present summary of SATNET stream set up concepts - Hoversten

12. Determine feasibility of public conference demonstration - Cerf

13. Document checksum algorithms for 10 & 11 - Reed

MEMOS DISTRIBUTED

The Catenet Model for Internetworking
Vint Cerf
IEN 48

Comments on TCP Checksum Function Design:
A Response to Internet Experimemt Note 45
P.L. Higginson
IEN 49 & INDRA Note 679

A Survey of End-to-End Retransmission Techniques
S.W. Edge & A.J. Hinchley
IEN 50

Types of Service in the Catenet
C.J. Bennett
IEN 51 & INDRA Note 680

Minimal Gateway Monitoring and Control Information

IEN Note Index
J. Postel

Internet Notebook Table of Contents
J. Postel

ATTENDEES

Vint Cerf	ARPA	CERF@ISIA
Richard Binder	BBN	BINDER@BBNE
Mike Brescia	BBN	BRESCIA@BBNE
Jon Cole	BBN	COLE@BBNE
John Davidson	BBN	DAVIDSON@BBNE
Radia Perlman	BBN	PERLMAN@BBN
William Plummer	BBN	PLUMMER@BBN
Virginia Strazisar	BBN	STRAZISAR@BBN
Kou-Mei Chuang	CCA	KOU-MEI@CCA
David L. Mills	COMSAT	MILLS@ISIE
Walt Roehr	DCEC	R740@BBN
Ray McFarland	DOD	MCFARLAND@ISIA
Danny Cohen	ISI	COHEN@ISIB
Jon Postel	ISI	POSTEL@ISIB
Jim Forgie	LINCOLN LAB	FORGIE@BBN
Estil Hoversten	LINKABIT	HOVERSTEN@ISIA
Noel Chiappa	MIT	JNC@MIT-AI
Steve Kent	MIT	SKENT@BBND
David Reed	MIT	DPR@MIT-ML
Yngvar Lundh	NDRE	YNGVAR@SRI-KA
Paal Spilling	NDRE	PAAL@SRI-KA
Ken Biba	SRI	BIBA@SRI-KL
Jim Mathis	SRI	MATHIS@SRI-KL
Andy Poggio	SRI	POGGIO@SRI-KL
Colin Bradley	UCL	KIRSTEIN@ISIA
Andrew Hinchley	UCL	UKSAT@ISIE
Peter Kirstein	UCL	KIRSTEIN@ISIA
John Shoch	XEROX	SHOCH@PARC