J. Postel
ISI
15 December 1978

IEN 70

Internet Meeting Notes - 4 December 1978

OBJECTIVES OF THE MEETING

The meeting was chaired by Jon Postel. Vint Cerf was unable to attend so the objectives of the meeting were taken to be the obvious ones of:

1) Determining the status of TCP implementations,

2) Gathering input for revision of the specification,

3) Demonstrating the TCP implementations.

Vint has indicated that TCP-4 is to be used in a major demonstration of the PRNET and ARPANET in April '79.

STATUS REPORTS

1. BBN - Bill Plummer

The problem of finding computer time for testing reported at earlier meetings have been easied to some extent. The current state is that TCP-2.5 is checked out, and work is under way on TCP-4. With respect to the schedule published in IEN 57, work is a little ahead of schedule. It is not currently possible for user written programs to access the internet level protocol or send datagrams. Bill is prepared to demonstrate TCP-2.5.

2. SRI - Jim Mathis

Jim reports that he has a preliminary version of TCP-4 on the LSI-11 available for testing. This version does not handle the rubber EOL or use URGENT. These features will be added before the program is distributed to others. The program may be ready for distribution in two to six weeks. Jim is ready to test his TCP-4 with others.

3. MIT - Dave Clark

Dave has a TCP-4 program in development and would like to test with others. Could provide an interface to the Datagram protocol. Dave has a toy user telnet program for testing purposes.

4. UCLA - Bob Braden

Bob has a TCP-4 program which is complete but still fragile. Bob has been able to communicate to his own program and to bounce packet through the testing-gateway provided by Bill Plummer. Bob's implementation does provide an interface to the datagram level. Bob is prepared to test his TCP-4 with other implementations.

5. DTI - Gary Grossman

Gary passed out a description of the TCP-3 produced by DTI. It will not interact with the TCP-4s since it is based on version 3 and has the AUTODIN II security features. Gary is prepared to demonstrate two instances of this program interacting.

6. BBN - Mike Wingfield

Mike has a TCP-4 written in C for the EDN. This program does not yet have the rubber EOL feature and does carry the AUTODIN II security features in the IN header options field. Mike also has a THP (AUTODIN II telnet) ready for testing (though it does no option processing). Mike is prepared to test his program against the others.

7. BBN - Jack Haverty

Jack's program is based on Jim Mathis'. Jack is working on the code that interfaces TCP to the rest of UNIX. An improvement in the performance of the UNIX interprocess communication (IPC) mechanism of five times has been achieved. Jack could demonstrate TCP-2.5.

8. FORD - Norman Abramovitz

Ford is in the early design stage for a TCP-4. Ford has no commitment to implement telnet.

9. MITRE - Anita Skelton

MITRE (Washington) is bringing up a cable-bus local net system. There will be a PDP-11/70 Unix host which may be a gateway between the local net and the ARPANET. Initially, the local net will have mini-hosts and may later have full hosts. The cable-bus system uses the RF band.

10. CCA - Thomas Lozenro

CCA is proceeding with the conversion of SRI TCP version 2.5 rather than wait for version 4. Estimated completion date is late February.

The Record and Relay (RAR) facility, as outlined at the last internet meeting, is running on CCA-SPEECH. The only current header transformation operates on ARPANET headers. It is expected that TCP version 2.5 and internet version X transformations will be up by the end of December. A current users' manual can be FTP'd from CCA-SPEECH using login name "RAR" (caps necessary), password "GUEST" and file-name "RARUSER.TXT".

ACTION ITEMS

1. Dave Clark was to write something about URGENT. Dave has some handwritten notes he will forward to Postel [ACTION: Clark].

2. Vint Cerf was to review the TCP spec and provide comments. Vint has some handwritten notes to review with Postel [ACTION: Cerf].

3. Information was supposed to be supplied by each implementor. Some have provided this information in the past. To meet this continuing need, we will start a biweekly report by implementors coordinated by Jim Mathis. Each implementor is to supply reports biweekly. The first report is due 15 December [ACTION: Mathis, Plummer, Clark, Braden, Wingfield, Haverty, Low, Stensby].

4. Bill Plummer was to supply some old notes on TCP scenarios. Bill distributed two memos "TCP Reset Mechanism" and "Sequence Number Arithmetic." Bill asked for comments when we had time to read the memos, and for the material to be considered in revising the TCP Specification.

TCP SPECIFICATION DISCUSSION

Jim Mathis suggested that TCP-4 may be an improvement on earlier versions since his TCP-4 is smaller than his TCP-2.5.

Jack Haverty brought up the problems of retransmission with rubber EOL's. Several people commented on this. Mainly one can't combine segments across a letter boundary. Some implementations save the formatted packet and others recompute the header.

Jack also wanted clarification on sending data outside the window, specifically what happens if a segment is partially acknowledged and

then the whole thing is retransmitted, is the data below the left edge legal? This seemed to be ok to everybody. One must also examine for control, especially Urgent, segments that are one octet to the above the right edge.

The time sequence of when to process options, as well as just what processing to do, should be described in the spec.

Gary Grossman has a collection of gripes and groans about the spec which he will SNDMSG to Postel [ACTION: Grossman]. In particular, Gary pointed to the state descriptions, the lack of specifics about the order of processing. Some comments about EOL being passed to the user being inconsistent. (Gary says EOL is not passed to the user.) Some problems with the specification of urgent, and the acceptability tests. The description of RESET on pages 21-24 was pointed out as being contradictory.

Several suggestions were made for improved specification format or style.

Gary distributed an excerpt from the specification of another protocol as an example. Carl Sunshine's note on specifications was distributed.

Jim Mathis suggested that example segments showing the values of various control fields could be useful.

Several people suggested that there be a way of indicating changes in the specification, e.g., change bars in the margin. Others suggested that a foreward or cover letter listing and explaining semantic changes would be more useful.

Dave Clark suggested that when several state/event/action descriptions are similar they be described once, with exceptions for the minor differences.

Dave also was confused about the open state as opposed to a listening state.

In the CLOSE state, it seems that at least one sequence of events would prevent the user from receiving all the data held at the receiving TCP.

The order of processing is confused or ambiguous at many points in the spec.

One suggestion was to "change state when this octet hits the left window edge."

DEMONSTRATIONS OF TCP IMPLEMENTATIONS

In the afternoon, we met at DCEC to test or demonstrate the TCP-4 implementations. The four programs that were in a state to attempt interconnections were Jim Mathis', Bob Braden's, Mike Wingfield's, and Dave Clark's. The first problems were to clarify the placement of the host and imp field in the internet header and the protocol version number. On getting these straightened out and checksumming turned on or off as needed, some messages were exchanged.

Wingfield's Report

Wingfield-Mathis

Mathis received Wingfield's SYN segment, forgetting to ignore the 4 byte internet option, and returned a RST segment with a bad ACK. This problem will be quickly fixed when Jim fixes his code.

Wingfield-Clark

A connection was established and data exchanged. However, because of a misinterpretation of the TCP spec concerning whether the ACK bit must always be set when the connection is established, the connection hung with Dave throwing away Mike's segments which did not have the ACK bit set.

Mike was interpreting the spec to mean that the ACK bit did not have to be on in new segments. This point needs to be resolved; either way is okay with me. This last experiment couldn't be completed because Multics went down for some time. The test was particularly interesting because Mike had set up a THP server and Dave was able to log in to Unix (almost).

[NOTE: There seems to be no advantage to send a TCP segment with out the ACK bit set and the current acknowledgment value once in the established state. There also seems to be no advantage in rejecting a TCP segment that has the ACK bit off, if it is otherwise ok. Thus, you are both wrong. In pratice all segments sent (once the established state has been reached) should carry valid acknowledgments and have the ACK bit on.]

Wingfield-Wingfield

Mike successfully demonstrated logging into EDN Unix from BBN Unix using a user and server THP on top of TCP.

Mathis' Report

Mathis-Braden

Exchanged SYNs and opened connections with UCLAs TCP. Correctly received data sent by UCLA, however, data from SRIs TCP did not make it to the user process. This test was run without checksums on UCLAs segments.

Mathis-Clark

Opened connection, exchanged data, and closed connection with MITs TCP twice. Also checked sending RSTs for non-existatn port numbers. This test was run without checksums.

Clark's Report

Clark-Mathis

Opened a connection sent data both ways and closed the connection. Generated RSTs by doing the wrong things, and thus checked out several error conditions. Opened and closed the same connection, and reused the same connection (at least the Multics end did).

Clark-Wingfield

Opened a connection passed data both ways, but hung on the problem of ACK bits being always or not always set. This was an especially interesting test as Dave got as far as an attempted login to Mike's Unix via Telnet. Unix said "login:" Multics responded "Wingfield" and Unix said "password:". At this point Multics rejected a segment from Unix.

Clark-Braden

Bob and Dave spent Tuesday following the meeting in a mutual debugging session at MIT. A connection was opened and reached the established state.

Braden's Report

No report is available from Bob at this time due to his being on vacation.

NEXT MEETING

The time and place of the next meeting is January 29 & 30 at ISI.

The agenda for the next meeting will be announced separately.

MEMOS DISTRIBUTED

-Agenda and Testing Plan - Postel

-DTI INFE TCP Status - Grossman

-The TCP Reset Mechanism - Plummer

-Sequence Number Arithmetic - Plummer

-Experps from DRAFT Host to Front End Protocol Specification version 3 - Grossman

-Specification of TCP from the User Point of View - Sunshine

ATTENDEES

Jack Haverty	BBN	JHAVERTY@BBND
Morris Kranc	BBN	MKRANC@BBND
William W. Plummer	BBN	PLUMMER@BBNA
Mike Wingfield	BBN	WINGFIELD@BBND
Tomas Lozeno	CCA	TLP@CCA
Ed Cain	DCA	DCEC-R850@BBNB
Ray McFarland	DOD	MCFARLAND@ISIA
Gary Grossman	DTI	grg@DTI
Norman Abramovitz	Ford Aero	BIBA@SRI-KL
Karl C. Kelley	Ford Aero	BIBA@SRI-KL
Danny Cohen	ISI	COHEN@ISIB
Jon Postel	ISI	POSTEL@ISIB
David Clark	MIT	DClark@MIT-Multics
Susan Poh	MITRE	POH@ISIA
Anita Skelton	MITRE	MITRE@BBN
Aage Stensby	NDRE	AAGE@SRI-KA
Jim Mathis	SRI	MATHIS@SRI-KL
Andy Poggio	SRI	POGGIO@SRI-KL
Bob Braden	UCLA-CCN	BRADEN@CCN
John Shoch	XEROX	SHOCH@PARC-MAXC