

RFC 683
NIC 32251

R. Clements
BBN
3 April 75

1

Introduction

2

In response to a long-known need for the ability to transfer TENEX paged files over the net via FTP, the TENEX FTP implementation has been extended.

2a

This implementation is an extension to the "OLD" protocol (RFC 354, ff). It was built after useful discussions with Postel, Neigus, et al. I do not mean to imply that they agreed that this implementation is correct, nor for that matter do I feel it is correct. A "correct" implementation will be negotiated and implemented in the "NEW" protocol (RFC 542), if funding ever appears for that task.

2b

The Problem(s)

3

This extension attacks two separate problems: Network reliability and TENEX disk file format's incompatibility with FTP. A checksummed and block-sequence-numbered transmission mode is seriously needed, in my opinion. This mode should also allow data compression.

3a

It is also necessary to handle paged, holey TENEX files. This latter problem, seriously needed for NLS, is the motivation for the current extension.

3b

The former problem requires a new MODE command, if done correctly; probably two MODES, to allow data compression in addition to checksumming. Actually, I think that is the tip of an iceberg which grows as 2^{*N} for additional sorts of modes, so maybe some mode combination system needs to be dreamed up. Cf the AN, AT, AC, EN, ET, EC TYPES. Also, one should be able to use MODE B and MODE C together (NEW protocol) to gain both the compression and restart facilities if one wanted.

3c

The second problem, TENEX files, are probably a new kind of STRUCTure. However, it should be possible to send a paper tape to a disk file, or vice versa, with the transfer looking like a paged file; therefore perhaps we are dealing with a data representation TYPE. This argument is a bit strained, though, so a paged STRUCTure is quite likely correct. I admit to feeling very unsure about what is a MODE, what is a TYPE and what is a STRUCTure.

3d

The (Incorrect) choices made

4

Having decided that new MODES and STRUctures were needed, I instead implemented the whole thing as a single new TYPE. After all, I rationalize, checksumming the data on the network (MODE) and representing the data in the processing system as a checksummed TYPE are really just a matter of where you draw the imaginary line between the net and the data. Also, a single new TYPE command reduced the size of the surgery required on the FTP user and server programs.

4a

Implementation details

5

The name of the new TYPE is "XTP". I propose this as a standard for all the Key Letter class of FTP commands: the "X" stands for "experimental" -- agreed on between cooperating sites. The letter after the "X" is signed out from the protocol deity by an implementor for a given system. In this case, "T" is for TENEX. Subsequent letter(s) distinguish among possibly multiple private values of the FTP command. Here "P" is "Paged" type.

5a

TYPE XTP is only implemented for STRU F, BYTE 36, and MODE S,

Information of TYPE XTP is transfered in chunks (I intentionally avoid the words RECORD and BLOCK) which consist of a header and some data. The data in a chunk may be part of the data portion of the file being transfered, or it may be the FDB (File Descriptor Block) associated with the file.

5b

Diversions: the TENEX Disk File

6

For those not familiar with the TENEX file system, a brief dissertation is included here to make the rest of the implementation meaningful.

6a

A TENEX disk file consists of four things: a pathname, a page table, a (possibly empty) set of pages, and a set of attributes.

6b

The pathname is specified in the RETR or STOR verb. It includes the directory name, file name, file name extension, and version number.

6c

The page table contains up to 2**18 entries. Each entry may be EMPTY, or may point to a page. If it is not empty, there are also some page-specific access bits; not all pages of a file need have the same access protection.

6d

A page is a contiguous set of 512 words of 36 bits each,

The attributes of the file, in the FDB, contain such things as creation time, write time, read time, writer's byte-size, end of file pointer, count of reads and writes, backup system tape numbers, etc,

6e

Note that there is NO requirement that pages in the page table be contiguous. There may be empty page table slots between occupied ones. Also, the end of file pointer is simply a number. There is no requirement that it in fact point at the "last" datum in the file. Ordinary sequential I/O calls in TENEX will cause the end of file pointer to be left after the last datum written, but other operations may cause it not to be so, if a particular programming system so requires,

6f

In fact both of these special cases, "holey" files and end-of-file pointers not at the end of the file, occur with NLS data files. These files were the motivation for the new TYPE,

6g

Meanwhile, back at the implementation,...

7

Each chunk of information has a header. The first byte, which is the first word (since TYPE XTP is only implemented for BYTE 36) of the chunk is a small number, currently 6, which is the number of following words which are still in the header. Next come those six words, and then come some data words,

7a

The six header words are:

Word 1: a checksum,

This is a one's complement sum (magnitude and end-around carry) of the six header words and the following data words (but not the leading "6" itself). The sum of all words including the checksum must come out + or = zero,

Word 2: A sequence number,

The first chunk is number 1, the second is number 2, etc,

Word 3: NDW,

the number of data words in this chunk, following the header. Thus the total length of the chunk is 1 (the word containing NHEAD) + NHEAD +NDW. The checksum checks all but the first of these.

Word 4: Page number.

If the data is a disk file page, this is the number of that page in the file's page map. Empty pages (holes) in the file are simply not sent. Note that a hole is NOT the same as a page of zeroes.

Word 5: ACCESS.

The access bits associated with the page in the file's page map. (This full word quantity is put into AC2 of an SPACS by the program reading from net to disk.)

Word 6: TYPE.

A code for what type of chunk this is. Currently, only type zero for a data page, and type -3 for an FDB are sent.

After the header are NDW data words. NDW is currently either 1000 octal for a data page or 25 octal for an FDB. Trailing zeroes in a disk file page will soon be discarded, making NDW less than 1000 in that case. The receiving portions of FTP server and user will accept these shortened pages. The sender doesn't happen to send them that way yet.

7b

Verification is performed such that an error is reported if either:

7c

The checksum fails,

The sequence number is not correct,

NDW is unreasonable for the given chunk type, or

The network file ends at some point other than immediately following the data portion of an FDB chunk.

Closing comments

8

This FTP server and user are in operation on all the BBN systems and at some other sites -- the user being more widely distributed since fewer sites have made local modifications to the user process.

8a

I believe the issues of checksumming and sequencing should be addressed for the "NEW" protocol. I hope the dissertation on TENEX files has been useful to users of other systems. It may explain my lack of comprehension of the "record" concept, for example. A TENEX file is just a bunch of words pointed to by a page table. If those words contain CRLF's, fine -- but that doesn't mean "record" to TENEX. I think this RFC also points out clearly that net data transfers are implemented like the layers of an onion: some characters are packaged into a line. Some lines are packaged into a file. The file is broken into other manageable units for transmission. Those units have compression applied to them. The units may be flagged by restart markers (has anyone actually done that?). The compressed units may be checksummed, sequence numbered, date-and-time stamped, and flagged special delivery. On the other end, the process is reversed. Perhaps MODE, TYPE, and STRU don't really adequately describe the situation. This RFC was written to allow implementors to interface with the new FTP server at TENEX sites which install it. It is also really a request for comments on some of these other issues.

8b

NWG/RFC# 683
FTPSRV == TENEX Extensions for Paged Files

RCC 8-APR-75 12:11 32251

(J32251) 8-APR-75 12:11;;; Title: Author(s): Robert C.
Clements/RCC; Sub=Collections: NWG NIC; RFC# 683; Clerk: JAKE;
Origin: < NETINFO, RFC683,NLS;2, >, 8-APR-75 11:15 JAKE ;;;;####;

USER SERVICES REPORT: COURSES AT GUNTER AFS, ALA, for users of
ARPA-NSW slot - MAR 19-21, 1975

1. MAR 19-21, 1975, [3 person-days] by JMB for personnel of Air Force Data Systems Design Center (DSDC)	1
2. Persons contacted:	2
Lt. Larry CRAIN	2a
Persons in courses:	2b
Sheila Roten (SHR using dir DSDC-SG)	2b1
Penny Glover (PAG using dir DSDC-SG)	2b2
Didn't attend second course	2b2a
Cynthia Pattillo (CFP using dir DSDC-PR)	2b3
Jo Wagner (JWV using dir DSDC-PR)	2b4
Ruth Woodley (RMW using dir DSDC-SC)	2b5
Didn't finish first course--she left during the afternoon of the second day--no second course.	2b5a
Raymond Uhler (using dir DSDC-SC; no ident)	2b6
3. COURSES:	3
BASIC COURSE completed	3a
SECOND COURSE as follows:	3b
Sections completed:	3b1
1. getting to nls	3b1a
2. organization of the system	3b1b
5. editing	3b1c
3. printing	3b1d
Sections NOT done:	3b2
4. addressing	3b2a
6. communicating	3b2b

USER SERVICES REPORT: COURSES AT GUNTER AFS, ALA, for users of
ARPA-NSW slot - MAR 19-21, 1975

7, trouble-shooting

3b2c

PRIMER distributed, but no assistance given, to those who attended
second course.

3c

JMB 2-APR-75 17:54 32253

USER SERVICES REPORT: COURSES AT GUNTER AFS, ALA, for users of
ARPA=NSW slot = MAR 19-21, 1975

(J32253) 2-APR-75 17:54;;; Title: Author(s): Jeanne M. Beck/JMB;
Distribution: /JMB([INFO-ONLY]) RH([INFO-ONLY]) SGR([INFO-ONLY
]) SLJ([INFO-ONLY]) JCN([INFO-ONLY]) RLL([INFO-ONLY]) DCE(
[INFO-ONLY]) JHB([INFO-ONLY]) ; Sub-Collections: SRI-ARC; Clerk:
JMB;

how do statement labels work?

i do not understand how to label statements.

1

as an example of a labeled statement, my initial file has two such, one is labeled "author" and the other "journal".

2

however, when i attempt to label a statement, for some reason it does not work. what magic is required to label my statement?? i want statement three of my initial file labeled "recieved" so that i can file away my recieved journal mail and not have to figgure out when i'm starting to see old stuff again.

3

thank you.

4

JGN 2-APR-75 17:59 32254

how do statement labels work?

(J32254) 2-APR-75 17:59;;; Title: Author(s): J. Gregory Noel/JGN;
Distribution: /FEEDBACK([ACTION]) FGB([INFO-ONLY]) SGR([
INFO-ONLY]) ; Sub-Collections: NIC FEEDBACK; Clerk: JGN;
Origin: < NAVIMP, TO-FEEDBACK,NLS;1, >, 2-APR-75 17:31 JGN ;;;;###;

Statement Names

Greg, 1

I am repeating an earlier message which probably was lost in
today's system crash. 1a

Basically, the first word in a statement is the statement name. 1b

Thus, to create a statement whose name is received 1c

Type [where (cr) = carriage return] 1d

is(cr)(cr)received(cr) 1d1

which echoes as 1e

insert C: Statement (to follow) A: L: 1e1

received 1e2

Base C: 1e3

At this point, your initial file JGN,NLS will probably look
something like: 1f

(Journal) 1f1

(Author) 1f2

received 1f3

These are all statement names. The default name delimiters (as in
"received") are NULL as LEFT DELIMITER and BLANK as RIGHT
DELIMITER. For reasons I won't go into now, the name delimiters in
the other two statements have been set to LEFT and RIGHT PARENS. 1g

If you have any questions, call me at AV 287-1618. 1h

Also, it's a good idea to update files (including your initial
file) before logging out. 1i

Regards, Frank 2

FGB 3-APR-75 17:26 32255

Statement Names

(J32255) 3-APR-75 17:26;;; Title: Author(s): Frank G. Brignoli/FGB;
Distribution: /JGN([ACTION]) ; Sub-Collections: NIC; Clerk: FGB;

ibm effort

Jim & Stoney, I n order to complete the SOW for this effort, I need to know specific identification of the items the government plans on furnishing. For IBM to become fully equated with all features of NLS, we have decided on furnishing them with a slot on Office-1 and have them access the ARPANET through NBS. I would like to see IBM get a CRT with mouse and keyset and line processor for accessing DNLS and a TI for accessing INLS and also to furnish them with a hardcopy capability). I don't know what we may need in the way of modems or communication lines. Ideaally, we would want the terminals to be at IBM (which is located a few blocks away from NBS) or if this cannot be arranged, we might set up the NLS stations at NBS. I would appreciate your help in obtaining this information, so we can move ahead with this effort. Thanks, Joe.

1

JpC 4-APR-75 12:18 32256

lbm effort

(J32256) 4-APR-75 12:18;;; Title: Author(s): Joe P. Cavano/JPC;
Distribution: /JCN([ACTION]) DLS([ACTION]) ; Sub-Collections:
RADC; Clerk: JPC;

JGN 4-APR-75 14:06 32257

watch for mail

did you get my sndmsg from nelc@isi??

1

JGN 4-APR-75 14:06 32257

watch for mail

(J32257) 4-APR-75 14:06;;; Title: Author(s): J. Gregory Noel/JGN;
Distribution: /FGB([ACTION]) ; Sub=Collections: NIC; Clerk: JGN;

IMM 5-APR-75 08:02 32259

This is one of many questions I will be forwarding to our new ident hoping we can get responses from anyone who knows about this.

(BEDFORD) 22-MAR-75 0944-PDT Experiments with the <>START Record and <>STOP Record commands.

Distribution: MATTIUZ, napke, bedford

Received at: 22-MAR-75 09:44:40

1

These are sort of neat, and you might want to experiment with them. When you <>START Record and name a file, all the subsequent interaction (including mistakes, backspaces, etc.) are recorded in a .CTL version of the named file.

When you <>STOP Record, the recording terminates.

When you <>PLAYback Record, the identical commands and text will be carried out.

I'm not sure yet what the difference between this and the <>PROCESS command is; <>PLAYback seems easier to build. One difference, perhaps trivial, is that when you build the PLAYback session, you are actually moving thru NLS routines in a real time environment; <>PROCESS commands aren't executed until a later time.

A further difference, again probably trivial, is that you can specify either "real" time or condensed time for the <>PLAYback sessions; that is, they can take place at the same speed at which you originally executed the commands, or at the fastest speed that the machine can work.

Inez, would you find out what further differences there are? I think we might be missing something something here.

1a

-mike

1b

(J32259) 5-APR-75 08:02;;; Title: Author(s): Inez M. Mattiuz/IMM;
Distribution: /AID([ACTION]) MTB([ACTION] this is the first
problem I am sending to AID as well as FEED) FEED([ACTION]) ;
Sub-Collections: BELL-CANADA AID; Clerk: IMM;

MIKE 5-APR-75 12:09 32260

this is a general comment; both imm and pan will see it.

this is a test of the individual comment feature we were talking
about.

1

MIKE 5-APR-75 12:09 32260

(J32260) 5-APR-75 12:09;;; Title: Author(s): Michael T.
Bedford/MIKE; Distribution: /IMM([ACTION]) PAN([INFO-ONLY]) ;
Sub-Collections: NIC; Clerk: MIKE;

DAP 7-APR-75 09:12 32261

New user=prog?

This is a request for development of a new user-program, similar to LETTER, but which produces memoes instead of letters.

New user=prog?

Dean, I think the Journal message I sent you last week on this topic got lost in last week's problems. If this is a duplicate (approximately), ignore it.

1

What I'd like is a program which does what user=prog LETTER does, but which generates memoranda instead. The format is much like this:

2

< POTTER, MEMO2,NLS;1, >, 18-MAR-75 20:09 DAP ;;;

2a

New user-prog?

Memorandum for: MR. MESSICK

2a1

cc: Mr. Elford
Mr. Forehand
Mr. Majetic
Mr. Manning
Mr. Maslow

2a2
2a3
2a4
2a5
2a6

Subject: Replacement for Dave Krathwohl Date: 8
APR 75

2a7

From: David A. Potter

2a8

Here are three possible replacements for Dave Krathwohl on the External Advisory Committee for the Job Analysis of Teaching Project, ordered in terms of what I must admit is a weak personal preference:

2a9

Phil Jackson

2a10

Don Medley

2a11

Bob Soar

2a12

Reactions? Other suggestions?

2a13

2a14

DAP 7-APR-75 09:12 32261

New user-prog?

(J32261) 7-APR-75 09:12;;; Title: Author(s): David A. Potter/DAP;
Distribution: /NDM([ACTION]) FEEDBACK([INFO-ONLY]) ;
Sub-Collections: NIC FEEDBACK; Clerk: DAP;

RMS2 7-APR-75 09:43 32262

Resetting directives

Can you offer any suggestion about this problem?

Resetting directives

There seems to exist a problem with the OUTPUT PROCESSOR, at least in the way I wish to use it. A document I am preparing has several pages that must have the OUTPUT PROCESSOR directives inserted in the text. This is not really unusual, however after the page have been typed I wish to return to the OUTPUT PROCESSOR directives as they existed prior to this page. The problem is to determine what the situation was prior to this page and then to reinstitute these directives. This can be rather difficult since I must change the page size, page length, top and bottom margin along with several other parameters. Once The page has been printed there is no need to retain these directives (in fact, they will interfere with the normal printouts), how then, can I return to the previous directives without ever knowing exactly what they were?

1

It would seem that there may be a need for a few more new commands in the OUTPUT PROCESSOR. Maybe something like a SAVE DIRECTIVES or/and a RESTORE DIRECTIVES command

2

Can you help with this Problem or can you offer any suggestions?

3

Resetting directives

(J32262) 7-APR-75 09:43;;; Title: Author(s): Robert M.
Sheppard/RMS2; Distribution: /FEEDBACK([ACTION]) AID([ACTION]) ;
Sub-Collections: NIC FEEDBACK AID; Clerk: RMS2; Origin: <
SHEPPARD, AID,NLS;3, >, 7-APR-75 09:39 RMS2 ;;;;###;

Network Performance Dialog, 31Jan-5Apr 75

2-APR-75 1023-EDT TOMLINSON at BBN-TENEXA: Delay measurements on
BBNB, Distribution: NET PERFORMANCE TECHNICAL GROUP
[BBN]<MCKENZIE>NPTG,TXT: Received at: 2-APR-75 08:26:03

1

As part of our effort to identify and correct the causes of poor performance that has been experienced using TENEX via the network, we have installed patches in the BBNB monitor to record the progress of NVT characters through the NCP. Below are the results of these measurements for the period from 1330 EDT to 1600 edt on Monday 31 March,

1a

The patches permit the time at which certain critical events occur to be recorded and for a user program to retrieve these times. The probing interval was 10 seconds unless delays required a longer period to permit the completion of one probe before the next was initiated. The times recorded are: 1) Time of arrival of 36-th bit of a message on the monitored host/link, 2) Time of arrival of the last bit of the same message, 3) Time at which the message was removed from the network input queue, 4) Time at which the character in the message was placed in the terminal input buffer (and its echo placed in the terminal output buffer), 5) Time at which the message containing the echo was generated, 6) Time at which that message was placed on the connection output queue, 7) Time at which the message was placed on the network output queue. This is different from the above only if there is an outstanding RFNM. No instances of this being different from the above actually occurred, 8) Time the message was removed from the output queue, 9) Time at which the 36-th bit of the message had been transferred to the IMP, 10) Time at which the last bit of the message was transferred to the IMP,

1b

The user program additionally recorded the time of day at which the BIN it was executing on that connection completed and the current 1 min load average. These times and other data were written on a file in raw form. The data presented below have been processed to extract the minimum, maximum, and average delays encountered between each of the points listed above plus the overall echo time and the overall response time. The overall echo time is just the difference between the time the last bit of the echo message was transferred to the IMP and the time the first 36 bits of the received message arrived from the IMP. The response time measurement is inferred by adding the interval from placing the input character in the terminal buffer to user program activation to the echo time,

1c

No conclusions have been drawn from this data yet, but the following observations are worth making. There are occasions when incoming messages spend an exceedingly long time on the network input queue (4.5 sec in one instance), since the process handling

Network Performance Dialog, 31Jan-5Apr 75

this queue has top priority in the system, this is difficult to understand. The delay in transmitting the echo following its being placed in the terminal buffers is also long on occasion (3.7 sec). The reason here may be the same as the reason for the preceding problem since the same process has this responsibility as well. There are also occasions when messages spend a long time on the network output queue. This can only be due to either the IMP blocking the host interface or an exceedingly large amount of output traffic. Since the IMP has been observed to block the interface for relatively long times, that is probably the reason for this large delay. Program activation delays are also fairly long on occasion, but this is mainly a function of program priorities and system loading which was clearly high during the experiment.

1d

The labels on the data below are a bit cryptic and are expanded here:

1e

RL=RI Interval from arrival of 36-th bit to arrival of last bit. (Time to transfer the message into TENEX from the IMP.) INQ Time the message spends on the input queue. CHI Time required to dispatch the character into the terminal input buffer. ECH=CHI Delay before the NCP fork notices the echo in the terminal buffer and generates the message containing the echo. CONQ Time the message spends on the connection queue. This is typically negative due the fact that the message never really spends any time there and that the points of measurement are actually reversed. OUTQ Time spent on the output queue. This reflects queueing delay due to prior traffic plus delay due to blocking of the interface on prior traffic. SND1 Time required to transfer the first word of the message to the IMP. SNDRST Time required to transfer the rest of the message into the IMP. PRGWAK Time required for the program to be activated from point of placing the character in the terminal input buffer until the program can execute a TIME JSYS. ECHO Time required to echo a character when no program activation is required. RESP Time required for a response character to be transmitted from the moment of arrival of an input character from the IMP to completion of transmission of the response character. LDAVG The 1 minute load average rounded to the nearest integer.

1f

The raw data may be found in the file [BBNB]<TOMLINSON>NET-DELAY-STATISTICS.BIN. Several files with this name exist for various periods. The write date should be referred to in order to determine which is the correct version.

1g

The file consists of a series of 13-word entries. Each entry contains: 0/ Time and date of the sample (GTAD format), 1/ Count of missed events (should be zero) 2-10/ The 9 times listed above

Network Performance Dialog, 31Jan-5Apr 75

(TIME format), 11/ Time of program wakeup in TIME format, 12/ Load average,

1h

The program for processing the data in the form below is [BBNB]<TOMLINSON>XXNCPD.SAV, It should be gotten and reentered, The source for XXNCPD.SAV is XXNCPD.MAC, It contains both the data taker program and the processor, When reentered, it processes the current version of NET-DELAY-STATISTICS.BIN, When started, it starts probing the NCP and generates a new version of NET-DELAY-STATISTICS.BIN, If not detached, it also processes the data at 10 minute intervals,

11

NCP DELAY STATISTICS FOR PERIOD FROM 31-MAR-75 13:26:28 TO 31-MAR-75 13:38:37, 60 SAMPLES,

1j

NAME	MIN	MAX	AVG	RL=R1	0	2	0 INQ
8 359	50 CHI	3	20	4 ECH=CHI	7	1089	
66 CONQ	-21	0	-1 OUTQ	0	5793	97 SND1	0
29 2 SNDRST		0	2	0 PRGWAK	52	1098	
503 ECHO	25	5889	221 RESP	99	6330	725	
LDAVG	6	15	10				

1k

NCP DELAY STATISTICS FOR PERIOD FROM 31-MAR-75 13:38:48 TO 31-MAR-75 13:50:21, 60 SAMPLES,

11

NAME	MIN	MAX	AVG	RL=R1	0	1	0 INQ
6 985	76 CHI	3	12	4 ECH=CHI	6	1803	
108 CONQ	-2	0	0 OUTQ	0	219	10 SND1	0
298 20 SNDRST		0	2	0 PRGWAK	62	952	
403 ECHO	26	2167	221 RESP	151	2708	624	
LDAVG	6	15	8				

1m

NCP DELAY STATISTICS FOR PERIOD FROM 31-MAR-75 13:50:32 TO 31-MAR-75 14:02:38, 60 SAMPLES,

1n

NAME	MIN	MAX	AVG	RL=R1	0	2	0 INQ
9 2810	169 CHI	3	24	4 ECH=CHI	8	479	
59 CONQ	-22	0	-1 OUTQ	0	510	15 SND1	0
648 21 SNDRST		0	2	0 PRGWAK	82	1277	
542 ECHO	29	2903	269 RESP	117	3212	812	
LDAVG	7	19	13				

1o

NCP DELAY STATISTICS FOR PERIOD FROM 31-MAR-75 14:02:50 TO 31-MAR-75 14:13:43, 60 SAMPLES,

1p

NAME	MIN	MAX	AVG	RL=R1	0	1	0 INQ
6 434	50 CHI	3	8	4 ECH=CHI	6	1700	
81 CONQ	-2	0	0 OUTQ	0	277	10 SND1	0
124 6 SNDRST		0	2	0 PRGWAK	21	1227	

Network Performance Dialog, 31Jan-5Apr 75

433 ECHO	33	1730	153 RESP	186	2461	586	
LDAVG	3	18	8				1q

NCP DELAY STATISTICS FOR PERIOD FROM 31-MAR-75 14:13:54 TO 31-MAR-75 14:25:05, 60 SAMPLES. 1r

NAME	MIN	MAX	AVG RL=R1	0	2	0 INQ	
6 1231	65 CHI	3	647	16 ECH=CHI	15	2083	
84 CONQ	=3	0	0 OUTQ	0	196	6 SND1	0
319 10 SNDRST		0	1	0 PRGWAK		41	1283
452 ECHO	34	2742	184 RESP	176	2928	636	
LDAVG	3	16	8				1s

NCP DELAY STATISTICS FOR PERIOD FROM 31-MAR-75 14:25:18 TO 31-MAR-75 14:36:31, 60 SAMPLES. 1t

NAME	MIN	MAX	AVG RL=R1	0	1	0 INQ	
4 371	48 CHI	3	54	5 ECH=CHI	14	705	
54 CONQ	=5	0	=1 OUTQ	0	139	4 SND1	0
720 19 SNDRST		0	1	0 PRGWAK		63	992
462 ECHO	36	914	134 RESP	117	1719	597	
LDAVG	7	14	9				1u

NCP DELAY STATISTICS FOR PERIOD FROM 31-MAR-75 14:36:42 TO 31-MAR-75 14:48:39, 60 SAMPLES. 1v

NAME	MIN	MAX	AVG RL=R1	0	3	0 INQ	
7 4557	207 CHI	3	40	4 ECH=CHI	11	965	
74 CONQ	=4	0	0 OUTQ	0	2	0 SND1	0
263 6 SNDRST		0	1	0 PRGWAK		46	934
466 ECHO	34	4594	293 RESP	119	5456	760	
LDAVG	9	20	15				1w

NCP DELAY STATISTICS FOR PERIOD FROM 31-MAR-75 14:48:50 TO 31-MAR-75 14:59:51, 60 SAMPLES. 1x

NAME	MIN	MAX	AVG RL=R1	0	1	0 INQ	
7 225	49 CHI	3	6	4 ECH=CHI	7	423	
54 CONQ	=3	0	0 OUTQ	0	217	6 SND1	0
176 15 SNDRST		0	2	0 PRGWAK		51	1149
465 ECHO	39	492	130 RESP	113	1360	596	
LDAVG	4	17	9				1y

NCP DELAY STATISTICS FOR PERIOD FROM 31-MAR-75 15:00:02 TO 31-MAR-75 15:11:28, 60 SAMPLES. 1z

NAME	MIN	MAX	AVG RL=R1	0	2	0 INQ	
7 1757	138 CHI	3	8	4 ECH=CHI	7	1729	
104 CONQ	=5	0	0 OUTQ	0	155	6 SND1	0

Network Performance Dialog, 31Jan-5Apr 75

701 30 SNDRST	0	1	0 PRGWAK	62	2175	
499 ECHO 37	3544	284 RESP	135	3849	784	
LDAVG 4	14	9				1ae

NCP DELAY STATISTICS FOR PERIOD FROM 31-MAR-75 15:11:39 TO 31-MAR-75 15:23:19, 60 SAMPLES, 1aa

NAME	MIN	MAX	AVG RL=R1	0	1	0 INQ
0 732	94 CHI	3	975	20 ECH=CHI	13	3707
148 CONQ	-19	0	-1 OUTQ 0	12	0 SND1	0
19 1 SNDRST		0	1	0 PRGWAK	56	1044
516 ECHO	36	4446	266 RESP	126	5092	783
LDAVG	10	16	13			1ab

NCP DELAY STATISTICS FOR PERIOD FROM 31-MAR-75 15:23:30 TO 31-MAR-75 15:34:36, 60 SAMPLES, 1ac

NAME	MIN	MAX	AVG RL=R1	0	2	0 INQ
7 333	52 CHI	3	23	4 ECH=CHI	6	390
36 CONQ	-19	0	-1 OUTQ 0	3	0 SND1	0
43 2 SNDRST		0	2	0 PRGWAK	34	932
421 ECHO	34	452	96 RESP 71	1035	518 LDAVG	
4 12	7					1ad

NCP DELAY STATISTICS FOR PERIOD FROM 31-MAR-75 15:34:48 TO 31-MAR-75 15:45:29, 60 SAMPLES, 1ae

NAME	MIN	MAX	AVG RL=R1	0	1	0 INQ
6 1103	56 CHI	3	62	6 ECH=CHI	6	233
36 CONQ	-11	0	-1 OUTQ 0	39	1 SND1	0
544 11 SNDRST		0	2	0 PRGWAK	21	1105
300 ECHO	28	1758	112 RESP	81	2099	412
LDAVG	2	9	5			1af

NCP DELAY STATISTICS FOR PERIOD FROM 31-MAR-75 15:45:39 TO 31-MAR-75 15:56:17, 60 SAMPLES, 1ag

NAME	MIN	MAX	AVG RL=R1	0	1	0 INQ
9 240	59 CHI	3	7	4 ECH=CHI	9	127
29 CONQ	-6	0	0 OUTQ 0	2	0 SND1	0
15 1 SNDRST		0	1	0 PRGWAK	59	860
405 ECHO	33	324	96 RESP 134	1042	501 LDAVG	
4 5	4					1ah

NCP DELAY STATISTICS FOR PERIOD FROM 31-MAR-75 15:56:27 TO 31-MAR-75 16:07:04, 60 SAMPLES, 1ai

NAME	MIN	MAX	AVG RL=R1	0	1	0 INQ
6 283	53 CHI	3	19	4 ECH=CHI	9	350

Network Performance Dialog, 31Jan-5Apr 75

32 CONG	-1	0	0 OUTQ	0	1	0 SND1	0
63 3 SNDRST		0	1	0 PRGWAK		48	1019
357 ECHO	30	383	94 RESP	80	1174	451 LDAVG	
3 6	4						

1aJ

NCP DELAY STATISTICS FOR PERIOD FROM 31-MAR-75 16:07:16 TO
31-MAR-75 16:14:10, 40 SAMPLES,

1aK

NAME	MIN	MAX	AVG RL=R1	0	2	0 INQ
6 235	48 CHI	3	7	4 ECH=CHI	6	324
60 CONG	-4	0	0 OUTQ	0	955	28 SND1
103 10 SNDRST		0	1	0 PRGWAK		66
443 ECHO	35	1097	151 RESP	123	1268	867
LDAVG	2	4	3			594

1aI

END OF FILE REACHED,

1aM

2-APR-75 1933=EDT IRBY at BBN-TENEXB: Some round trip character
times under low load conditions at Office-1 Distribution: NET
PERFORMANCE TECHNICAL GROUP [BBN]<MCKENZIE>NPTG.TXT: Received at:
2-APR-75 16:13:29

2

We have been periodically running Jim White's two-terminal
watchdog program at Office-1 to try to detect low load -- high
round trip character times. The program sends characters (5 per
minute) to both a local and a TIP looped-back terminal and records
the average round trip time for the two. In this case, the TIP is
TYMSHARE TIP (thus, 0 hops). The following are some interesting
sample numbers we have been observing. Times are in milliseconds.
Monday 3/31/75 TIME LOAD LOCAL REMOTE

----- 16:13:35				0.72	50
220 16:14:35	0.78	60	300 16:15:36	0.78	
30 420 16:16:36	0.80	30	390 16:17:37		
0.81 50 480 16:18:37	0.80	50	300 16:19:38		
0.86 40 250 16:20:39	0.90	90	380 Tuesday		

4/1/75 TIME LOAD LOCAL REMOTE				----- 08:57:11		1.64	90
350 08:58:12	1.56	90	1000 08:59:13	1.50			
80 800 09:00:16	1.44	70	920 09:01:17				
1.33 110 700 09:02:17	1.21	50	870 09:03:19				
1.29 60 860 09:04:20	1.19	80	710 09:05:22				
1.29 120 920 09:06:23	1.20	70	630 09:07:23				
1.50 120 750 09:08:24	2.24	200	1260 09:09:26				
2.27 130 430 09:10:28	2.54	170	590 09:11:30				
2.82 170 440 09:12:30	2.69	70	490 09:13:31				
2.29 70 660 09:14:32	2.25	190	560 09:15:32				
2.16 110 490 09:16:33	2.06	90	320 09:17:34				
1.89 60 320							

2a

Network Performance Dialog, 31Jan=5Apr 75

10:11:01	1.72	110	740	10:12:03	1.69	80
1760 10:13:05	1.67	190	1110	10:14:06	2.43	
200 470 10:15:08	2.52	180	430	10:16:08		
2,29 100	300 10:17:09	1.98	40	310 10:18:10		
1,82 60	400 10:19:10	1.63	60	370 10:20:10		
1,47 60	290 10:21:11	1.38	60	310 10:22:11		
1,29 70	280 10:23:12	1.40	170	490 10:24:13		
1,56 170	510 10:25:15	1.85	380	880 10:26:15		
1,69 40	610 10:27:16	1.66	140	710 10:28:17		
1,59 100	410 10:29:17	1.59	110	300 10:30:18		
1,52 70	390 10:31:20	2.13	330	600 10:32:23		
3,24 530	1130 10:33:24	3.45	310	520 10:34:27		
3,69 340	1130 10:35:28	3.44	160	500 10:36:29		
3,53 280	590 10:37:31	4.18	240	550 10:38:33		
4,57 320	370 10:39:34	4.27	410	510 10:40:36		
4,26 280	460					

2b

-- Charles.

2c

3-APR-75 1545=EDT CLEMENTS at BBN=TENEXA: OPEN SCANNER LINES - REF
IRBY'S REPORT OF TODAY. Distribution: NET PERFORMANCE TECHNICAL
GROUP [BBN]<MCKENZIE>NPTG.TXT: Received at: 3-APR-75 13:12:19

3

OPEN SCANNER LINES ON A DC10 SCANNER SHOULD NOT BE A PROBLEM (BUT CAN BE). THE FIX IS DOCUMENTED IN BOTH DEC'S FIELD SERVICE TECH TIPS AND IN THE MAINTENANCE MANUAL FOR THE DC10. THERE IS A JUMPERING ARRANGEMENT ON THE W706 TTY RECEIVER CARD WHICH CAUSES AN OPEN LINE TO GENERATE ONE AND ONLY ONE NULL CHARACTER, UNLESS TRANSITIONS FROM SPACE BACK TO MARK ARE OCCURRING. WHEN A LINE BECOMES OPEN, IT CHANGES FROM MARKING (IDLE) TO SPACING. THIS GENERATES ONE NULL, UNTIL A TRANSITION BACK TO MARK, I.E., UNTIL THE LINE IS POWERED BACK UP OR PLUGGED BACK IN, NO FURTHER INTERRUPTS ARE GENERATED. /RCC

3a

3-APR-75 1102=EDT MILLSTEIN at BBN=TENEXB: revised document
Distribution: BAGGIANO AT ISI, BALZER AT ISIB, BTHOMAS AT BBN,,
BURCHFIEL AT BBN, CARLSON AT ISI, CARLSTROM AT ISI,, CRAIN AT ISI,
CROCKER AT ISIB, IRBY AT BBNB, JACOBS AT BBN,, LAWRENCE AT OFFICE-1,
LEHTMAN AT BBNB, LLOYD AT ISI,, MAYHAN AT ISI, MICHAEL AT BBNB,
MILLSTEIN AT BBNB,, POSTEL AT BBNB, RETZ AT ISI, RIDDLE AT OFFICE-1,,
SCHAFFNER AT BBNB, SCHANTZ AT BBN, STONE AT OFFICE-1,, STUBBS AT BBN,
TRIOLO AT BBNB, UHLIG AT OFFICE-1,, WAAL AT BBNB, WALKER AT BBNB,
WARSHALL AT BBNB,, WATSON AT BBNB, WEEKS AT OFFICE-1, WHITE AT BBNB,,
WINGFIELD AT OFFICE-1, HOLG AT ISIB, BRADEN AT CCN,, POGRAN AT
MIT-MULTICS, BOLDUC AT BBNB, SATTLEY AT BBNB Received at: 3-APR-75
13:12:44

4

there is a new version of <MILLSTEIN>WM-PROCEDURES.TXT at BBNB

Network Performance Dialog, 31Jan=5Apr 75

containing the definitions of two new variables maxitem and
Maxlist and a new paragraph at the end of Section 1,

4a

we invite your comments.

4b

steve warshall/bob millstein

4c

3=APR=75 1736=PDT ENGELBART: To Steve Walker re plans to improve
network Distribution: NETWORK PERFORMANCE GROUP
[OFFICE=1]<ENGELBART>NPG.TXT: Received at: 3=APR=75 17:36:33

5

Steve: We at ARC must make decisions soon regarding NLS-service
computers to support ARC workers and new Utility clients. The
intense studies over the past month reveal quite clearly some
problems in the Network that prevent bandwidth that is adequate to
support DNLS. Reconfiguring to reduce the IMP-hop distances, and
improving host-host and Host=Imp protocols to avoid blocking
problems and if possible to allow message streaming without RFNM
waits, seem critical. Effective bandwidths otherwise appear
ludicrously low in view of sub-net capability.

5a

We need to know what commitments we can count on regarding ARPA or
DCA improving network performance. I would appreciate your
informing us about what you now know or intend. We soon have to
commit ourselves to two TENEXs worth of resources for next summer,
and Net performance is a serious issue. (Note: We are also
intensely involved in a parallel assessment of TENEX's part of the
NLS-service problems.)

5b

I would like to point out that the NSW and VLDB programs most
certainly will be in trouble if the network performance isn't
brought up to par,

5c

Regards, Doug

5d

P. S. We notice that a local (SF Bay Area) change in Net
connections could put Office-1 just one hop from ISI -- reducing
from 13 to 7 the hops from ARPA, and increasing no other O-1
user's "distance." Does that interest you?

5e

3=APR=75 1719=EDT MCKENZIE at BBN-TENEX: Your Recent Summary
Distribution: IRBY AT BBNB, net performance technical group
[bbnl]<mckenzie>nptg.txt: Received at: 3=APR=75 20:45:25

6

Charles, Somehow or other some wires seem to have gotten crossed.
I am referring specifically to the portion of your recent summary
which refers to the "BREAK CHARACTERS TO TIP" problem. The
following note, which I thought went to everyone on 21 March,
describes the status of this particular problem. In fact, NOTHING

Network Performance Dialog, 31Jan-5Apr 75

has been done to the TIP since that time, so both your statement that the reason was "unknown" and your statement that "the TIP has been modified to prevent it" are probably incorrect. Regards, Alex
Previous message follows:

6a

There has been much discussion recently about the problems which the Tymshare TIP (and perhaps other TIPs) has been having with the 4800 baud modems used by the "lineprocessors". This is an attempt to clear up a few misconceptions about what does and does not happen.

6b

The modem/lineprocessor combination, as currently implemented, does not use any "data terminal ready" signal to tell the TIP when to listen to the terminal. Apparently the modem could convey this signal to the TIP, but the lineprocessor was not implemented with this in mind; therefore the modem at the terminal end has the "data terminal ready" signal wired on. We have asked the SRI people to correct this design, and we believe they are in the process of beginning to do so.

6c

The Tymshare TIP was recently fixed (as part of a general TIP retrofit program) to ignore a steady stream of "breaks" coming from a terminal. This fix, however, was designed under the assumption that "externally clocked" devices would only send constant "breaks" if they really meant them. Unfortunately, the modems used with the lineprocessor fall into the category of "externally clocked" devices. We are now reviewing our design to try to find a way of discarding (in hardware) breaks from such devices.

6d

Even in the current case, with a somewhat deficient design of the TIP hardware to discard multiple breaks, and a somewhat deficient design of the lineprocessor to not generate "data terminal ready", we believe that essentially no TIP bandwidth is consumed by the breaks IFF there is no open network connection (to a Host) from such a device. We believe that a significant fraction of TIP bandwidth is consumed in the case where the user fails to close (@c <cr>) his TIP connection before turning off his lineprocessor. We suspect that this was happening frequently at Tymshare, since the reason this problem first came to light was Postel's observation that traffic generated by the Tymshare TIP seemed abnormally high.

6e

It appears that the lineprocessor(s) connected to the ARPA TIP are not using modems and do provide a "data terminal ready" signal to the TIP. The TIP examines this signal every few seconds. Thus, when a lineprocessor is turned off, the TIP sees a constant stream of breaks (the lineprocessor is still an "externally clocked" device) until the next time the TIP probes the "ready" signal, at

Network Performance Dialog, 31Jan-5Apr 75

which point further input is ignored. This is consistent with the findings Hardy reported in his message of 20 March 1946-EDT,

6f

3-APR-75 1926-EDT VICTOR at BBN-TENEXB: differences between ray tomlinsons work and that of jim white and ken victor Distribution: NET PERFORMANCE TECHNICAL GROUP [BBN]<MCKENZIE>NPTG.TXT: Received at: 3-APR-75 20:50:08

7

ray tomlinsons work traced an nvt character from its arrival at the tenex imp interrupt routine to a user program and then back out to the imp interrupt routine (as an echoed character), this work provided a good measure of the ncp fork inside tenex and pointed up some possible delays at the host-imp interface. The measurements made by j. white and me traced a character that started in a user program and went through the ncp (without the overhead of going through an nvt terminal) and out through the imp into the subnet to a tip that then looped back the character which was then traced in its path back to a user program, this work also provided a measure of tenex handling of the ncp (although not as detailed as rays work), however, this work also provided a measure of subnet times and showed that at times the subnet was experiencing imp hop times of as large as 200ms, the modifications i made to the ai tenex also provide for metering the number and sizes of messages sent to and from tenex and the imp, as yet, we have no user programs to analyze this additional information (specifically delays associated with multi-packet messages), we (arc) currently have no plans to write such a program as we believe it belongs as an integral part of a generalized measurement package in tenex. We are in the process of putting together a document describing in detail what we feel this general measurement package should contain, until such a package is implemented, i would suggest that rays work be upgraded to be able to measure sub-net times in addition to what it currently measures,

7a

3-APR-75 1339-EDT IRBY at BBN-TENEXB: Summary of Findings of ARC Network Performance task force Distribution: NET PERFORMANCE TECHNICAL GROUP [BBN]<MCKENZIE>NPTG.TXT:, [bbnb]<postel>nsw-steering-committee,list: Received at: 3-APR-75 20:57:44

8

This document attempts to summarize the findings of the ARC Network Performance task force. Recommendations are also made where appropriate. In addition time estimates are given for length of time until a problem is fixed. As I learn more about the problems, I will update this file,

8a

BREAK CHARACTER TO TIP

8b

Network Performance Dialog, 31Jan-5Apr 75

The 208 modems connected to Line Processors were sending continuous breaks to the TIP whenever the Line Processor was powered down. This was because the Line Processor-modem connection had the "data ready" bit wire to true all the time, since the Line Processor did not provide such a signal. The TIP had been modified to detect this sort of thing and disable it. The TIP modification did not work for some unknown reasons.

8c

This is supposedly not happening anymore. The TIP has been modified to prevent it and the Line Processor is being changed to prevent it. Martin Hardy should make sure this is not happening to any other TIPs we are associated with.

8d

This caused a very significant degradation of service to users using local host if TIP connect was to the local host and was not closed.

8e

BREAK CHARACTERS TO PDP-10

8f

An open line to the PDP-10's Data Line scanner was found and fixed. May have accounted for a 1% to 2% load on system (this guess is based on our experience that an open line at 9600 baud caused a 20% load on our local PDP-10). The open line at Office=1 was at 1200 baud.

8g

VDH interference

8h

The Very Distant Host interface code and buffers now being installed in the TYMSHARE TIP is highly suspect but no conclusive proof as yet. It appears to be causing IMP crashes and strange behavior.

8i

I expect this will get fixed within a month. I would recommend moving the VDH to another IMP since it eats up 50% of the IMP's buffer space and seems flacky.

8j

HOST blocking imp-host interface

8k

It was discovered that TENEX is blocking IMP from inputting on 1,31 TENEX. Office=1 is moving to 1,33 to fix this and other network problems. I expect this will be up within three weeks. This may not entirely alleviate the problem and we should continue to watch it.

8l

imp blocking host-imp interface

8m

This seems to be happening quite a lot and probably accounts for the 15 second periods users have observed when there is no interaction with the servicing host. There are two causes known

Network Performance Dialog, 31Jan-5Apr 75

to me to date. Hopefully the bbn-network people will clarify this area for us,

8n

multi-packet messages

8o

Can cause the interface to be blocked for at least the round trip time to the receiving imp. This causes the host to be unable to send data to any other host for a period of at least 1/2 second for the ARPA-TIP -- Office-1 and ARC-ELF -- BBNB situations. Fixing this requires a change to host-imp protocol which I recommend ARPA should support. This will require many months to implement for all hosts but could be fixed in all TENEX's more quickly.

8p

receiving host slow in taking data

8q

This may account for long (15 second) pauses where a host does not respond. The subnet blocks a host from sending too much data to a receiver host that is not taking data quickly enough. If the host is going down or faulting, the delay in taking data becomes the timeout period or 15 (or is it 40) seconds.

8r

Hopefully the bbn-network people will clarify this one for us shortly. The fix here is also a change in imp-host protocol or perhaps a change to sending hosts to monitor their outstanding messages more closely and to anticipate the imp blockage before it occurs.

8s

general tenex overhead -- NCP and scheduler

8t

Initial measurements from Jim White's and Ken Victor's work as well as Ray Tomlinson's recent memo indicate that general TENEX NCP/scheduler overhead often exceeds network transmission times for single character interactions. It is perhaps possible to tune TENEX so that this overhead is reduced. BBN-TENEX people are currently looking into these problems and hopefully will come up with some answers.

8u

It almost goes without saying that TENEX load has a significant effect on responsiveness. However, the pie slice scheduler does provide some insulation from this.

8v

subnet performance

8w

The subnet performance seems about according to specs (about 50 MS per hop). However we have observed times when it appeared to be more like 200 MS per hop. Long network paths seem to cause some problems because the net was designed assuming a max of 6 hops and the average length is now more like 10 hops. This is

Network Performance Dialog, 31Jan=5Apr 75

very serious and requires either changes in network topology or network design and protocols to accomodate this deviation from initial design assumptions,

8x

The difference in performance of printers and high speed terminals when zero and one or three hops away is still unexplained. We have experienced differences of a factor of two in throughput to these devices depending on whether they were zero or more hops away (going from zero to one hop for the printer and zero to three for the terminal -- we ran two terminals side by side, one zero hops and one three),

8y

ELF problems

8z

ELF crashes, hangs, and character dropping has been a very serious problem for us. There are now three people working on various of the problems. My guess is that most of our serious elf problems will be solved within a week if Dave Retz pays attention to them. Otherwise, I would guess three weeks,

8a@

Mouse keyset handling

8aa

The reliability of mouse/keyset usage has dropped considerably. It was decided that this would be corrected within NLS rather than changing the Line Processor (which we now consider to be frozen except for bug fixes),

8ab

I expect the fixes will be in the running BBNB NLS within a few days

8ac

suspended connections

8ad

This continues to be a problem for me. My connection tends to be broken to BBNB about four times a day. I have never been able to type in the daily status report without at least 2 breaks in my work, some of which are auto hangup of the data set at the ames tip,

8ae

single vs multi-packet messages

8af

It was learned that sending multi-packet messages to tip terminals over long network paths did not win. We have changed tip/elf buffers to be smaller than a packet to avoid this. However, file transfers and other network uses can still cause problems.

8ag

-- Charles,

8ah

3-APR-75 1350=EDT IRBY at BBN=TEEXB: Summary of Findings of ARC

Network Performance Dialog, 31Jan-5Apr 75

Network Performance task force Distribution: NET PERFORMANCE
 TECHNICAL GROUP [BBN]<MCKENZIE>NPTG.TXT;
 [bbnb]<postel>nsw-steering-committee.list: Received at: 4-APR-75
 01:03:28

This document attempts to summarize the findings of the ARC Network performance task force. Recommendations are also made where appropriate. In addition time estimates are given for length of time until a problem is fixed. As I learn more about the problems, I will update this file.

BREAK CHARACTER TO TIP

The 208 modems connected to Line Processors were sending continuous breaks to the TIP whenever the Line Processor was powered down. This was because the Line Processor-modem connection had the "data ready" bit wire to true all the time, since the Line Processor did not provide such a signal. The TIP had been modified to detect this sort of thing and disable it. The TIP modification did not work for some unknown reasons.

This is supposedly not happening anymore. The TIP has been modified to prevent it and the Line Processor is being changed to prevent it. Martin Hardy should make sure this is not happening to any other TIPs we are associated with.

This caused a very significant degradation of service to users using local host if TIP connect was to the local host and was not closed.

BREAK CHARACTERS TO PDP-10

An open line to the PDP-10's Data Line scanner was found and fixed. May have accounted for a 1% to 2% load on system (this guess is based on our experience that an open line at 9600 baud caused a 20% load on our local PDP-10). The open line at Office-1 was at 1200 baud.

VDH interference

The Very Distant Host interface code and buffers now being installed in the TYMSHARE TIP is highly suspect but no conclusive proof as yet. It appears to be causing IMP crashes and strange behavior.

I expect this will get fixed within a month. I would recommend moving the VDH to another IMP since it eats up 50% of the IMP's buffer space and seems flacky.

Network Performance Dialog, 31Jan-5Apr 75

HOST blocking imp-host interface

9K

It was discovered that TENEX is blocking IMP from inputting on 1.31 TENEX. Office-1 is moving to 1.33 to fix this and other network problems. I expect this will be up within three weeks. This may not entirely alleviate the problem and we should continue to watch it,

9L

imp blocking host-imp interface

9M

This seems to be happening quite a lot and probably accounts for the 15 second periods users have observed when there is no interaction with the servicing host. There are two causes known to me to date. Hopefully the bbn-network people will clarify this area for us,

9N

multi-packet messages

9O

Can cause the interface to be blocked for at least the round trip time to the receiving imp. This causes the host to be unable to send data to any other host for a period of at least 1/2 second for the ARPA-TIP -- Office-1 and ARC-ELF -- BBNB situations. Fixing this requires a change to host-imp protocol which I recommend ARPA should support. This will require many months to implement for all hosts but could be fixed in all TENEX's more quickly,

9P

receiving host slow in taking data

9Q

This may account for long (15 second) pauses where a host does not respond. The subnet blocks a host from sending too much data to a receiver host that is not taking data quickly enough. If the host is going down or faulting, the delay in taking data becomes the timeout period or 15 (or is it 40) seconds,

9R

Hopefully the bbn-network people will clarify this one for us shortly. The fix here is also a change in imp-host protocol or perhaps a change to sending hosts to monitor their outstanding messages more closely and to anticipate the imp blockage before it occurs,

9S

general tenex overhead -- NCP and scheduler

9T

Initial measurements from Jim White's and Ken Victor's work as well as Ray Tomlinson's recent memo indicate that general TENEX NCP/scheduler overhead often exceeds network transmission times for single character interactions. It is perhaps possible to tune TENEX so that this overhead is reduced. BBN-TENEX people are

Network Performance Dialog, 31Jan-5Apr 75

currently looking into these problems and hopefully will come up with some answers,

9u

It almost goes without saying that TENEX load has a significant effect on responsiveness. However, the pie slice scheduler does provide some insulation from this.

9v

subnet performance

9w

The subnet performance seems about according to specs (about 50 MS per hop). However we have observed times when it appeared to be more like 200 MS per hop. Long network paths seem to cause some problems because the net was designed assuming a max of 6 hops and the average length is now more like 10 hops. This is very serious and requires either changes in network topology or network design and protocols to accomodate this deviation from initial design assumptions.

9x

The difference in performance of printers and high speed terminals when zero and one or three hops away is still unexplained. We have experienced differences of a factor of two in throughput to these devices depending on whether they were zero or more hops away (going from zero to one hop for the printer and zero to three for the terminal -- we ran two terminals side by side, one zero hops and one three).

9y

ELF problems

9z

ELF crashes, hangs, and character dropping has been a very serious problem for us. There are now three people working on various of the problems. My guess is that most of our serious elf problems will be solved within a week if Dave Retz pays attention to them. Otherwise, I would guess three weeks.

9a@

Mouse keyset handling

9aa

The reliability of mouse/keyset usage has dropped considerably. It was decided that this would be corrected within NLS rather than changing the Line Processor (which we now consider to be frozen except for bug fixes).

9ab

I expect the fixes will be in the running BBNB NLS within a few days

9ac

suspended connections

9ad

This continues to be a problem for me. My connection tends to be broken to BBNB about four times a day. I have never been able to type in the daily status report without at least 2 breaks

Network Performance Dialog, 31Jan-5Apr 75

in my work, some of which are auto hangup of the data set at the
ames tip,

9ae

single vs multi-packet messages

9af

It was learned that sending multi-packet messages to tip
terminals over long network paths did not win. We have changed
tip/elf buffers to be smaller than a packet to avoid this.
However, file transfers and other network uses can still cause
problems.

9ag

-- Charles,

9ah

4-APR-75 0741=EDT WALDEN at BBN-TENEX: DIFFERENCES BETWEEN
TOMLINSON AND WHITE/VICTOR WORK Distribution: VICTOR AT BBNB, net
performance technical group [bbn]<mckenzie>nptg.txt: Received at:
4-APR-75 04:41:20

10

IF I UNDERSTAND CORRECTLY WHICH FIGURES YOU ARE TALKING ABOUT WHEN
YOU CONSIDER THE WHITE/VICTOR FIGURES, THEN I THINK THE FACT THAT
THEY LUMP NCP+TENEX TIMES TOGETHER SO ONE CAN'T SEE WHICH IS
CONTRIBUTING WHAT IS NOT ALL THAT USEFUL. I FOUND RAY'S FIGURES
VERY USEFUL BECAUSE OF THE CAREFUL BREAKDOWN ON DELAY
CONTRIBUTIONS. I THINK IT WOULD BE WONDERFUL IF RAY'S PROGRAM WAS
EXPANDED TO INCLUDE NET DELAY, BUT ONLY IF DONE IN A WAY WHICH
CAREFULLY SEPARATES THE CONTRIBUTORS.

10a

4-APR-75 0749=EDT WALDEN at BBN-TENEX: LOST CONNECTIONS
Distribution: IRBY, net performance technical group
[bbn]<mckenzie>nptg.txt: Received at: 4-APR-75 04:59:48

11

IN YOUR SUMMARY OF YESTERDAY, YOU MENTION THAT YOUR CONNECTION IS
GETTING BROKEN OR THE TIP IS HANGING UP YOUR DATA SET OR SOMETHING
SEVERAL TIMES A SESSION. CAN YOU EXPAND ON WHAT YOU THINK IS
HAPPENING. WE HAVEN'T BEEN WORKING VERY MUCH ON THIS PROBLEM NOT
KNOWING THAT IS WAS A SERIOUS PROBLEM. IS IT A SERIOUS PROBLEM,
DO YOU THINK WE ARE WORKING ON IT.

11a

IF THE PROBLEM IS YOU LOCAL TIP HANGING UP ITS DATASET, WE SHOULD
LOOK INTO THAT. IF THE PROBLEM IS THAT TENEX IS JUST CLOSING THE
CONNECTION EVERY NOW AND THEN, THEN THE TENEX GUYS MUST LOOK INTO
THE REASONS WHY TENEX CLOSES A CONNECTION WHICH IS OPEN. I'D
ENCOURAGE TENEX BEING FIXED TO LOG EVERY CONNECTION CLOSE AND THE
REASON FOR CLOSING WITH A TIME STAMP AS THE WAY TO TRACK THIS
DOWN. IF IT IS SOMETHING IN BETWEEN, THEN I SUPPOSE WE WILL ALL
HAVE TO LOOK. ONCE AGAIN, IT WILL HELP IF YOU KNOW EXACTS TIMES AS
THESE CAN BE CORRELATED WITH THE NCC AND TENEX LOGS, DAVE

11b

4-APR-75 0932=EDT MCKENZIE at BBN-TENEX: Broken Connections

Network Performance Dialog, 31Jan-5Apr 75

Distribution: WALDEN, net performance technical group
 [bbn]<mckenzie>nptg.txt:,, irby Received at: 4-APR-75 06:40:52

12

Dave, I think another possibility to be investigated is that Irby's terminal and/or modem is failing to maintain the signal that tells the modem at the TIP end that the terminal is still there. It seems to me we once saw this type of result with an acoustically-coupled terminal, although I can't remember where or when. Perhaps Charles can tell us if he knows that lots of people are seeing this problem or only himself; whether he sees it at both the AMES TIP and the Tymshare TIP or only one of them; whether he works with any service Host other than BBNB and if connections get lost there as well as at BBNB, etc. Since we have had no noticable complaints about this problem before, I am inclined to suspect the terminal a bit more than I suspect other parts of the system. Regards, Alex

12a

4-APR-75 1904=EDT VICTOR at BBN-TENEXB: a tenex measurement package Distribution: NET PERFORMANCE TECHNICAL GROUP
 [BBN]<MCKENZIE>NPTG.TXT: Received at: 4-APR-75 16:05:17

13

< ANDREWS, MEASURE133,NLS;6, >, 4-APR-75 18:52 KEY ;;;; This is a list of measurement capabilities that we at SRI-ARC would like to see added to the current capabilities in TENEX 1.33. Part of these capabilities existed in SRI-ARC's 131 system. We may install part or all of these capabilities in SRI-AI TENEX 1.33. Motivation Experiences with SRI-ARC and OFFICE-1 measurement system (monitor changes plus superwatch user program) have shown us that such a system is very desirable. An instance of superwatch running at a low frequency provides a complete record of all up-time performance and is valuable. Recent experiences of network performance problems coupled with host performance problems indicate to us that performance records as well as high frequency probing capabilities are important now to identify problem areas. But such tools will be CRUTIAL in beating the performance problems out of the NSW complex. Definitions A COUNTER is a word in monitor memory that is incremented every time an event (which it is counting) occurs. Example: number of clock ticks. A COUNT is a word in monitor memory that contains an integer of some particular meaning, and it is correct at any instant. Example: number of user pages available in system. A METER is a word in monitor memory that contains a sum or integral. Strictly speaking, a counter is a meter. Example: integral of balances set jobs dt. Example: Sold time since system was started. A PARAMETER is a significant variable that is computed by a user program from COUNTERS, COUNTS and METERS, all lifted from the monitor space in a sufficiently short period of time. Example: per cent of sold time for the last minute of real tme. Clock interrupt sampling

13a

Network Performance Dialog, 31Jan-5Apr 75

Some of the meters we would like to have available are best maintained by clock-driven sampling. The SRI-ARC 1.31 monitor did its sampling every 50 ms., when the two clocks were synchronized. The sampling operation is fast enough that it does not introduce any appreciable overhead. Swapping/memory management measurements REAL drum measurements (may be omitted when no physical drum) drum queue length This is a count of the number of entries in the physical drum queue (0=empty), drum busy meter This records the number of 50ms samples for which the drum queue length was non-zero, drum queue sum This is a meter containing the sum of the non-zero drum queue samples (yields average drum queue length given number of samples), drum transfer counter This counts all reads and writes to the physical drum, VIRTUAL drum measurements (may be omitted when virtual drum = real drum) NOTE: There may be both real and virtual drums, as in the case of OFFICE-1. In that case the real drum is a subset of the virtual drum, V-drum queue length V-drum busy meter V-drum queue sum V-drum transfer counter REAL disk measurements NOTE: Virtual drum transfers going to/from the disk are included here, disk queue length disk busy meter disk queue sum disk transfer counter Virtual drum cleanup activity A counter recording the number of page transfers from virtual drum to disk, Pages regained without writing on the disk are not counted, Drum usage counts count of free pages on virtual drum (exists but not generally accessible)

13b

count of free pages on physical drum (" ") Core cleanup activity number of transfers from core to (virtual) drum (counter) number of transfers from core to disk (counter) Working set management activity NOTE: We have a poor understanding of the TENEX memory management and would like to see these measurements to aid our understanding but also to find out how TENEX treats an NLS working set under different conditions. In the long term, these may not be valuable things to have in the measurement package.

number of calls to routine XGC (counter) number of pages 'removed' by XGC (counter) number of times routine GCPC is called (counter) Balance Set measurements number of balance set jobs A meter containing the sum of number of BS jobs for each 50 ms. sample, (somewhat redundant but necessary to compute parameters from following meters), number of balance set jobs runnable A meter containing the sum of number of BS jobs that are not in page wait state (for each 50 ms. sample) number of pages held for balance set jobs A meter containing the sum of FKNR for each BS job (for each 50 ms. sample) number of pages held for other reasons A meter similar to above, Maintained every 50 ms. sample, Scheduler/system measurements These are super as of TENEX 1.33, One simple addition give an important statistic: time in system mode A counter of the number of 50 ms. samples taken with the user bit off in the interrupt word (i.e. it was in system mode when

Network Performance Dialog, 31Jan-5Apr 75

interrupted), Subsystem measurements we would like the monitor to collect more information on

13c

subsystem operation and the most reasonable thing seems to be to provide it with a subsystem name and have it collect extra information for forks running that subsystem. Here is what we would like to see: time spent executing on each queue A meter for each queue, containing the sum of execute time, average working set size NOTE: current WS size numbers are size at blocking time. We would like a distribution of WS size obtained from a 50 ms sample of processes running the given subsystem and in the balance set. Perhaps a reasonable way to do this is to have two user-specified values that would define three regions ($ws < p1$, $p1 \leq ws < p2$, $ws \geq p2$). The monitor would maintain three meters counting the occurrences of forks with working set in each region. In general the monitor could maintain N regions, where N is determined at assembly time (?). The previous meters would necessitate another containing the sum of the number of balance set forks running the given subsystem, for each 50 ms. sample, Integral of CAPT dt This is a meter containing the sum of CAPT dt for each fork running the given subsystem. CAPT dt is added to it every time CAPT is changed for a fork. The previous meter necessitates a meter containing the sum of the dt factors every time the previous meter is changed. This will allow computation of the average cutoff age for the given subsystem. Alternatively, this could be estimated by a meter controlled by the 50 ms. sampling code. If CAPT for every balance set process were in a monitor table, the sampling code could sum the CAPTs for each process running the subsystem in question

13d

and accumulate the sums in the meter. Integral of IFAV dt Similar to the CAPT integral. Will allow computation of the average IFAV for a subsystem. This could be estimated by the 50 ms. sampling code as for CAPT.

NCP measurements With respect to measuring the NCP it seems appropriate to have two (and possibly three) modes of operation. The first mode is to measure all network traffic; the second mode would be to measure only traffic on specific host-link pairs (for input and output) or for a specific NVT. (The third mode is to turn off all measurement due to the overhead associated with collecting information at interrupt time.) These measurements should be resettable so that when it is desired to start measuring a specific host-link pair (as opposed to all or no network traffic), the measurements will be consistent. In addition to the time meters mentioned below, when the measurements are tracing a specific host-link pair, the discrete times of state changes (e.g. message moved to connection queue, etc.) should also be available to user programs analyzing the data. Some of the measurements listed below may already be collected, however, they are listed

Network Performance Dialog, 31jan-5Apr 75

anyway for the sake of completeness, Counts; number of input buffers available for receiving messages from the imp
 length of the imp output queue
 Counters: number of times ncp fork is run
 number of times ncp fork is run because impiflg was raised
 number of messages sent/received
 number of impbugs received
 number of high priority messages sent/received
 number of messages sent/received on link 0
 number of multi-packet messages sent/received on link 0

13e

length (in bits) of non multi-packet messages sent/received on link 0
 length (in bits) of multi-packet messages sent/received on link 0
 number of non link 0 regular messages the following counters only apply to non link 0 regular messages
 number of multi-packet messages sent/received length (in bits) of multi-packet messages sent/received length (in bits) of non multi-packet messages sent/received
 Meters: time to send/receive single/multi-packet link 0 messages to/from imp time to send/receive single/multi-packet non link 0 regular messages to/from imp time a single/multi-packet message spends on the input/output queues

rfnm wait times for single/multi-packet non link 0 regular / link 0 messages time a message spends on the connection queues time to move a message from/to the connection queue to/from its destination/source (tty buffers or file windows or imp i/o queues)

time between invocations of the ncp fork time between invocations of the ncp fork for when the ncp fork is run due to the raising of impiflg
 Access to measurements by user programs
 All the above counts and meters, and also the "tasktb" meters should be available to a user program of non-wheel status. The execution time required to get them should be very small. THE SYSGT-GETAB method is not fast enough. PEEK would be satisfactory if (1) there were a way to get an address given a symbol describing what the program wanted and (2) PEEK were not privileged. A JSYS similar to GETAB would be satisfactory if it would BLT a whole table into user space rather than return one

13f

word at a time. The user program should be able to make probes at a high frequency (every one or two seconds) without presenting much load to the system. We recognize that most every one of these "hooks" has been put into some TENEX somewhere at sometime by a proficient hacker. But a standard TENEX measurement package would be invaluable in comparing configurations, different kinds of loads, KA vs. KI etc., as well as in identifying bottlenecks, evaluating changes in subsystems, monitors, and hardware.

13g

Network Performance Dialog, 31jan-5Apr 75

4-APR-75 2049=EDT IRBY at BBN-TENEXB: Broken connections
 Distribution: WALDEN, MCKENZIE, net performance technical group
 [bbn]<mckenzie>nptg.txt: Received at: 4-APR-75 17:56:59

14

1) auto data set hangups: I have been getting this on Ames TIP 964-8997. Am using same terminal/coupler I have used for a long time with no similar problems before I started using Ames TIP. Hasn't happened this week at all. About once per day before.

14a

2) disconnects: These seem to originate from several sources. The TIP can sometimes restore the connection (ELF doesn't even try). BBNB usually does not crash, since the jobs are usually detached when we re-connect. "host not responding" is very common TIP message. We understand that Ray Tomlinson was trying things yesterday and inadvertently caused several breaks in about one hour's time. However, it does seem to happen nearly every morning around 10:00 or 11:00 PDT and one, two, or three other times each day. Following is an informal log kept at ARC (we do not pretend that it is complete)

14b

DATE REPORTER	TIME DOWN	TIME UP	COMMENTS
(I jus lostt my connection -- TIP resttored 10:15 our time) 3/27			
12:15	13:05	Host Rejecting (BBNB)	Hopper 3/28
09:08	?	BBN IMP DEAD	Postel I just
lost my connection twice entering the above line -- time above			
should be 9:45 not 10:15, 3/28	18:45	?	Host
not responding (BBNB)	Postel 3/30	13:00	1440 net
down then BBNB down Hopper 3/31	0735	0800	host
not responding (BBNB)	Hopper 3/31	08:40	?
Host dead	Ehard 4/1	1400	? BBNB
down	Postel 4/1	2049	? BBNB
down	Postel 4/2	10:10	? BBN imp down
then BBNB down	Maynard 4/2	11:55	12:08 BBN
imp down	Postel 4/2	12:52	13:05 BBN imp down
Postel 4/3 08:30	08:35	host not responding (BBN)	
Posttel 4/3 08:50	9:00	BBNB down	Watson
4/3 10:08	10:14	host dead	Hopper 4/3 10:22
? host dead	Lieberman 4/3	10:30	10:31
Bug halt at bbnb	Maynard 4/3	13:58	13:59 host
dead	Posttel 4/4	9:45	9:50 host
not responding (BBNB)	Postel 4/4	10:00	10:01 host
not responding (BBNB)	Postel 4/4	10:01	10:03 host
not responding (BBNB)	Postel (TIP resored on last three, ELF users		
had to reconnect) 4/4	10:22	10:30	host not
responding(BBNB)	Irby 4/4	10:55	10:57 Host
not responding(BBNB)	Postel 4/4	11:10	?
Host not responding (BBNB)	Postel 4/4	14:05	?

Network Performance Dialog, 31Jan=5Apr 75

Host not responding (BBNB) Postel 4/4 14:35 14:38 14c
 Net trouble Lieberman

-- Charles, 14d

4=APR=75 2130=EDT WALDEN at BBN=TENEX: broken connections
 Distribution: IRBY AT BBNB, net performance technical group
 [bbn]<mckenzie>nptg.txt: Received at: 4=APR=75 18:35:34 15

charles, 15a

it seems to me that your informal log shows several different
 problems and they should be treated separately, 15b

1. tip/data-set hanging up on you, the ncc staff can help try to
 track this one down if it reoccurs, 15c

2. net trouble or bbn imp dead, no fancy isolation diagnosis
 needed here, just have to keep trying to improve imp reliability.
 we recently got one bug out of the tip which was crashing the bbn
 tip about once per day, and we are continuing to look for such
 bugs, we postmortem every crash carefully, 15d

3. tenex crash, same story as with the tip crashes, the tenex
 guys try very hard to get the bugs out of tenex and i am sure
 will continue to try, 15e

4. tenex bug hlt's, the tenex guys argued very convincingly in an
 rfc some time ago that some bug hlt's are unavoidable and the best
 that could be done is to have the tip and tenex try to put the
 connection back together to minimize the users pain, we have had
 some trouble in the past making this reconnection mechanism work,
 both on the tenex side and the tip side, and if you consistently
 see the tip fail to restore suspended connections, please let us
 know so we can debug the mechanism some more, of course, in some
 cases tenex gives up and restarts rather than continuing and the
 connection can not get restored for this reason, there is nothing
 to debug in this situation, if there is any confusion over this
 mechanism, i will be glad to point you to documentation or to try
 to write better documentation, we do encourage elf to adopt the
 tip/tenex connection restoration mechanism and will be glad to
 supply documentation on the protocols, 5, there can be another
 kind of connection breakage which your message does not lead me to
 believe is happening but which you should be aware of all the
 same, if tenex doesn't take traffic from the imp fast enough, the
 imp will declare it tardy, we have not been seeing this much at
 all with tenex 1.33 at bbn of late, another, problem which your
 message does not indicate is happening but which we should be
 notified of if it begins to happen is tenex just closing (not

Network Performance Dialog, 31Jan-5Apr 75

breaking) the connection out of the blue, if you see any of this, we want to know as will the tenex guys i am sure, this would undoubtedly require joint diagnostic efforts,

15f

while the number of broken connections or dead systems is, i am sure, a pain for you, i am encourage by your message as it doesn't indicate anything mysterious happening, just system crashes, we can all understand the latter and try to make them happen less, this would not be true if you were seeing many mysterious "hung connections", for instance,

15g

regards, dave

15h

4-APR-75 2149-EDT WALDEN at BBN-TENEX: irby's log Distribution: MCKENZIE, MALMAN, net performance technical group
[bbn]<mckenzie>nptg,txt: Received at: 4-APR-75 18:58:29

16

joel, will you please check over the autodialler logs for the past several weeks to see if any trouble with 964-8997 at the ames tip has been apparent, data-set hang-up has been reported on this port,

16a

alex, would it be possible for one of the ncc operators (for instance) to check back over the ncc log for the times mentioned in irby's log of connection breaks to see what can be learned?

16b

dave

16c

26-MAR-75 1823-EDT IRBY at BBN-TENEXB: NCP/TENEX measurements Distribution: CLEMENTS, net performance technical group
[bbn]<mckenzie>nptg,txt: Received at: 26-MAR-75 15:37:41

17

Bob, I think I have as strong an aversion to re-inventing wheels as you do. Why, then, when Ken sent out Xn announcement that he was going to collect NCP measurements did you and Tomlinson not tell us you were already doing that??? We could certainly have found better ways to spend our time! From our vantage point you TENEX guys have been surprisingly quiet thruout this whole affair,

17a

I don't mean to seem too harsh. We are just trying to do our jobs,

17b

-- Charles

17c

27-MAR-75 0112-EDT POSTEL at BBN-TENEXB: NLS and Multi-Packet Messages Distribution: NETWORK=PERFORMANCE=GROUP:, naylor at isi
Received at: 27-MAR-75 03:31:45

18

21-MAR-75 21:54:00=EDT,545:000000000000 Mail from USC-ISI rcvd at

Network Performance Dialog, 31Jan-5Apr 75

21-MAR-75 2153-EDT Date: 21 MAR 1975 1855-PDT From: NAYLOR at
 USC-ISI Subject: HOSTS WITH NLS To: POSTEL at BBNB cc: NAYLOR

18a

AS YOU ALREADY KNOW, THE HOST RUNNING NLS SEND A SIGNIFICANTLY
 HIGHER NUMBER OF MULTIPACKET MESSAGES. OUR EARLY RESULTS SHOW
 THAT: 3.1% OF THE PACKETS SENT BY ISI WERE FROM MULTIPACKET
 MESSAGES, WHILE OFFICE-1 SENT AT A RATE OF 15.0% AND BBNB AT
 12.6%. WE'RE PLANNING FURTHER EXPERIMENTS AND WILL GIVE YOU THOSE
 RESULTS AS THEY BECOME AVAILABLE.

18b

REGARDS,... BILL ----- --jon

18c

26-MAR-75 1701-EDT CLEMENTS at BBN-TENEX: VICTOR'S MESSAGE ON MODS
 TO TENEX Distribution: NET PERFORMANCE TECHNICAL GROUP
 [BBN]<MCKENZIE>NPTG.TXT: Received at: 26-MAR-75 14:30:04

19

I FEEL THAT THIS WHOLE EFFORT IS RUNNING BADLY OUT OF CONTROL. KEN
 VICTOR HAS WRITTEN A SET OF CHANGES TO THE IMP DRIVER. CHARLES
 IRBY SAYS "IT WAS DECIDED THAT KEV WOULD..." [DECIDED BY WHO?
 AFTER TALKING IT OVER WITH WHO?] MEANWHILE, HERE AT BBN, WE ARE
 ACTIVELY LOCKING AT THE PROBLEM. RAY TOMLINSON, THE MAIN AUTHOR OF
 THE IMP CODE, IS DOING A SET OF MEASUREMENTS WHICH SEEM TO BE
 QUITE SIMILAR TO WHAT KEN IS DOING. THE WHEEL IS BEING INVENTED
 MANY TIMES, AND THE MEASURE- MENT PROGRAMS MAY WELL PRINT OUT
 APPLES ON THE WEST COAST AND ORANGES ON THE EAST COAST.

19a

FOR THE MOMENT, OUR PLAN IS AS FOLLOWS: WE WILL CONTINUE TO
 MEASURE WITH OUR METERS AND OUR ANALYSIS PROGRAMS. [KEN'S AND
 CHARLES'S MAY BE JUST AS GOOD, BUT WE HAVE OURS GOING AND WE SEE
 NO POINT IN STOPPING OUR EFFORT JUST BECAUSE WE WEREN'T TOLD OF
 THE OTHER TOOLS BEFORE THEY WERE GIVEN TO US AS A FAIT ACCOMPLI.]

19b

WE WILL CHECK OVER KEN'S CHANGES AND ADD ANYTHING TO OURS WHICH HE
 CAUGHT BUT WE MISSED.

19c

IF OUR EXPERIMENTS ARE INCONCLUSIVE, WE WILL CONSULT WITH THE
 APPROPRIATE WIZARDS TO DETERMINE WHAT TO EXAMINE NEXT.

19d

/RCC

19e

27-MAR-75 1718-EDT MCKENZIE at BBN-TENEX: SUMEX VDH code in
 Tymshare TIP Distribution: NET PERFORMANCE TECHNICAL GROUP
 [BBN]<MCKENZIE>NPTG.TXT: Received at: 27-MAR-75 15:05:14

20

After doing as much diagnosis as possible on one Tymshare TIP
 crash Tuesday, 3 Wednesday, and a few today (Thursday) the NCC
 programmers have reached a preliminary and tentative conclusion
 that the IMP is frequently out of buffers. We think that there
 may be some flaws in the book-keeping associated with buffer

Network Performance Dialog, 31Jan-5Apr 75

management which are seen only when the IMP is totally devoid of free buffers for extended periods and are currently studying this possibility. Recall that Tuesday was the first day when the VDH code was left in the IMP during working hours and that the VDH code usurps almost 50% of the IMP's buffer space. Please note that this report is quite tentative and that some problem completely unrelated to the VDH code may yet be found. Alex McKenzie

20a

27-MAR-75 1735=EDT BUTTERFIELD at BBN=TENEX: Host Blocking Distribution: NET PERFORMANCE TECHNICAL GROUP
[BBN]<MCKENZIE>NPTG,TXT: Received at: 27-MAR-75 15:21:43

21

On March 20 from one to five o'clock we patched the ARPA, TYMSHARE, and BBN10X TIPS to report to the NCC whenever a host had to wait for longer than some clip before it was allowed to send to the IMP. Afterwards I totalled the reports and sent Ben Barker a message with the results. Here they are preceded with a bit more explanation of the reasons. The MSGx reason for the host having to wait applies to both one and eight (more than one) packet messages and can be at the beginning of a host/host conversation while the message blocks are being set up or it can be whenever eight messages are waiting for RPNMs. The PPT and PLT apply respectively to one and eight packet messages and are tables that the hosts on an IMP share to keep track of the messages as requests for destination IMP storage are made. Only eight packet messages have to wait for the IMP to receive the allocation which is what ALL8 is. I make some attempt to explain the TASK and FREE reasons, but basically they are the subnet or storage waits which no doubt occur while the timer runs out after having spent a large amount of time at any of the others. Note that this can also be the case for the PPT and PLT reasons.

21a

Ben, here are the experiment's (3/20 1300-1700) totals for host blocking. The three TIPS had different clips: ARPA at 1/4 sec, TYMSHARE at 1/2, and BBN10X at 5/8. The first five reasons are pretty self-explanatory: all message number stuff. "TASK" means that HI was waiting for TASK to take the packet; HI had probably waited before for some time at one of the five. The one "FREE" wait for a buffer to input into also probably had waited at one of the five. (The PDP-15 at ARPA also only had to wait at MSG1 and did so 71 times.)

21b

ARPA-TIP OFFICE=1 TYMSHARE-TIP BBN-TENEXB BBN10X-TIP
BBN-TENEXC (1/4 sec) (1/2 sec) (1/2 sec) (5/8 sec) (5/8 sec) (5/8 sec)

21c

MSG1: 104	260	5	166	48
172 PPT: 2	62	45	107	18 MSG8: 3
PLT:				

Network Performance Dialog, 31Jan-5Apr 75

ALL8: 406 1 104 3 86 TASK: 56
 6 25 3 5 FREE: 1 21d
 Steve ----- 21e

27-MAR-75 1718-EDT MCKENZIE at BBN-TENEX: SUMEX VDH code in
 Tymshare TIP Distribution: NET PERFORMANCE TECHNICAL GROUP
 [BBN]<MCKENZIE>NPTG,TXT: Received at: 27-MAR-75 15:05:14 22

After doing as much diagnosis as possible on one Tymshare TIP
 crash Tuesday, 3 Wednesday, and a few today (Thursday) the NCC
 programmers have reached a preliminary and tentative conclusion
 that the IMP is frequently out of buffers. We think that there
 may be some flaws in the book-keeping associated with buffer
 management which are seen only when the IMP is totally devoid of
 free buffers for extended periods and are currently studying this
 possibility. Recall that Tuesday was the first day when the VDH
 code was left in the IMP during working hours and that the VDH
 code usurps almost 50% of the IMP's buffer space. Please note that
 this report is quite tentative and that some problem completely
 unrelated to the VDH code may yet be found, Alex McKenzie 22a

27-MAR-75 1210-P CHI: NPTG-ARC findings for 26-Mar-75
 Distribution: CHI SRI-ARC DLS
 Received at: 27-MAR-75 15:45
 Location: (JJOURNAL, 25634, 1:w) 23

25-MAR-75 1344-P CHI: NPTG Findings: Daily Report 24-mar-75
 Distribution: CHI DCE SRI-ARC DLS
 Received at: 26-MAR-75 12:54
 Location: (JJOURNAL, 25622, 1:w) 24

Comments: Sent via SNDMSG to Network Performance Technical Group
 (NPTG), 24a

27-MAR-75 1340-EDT CLEMENTS at BBN-TENEX: CHARLES'S REPORT,
 TODAY'S DATE Distribution: NET PERFORMANCE TECHNICAL GROUP
 [BBN]<MCKENZIE>NPTG,TXT: Received at: 27-MAR-75 10:49:01 25

I WISH TO CORRECT CHARLES'S DESCRIPTION OF MY NOTE YESTERDAY. I
 CERTAINLY DID NOT CRITICIZE ANYONE FOR MAKING MEASUREMENTS. I WAS
 TRYING TO POINT OUT A LACK OF COORDINATION, IN PARTICULAR,
 BBN-TENEX AND SRI-ARC HAVE SOME GROUND RULES FOR OUR RESPONDING TO
 THEIR REQUESTS FOR CHANGES IN OUR OPERATING SYSTEM ON OUR SERVICE
 MACHINES, AND THESE WERE NOT FOLLOWED IN THIS INSTANCE. PERHAPS I
 DIDN'T MAKE THIS CLEAR ENOUGH IN MY NOTE; BUT I DON'T SEE HOW MY
 NOTE CAN BE READ AS A CRITICISM OF MAKING MEASUREMENTS, /RCC 25a

Network Performance Dialog, 31Jan-5Apr 75

27-MAR-75 1249-EDT IRBY at BBN-TENEXB: Findings for 26-Mar-75
 Distribution: NET PERFORMANCE TECHNICAL GROUP
 [BBN]<MCKENZIE>NPTG.TXT: Received at: 27-MAR-75 10:12:30

26

Daily report of findings for 26-Mar-75

26a

ARPA-TIP HUNG TERMINALS

26b

Several instances of hung terminals were reported yesterday by users trying to access Office-1 and BBNB. The symptom was that they would suddenly be unable to get characters to/from the servicing host. In many cases, the only remedy was to disconnect and reconnect. In other cases, the situation would clear itself after a few minutes. Ben Barker looked into one such hung terminal and found a 0 bit allocation and was able to clear up another user's problem by flushing the TIP buffer. He should explain the nature of these problems to us.

26c

It may be just a coincidence, but the VDH software has been up and down alot in the TYMESHARE TIP since Tuesday, and there have been quite a few anomalies during that period of time, including retransmission of messages that actually got delivered, resulting in duplicate output on terminals.

26d

The Test on the 25th did not happen

26e

Ben also reported that because of a number of difficulties, the test they were trying to run on Tuesday did not actually happen. However, some testing and debugging of data collection mechanisms did take place on Wednesday.

26f

White/Victor Measurements

26g

Jim will release some measurements he and Ken collected regarding TENEX/NCP/NETWORK measurements as described in my memo of yesterday. It should be pointed out that Bob Clements criticized us for making these measurements and that I criticized him for not telling us that they were making similar measurements when we put out an announcement of our intentions a week before we did it. Ken has discussed the situation with Ray Tomlinson. Ray agreed to send us his changes and to look at ours. We would very much like to run Ray's changes since they should be much more complete than ours. It should also be pointed out that Jim White will no longer be working on performance issues but rather is returning to the implementation of the Distributed Programming System. Ken Victor and Jan Kramers will be taking on his measurement work. Martin Hardy and Rene Ortega will be trained to run the measurement programs that now exist.

26h

Network Performance Dialog, 31Jan-5Apr 75

LEVIN's memo reporting problems in the network during a 35 minute period,

261

In a conversation with Ben Barker, I asked if the problems reported in Levin's message were fairly typical. He said they were. This somewhat alarmed me, so I probed some of the events reported more closely. Ben said he should study the message and talk to me later or send out a clarifying message. However, he did explain to me that at least some IMps have been patched to report to NCC whenever they turn their host interface off for more than 1/2 second. This appears to be happening quite a lot. So far, I understand that there are at least two reasons for turning this interface off. One is the sending of multipacket messages that require a subnet allocate and the other is that the host is trying to send a message to another host that is not taking messages fast enough.

26j

This latter event can cause the host-imp interface to be blocked for 15 or more seconds at a time. This was added to the subnet flow control when it was learned (some time ago) that the RFNM mechanism was not adequate because it only controlled things at the link level. Since two hosts could have many links between them, the subnet had to take additional measures to prevent the net from getting clogged and crashing. This is very important since it prevents the sending host from sending to any other host for significant periods of time. The 15 second periods of no response correlates well with user experience of service going away for such periods of time and then catching up suddenly. As I understand it, a change in the host-imp protocol would be required to fix this one. Perhaps the BBN=Network guys could address this in a separate message.

26k

MULTI-PACKET MESSAGES TO TIP/ELF

26l

We have changed buffer sizes in Tips and in our ELF to avoid problems associated with servicing hosts sending multipacket messages to terminals over long network paths. This only effect NLS display terminals, of course, and does not effect XGP's, printers, and other high-baud-rate "terminals" attached to TIPS,

26m

-- Charles,

26n

26-MAR-75 1324-EDT IRBY at BBN=TENEXB: Network Performance --
Findings for 25-Mar-75 Distribution: STONE AT OFFICE-1, MARTINEZ AT
OFFICE-1 Received at: 26-MAR-75 11:33:24

27

Findings for 25-mar-75:

27a

Office-1:

27b

Network Performance Dialog, 31Jan-5Apr 75

It was decided that KEV would not modify his NCP monitoring package for Office-1's 1.31 monitor. Instead we would wait for them to come up on 1.33. In addition, DIA and Dan Lynch should put off changes to the drum/disk stuff until after 1.33 is up at Office-1.

27c

We should put a hardware probe on the IMP-READY-FCR-NEXT-HOST-BIT light on the IMP interface so we can give BBN some REAL numbers for how often it is off. I will discuss this with MEH.

27d

More testing:

27e

BBN ran some more tests yesterday from 10:00 to 2:00 PDT. We have not sent them our superwatch statistics yet and have heard nothing from them. Jim White could not run the two-terminal watchdog program in parallel with the test because the loop-backed terminals were being used in trying to get his program to run on SRI-AIC with KEV's NCP measures. We still need the additional time measure which DIA is to provide, which will tell us when the fork is made runnable, as opposed to when it gets to run. This combination of programs gives us by far the most complete measure of total system performance.

27f

The program sends a character out to a terminal attached to a TIP. This terminal is looped-back so the character comes right back to the job that sent it. On each round trip, we can get the real times that the character is handed to the monitor, when it is handed to the IMP (start and end of transmission to IMP), when NCP gets it back from IMP, when scheduler puts it into line buffer for the program (and sets it runnable), and when the program reads the character. The results of this should be forthcoming in the next day or so.

27g

LEVEN from BBN-Network sent out an interesting log of troubles for the ARPA-TIP in response to complaints from users. There were many problems during the 35 minute period being examined! I will send a copy of the message to anyone who has not seen it and wants to.

27h

There is no new data available for Office-1 behavior. We are turning our attention to BBNB for the time being (until something new happens or 1.33 is up there).

27i

BBNB:

27j

BENCHMARK for NLS

27k

We talked with Dan Lynch to find out how he does

Network Performance Dialog, 31jan-5Apr 75

benchmarks for his lisp users. He uses RUNFIL and does his testing under low-load conditions. It seems we can do the same for NLS but will have to run several simultaneous copies of the benchmark because BBNB does not perform as well as a TENEX with a drum under low load conditions (the claim is that it is competitive under medium to high load conditions).

27l

NCP measurement package

27m

KEV sent off his package to Clements to add to BBNB and perhaps other BBN TENEX's. In addition, I sent a note to Bob asking for his close cooperation in this effort (and for info on other 1.33 measurements that we might be able to use in analyzing NLS behavior at bbnb).

27n

Whether or not this approach works out will determine the need to fly Bob out here.

27o

Drum/Disk measures

27p

Dan Lynch suggested some interesting meters to add to BBNB. I sent them to Clements asking him to add them.

27q

Printer behavior:

27r

BBN is looking into our report that our printer does not run full speed when two IMPs are between it and the sending host.

27s

OTHER:

27t

ELF stability:

27u

Dave Retz, Frank Walden (ADR), and several people from ARC met yesterday to formulate a plan for increasing ELF stability. Joe Ehardt outlined what he believed the main problems were. Dave agreed to spend the next couple of days working on the serious one, Frank agreed to tackle some other problems (with Dave's guidance), and some will be handled by telling users how to avoid them. I left the meeting with the feeling that the crashes and hangs that we now see in ELF will not be occurring in two or three days time. This does not mean that new problems will not arise, however.

27v

Mouse/keyset problems:

27w

The reliability of mouse/keyset interactions, especially in type-ahead, has declined markedly. I personally find that working ahead of the system to be very dangerous now. I believe this to be a combination of two things (at least): 1) a software

Network Performance Dialog, 31Jan-5Apr 75

bug in NLS input routines or in the Line Processor that is causing "no such marker" and other such messages (this could be hardware or buffer overflow also) and 2) the difference in resolution of the mouse on alpha-numeric displays,

27x

This second points warrents further explanation. The coordinates of the mouse are input on both the down and up stroke of the mouse buttons. On a bug-selection, then, only the coordinates of the up stroke are used. On Tasker and IMLAC displays this made no difference since a slight difference in the coordinate value did not generally yield an erroneous selection (each character was several units wide). On the alpha-numeric display, however, characters are one unit wide. The mouse traversal is about the same as for Tasker and IMLAC displays but the user has no feedback to indicate that he is near a character boundary. Thus, any error in the coordinates being input can yield an erroneous selection. It is my belief that the coordinates of the down stroke should be used for selections, since that is the point in time when the user is being precise about the cursor location. This can be achieved in two ways: 1) don't send coordinates with up strokes (this greatly reduces the amount of traffic generated by the user) or 2) fix NLS to use only down-stroke coordinates for selections. To my knowledge, the only functional capability that would be lost by discarding up-stroke coordinates is (I don't think anyone uses this) that when changing viewspecs for multi-window situations, the user would not be able to pick up the viewspecs from one window, modify them slightly and deposit them in another window. This is currently possible because we get the coordinates with the down stroke(s) which tells NLS which window to use for initial viewspecs and we get the coordinates on the up stroke which

27y

tells NLS which window to store the updated viewspecs in.

27z

I would very much like to see RLB2 investigate and remedy this problem at once,

27a@

-- Charles,

27aa

25-MAR-75 2002=EDT VICTOR at BBN-TENEXB; mods to tenex
Distribution: NET PERFORMANCE TECHNICAL GROUP
[BBNB]<ENGELBART>NPTG.TXT: Received at: 25-MAR-75 17:06:05

28

bob clements: the file [sri=ail<133=tenex>impdv.kev contains a measurement package that has been running on the ai ten for a couple of days now, we have just started running jim whites program to extract some of the info, (i suspect jim will be releasing data soon), it would be very nice if we could get these mods implemented at bbnb and perhaps bbna or some other places for

Network Performance Dialog, 31Jan-5Apr 75

some comparison studies, note the addition of two jsies. (this package is a quick and dirty one that works, i feel that a one-two man-month effort would be invovled to implement a clean complete package, i also expect to expand this package for the next coupel of days (week? or so?) but it would be nice if you could pick up this initiail implementation.)

28a

kev

28b

25-MAR-75 1801-EDT LEVIN at BBN-TENEX: ARPA-TIP / OFFICE-1
difficulties Distribution: NET PERFORMANCE TECHNICAL GROUP
[BBN]<MCKENZIE>NPTG,TXT: Received at: 25-MAR-75 15:01:23

29

Checking over the log between 9:15 and 9:50 (EDT) this morning following a complaint by TIP users at ARPA, I found the followings:

29a

1. ARPA was reporting from 0 to 3 traps per minute (average around 1 per minute) that a host was blocked waiting for message number or block for at least 1/2 second. These seemed randomly distributed between the TIP host and the DMS host (during the period that DMS was up) and appeared fairly normal and consistent with the assumption that users were coming on and using varous hosts in the net.

29b

2. Utah's IMP was pretty much screwed up from around 9:40 to 9:50 in its allocation system, with its not going tardy at least once during that period. The OFFICE-1 host appeared to be blocked for 15 seconds several times trying to talk to Utah. Of course such items as allocates to ARPA terminals or echo characters would have to wait each time.

29c

3. Also, OFFICE-1 had incomplete transmission interactions with 37 (RML) and 50 (BBN), and was blocked 15 seconds once talking to RML-TIP. Around this time a line (25) and possibly an IMP (48) in that path were down momentarily.

29d

Fax JBL

29e

25-MAR-75 1615-EDT HARDY at BBN-TENEXB: Three Messages discussing the possibility of the TIP dropping Lineprocessor characters, causing protocol errors, because of LP to TIP input rate. Distribution: IRBY, network-performance-group: Received at: 25-MAR-75 13:19:04

30

24-MAR-75 13:55:17-EDT,774;000000000000 Date: 24 MAR 1975 1355-EDT
From: HARDY Subject: FINBUF To: BARKER at BBN cc: HARDY, IRBY

30a

BEN: LATE LAST WEEK I SENT A MESSSAGE INTENDED FOR YOU SPECIFICALLY TO THE NETWORK PERFORMANCE GROUP, PERHAPS IT DID NOT CATCH YOU EYE.

30b

Network Performance Dialog, 31Jan-5Apr 75

I NOTICED IN TIP MANUAL 1822 IT RECOMMENDED TAYLORING A PORT TO USE "FINBUF" IF THE TERMINAL OPERATED AT INPUT SPEEDS GREATER THAN 3KBPS. OUR LINEPROCESSOR OPERATES AT 4.8KB, INTO TIPS HEAVILY LOADED WITH MANY TERMINALS (ARPA, MITRE, RUTGERS, ETC). IN THESE CASES 1822 IMPLIES WE MAY BE LOSING CHARACTERS, OR GARBING THEM, NOW AND THEN. MY QUESTION IS: IS "FINBUF" OPTION REALLY AVAILABLE? (MALMAN DOES NOT KNOW), AND IF SO, AND WE USE IT, DO SEE ANY ADVERSE EFFECTS ON THE OTHER TIP TERMINALS? ...MARTIN...

HARDY@BBNB ----- 24-MAR-75 14:59:17-EDT,581;000000000000 Mail from BBN-TENEX rcvd at 24-MAR-75 1459-EDT Date: 24 MAR 1975 1454-EDT From: BARKER at BBN-TENEX Subject: FINBUF To: HARDY at BBNB cc: MALMAN, COSELL, MIMNO, BARKER

30c

BASICALLY NO, IT IS NOT AVAILABLE. I WOULD CERTAINLY BE INTERESTED IN WHERE YOU SAW IT RECOMMENDED. THE TIP DOES NOT CURRENTLY AND NEVER HAS USED THAT BIT. TO IMPLEMENT IT WOULD FURTHER DRAIN THE TINY TIP SPACE WHICH IS STILL AVAILABLE, AND PERHAPS MAKE IT EVEN MORE UNUSEABLE. IF THIS IS A PROBLEM, THE SYMPTOM WOULD CLEARLY BE SIMPLY DROPPED CHARACTERS ARE YOU SEEING THAT AS SUCH? -----

25-MAR-75 14:08:02-EDT,582;000000000000 Date: 25 MAR 1975 1408-EDT From: HARDY Subject: FINBUF To: BARKER at BBN cc: HARDY

30d

BEN: THE FINBUFF REFERENCE IS IN 2184 INPUT SECTION 1.4.1 (NOT 1822 AS REFERENCED EARLIER). WE HAVE NO POSITIVE RECORDS OF CHARACTERS BEING DROPPED BY THE TIP, HOWEVER, WE HAVE HAD SEVERAL INSTANCES, USING ARPA AND TYMSHARE TIP, OF PROTOCOL ERRORS. WE SUSPECT DUE TO DROPPED OR GARBLED CHARACTERS SENT TO THE HOST. IT IS NOT CLEAR IF WHO IS AT FAULT: LINEPROCESSOR CODE, NOISY LINE, OR TIP. I WILL PURSUE IN MORE DETAIL AND LET YOU KNOW WHAT I FIND. ...MARTIN..

30e

25-MAR-75 1150-EDT BARKER at BBN-TENEX: EXPERIMENTS Distribution: NET PERFORMANCE TECHNICAL GROUP [BBN]<MCKENZIE>NPTG.TXT: Received at: 25-MAR-75 09:08:13

31

ON THURSDAY, MARCH 20, WE CONDUCTED AN EXPERIMENT TO TRY TO GATHER DATA ABOUT VARIOUS SYSTEMS WHICH MIGHT BE CONTRIBUTING TO GENERAL NETWORK UGLINESS. WE HOPED TO CORRELATE THIS DATA WITH USERS' SUBJECTIVE EVALUATION OF WHEN IT WAS BAD.

31a

MANY PEOPLE GATHERED AND REPORTED USEFUL DATA. WE WOULD LIKE TO EXPRESS OUR APPRECIATION FOR THE EFFORTS OF ALL.

31b

DATA WAS COMPILED ON:

31c

1) THE BUSY-NESS OF VARIOUS TIPS AND IMPS, IN TERMS OF TIME SPENT IN BACKGROUND AS WELL AS NUMBER OF TIP USERS.

31d

2) THE LOAD AVERAGES OF VARIOUS TENEX HOSTS IN THE NET.

31e

Network Performance Dialog, 31Jan-5Apr 75

3) USER-PROGRAM TO USER-PROGRAM ROUND TRIP TIMES BETWEEN
VARIOUS TENEX HOSTS. 31f

4) CAUSES AND FREQUENCIES OF IMPS HOLDING OFF HOST INPUT
LONGER THAN CERTAIN CLIPS. 31g

5) WATCH DATA ON TENEX MONITOR PERFORMANCE 31h

6) ECHO TIME COMPARISONS BETWEEN CHARACTERS LOOPED THROUGH THE
NET AND LOOPED THROUGH A SCANNER. 31i

7) SOME ROUND TRIP TIME SAMPLE STATISTICS 31j

UNFORTUNATELY, WE DID NOT RECEIVE A GREAT DEAL OF USER SUBJECTIVE
EVALUATION INPUT. WHAT LITTLE WE DID RECEIVE DID NOT CORRELATE
PARTICULARLY WELL WITH ANY INTERESTING DATA IN OUR STATISTICS. 31k

CERTAINLY THE MOST NOTICEABLE CORRELATIONS FROM OUR DATA ARE THE
INCREASE IN NET ECHO TIME ABOVE AND BEYOND SCANNER ECHO TIME WITH
THE INCREASE IN SCANNER ECHO TIME AND HOST LOAD AVERAGE, NOTICING
THE HIGH LOAD AVERAGE ON BBNB AND OFFICE-1 31l

OUR CURRENT GUESS IS THAT THE IMP SUBNET IS DOING REASONABLY WELL
GIVEN THE MEMORY RESTRICTIONS, HOST PROTOCOL, AND NETWORK TOPOLOGY
WHICH CURRENTLY EXIST. 31m

THERE ARE KNOWN PROBLEMS WHICH EXISTED AT THE TIME OF THE
EXPERIMENT ON BOTH BBNB AND OFFICE-1 WHICH MAY HAVE CONTRIBUTED TO
SLOWNESS. IN ADDITION TO HIGH LOAD AVERAGES, OFFICE-1 WAS (AND IS)
RUNNING TENEX 1.31, WHICH HAS KNOWN TROUBLE TAKING DATA FROM THE
NET. BBNB WAS RUNNING WITH 64K OF CORE DOWN, AND A VIRTUAL DRUM
WHICH MAY HAVE BEEN TOO SMALL. BOTH OF THE BBNB PROBLEMS ARE FIXED
BY NOW (I THINK). 31n

WE WILL BE CONDUCTING ANOTHER EXPERIMENT THIS AFTERNOON. TOMLINSON
IS ADDING SOME MEASUREMENT TOOLS TO THE TENEX MONITOR. WE WILL
GENERATE TIP TRAFFIC, MEASURE IT THROUGH THE SUBNET, IN AND OUT OF
THE TENEX HOST, AND BACK THROUGH THE SUBNET TO THE TIP, HOPEFULLY
AN ANALYSIS OF THE DATA YIELDED BY THAT EXPERIMENT WILL TELL US
EXACTLY WHAT DELAYS ARE BEING INSERTED BY WHAT COMPONENTS. 31o

ONCE AGAIN, SUBJECTIVE EVALUATIONS ARE IMPORTANT. DATA OF THE FORM
"IT WASN'T SO BAD THIS AFTERNOON" WOULD BE APPRECIATED, DATA OF
THE FORM "AT 1531 I GOT 5 40-SECOND DELAYS IN A ROW" WOULD BE
TERRIFIC. 31p

WE WILL BE MEASURING DELAYS THROUGH BBNB FROM VARIOUS TIPS. THE
EXPERIMENT WILL RUN 1300-1700 TODAY, MARCH 25. 31q

Network Performance Dialog, 31Jan-5Apr 75

/BEN BARKER

31r

25-MAR-75 0950-EDT BURCHFIEL at BBN-TENEXA; BBNB Service
 Interruptions Distribution: WATSON AT BBNB, STROLLO, CHIPMAN, ALLEN,
 net performance technical group [bbn]<mckenzie>nptg.txt:,, burchfiel
 Received at: 25-MAR-75 07:03:56

32

Dick, At the PI meeting, you mentioned problems with service
 interruptions on BBNB in addition to response problems. I've
 collected the following data from the Network Control Center for
 the first half of March:

32a

SYSTEM B SERVICE INTERRUPTIONS MEASURED BY NCC - MARCH 1019, 1975

32b

MARCH	TIME(EDT)	DURATION(min)	COMMENT
3	17:58	10	Reload of IMP
4	4:02	5:24	Scheduled P.M. 18:02
11	TENEX software work		
6	12:14	1	Unknown - unnoticed 16:37
1:38		Reload of IMP 7	11:36 1
Reload of IMP 9	10:46	?	Unknown -
Unnoticed 1-	19:22	27	Reload of IMP
19:58	17	Reload of IMP 11	4:02
3:55	Scheduled P.M. 14	12:42	1
Unknown - unnoticed 17	6:35	3:25	
Installation of 1 Msec clock 18		4:01	4:00
Scheduled P.M. 19	10:17	19	Spurious power
clear of processor			

32c

32d

32e

32f

I would like to hear how these measurements compare with your
 records,

32g

Thanks, Jerry

32h

25-MAR-75 0946-PDT TRAVIS at USC-ISIB: Office=1 Experiment
 Distribution: ENGLEBART AT OFFICE=1, NORTON AT OFFICE=1,, WATSON
 AT OFFICE=1, STONE AT OFFICE=1, HEART AT BBN,, WALDEN AT BBN,
 MCKENZIE AT BBN, SUTHERLAND AT BBN,, BURCHFIEL AT BBN, WALKER AT ISI,
 MCLINDON AT ISI,, RUSSELL AT ISI
 Received at: 25-MAR-75 09:58:51

33

I am sorry to say that the following information is not as you
 wished it to be presented. I have been compiling data as
 requested
 as to the performance of OFFICE=1, but was not aware that it was
 to be sent daily. I hope same is still of some use to you,

33a

Network Performance Dialog, 31Jan-5Apr 75

Sincerely,
Nancy Travis

33b

Date String	Time	Load Average	Jobs	Speed, Char.	
3/10	10:10	5.12, 5.84, 4.95	20	Moderate	
	1:30	8.00, 5.53, 3.88	23	Slow	
3/13	8:40	9.90, 8.34, 7.78	23	Moderate	
	10:00	1.90, 2.15, 2.62	18	Very Slow	
	11:30	5.13, 5.94, 5.30	22	Very Slow	
	2:45	2.13, 2.63, 3.71	18	Moderate	
	5:15	2.49, 2.56, 2.58	19	Very Slow	
3/14	9:15	5.07, 4.17, 5.28	24	Moderate	
	10:02	2.04, 2.63, 1.95	17	Moderate	
	11:45	1.96, 5.00, 5.97	28	Slow	
	3:30	unable to log on, all in use			
	4:10	1.01, 0.64, 0.62	7	Moderate	
3/17	9:30	3.64, 3.40, 3.84	17	Very Slow	
	10:45	unable to log on, all in use			
	11:30	2.90, 2.61, 1.95	19	Slow	
	2:15	3.94, 5.44, 6.74	19	Moderate	
	4:20	3.40, 2.78, 2.07	15	Moderate	
3/18	9:45	7.86, 5.33, 4.49	23	Moderate	
	12:00	unable to log on, all in use			
	1:15	5.31, 6.48, 5.88	24	Moderate	
	4:30	1.43, 1.20, 1.26	14	Very Slow	
	5:30	unable to log on, all in use			
3/19	9:30	2.62, 3.23, 4.22	19	Slow	
	11:15	10.40, 6.44, 5.25	28	Very Slow	
	2:45	0.89, 0.84, 1.08	16	Very Slow	
	4:00	0.31, 0.28, 0.31	11	Moderate	
	5:00	0.79, 0.74, 0.73	17	Very Slow	
3/20	8:45	3.11, 2.39, 2.76	25	Very Slow	
	11:15	4.68, 8.02, 9.24	32	Moderate	
	12:15	9.28, 10.06, 10.30	31	Moderate	
	2:30	3.41, 2.48, 2.82	18	Very Slow	
	4:15	0.70, 0.44, 0.74	16	Moderate	
3/21	9:45	10.05, 8.52, 8.43	22	Moderate	
	10:45	11.57, 8.77, 7.49	26	Very Slow	
	3:15	2.17, 1.46, 2.09	14	Fast	
	4:45	0.84, 0.97, 1.28	12	Very Slow	33d

Again, I am very sorry this is late, I hope you can still find use for the above information.

33e

Nancy

33f

24-MAR-75 0718-PDT WALKER at USC-ISI: SUMMEX VDH

Network Performance Dialog, 31Jan-5Apr 75

Distribution: NETWORK PERFORMANCE GROUP:
 Received at: 24-MAR-75 07:18:06

34

The Summex VDH will go in at the TYMSHARE TIP on 25 March.
 As soon as arrangements can be made an IMP will be installed
 at Summex. Until that time, let us carefully observe the impact of
 the
 VDH on performance at Office-1.

34a

Steve

34b

21-MAR-75 1738-EDT BARKER at BBN-TENEX: LOAD AVERAGES DURING THE
 EXPERIMENT

Distribution: NET PERFORMANCE TECHNICAL GROUP
 [BBN]<MCKENZIE>NPTG.TXT:
 Received at: 21-MAR-75 14:38:52

35

FIRST IN A TABLE, THEN IN A GRAPH.
 NOTE THAT BBNB AND OFFICE-1 ARE CONSISTANTLY HIGH WHERE OTHERS ARE
 CONSISTANTLY LOW.

35a

DATA STARTING THU 20 MAR 75 12:40

35b

---- INDICATES SITE DOWN
 -ND- INDICATES SITE UP/NO DATA

35c

	BBNA	BBNB	BBN	BBND	OFF-1	ISIA	ISIB	ECL
SRI-AI								
1240								
1,83	9,78	1,61	,41	3,46	,48	-ND-	-ND-	
2,02								
1,83	9,78	1,61	,41	3,46	,48	-ND-	-ND-	
2,02								
1,04	6,48	1,82	,28	3,70	1,30	-ND-	-ND-	
2,45								
1,04	6,48	1,82	,28	3,70	1,30	-ND-	-ND-	
2,45								
1,04	6,48	1,82	,28	3,70	1,30	-ND-	-ND-	
2,45								
,99	5,37	1,64	,60	2,75	1,09	-ND-	-ND-	
2,07								
,99	5,37	1,64	,60	2,75	1,09	-ND-	-ND-	
2,07								
,99	5,37	1,64	,60	2,75	1,09	-ND-	-ND-	
2,07								
1,52	5,37	1,67	,70	2,75	,71	-ND-	-ND-	
3,18								
1,52	5,37	1,67	,70	2,75	,71	-ND-	-ND-	

35d

Network Performance Dialog, 31Jan-5Apr 75

3,18								
1300								
1,52	5,37	1,67	,70	2,75	,71	-ND-	-ND-	
3,18								
1,52	5,37	1,67	,70	2,75	,71	-ND-	-ND-	
3,18								
1,46	7,91	1,98	,75	2,96	,53	-ND-	-ND-	
6,45								
1,46	7,91	1,98	,75	2,96	,53	-ND-	-ND-	
6,45								
1,46	7,91	1,98	,75	2,96	,53	-ND-	-ND-	
6,45								
1,89	9,68	1,76	,88	3,40	,49	-ND-	-ND-	
3,42								
1,89	9,68	1,76	,88	3,40	,49	-ND-	-ND-	
3,42								
1,89	9,68	1,76	,88	3,40	,49	-ND-	-ND-	
3,42								
1,12	6,09	1,67	,82	3,40	,62	-ND-	-ND-	
3,41								
1,12	6,09	1,67	,82	4,81	,62	-ND-	-ND-	
3,41								
1320								
1,12	6,09	1,67	,82	4,81	,62	-ND-	-ND-	
3,41								
1,12	6,09	1,67	,82	4,81	,62	-ND-	-ND-	
3,41								
,97	3,79	1,76	,29	4,52	1,12	-ND-	-ND-	
5,55								
,97	3,79	1,76	,29	4,52	1,12	-ND-	-ND-	
5,55								
,97	3,79	1,76	,29	4,52	1,12	-ND-	-ND-	
5,55								
,82	4,63	2,53	,30	4,52	,68	-ND-	-ND-	
7,45								
,82	4,63	2,53	,30	4,52	,68	-ND-	-ND-	
7,45								
,82	4,63	2,53	,30	4,52	,68	-ND-	-ND-	
7,45								
2,15	6,50	2,51	,19	4,52	,43	-ND-	-ND-	
4,30								
2,15	6,50	2,51	,19	12,18	,43	-ND-	-ND-	
4,30								
1340								
2,15	6,50	2,51	,19	12,18	,43	-ND-	-ND-	
4,30								
2,15	6,50	2,51	,19	12,18	,43	-ND-	-ND-	
4,30								

35e

Network Performance Dialog, 31Jan-5Apr 75

2,51	9,24	2,48	,57	10,51	,52	-ND-	-ND-
3,12							
2,51	9,24	2,48	,57	10,51	,52	-ND-	-ND-
3,12							
2,51	9,24	2,48	,57	10,51	,52	-ND-	-ND-
3,12							
3,24	9,24	1,83	,60	7,18	1,15	-ND-	-ND-
2,63							
3,24	9,24	1,83	,60	7,18	1,15	-ND-	-ND-
2,63							
3,24	9,24	1,83	,60	7,18	1,15	-ND-	-ND-
2,63							
3,24	9,24	2,31	,60	7,18	1,15	-ND-	-ND-
3,03							
2,13	7,10	2,31	,43	7,80	1,55	-ND-	-ND-
3,03							
1400							
2,13	7,10	2,31	,43	7,80	1,55	-ND-	-ND-
3,03							
2,13	7,10	2,31	,43	7,80	1,55	-ND-	-ND-
3,03							
2,63	7,10	2,38	,26	7,25	2,20	-ND-	-ND-
2,96							
2,63	7,10	2,38	,26	7,25	2,20	-ND-	-ND-
2,96							
2,63	7,10	2,38	,26	7,25	2,20	-ND-	-ND-
2,96							
2,70	7,98	2,58	,55	11,86	3,23	-ND-	-ND-
6,21							
2,70	7,98	2,58	,55	11,86	3,23	-ND-	-ND-
6,21							
2,70	7,98	2,58	,55	11,86	3,23	-ND-	-ND-
6,21							
2,70	7,98	3,52	,55	11,86	3,23	-ND-	-ND-
4,74							
2,03	8,46	3,52	,47	11,39	2,17	-ND-	-ND-
4,74							
1420							
2,03	8,46	3,52	,47	11,39	2,17	-ND-	-ND-
4,74							
2,03	8,46	3,52	,47	11,39	2,17	-ND-	-ND-
4,74							
1,32	9,62	2,58	,38	10,56	1,46	-ND-	-ND-
2,47							
1,32	9,62	2,58	,38	10,56	1,46	-ND-	-ND-
2,47							
1,32	9,62	2,58	,38	10,56	1,46	-ND-	-ND-
2,47							

35f

Network Performance Dialog, 31Jan-5Apr 75

	.89	9.95	4.19	.25	7.19	2.24	-ND-	-ND-
2,74	.89	9.95	4.19	.25	7.19	2.24	-ND-	-ND-
2,74	.89	9.95	4.19	.25	7.19	2.24	-ND-	-ND-
2,74	.89	9.95	4.19	.25	7.19	2.24	-ND-	-ND-
2,30	.89	9.95	4.19	.25	7.19	2.24	-ND-	-ND-
2,30	1.51	8.13	3.53	1.30	9.81	2.44	-ND-	-ND-
1440								
2,30	1.51	8.13	3.53	1.30	9.81	2.44	-ND-	-ND-
2,30	1.51	8.13	3.53	1.30	9.81	2.44	-ND-	-ND-
2,30	.89	4.65	2.61	1.65	12.24	3.39	-ND-	-ND-
2,52	.89	4.65	2.61	1.65	12.24	3.39	-ND-	-ND-
2,52	.89	4.65	2.61	1.65	12.24	3.39	-ND-	-ND-
2,52	.89	4.65	2.61	1.65	12.24	3.39	-ND-	-ND-
1,65	1.89	3.76	2.61	1.69	11.19	2.55	-ND-	-ND-
1,65	1.89	3.76	2.61	1.69	11.19	2.55	-ND-	-ND-
1,65	1.89	3.76	2.61	1.69	11.19	2.55	-ND-	-ND-
1,65	1.89	3.76	2.61	1.69	11.19	2.55	-ND-	-ND-
1,65	1.89	3.76	2.61	1.69	11.19	2.55	-ND-	-ND-
1,65	3.07	4.19	2.34	.79	8.79	2.30	-ND-	-ND-
1,46								
1500	3.07	4.19	2.34	.79	8.79	2.30	-ND-	-ND-
1,46	3.07	4.19	2.34	.79	8.79	2.30	-ND-	-ND-
1,46	3.25	5.36	2.09	.41	8.79	1.71	-ND-	-ND-
1,91	3.25	5.36	2.09	.41	8.79	1.71	-ND-	-ND-
1,91	3.25	5.36	2.09	.41	8.79	1.71	-ND-	-ND-
1,91								
	3.25	5.16	1.96	.35	8.79	2.08	-ND-	-ND-
1,63	3.25	5.16	1.96	.35	11.58	2.08	-ND-	-ND-
1,63	3.25	5.16	1.96	.35	11.58	2.08	-ND-	-ND-
1,63								

35g

Network Performance Dialog, 31Jan-5Apr 75

3,25	5,16	1,96	,35	11,58	2,08	-ND-	-ND-
1,63							
3,25	5,10	2,87	,30	10,36	2,66	-ND-	-ND-
1,27							
1520							
3,25	5,10	2,87	,30	10,36	2,66	-ND-	-ND-
1,27							
3,25	5,10	2,87	,30	10,36	2,66	-ND-	-ND-
1,27							
3,25	4,85	2,32	,17	10,35	1,88	-ND-	-ND-
2,08							
3,25	4,85	2,32	,17	10,35	1,88	-ND-	-ND-
2,08							
3,25	4,85	2,32	,17	10,35	1,88	-ND-	-ND-
2,08							
3,25	4,33	2,77	,14	10,35	1,88	-ND-	-ND-
1,25							
3,25	4,33	2,77	,14	8,67	1,69	-ND-	-ND-
1,25							
3,25	4,33	2,77	,14	8,67	1,69	-ND-	-ND-
1,25							
3,25	4,33	2,77	,14	8,67	1,69	-ND-	-ND-
1,25							
3,25	4,33	3,39	,15	13,00	,73	-ND-	-ND-
1,54							
1540							
3,25	4,33	3,39	,15	13,00	,73	-ND-	-ND-
1,54							
3,25	4,33	3,39	,15	13,00	,73	-ND-	-ND-
1,54							
,89	6,19	5,02	,17	13,00	,87	-ND-	-ND-
1,10							
,89	6,19	5,02	,17	10,27	,87	-ND-	-ND-
1,10							
,89	6,19	5,02	,17	10,27	,87	-ND-	-ND-
1,10							
,89	6,19	5,02	,17	10,27	,87	-ND-	-ND-
1,10							
,89	6,95	5,30	,36	9,53	1,44	-ND-	-ND-
1,62							
,89	6,95	5,30	,36	9,53	1,44	-ND-	-ND-
1,62							
,89	6,95	5,30	,36	9,53	1,44	-ND-	-ND-
1,62							
1,08	5,53	3,29	,85	9,76	1,42	-ND-	-ND-
1,73							
1601							
1,08	5,53	3,29	,85	9,76	1,42	-ND-	-ND-

35h

Network Performance Dialog, 31Jan-5Apr 75

1,73								
1,08	5,53	3,29	,85	9,76	1,42	-ND-	-ND-	
1,73								
1,08	5,53	3,29	,85	9,76	1,42	----	-ND-	
1,73								
2,85	7,28	2,44	1,67	7,23	1,14	-ND-	-ND-	
1,30								
2,85	7,28	2,44	1,67	7,23	1,14	-ND-	-ND-	
1,30								
2,85	7,28	2,44	1,67	7,23	1,14	-ND-	-ND-	
1,30								
2,85	5,15	3,51	1,96	5,42	,64	-ND-	-ND-	
,79								
2,85	5,15	3,51	1,96	5,42	,64	-ND-	-ND-	
,79								
2,85	5,15	3,51	1,96	5,42	,64	-ND-	-ND-	
,79								
2,85	5,11	4,46	1,96	5,42	,64	-ND-	-ND-	
1,86								
1621								
2,85	5,11	4,46	2,73	5,42	,40	-ND-	-ND-	
1,86								
2,85	5,11	4,46	2,73	5,42	,40	-ND-	-ND-	
1,86								
2,85	5,11	4,46	2,73	5,42	,40	-ND-	-ND-	
1,86								
2,85	5,50	4,36	1,49	10,42	,40	-ND-	-ND-	
1,17								
2,85	5,50	4,36	1,49	10,42	,40	-ND-	-ND-	
1,17								
2,85	5,50	4,36	1,49	10,42	,40	-ND-	-ND-	
1,17								
2,85	5,50	4,29	1,34	11,03	,81	-ND-	-ND-	
1,68								
2,85	5,50	4,29	1,34	11,03	,81	-ND-	-ND-	
1,68								
2,85	5,50	4,29	1,34	11,03	,81	-ND-	-ND-	
1,68								
2,85	5,50	3,00	1,34	11,03	,81	-ND-	-ND-	
,92								
1641								
2,85	4,31	3,00	,77	11,03	1,01	-ND-	-ND-	
,92								
2,85	4,31	3,00	,77	11,03	1,01	-ND-	-ND-	
,92								
2,85	4,31	3,00	,77	11,03	1,01	-ND-	-ND-	
,92								
3,31	5,39	3,26	,24	10,83	2,15	-ND-	-ND-	
1,17								

Network Performance Dialog, 31Jan-5Apr 75

1,17	3,31	5,39	3,26	,24	10,83	2,15	-ND-	-ND-	351
1,17	3,31	5,39	3,26	,24	10,83	2,15	-ND-	-ND-	
1,56	3,85	5,39	3,36	,22	7,28	1,73	-ND-	-ND-	
1,56	3,85	5,39	3,36	,22	7,28	1,73	-ND-	-ND-	
1,56	3,85	5,39	3,36	,22	7,28	1,73	-ND-	-ND-	
1,56	3,85	5,39	3,36	,22	7,28	1,73	-ND-	-ND-	
1,56	3,85	5,39	3,36	,22	7,28	1,73	-ND-	-ND-	
1,69	4,55	5,39	3,06	1,10	5,36	,94	-ND-	-ND-	
1,69	4,55	5,39	3,06	1,10	5,36	,94	-ND-	-ND-	
1,69	4,55	5,39	3,06	1,10	5,36	,94	-ND-	-ND-	
1,69	4,55	5,39	3,06	1,10	5,36	,94	-ND-	-ND-	
1,41	4,22	8,00	2,41	,96	3,27	,57	-ND-	-ND-	
1,41	4,22	8,00	2,41	,96	3,27	,57	-ND-	-ND-	
1,41	4,22	8,00	2,41	,96	3,27	,57	-ND-	-ND-	
1,41	4,22	8,00	2,41	,96	3,27	,57	-ND-	-ND-	
3,75	2,10	4,97	2,21	,66	4,16	1,65	-ND-	-ND-	
3,75	2,10	4,97	2,21	,66	4,16	1,65	-ND-	-ND-	
3,75	2,10	4,97	2,21	,66	4,16	1,65	-ND-	-ND-	
3,75	2,10	4,97	2,21	,66	4,16	1,65	-ND-	-ND-	
1721	3,87	5,38	4,40	,35	4,16	1,66	-ND-	-ND-	
4,75	3,87	5,38	4,40	,35	4,16	1,66	-ND-	-ND-	
4,75	3,87	5,38	4,40	,35	4,16	1,66	-ND-	-ND-	
4,75	3,87	5,38	4,40	,35	4,16	1,66	-ND-	-ND-	
4,10	3,87	5,40	4,95	,61	3,50	1,36	-ND-	-ND-	
4,10	3,87	5,40	4,95	,61	3,50	1,36	-ND-	-ND-	
4,10	3,87	5,40	4,95	,61	3,50	1,36	-ND-	-ND-	
2,71	6,01	4,02	4,27	,62	3,50	1,36	-ND-	-ND-	

Network Performance Dialog, 31Jan-5Apr 75

6,01	4,02	4,27	,62	2,60	1,20	-ND-	-ND-
2,71							
6,01	4,02	4,27	,62	2,60	1,20	-ND-	-ND-
2,71							
6,01	4,02	4,27	,62	2,60	1,20	-ND-	-ND-
2,71							
1741							
4,79	3,76	4,08	,46	1,95	1,74	-ND-	-ND-
1,69							
4,79	3,76	4,08	,46	1,95	1,74	-ND-	-ND-
1,69							
4,79	3,76	4,08	,46	1,95	1,74	-ND-	-ND-
1,69							
6,80	4,13	4,03	,37	2,33	1,34	-ND-	-ND-
2,02							
6,80	4,13	4,03	,37	2,33	1,34	-ND-	-ND-
2,02							
6,80	4,13	4,03	,37	2,33	1,34	-ND-	-ND-
2,02							
6,95	3,03	3,80	,54	2,33	1,34	-ND-	-ND-
4,68							
6,95	3,03	3,80	,54	2,34	1,42	-ND-	-ND-
4,68							
6,95	3,03	3,80	,54	2,34	1,42	-ND-	-ND-
4,68							
6,95	3,03	3,80	,54	2,34	1,42	-ND-	-ND-
4,68							
1801							
4,60	4,30	3,24	,36	1,87	2,29	-ND-	-ND-
4,30							
4,60	4,30	3,24	,36	1,87	2,29	-ND-	-ND-
4,30							
4,60	4,30	3,24	,36	1,87	2,29	-ND-	-ND-
4,30							
4,14	4,04	2,50	,21	1,61	2,16	-ND-	-ND-
3,98							
4,14	4,04	2,50	,21	1,61	2,16	-ND-	-ND-
3,98							
4,14	4,04	2,50	,21	1,61	2,16	-ND-	-ND-
3,98							
4,07	4,42	1,91	,21	1,69	2,21	-ND-	-ND-
4,37							
4,07	4,42	1,91	,21	1,69	2,21	-ND-	-ND-
4,37							
4,07	4,42	1,91	,21	1,69	2,21	-ND-	-ND-
4,37							
4,07	4,42	1,91	,21	1,69	2,21	-ND-	-ND-
4,37							

351

Network Performance Dialog, 31Jan-5Apr 75

1821	4.21	4.28	2.66	.97	1.84	2.62	-ND-	-ND-
5.84	4.21	4.28	2.66	.97	1.84	2.62	-ND-	-ND-
5.84	4.21	4.28	2.66	.97	1.84	2.62	-ND-	-ND-
5.84	4.98	4.23	1.47	.53	.89	1.28	-ND-	-ND-
7.78	4.98	4.23	1.47	.53	.89	1.28	-ND-	-ND-
7.78	4.98	4.23	1.47	.53	.89	1.28	-ND-	-ND-
7.78	4.98	4.23	1.47	.53	.89	1.28	-ND-	-ND-
6.24	6.03	3.07	1.99	.95	1.00	.85	-ND-	-ND-
6.24	6.03	3.07	1.99	.95	1.00	.85	-ND-	-ND-
6.24	6.03	3.07	1.99	.95	1.00	.85	-ND-	-ND-
6.24	6.03	3.07	1.99	.95	1.00	.85	-ND-	-ND-

35K

351

DATA STARTING THU 20 MAR 75 12:40


35m

```

+ BBNA
X BBNB
. BBN
* BBND
# OFF=1
= ISIA
= ISIB
! ECL
- SRI=AI

```

35n



350

35p

35r

36

36a

36b

36c

37

37a

Network Performance Dialog, 31Jan-5Apr 75

21-MAR-75 1044-EDT MCKENZIE at BBN-TENEX: Comments on
 Lineprocessor/TIP Interaction Distribution: NET PERFORMANCE
 TECHNICAL GROUP [BBN]<MCKENZIE>NPTG,TXI: Received at: 21-MAR-75
 07:44:47

38

There has been much discussion recently about the problems which the Tymshare TIP (and perhaps other TIPs) has been having with the 4800 baud modems used by the "lineprocessors". This is an attempt to clear up a few misconceptions about what does and does not happen.

38a

The modem/lineprocessor combination, as currently implemented, does not use any "data terminal ready" signal to tell the TIP when to listen to the terminal. Apparently the modem could convey this signal to the TIP, but the lineprocessor was not implemented with this in mind; therefore the modem at the terminal end has the "data terminal ready" signal wired on. We have asked the SRI people to correct this design, and we believe they are in the process of beginning to do so.

38b

The Tymshare TIP was recently fixed (as part of a general TIP retrofit program) to ignore a steady stream of "breaks" coming from a terminal. This fix, however, was designed under the assumption that "externally clocked" devices would only send constant "breaks" if they really meant them. Unfortunately, the modems used with the lineprocessor fall into the category of "externally clocked" devices. We are now reviewing our design to try to find a way of discarding (in hardware) breaks from such devices.

38c

Even in the current case, with a somewhat deficient design of the TIP hardware to discard multiple breaks, and a somewhat deficient design of the lineprocessor to not generate "data terminal ready", we believe that essentially no TIP bandwidth is consumed by the breaks IFF there is no open network connection (to a Host) from such a device. We believe that a significant fraction of TIP bandwidth is consumed in the case where the user fails to close (i.e. <cr>) his TIP connection before turning off his lineprocessor. We suspect that this was happening frequently at Tymshare, since the reason this problem first came to light was Postel's observation that traffic generated by the Tymshare TIP seemed abnormally high.

38d

It appears that the lineprocessor(s) connected to the ARPA TIP are not using modems and do provide a "data terminal ready" signal to the TIP. The TIP examines this signal every few seconds. Thus, when a lineprocessor is turned off, the TIP sees a constant stream of breaks (the lineprocessor is still an "externally clocked" device) until the next time the TIP probes the "ready" signal, at

Network Performance Dialog, 31Jan-5Apr 75

which point further input is ignored. This is consistent with the findings Hardy reported in his message of 20 March 1946-EDT,

38e

Regards, Alex

38f

21-MAR-75 1017-EDT BURCHFIEL at BBN-TENEXA: Potential source of long blockages of host-IMP interface Distribution: ENGELBART AT OFFICE-1, NORTON AT OFFICE-1,, WATSON AT OFFICE-1, HEART AT BBN, WALDEN AT BBN, SUTHERLAND,, BURCHFIEL, MCLINDON AT ISI, CARLSON AT ISI, BLUE AT ISI,, STUBBS AT ISI, LICKLIDER AT ISI, RUSSELL AT ISI,, WALKER AT ISI, STONE AT OFFICE-1, STROLLO Received at: 21-MAR-75 07:36:35

39

In the past, we have observed TENEX blocking the input interface from the IMP for long (multi-second) periods. This used to happen when the NCP fork would hang on the lock IDVLCK while the fork which set the lock is low queued by the scheduler,

39a

Patches were distributed for TENEX system 1.32 and TENEX 1.33 correcting this problem by high queueing IDVLCK lockers. This symptom has not been observed on BBN System B. Office-1 should be examined for this problem, and if it is occurring, the above upgrade of the system code should be considered.

39b

Thanks, Jerry (BURCHFIEL@BBN)

39c

21-MAR-75 1014-PDT PETERS at OFFICE-1: Office-1 statics during NET testing Distribution: WALKER AT ISI, MCLINDON AT ISI, CARLSON AT ISI,, KLEINROCK AT ISI, STONE AT OFFICE-1, NORTON AT OFFICE-1,, HEART AT BBN, WALDEN AT BBN, MCKENZIE AT BBN,, BURCHFIEL AT BBN, STROLLO AT BBN, CLEMENTS AT BBN,, ENGELBART AT BBNB, WATSON AT BBNB, IRBY AT BBNB Received at: 21-MAR-75 10:15:11

40

Some pertinent OFFICE-1 statistics are available at [OFFICE-1]<NICGUEST>3/20/afternoon.txt - these were collected while BBN and ARC were collecting information on net performance, blocking etc. More info can be extracted from the same data base if desired. Column headings and meanings are: IDL: per cent of time in idle GU: ave. number of "go jobs" (load, but NOT taken from exponential ave) %U: per cent used time IQW: per cent I/O wait time %DB: per cent of time IMPNFI>0 based on 50 ms. samples (imp buffers available) DQL: ave. size of IMPNFI when it WAS GREATER THAN ZERO. DKR, DKW: Disk reads and writes per second, NOT ACCURATE (too small) %KB: per cent of time disk queue length greater than zero KPT: disk page transfer time (ms). NOT ACCURATE (too large) KQL: ave. disk queue length when it was GREATER THAN ZERO HR, MN, SC: PDT according to OFFICE-1 TENEX

40a

Don Andrews (ANDREWS at BBNB)

40b

Network Performance Dialog, 31Jan-5Apr 75

20-MAR-75 1959-EDT HARDY at BBN-TENEX: TIP "FAST INPUT BUFFER" FOR LINEPROCESSORS, AND ARPA TIP LP BREAK TESTS. Distribution: NETWORK-PERFORMANCE-GROUP: Received at: 20-MAR-75 16:59:44

41

BEN: IN READING BBN REPORT 2184 (TIP HARDWARE MANUAL) SECTION 1.4.1 (INPUT), IT APPEARS THAT WE SHOULD HAVE EACH LINEPROCESSOR PORT TAYLORED TO USE "FINBUF" (FAST INPUT BUFFER) INSTEAD OF "INBUF" AS THEY DO NOW. LINEPROCESSOR CONNECTIONS ARE TYPICALLY 4800 BAUD. THE TIP MANUAL RECOMMENDS USING FINBUF FOR CONNECTION WITH LINE SPEEDS GREATER THAN 3000 BPS TO INSURE TRANSFER TO 316 MEMORY BEFORE NEXT CHARACTER IS RECIEVED. IS THIS A REASONABLE THING TO DO? USING INBUF FOR LINEPROCESSOR CONNECTIONS AT TIPS WITH MANT TERMINALS CONNECTED COULD EXPLAIN SOME OF THE NOISE CHARACTERS RECIEVED BY THE HOST.

41a

ARPA TIP LP BREAK TESTS

41b

ALL THE LINEPROCESSORS INSTALLED AT ARPA USE THE DTR (DATA TERMINAL READY) FUNCTION. TODAY A TEST WAS PERFORMED BY KEYDATA (COMPANY WHO MAINTAINS ARPA TERMINAL SYSTEMS). TEST SHOWED THE FOLOWING: 1) WITH DTR ON, AND WITH DTR OFF POWER OFF ON LP CAUSED TIP BRK LIGHT TO COME ON FOR 2 TO 5 SEC.

41c

2) DTR OFF FORCED CONNECTION SUSPENSION FROM THE HHST (LEAVING A DETACHED JOB). P.S ALL LINEPROCESSORS AT ARPA ARE CONNECTED TO THE TIP HARDWIRED (NO MODEM). ..MARTIN.. HARDY

41d

20-MAR-75 1952-EDT WALDEN at BBN-TENEX: TENEX MEASUREMENTS Distribution: VICTOR AT BBNB, net performance technical group [bbn]<mckenzie>nptg.txt: Received at: 20-MAR-75 17:35:49

42

YOUR SUGGESTED MEASUREMENTS ONLY LOOK AT THE NCP'S INTERFACE TO THE NETWORK. I SUGGEST YOU ALSO ADD MEASUREMENTS TO LOOK AT THE NCP INTERFACE TO THE REST OF TENEX AND TO LOOK AT THE INTERNALS OF THE NCP. FOR INSTANCE, THE FOLLOWING SORT OF INFORMATION WOULD SEEM USEFUL TO ME:

42a

-WHAT IS THE AVERAGE SEARCH TIME THROUGH THE CONNECTION HASH TABLE -- IS IT SOMETIMES VERY BIG? -- ARE HASH TABLE SEARCHES SOMETIMES LOCKED OUT FOR A PERIOD AS WHEN THE HASH TABLE IS BEING GARBAGE COLLECTED? -- WHAT IS THE AVERAGE NUMBER OF ENTRIES IN THE HASH TABLE?

42b

-HOW MANY COMMANDS DOES TENEX NORMALLY PACK INTO ONE MESSAGE ON THE CONTROL LINK?

42c

-WHAT SHARE OF THE MACHINE DOES THE NCP GET?

42d

Network Performance Dialog, 31Jan-5Apr 75

-HOW MUCH OF THE TIME IS THE NCP WAITING FOR THE PROCESSOR OR FOR THE DRUM OR FOR PAGES? 42e

-CAN TOO MANY CALLS ON THE NCP BY USERS PROCESSES CAUSE IT TO SPIN ITS WHEELS? 42f

-THE AVERAGE MESSAGE SIZE TENEX SENDS TO THE NETWORK AND THE AVERAGE MESSAGE SIZE USERS SEND TO THE NCP? 42g

-WHAT PROPORTION OF THE TIME DOES TENEX HAVE A ZERO MESSAGE OR BIT ALLOCATE FROM A CONNECTION. 42h

-WHAT IS THE MIX OF OTHER HOSTS TO WHICH TENEX IS COMMUNICATING? 42i

-HOW OFTEN IS TENEX GETTING "IMBUGS". 42j

- HOW MUCH OF THE TIME IS TENEX TRYING TO RETRANSMIT A MESSAGE AND DOES THIS BLOCK OTHER WORK? 42k

-DOES PRINTING AN IMBUG HANG UP THE NCP SO I CAN'T DO OTHER WORK AND IF IT IS THOUGHT THAT IIS DOES NOT, WHY HAS IT LOOKED LIKE IT DID EVVRY TIME I HAVE LOOKED IN THE PAST? 42l

-IS THERE SOME SIZE OR FREQUENCY OF MESSAGES RECEIVED WHICH CAN DEGRAD THE NCP'S PERFORMANCE, FOR INSTANCE TOO MNAY SHORT MESSAGE FRAGMENTING ITS INPUT BUFFER SPACE. 42m

-HOW MUCH OF THE TIME CAN'T TENEX FIND A BUFFER THE RIGHT SIZE INTO WHICH TO PUT AN INCOMING MESSAGE. 42n

-IF INCOMING MESSAGES MUST EVENTUALLY BE PUT INTO USER SPACE, HOW MUCH TIME IS SPENT WAITING FOR THIS SPACE TO BECOME AVAILABLE. 42o

-HOW MUCH OF THE TIME DOES TENEX'S ALLOCATE FOR OTHER HOSTS GO TO ZERO? 42p

-ETC. 42q

SOME OF THE ABOVE SUGGESTIONS MAY NOT MAKE MUCH SENSE BUT THAT IS BECAUSE I DON'T REALLY UNDERSTAND HOW TENEX OR ITS NCP WORKS? NONETHELESS, ASKING YOURSELF THAT SORT OF QUESTIONS AND TRYING TO DO MEASUREMENTS TO GET THE ANSWERS WOULD SEEM USEFUL TO ME. 42r

REGARDS, DAVE 42s

20-MAR-75 1934-EDT WALDEN at BBN-TENEX: OFFICE-1 HOST TARDY TAKING DATA FROM IMP Distribution; NET PERFORMANCE TECHNICAL GROUP [BBN]<MCKENZIE>NPTG.TXT: Received at: 20-MAR-75 17:13:22

43

Network Performance Dialog, 31Jan-5Apr 75

TO IRBY:

43a

SORRY CHARLES, BUT I MOMENTARILY FORGOT YOUR MENTION OF THE TENEX 1.31 INPUT BUFFER PROBLEM IN YOUR LONG SUMMARY MESSAGE. WE ARE, OF COURSE, WILLING TO LOOK INSIDE IMP WHILE YOU LOOK INSIDE TENEX, LET US KNOW WHEN YOU WANT US TO LOOK, DAVE

43b

20-MAR-75 1545-EDT VICTOR at BBN-TENEXB: tenex measurements
Distribution: NETWORK PERFORMANCE GROUP; Received at: 20-MAR-75
12:46:19

44

i am in the process of specifying changes to tenex that will enable us to get statistics with respect to host-imp and imp-host interactions, i plan to bring up these changes at sri-ai in the next couple of days; at office-1 when it comes up on 133; and hopefully on bbnb if the measurements prove interesting,

44a

this memo specifies what measurements i plan on taking with regard to host-imp interactions, i would appreciate any feedback as to whether people think these are the right measurements to be taking, in particular if anyone with tenex knowledge knows if any of these measurements are currently being made, or if there is a better way of making than the way i intend to use (get in touch with me for details) i would like to be so informed,

44b

we will be able to measure all host-imp messages, no host-imp messages, or host-imp messages for a specific host-link pair,

44c

the following meters will be kept (a meter is a constantly increasing number; statistics can and will be made by sampling these meters):

44d

number of messages sent

44e

number of high priority messages sent

44f

number of multi-packet messages sent

44g

number of regular messages sent

44h

number of regular messages sent on link 0

44i

number of regular messages with byte size other than 8 bits

44j

number of multi-packet regular messages with byte size other than 8 bits

44k

number of regular messages with 8-bit byte size

44l

Network Performance Dialog, 31Jan-5Apr 75

number of 8-bit byte size regular messages with byte count = 1 44m

number of 8-bit byte size multi-packet regular messages 44n

number of bits divided by 36 sent 44o

time to send the 32-bit leader 44p

time to send entire messages 44q

time from sending end of output until sending leader 44r

in addition to the above meters i will bump a counter every time a message is place on theimps output queue and decrement this counter when tenex sends the leader of a message. 44s

(note that at times it appears that tenex has a message complete, but does not place it on theimps output queues because there are still outstanding rfnms on the particular link. i dont fully understand this yet, but am working on it.) 44t

20-MAR-75 1317-PDT LYNCH at SRI-AI: TENEX TUNING PARAMETERS WRITEUP

Distribution: NET PERFORMANCE TECHNICAL GROUP

[BBN]<MCKENZIE>NPTG,TXT; Received at: 20-MAR-75 13:16:20 45

I have prepared a 12 page document on what I have been able to find in 133 TENEX that can be considered to be ways to change the performance of TENEX as it treats it various resources and demands, it is called [sri-ai]<lynch>tenex . The anonymous FTP login feature works here, so any one can get the file to read it, I would appreciate and feedback/corrections/additions that any of you may have. I would like to have the document grow in content and usefulness, Dan Lynch 45a

20-MAR-75 1246-EDT IRBY at BBN-TENEXB: host tardy in taking data from imp Distribution: WALDEN AT BBN, walker at isi, mclindon at isi, carlson at isi, kleinrock at isi, heart at bbn, walden at bbn,, mckenzie at bbn, burchfiel at bbn, strollo at bbn,, clements at bbn, irby at bbn, hardy at bbn, stone at office-1,, engelbart at office-1, norton at office-1, watson at office-1,, martinez at office-1, lynch at sri-aic, victor at sri-aic Received at: 20-MAR-75 10:21:48 46

Dave, yes we have investigated this. As I pointed out in my "findings" memo the TENEX NCP at Office-1 sometimes runs out of input buufers for the NET and turns the IMP off. Att times this was observed to happen alot (like 50% of the time), BBN fixed this in 1.33 TENEX by tripling the number of buffers reserved for this purpose and running the NCP more frequently. 46a

Network Performance Dialog, 31Jan-5Apr 75

If you are willing to help, we would like to try patching Office-1's 1.31 TENEX to do what 1.33 does and see if the problem goes away. We can measure inside the NCP to see if it is still turning the IMP off while you guys are measuring from the IMP side.

46b

-- Charles, p.s. -1 will move to 1.33 within the next few weeks,

46c

20-MAR-75 1032-EDT MCKENZIE at BBN-TENEX: Mailing List
Distribution: NET PERFORMANCE TECHNICAL GROUP
[BBN]<MCKENZIE>NPTG,TXT: Received at: 20-MAR-75 10:29:53

47

The title gives the pathname for a composite list including all addressees suggested by Engelbart and all those suggested by Walden. The address NAC@BBN is used for Mario Gerla, Alex McKenzie

47a

19-MAR-75 2130-EDT ENGELBART at BBN-TENEXB: Distribution for Net Performance Technical Notes Distribution: NET PERFORMANCE TECHNICAL GROUP [BBNB]<ENGELBART>NPTG,TXT: Received at: 19-MAR-75 18:30:19

48

I would like to propose that Walker make "official" a distribution list for technical-note exchanges on the current multi-party effort regarding Net performance. I propose the following list for his consideration:

48a

@isi, walker, mclindon, carlson, kleinrock, @bbn, heart, walden, mckenzie, burchfiel, strollo, clements, irby, hardy, @office-1, stone, engelbart, norton, watson, martinez, @SRI-AIC, lynch

48b

This alters the original list in the following way:

48c

Removes Licklider, Russell, Blue and Stubbs, and adds Kleinrock, as per Walker's 16 March memo;

48d

Removes Sutherland who told me that he is phasing out at BBN, and adds Strollo and Clements with whom we've been communicating fairly heavily on BBNB-TENEX (and general-TENEX) matters;

48e

Adds Irby and Hardy at SRI ARC here who respectively are coordinating software studies and hardware (plus interface to Martinez);

48f

Adds Martinez at TYMSHARE who is responsible for OFFICE-1 operations there;

48g

And adds Lynch at SRI-AIC, their TENEX king who has become heavily involved with us in these studies,

48h

Network Performance Dialog, 31Jan-5Apr 75

Also, it would be nice if more of us learned how to use the SNDMSG distribution-specification feature that provides what Walker's 16-March memo did, i.e. where the recipient got only "Distribution: NETWORK PERFORMANCE GROUP:" (That's what I'm going to try on this note, to the above list,)

481

19-MAR-75 1722-EDT WALDEN at BBN-TENEXB; TIP ALLOCATES STRATEGY AND BUFFERING Distribution; WALKER AT ISI, MCLINDON AT ISI, CARLSON AT ISI,, KLEINROCK AT ISI, HEART AT BBN, WALDEN AT BBN,, MCKENZIE AT BBN, SUTHERLAND AT BBN, BURCHFIELD AT BBN,, STROLLO AT BBN, STONE AT OFFICE-1, ENGELBART AT OFFICE-1,, NORTON AT OFFICE-1, WATSON AT OFFICE-1, IRBY AT BBNB,, LEVIN AT BBNB, BARKER AT BBNB, SANTOS AT BBNB,, BUTTERFIELD AT BBNB, COSELL AT BBNB, MIMNO AT BBNB, mcquillan Received at: 19-MAR-75 14:29:38

49

CHARLES IRBY,

49a

1. WE DO THE TIP ALLOCATION THE WAY WE DO BECAUSE IT IS EASY THIS WAY -- DOUBLE BUFFERING IS SOMEWHAT SIMPLER THAN A CIRCULAR BUFFER WHICH IS WHAT YOU HAVE IN MIND AS THE ALTERNATIVE.

49b

2. OUR ASSUMPTION IS THAT THE SENDING HOST, IF TRYING FOR HIGH THROUGHPUT, WHICH IS WHAT YOU WANT WITH YOUR DISPLAYS, WILL SEND MESSAGES FULL TO THE ALLOWABLE BIT ALLOCATIONS -- THIS SHOULD BE NO BIG PROBLEM GIVEN THE SMALL SIZE OF THE TIP BUFFERS (I REALIZE THAT TENEX MAY NOT ALLOW THE USER TO CONTROL HOW FULL MESSAGES ARE AND TO THIS EXTENT OUR ASSUMPTION IS NOT REALISTIC).

49c

3. TO GO TO A CIRCULAR BUFFER SCHEME PROBABLY ADDS SOME WORDS OF CODE (MAYBE LESS THAN 100). HOWEVER, SOME OF THESE WORDS ARE IN THE INNER LOOP (I.E., END TESTS ARE HARDER IN A CIRCULAR BUFFER SCHEME WHERE THE BOUNDARIES BETWEEN MESSAGES VARIES AROUND IN THE BUFFER AND WHERE ONE HAS TO BE CONCERNED ABOUT WRAPING AROUND THE END OF THE BUFFER. WORSE, GOING TO A CIRCULAR BUFFER SCHEME MEANS ONE HAS REMEMBER EXACTLY HOW MANY OUTSTANDING ALLOCATES SO THERE ARE RETRANSMISSION OF ALLOCATES CAN BE DONE CORRECTLY -- THIS WILL REQUIRE AT LEAST A NEW 64 WORDS TIMES 11 BITS PER WORD TABLE TO HOLD THE ADDITIONAL BIT ALLOCATE MEMORY.

49d

4. GIVEN THE SMALL SIZE OF THE TIP BUFFERS ALREADY, IT IS NOT CLEAR THAT THE ADDITIONAL LOSS OF MEMORY TO GO TO A CIRCULAR BUFFER SCHEME WOULD BE COMPENSATED FOR BY THE LESS BREAKAGE THAT A CIRCULAR BUFFER SCHEME OFFERS OVER A DOUBLE SCHEME.

49e

5. FINALLY, GIVEN SUFFICIENT BUFFERING, IN THE TIP AND A HOST PUSHING AS HARD AS ALLOWED, DOUBLE BUFRING SHOULD BE ABLE TO COVER ANY NETWORK (AND OTHER) DELAYS. I THINK THE FUNDAMENTAL PROBLEM, GIVEN THE LARGE DELAYS ACROSS THE NET AND THROUGH THE HOSTS, IS LACK OF MEMORY, NOT THE ALLOCATION STRATEGY. ANY

Network Performance Dialog, 31Jan-5Apr 75

ALLOCATION STRATEGY WOULD WORK POORLY WITH THE SMALL AMOUNT OF BUFFERING AVAILABLE ON MANY TIPS. ONCE AGAIN I URGE YOU TO GET MORE MEMORY FOR THE TYMSHARE TIP AND FOR ARPA (OR THE TIP OWNERS) TO GET MORE MEMORY FOR THE REST OF THE TIPS. ADDITIONALLY, WE MUST ALL CONTINUE OUR QUEST TO REDUCE THE HOST/NET/TIP/ELF/IMP DELAYS. NONE OF OUR SYSTEMS WERE OPTIMIZED TO WORK WITH SUCH DELAYS.

49f

REGARDS, DAVE

49g

19-MAR-75 1446-EDT IRBY at BBN-TENEX: TIP allocation policy question Distribution: WALDEN AT BBNA, walker at isi, mclindon at isi, carlson at isi,, Kleinrock at isi, heart at bbn, walden at bbn,, mckenzie at bbn, sutherland at bbn, burchfiel at bbn,, strollo at bbn, stone at office-1, engelbart at office-1,, norton at office-1, watson at office-1, irby at bbn Received at: 19-MAR-75 11:55:12

50

Dave, If we understand the TIP allocation policy correctly, it seems highly likely that TENEX will use up the message allocate before the bit allocation and thus be unable to output to the terminal until the TIP switches buffers and sends a new allocate. We wanted to make sure we understood the policy so we know how seriously to pursue timing measurements, etc. to see if this is hurting us for long net paths and/or high speed terminals. We are not accusing you of doing anything in a stupid or sub optimal way. We know you have thought about this alot. We are just asking the obvious question to clarify our understanding and perhaps to uncover some other design assumptions that are getting in the way of our users.

50a

-- Charles,

50b

19-MAR-75 0948-EDT WALDEN at BBN-TENEX: ADDITIONAL COMMENT Distribution: NETWORK PERFORMANCE GROUP:, ENGELBART AT OFFICE-1,, NORTON AT OFFICE-1, WATSON AT OFFICE-1, STONE AT OFFICE-1,, HEART, WALDEN, MCKENZIE, SUTHERLAND, BURCHFIELD, STROLLO,, MCLINDON AT ISI, CARLSON AT ISI, KLEINROCK AT ISI,, WALKER AT ISI Received at: 19-MAR-75 09:57:52

51

IN IRBY'S MESSAGE OF 17 MARCH (0141 EDT), HE DISCUSSED THE 1/4 SECOND TIMEOUT OF THE IMP LEVEL ALLOCATE AT THE SOURCE IMP. IT HAS BEEN POINTED OUT TO ME BY MCQUILLAN THAT CHARLES DESCRIBED THE MECHANISM SLIGHTLY INCORRECTLY. THE TIMEOUT COUNTER STARTS RUNNING AT THE TIME THE RFNM (CARRYING THE PIGGYBACKED ALLOCATE) GETS TO THE SOURCE IMP, NOT AT THE TIME THE RFNM (AND ALLOCATE) ARE GENERATED BY THE DESTINATION IMP. THUS, THE SOURCE HOST DOES NOT HAVE TO LOSE ALL THE TIME IN THE FACE OF 1/2 SECOND NETWORK TRANSIT TIMES AS CHARLES SUGGESTS. SORRY I DIDN'T PICK UP THIS POINT WHEN I READ THE MESSAGE THE FIRST TIME.

51a

Network Performance Dialog, 31Jan-5Apr 75

DAVE

51b

19-MAR-75 0836-EDT WALDEN at BBN-TENEX: IRBY'S QUESTIONS ABOUT WHICH BIT IMP BLOCKS ON Distribution: NETWORK PERFORMANCE GROUP:, ENGELBART AT OFFICE-1,, NORTON AT OFFICE-1, WATSON AT OFFICE-1, STONE AT OFFICE-1,, HEART, WALDEN, MCKENZIE, SUTHERLAND, BURCHFIELD, STROLLO,, MCLINDON AT ISI, CARLSON AT ISI, KLEINROCK AT ISI,, WALKER AT ISI Received at: 19-MAR-75 09:58:22

52

I cannot answer the questions about what bit the IMP stops taking data at without help from a hardware person. That is, I can explain when the IMP holds off on accepting input from the interface, but whether the interface then accepts a register's worth of bits after that I do not know.

52a

In the case of a block waiting for a multi-packet allocate, when the IMP has received a full 63 words an interrupt occurs. At that time the IMP checks to see if the last bit stored in memory was accompanied by a Last-Host-Bit signal. If not, input is deferred until an allocate becomes available.

52b

In the case of single packet messages with eight messages already outstanding, the entire message including its last bit has been received. Thus the line is blocked at some point after complete transmission of the message.

52c

----- THE ABOVE IS JOEL LEVIN'S ANSWER TO IRBY'S QUESTION. CHARLES, WILL YOU PLEASE PASS THIS ON TO THE REST OF YOUR CIRCULATION LIST?

52d

DAVE

52e

18-MAR-75 2103-EDT VICTOR at BBN-TENEX: count of host-imp messages Distribution: WALKER AT ISI, MCLINDON AT ISI, CARLSON AT ISI, BLUE AT ISI,, STUBBS AT ISI, LICKLIDER AT ISI, RUSSELL AT ISI,, HEART AT BBN, WALDEN AT BBN, MCKENZIE AT BBN,, SUTHERLAND AT BBN, BURCHFIELD AT BBN, STONE AT OFFICE-1,, ENGELBART AT OFFICE-1, NORTON AT OFFICE-1, WATSON AT OFFICE-1
Received at: 18-MAR-75 18:03:21

53

I made the following measurements today at office-1, sri-al, and isi-ka-tenex. the measurements were made by placing a patch in the tenex monitor. the meaning of the individual rows is as follows:
TOTAL # MSG - number of messages sent from the host to the imp from the time the patch was installed until the data was examined
TOTAL # WORDS - the total number of words in all the messages that

Network Performance Dialog, 31Jan-5Apr 75

were measured,

the first two words of each message is part of the host-host protocol, the remaining words are data with most likely four characters per word

(my counts may include one extra word per message since i am still learning about the way tenex deals with the net,

however,

the relationships found should still be valid,)

AVERAGE MSG = the average number of words per message

MAX MSG = the largest message sent in the time measured

MULTI PACKET MSG = the number of messages that would correspond

to

multi-packet messages being sent from imp to imp (as determined

by a

message length of greater than 28 words)

% MULTI PACKET MSG = the percentage of multi-packet messages

	AI	OFF-1	ISI-KA
TOTAL # MSGS	16878	32168	27560
TOTAL # WORDS	87776	314617	179845
AVERAGE MSG	5	9	6
MAX MSG	252	237	240
MULTI PACKET MSG	43	1950	269
% MULTI PACKETS MSG	.25	6.06	.98

53a

18-MAR-75 1747-EDT IRBY at BBN-TENEXB: Message Blocking by IMP Question.

Distribution: WALDEN AT BBN, walker at isi, mclindon at isi, carlson at isi, blue at isi,, stubbs at isi, licklider at isi, russell at isi, heart at bbn,, walden at bbn, mckenzie at bbn, sutherland at bbn,, burchfiel at bbn, strollo at bbn, tomlinson at bbn,, clements at bbn, allen at bbn, kleinrock at isi,, stone at office-1, engelbart at office-1,, norton at office-1, watson at office-1, lynch at sri-ai,, postel at bbnb, irby at bbnb, hardy at bbnb, jwhite at bbnb,, andrews at bbnb, victor at bbnb, hopper at bbnb,, ehardt at bbnb

Received at: 18-MAR-75 14:47:05

54

Dave:

When an IMP blocks the host to IMP line when it needs to get an (IMP level) allocate from the destination IMP for a multi-packet message the line is blocked on the bit that is one past the last bit that fits into the first packet. (how many bits is that exactly ?)

But if there are (now) 8 single packet messages outstanding between the pair of hosts that this message is to, and the line is

blocked for this reason after which bit is the line stopped ?

--jon.

54a

Network Performance Dialog, 31Jan-5Apr 75

18-MAR-75 1712-EDT IRBY at BBN-TENEXB: TIP Buffering/Allocation Strategy Question

Distribution: WALDEN AT BBN, walker at isi, mclindon at isi, carlson at isi, blue at isi,, stubbs at isi, licklider at isi, russell at isi,, heart at bbn, walden at bbn, mckenzie at bbn,, sutherland at bbn, burchfiel at bbn, strollo at bbn,, tomlinson at bbn, clements at bbn, allen at bbn,, kleinrock at isi, stone at office-1, engelbart at office-1,, norton at office-1, watson at office-1, lynch at sri-ai,, postel at bbnb, irby at bbnb, hardy at bbnb,, jwhite at bbnb, andrews at bbnb, victor at bbnb,, hopper at bbnb, ehardt at bbnb

Received at: 18-MAR-75 14:19:00

55

Dave our understanding of the TIP buffer allocation policy is as follows:

55a

The TIP has a pair of output buffers of equal size, say 800 bits (100 characters) for each terminal. The TIP initially allocates to the sending host 1 message and 800 bits. One of the buffers is always being used to output to the terminal while the other is used to accept data from the sending host.

55b

Lets call the buffer currently pointing at the terminal the Tbuf and the buffer currently pointing at the network the Nbuf.

55c

When the first message arrives from the network that data is put into Nbuf.

55d

The buffers are toggled and a new allocation is sent of 1 message and "1" bits where "1" is the length of the previous message. The buffer just filled, now Tbuf, is output to the terminal. When the next message arrives from the network it is put into Nbuf.

55e

When Tbuf is empty the buffers are toggled and a new allocate is sent of 1 message and "1" bits.

55f

We do not understand why the message allocation is limited to one at a time. It seems to us that it would be possible to allocate several messages and append the data that arrives

Network Performance Dialog, 31Jan-5Apr 75

to
the data already in Nbuf,

55g

18-MAR-75 1526-EDT HAMILTON at BBN=TENEXB: Lineprocessor Power Off Problem (TIP Break Test),

Distribution: WALKER AT ISI, MCLINDON AT ISI, CARLSON AT ISI, BLUE AT ISI,, STUBBS AT ISI, LICKLIDER AT ISI, RUSSELL AT ISI,, HEART AT BBN, WALDEN AT BBN, MCKENZIE AT BBN,, SUTHERLAND AT BBN, BURCHFIELD AT BBN, STROLO AT BBN,, TOMLINSON AT BBN, CLEMENTS AT BBN, ALLEN AT BBN,, KLEINROCK AT ISI, STONE AT OFFICE=1, ENGELBART AT OFFICE=1,, NORTON AT OFFICE=1, WATSON AT OFFICE=1, LYNCH AT SRI-AI,, POSTEL AT BBN, IRBY AT BBNB, HARDY AT BBNB,, JWHITE AT BBNB, ANDREWS AT BBNB, VICTOR AT BBNB,, HOPPER AT BBNB, EHARDT AT BBNB, hardy, hardware

Received at: 18-MAR-75 12:25:59

56

< HAMILTON, LPPPOWEROFF,NLS;2, >, 18-MAR-75 14:55 JOAN ;;;;

56a

1 A more detailed look shows that turning off power to a Lineprocessor connected to the Tymshare TIP and Office=1 via a 208A modem has two effects depending on user logo procedure,

56b

1A 1) If the user DOES NOT close his Network connection (type @c

<CR>) when he powers off the Lineprocessor the TIP sends repeating

breaks (or repeating something) to Office=1 eating up Host-IMP bandwidth severely loading the NCP and limiting other user response.

This loading is not reflected in load average statistics (CTRL T).

The effect is so bad that if two Lineprocessors using modems are

turned off simultaneously the Office=1 NCP crashes,

56c

1B 2) If the user closes his Network connection, the TIP still tries

to send repeating breaks. However, there is no host to send it to,

therefore only other Tymshare TIP users are effected. It could be

that the TIP becomes over loaded which causes the IMP to crash in a

funny way,

56d

2 These two effects could very well explain a lot of the random slow

response users have been experiencing at Office=1 independent of

Network Performance Dialog, 31Jan-5Apr 75

load
average,

56e

3 The cause appears to be that the Tymshare TIP is not inhibiting from sending multiple breaks as it should. (It is not as we earlier suspected that the Lineprocessor modem is sending them.) It may turn out that the cause is conditional depending on port configuration, or that there is a hardware malfunction in the Tymshare TIP alone. Ben Barker and Earle Hiscox are investigating in more detail.

56f

4 We have run the TIP load average test suggested by Walden, and Ben and Ernie should have useful data to look at. An extension to the test was to test the effects of turning power off to a Lineprocessor connected to BBNB.

56g

.....mmartin..
.....HARDY@SRI-ARC

56h

18-MAR-75 0730-EDT WALDEN at BBN-TENEX: CONTINUOUS BREAKS AT TYMSHARE

Distribution: NETWORK PERFORMANCE GROUP:, ENGELBART AT OFFICE-1,, NORTON AT OFFICE-1, WATSON AT OFFICE-1, STONE AT OFFICE-1,, HEART, WALDEN, MCKENZIE, SUTHERLAND, BURCHFIEL, STROLLO,, MCLINDON AT ISI, CARLSON AT ISI, KLEINROCK AT ISI,, WALKER AT ISI

Received at: 18-MAR-75 12:08:59

57

ALTHOUGH SENDNING CONTINUOUS BREAKS TO THE TIP WITH A CONNNECTION OPEN LOADS DOWN THE TIP BADLY AND WASTES 10% OF OFFICE-1 TENEX, LOOK AT THE GOOD SIDE: IT HELPS THE NETWORK TRAFFIC STATISTICS LOOK BIGGER.

57a

18-MAR-75 0113-EDT IRBY at BBN-TENEXB: Progress in Office-1 Performance Investigation Distribution: WALKER AT ISI, MCLINDON AT ISI, CARLSON AT ISI,, BLUE AT ISI, STUBBS AT ISI, LICKLIDER AT ISI,, RUSSELL AT ISI, HEART AT BBN, WALDEN AT BBN, MCKENZIE AT BBN,, SUTHERLAND AT BBN, BURCHFIEL AT BBN, STONE AT OFFICE-1,, ENGELBART AT OFFICE-1, NORTON AT OFFICE-1,, WATSON AT OFFICE-1 Received at:

17-MAR-75 22:45:34

58

At 5:00 PDT members of the ARC staff worked with NCC staff to determine the extent of loading the open modem lines caused on the TYMSHARE TIP. The results were very interesting. If the connection from the TIP to Office-1 was closed before the terminal

Network Performance Dialog, 31Jan-5Apr 75

was powered down, there was little effect on other Office-1 users because of the continuous break characters sent by the modem. However, if the terminal was powered down without closing the connection, the effect was devastating to other users accessing the Office-1 system through the net (ARC-ELF and SRI-AI). We tried powering the terminals up and down with very consistent result: fine performance when terminal on and horrendous (many seconds to echo characters) response when terminal off and modem sending breaks. We tried connecting to BBNB and SRI-AI and then causing breaks to be sent by modem. Little visible degradation. This may be because 1.33 TENEX breaks connection if login has not yet occurred or the long wait for RFNM may have lessened the effect. BBN is now trying to make there earlier fix to prevent this break nonsense work (it wasn't that the modem was sending start-stop bits, the fix just plain did not work).

58a

-- Charles, P.S. The amount of time spent in job zero in Office-1 TENEX jumped from about 3% to about 15% while the breaks were being input.

58b

17-MAR-75 2131-EDT STROLLO at BBN-TENEXB: 1 MS CLOCK ON BBN SYSTEM
B Distribution: WATSON, NORTON, engelbart at office-1, stone at
office-1, heart, walden,, mckenzie, sutherland, burchfiel, mclindon
at isi,, carlson at isi, walker at isi, kleinrock at isi,, chipman at
bbn, depesa, clements, allen, tomlinson Received at: 17-MAR-75
18:36:53

59

IT IS NOW CLAIMED TO BE FUNCTIONING, TED

59a

17-MAR-75 1243-PDT WALKER at USC-ISI: SUMEX VDH CONNECTION TO
TYMSHARE TIP Distribution: NETWORK PERFORMANCE GROUP: Received at:
17-MAR-75 16:19:11

60

BASED ON THE PROGRESS BEING MADE IN SHAKING OUT NETWORK
DIFFICULTIES, THE CONNECTION OF THE SUMMEX VDH TO THE TYMSHARE TIP
DURING DAYTIME HOURS IS POSTPONED UNTIL 25 MARCH. ALEX, PLEASE SO
INFORM THE SUMMEX PEOPLE, THANKS.

60a

STEVE

60b

17-MAR-75 0734-EDT WALDEN at BBN-TENEX: REPONSE TO IRBY'S MESSAGE
Distribution: NETWORK PERFORMANCE GROUP:, ENGELBART AT OFFICE-1,,
NORTON AT OFFICE-1, WATSON AT OFFICE-1, STONE AT OFFICE-1,, HEART,
WALDEN, MCKENZIE, SUTHERLAND, BURCHFIEL, STROLLO,, MCLINDON AT ISI,
CARLSON AT ISI, BLUE AT ISI, STUBBS AT ISI,, LICKLIDER AT ISI,
RUSSELL AT ISI, WALKER AT ISI Received at: 17-MAR-75 05:22:00

61

RE THE OPEN LINES ON THE TYMSHARE TIP, I WISH TO MEASURE THE
EFFECTS OF THIS IN A CONTROLLED TEST TODAY. WE HAVE A PATCH IN

Network Performance Dialog, 31Jan-5Apr 75

THE TYMSHARE IMP NOW TO MEASURE ITS "LOAD AVERAGE" AND I WOULD LIKE TO SEE THE EFFECTS ON THE LOAD AVERAGE OF BOTH ONE AND TWO 4800BAUD MODEMS RUNNING OPEN WITH AND WITHOUT AN OPEN CONNECTION. BEN BARKER AND ERNIE HISCOX OF BBN WITH TAKE THE INITIATIVE IN THESE MEASUREMENTS TODAY AND I ASSUME THE HELP OF SOMEONE AT OFFICE-1 TO LOOK AT THE MLC LIGHTS AND SOMEONE AT SRI TO TURN POWER ON AND OFF ON THE LINE PROCESSOR.

61a

RE THE IMP READY FOR NEXT HOST BIT LIGHT GOING OFF, WE HAVE WORKED UP A PATCH WHICH IS THE IMP TO WHICH BBNB IS CONNECTED WHICH DISPLAYS IN THE IMP LIGHTS 16 REASONS WHY THE IMP MIGHT BE BLOCKING THE HOST FOR A HALF-SECOND EVERY SO OFTEN. OUR OBSERVATION (BRIEFLY OVER THIS PAST WEEKEND) IS THAT THE IMP IS READY TO TAKE FROM THE HOST THE VAST MAJORITY OF THE TIME, BUT IT DOES NOT TAKE FOR THE ONE-HALF SECOND NOTED BY HARDY EVERY SEVERAL SECONDS (E.G., MAYBE 5% OF THE TIME). WE WILL INVESTIGATE THIS FURTHER TODAY AND MAY WELL WANT SOMEBODY TO LOOK AT THE LIGHTS AT TYMSHARE WHILE WE DO THE SAME THINGS.

61b

CHARLES' UNDERSTANDING OF PACKETS VS. MESSAGES IS PRETTY MUCH CORRECT (NOT PERFECT HOWEVER). WE HAVE LONG WANTED TO PUT IN THE HOST TO IMP CONTROL MESSAGE CHARLES SUGGESTS LETTING THE HOST ASK THE IMP FOR THE NECESSARY RESOURCES SO THE INTERFACE DIDN'T HAVE TO GET BLOCKED, BUT WE HAVE ALWAYS BEEN SERVICING A HIGHER PRIORITY INTERRUPT. WE HAVE ALL KNOWN FOREVER THAT HOSTS SHOULD BE ABLE TO HAVE MORE THAN ONE MESSAGE OUTSTANDING AT ONE TIME, BUT THIS WOULD REQUIRE A HOST/HOST PROTOCOL CHANGE EVERYWHERE. NOTE, THE TIP MAG TAPE OPTION DOES USE MULTIPLE MESSAGE OUTSTANDING AT A TIME AND A SPECIAL HOST LEVEL RETRANSMISSION PROTOCOL ON TOP OF HOST-TO-HOST PROTOCOL AND TELNET. BETWEEN TENEXS, FOR INSTANCE, A SIMILAR AD HOC PROTOCOL COULD BE USED FOR FILE TRANSFERS. (NOTICE, THE LINK ZERO BLOCKING PROBLEM WAS A PROBLEM WITH NOT ENOUGH MESSAGES OUTSTAND AT A TIME -- WE WERE GOING TO USE THOSE EXTRA FOUR MESSAGE ID BITS TO KEEP THINGS IN ORDER, ALTHOUGH IN THE CASE OF ALLOCATE, THIS CAN BE DONE WITHOUT THE KNOWLEDGE OF THE RECEIVER,) BY THE WAY, IF WHAT IS BEING DONE IS JUST TERMINAL OUTPUT (E.G., PAINTING A SCREEN), THEN WHY NOT JUST MODIFY TENEX SO UNDER THESE CIRCUMSTANCES IT DOESN'T WAIT FOR A BLOCKED LINK TO UNBLOCK BEFORE SENDING THE NEXT MESSAGE. THE SUBNET DOESN'T LOSE MESSAGE VERY OFTEN. DON'T WORRY ABOUT RETRANSMISSION IN THE RARE CASE IT DOES TO CRT TERMINALS. JUST LET THE SCREEN GET PAINTED OVER AGAIN CORRECTLY THE NEXT TIME. THE NETWORK TOPOLOGY SHOULD BE CHANGE SO THE GREATEST DISTANCE IS SAY 6 HOPS. ALL THE NETWORK ALGORITHMS WERE DESIGNED ORIGINALLY WITH SUCH AN ASSUMPTION IN MIND. THERE IS NOT ENOUGH STORAGE IN THE IMPs, THE TIMING ARE ALL WRONG (E.G., THE 1/4 SECOND ALLOCATE GIVE BACK TIME), ETC. WITH THE PRESENT NETWORK TOPOLOGY, (I DON'T KNOW WHAT CHARLES THINKS THE DIFFERENCE BETWEEN "HOST MESSAGES" AND "IMP MESSAGES"

61c

Network Performance Dialog, 31Jan-5Apr 75

IS; IN THE STORAE AND FORWARD NETWORK THEY ARE ALL BROKEN DOWN INTO PACKETS AND PACKET TAKE 50MSEC PER PACKET! TO TRAVERSE EACH HOP.)

61d

RE PAGE TRANSFER TIMES AT OFFICE-1, ARE THERE ANY SIMILAR MEASUREMENTS ON BBNB. REMEMBER MY SPECULATION OF A FEW DAYS AGO THAT THE BBNB VIRTUAL DRUM WAS OVERFLOWING.

61e

RE TIMING OF CHARACTER ECHOING AND TIMING OR TENEX TO IMP TO TENEX, LET'S RUN THE EXPERIMENTS AGAIN WITH US TRACING WHAT IS GOING ON IN THE IMPS IN QUESTION.

61f

REGARDS, DAVE

61g

17-MAR-75 0124-P CHI: Findings to date in Office-1 and BBNB performance problems investigation

Distribution: CHI SRI-ARC WEC DLS

Received at: 17-MAR-75 17:49

Location: (IJOURNAL, 25569, 1:w)

62

Comments: Sent via SNDMSG to distribution list in <engelbart>slst.

62a

16-MAR-75 2308-PDT LYNCH at SRI-AI: MORE TIMING TESTS Distributions: LICKLIDER AT ISI, STUBBS AT ISI, BLUE AT ISI,, CARLSON AT ISI, WALKER AT ISI, RUSSEL AT ISI,, MCLINDON AT ISI, DALE AT ISIB, HEART AT BBN, WALDEN AT BBN,, MCKENZIE AT BBN, SUTHERLAND AT BBN, BURCHFIELD AT BBN,, STROLLO AT BBN, HGM AT CCA, ENGELBART AT OFFICE-1,, NORTON AT OFFICE-1, WATSON AT OFFICE-1, STONE AT OFFICE-1,, TOMLINSON AT BBN, POSTEL AT BBN, VICTOR AT BBN,, ELLISON AT UTAH-10, PEPIN AT ECL, LYNCH AT SRI-AI Received at: 16-MAR-75 23:22:12

63

The following data were all collected on Sunday evening in order to try to ascertain what the "maximum" unloaded throughput that each TENEX site could provide in its current mode of operation. A few comments are in order:

63a

1) All sites were very lowly loaded. By that I mean that the load averages were below 1.0 in all cases except for CCA,

63b

2) The ?? marks for ISIB are due to the silly times received back from the system that gives the process accumulated runtime.

63c

3) ISIB and ECL are KI-TENEX machines, running KI-TENEX version 131 as modified by Rainer Schultz.

63d

4) Running 133 TENEX were BBNA, BBNB, BBNC, SRI-AI and ISIC,

63e

5) Running 132 TENEX were ISIA and CCA.

63f

Network Performance Dialog, 31Jan-5Apr 75

6) Running 131 TENEX were OFFICE, UTAH and the two KI systems, 63g

An explanation of the program that generated these data are in order! 63h

The program is a short assembly language program that was written by Hal MURRAY at CCA and modified (trivially) by Dan Lynch at SRI-AI to be able to run at any site. The program simply shoves data at itself through its own IMP and times all the events. It is quite a flexible program and may be obtained as: [SRI-AI]<LYNCH>WHIZZ.MAC. The ANONYMOUS login will work at SRI-AI for anyone who wants to look at the program. Also, All the raw data taken tonight exists in a file called [SRI-AI]<LYNCH>TELNET.15-MAR-75-NET=MEASUREMENTS. 63i

An explanation of the columns in the table is in order! 63j

1) The bits/sec (realtime) column is the true number of bits that travelled from the program out to its own IMP and back to the program divided by the amount of real time it took to do it. Thus it is a measure of human being time, 63k

2) The cpu-time (ms) is the number of milliseconds that TENEX charged the program for doing all the work. It does not include the very low level IMP interrupt code time (which is never measured [for the user] by TENEX), 63l

3) the cpu% column is simply the ratio of time that the program got while doing the transfers divided by real time. Thus it is a measure of both the "efficiency" of the Network Control Program and the ability of individual Schedulers to rapidly switch to and from "high priority" programs, 63m

4) The bits/cpu-sec (process-time) is really a measure of the efficiency of the NCP program and the speed of the CPU, 63n

SITE NAME	BITS/SEC REALTIME	CPU-TIME MS	CPU%	BITS/CPU-SEC PROCESS-TIME
-----------	-------------------	-------------	------	---------------------------

63o

TEST ONE == BIG TRANSFERS USING STRING I/O FOR 220 36 BIT WORDS 63p

BBNA	59000	4200	24	240000	BBNB
39000		3700	14	275000	BBNC
3800		16	260000	20500	3500
7	280000	20000	4300	9	235000
SRI-AI	27500	4300	11	234000	ISIA
20500		5300	10	190000	ISIB
??	??	??	ISIC	25000	3800
					9

Network Performance Dialog, 31Jan-5Apr 75

260000 ECL	24000	2700	6	375000 CCA	
44000	3500	15	288000		63q

TEST TWO == SINGLE CHARACTER THROUGHPUT AS IN TELNET OR TIP
USE 63r

BBNA	225	16950	38	590 BBNB	
118	11000	12	900 BBNC	80	
10300		8	970 OFFICE	35	5800
2	1700 UTAH	60	12200	7	820
SPI=AI	200	16300	33	610 ISIA	
100	25550	24	390 ISIB	175	
??	??	?? ISIC	100	9300	9
1125 ECL	75	11500	9	850 CCA	
125	9500	11	1060		63s

16-MAR-75 0904=PDT WALKER at USC=ISI: Thanks and Move on
Distribution: NETWORK PERFORMANCE GROUP: Received at: 16-MAR-75
14:42:47 64

I would like to acknowledge all the hard work being poured into the network response problems and the cooperative atmosphere which prevails. Things appear to be going quite well even if in a slightly unstructured way. As long as progress continues I cannot suggest any changes. If anyone feels specific actions should be taken please advise me. 64a

In the interests of not overloading the net with mail traffic and not saturating individuals mailboxes, I request that Licklider, Russell, Stubbs, and Blue not be included in the detailed reporting messages. They should be included when actions are requested or conclusions drawn. Does anyone else want out? 64b

I would like to add Len Kleinrock (@ISI) to our deliberations. 64c

Thanks for all your efforts. 64d

Steve 64e

15-MAR-75 0743=EDT WALDEN at BBN=TENEX: ANOTHER DISCOVERY IN THE OFFICE-1/NET EFFORT Distribution: NETWORK PERFORMANCE GROUP:, ENGELBART AT OFFICE-1,, NORTON AT OFFICE-1, WATSON AT OFFICE-1, STONE AT OFFICE-1,, HEART, WALDEN, MCKENZIE, SUTHERLAND, BURCHFIELD, STROLLO,, MCLINDON AT ISI, CARLSON AT ISI, BLUE AT ISI, STUBBS AT ISI,, LICKLIDER AT ISI, RUSSELL AT ISI, WALKER AT ISI Received at:
16-MAR-75 14:36:41 65

MARTIN HARDY HAS PROMISED TO WRITE THIS UP IN GREATER DETAIL, BUT

Network Performance Dialog, 31Jan-5Apr 75

I WON'T WAIT FOR HIM, BEFORE I GIVE YOU ALL A HINT WHAT HAS BEEN DISCOVERED:

65a

IT HAS BEEN FOUND THAT WHEN THE LINE PROCESSORS AT SRI CONNECTE BY TWO 4800BAUD MODEMS TO THE TYMSHARE TIP ARE TURNED OFF, THE MODEMS GENERATE A STREAM OF BREAK CHARACTERS TO THE TIP PUTTING A VISIBLE BURDEN ON THE TIP. CLEVERLY, THE MODEMS DON'T RUN OPEN, WHICH THE TIP WAS JUST MODIFIED TO DETECT (AS I STATED YESTERDAY), BUT THEY GENERATE NICELY FORMATTED CHARACTERS OF BREAK (I.E., EACH WITH ITS OWN STOP BIT) WHICH THE TIP THINKS IS A USER REALLY TYPING BREAKS ON HIS TERMINAL. NOW, 1) MARTIN WILL EXPLAIN THE DISCOVERY IN GREATER DETAIL AND ANY POSSIBLE RAMIFICATIONS HE SEES; AND 2) WE SHOULD RUN CONTROLLED TESTS NEXT WEEK OF THE BREAKS BEING POURED ONTO THE TIP TO UNDERSTAND HOW LOADED DOWN THE TIP GETS IN VARIOUS POSSIBLE SITUATIONS (E.G., A CONNECTION OPEN, NO CONNECTION OPEN) TO SEE IF ANY ADDITIONAL HARDWARE OR SOFTWARE FIXES ARE IN ORDER, OR IF THERE ARE ANY INCIDENTAL CLUES TO THE WAYS THE TIP CAN BE OVERLOADED.

65b

INCIDENTALLY, WE HAD OUR FIELD ENGINEER AT TYMSHARE LOOKING AT THE TIP LIGHTS WHILE THE SRI GUYS WERE RUNNING THEIR EXPERIMENTS AND THIS APPEARS TO HAVE BEEN HELPFUL, ALTHOUGH ANY OTHER PAIR OF VISUAL INTEGRATORS (I.E., EYES) COULD HAVE DONE THE JOB AND I BELIEVE MARTIN HIMSELF WENT DOWN TO CUPERTINO THE DAY BEFORE YESTERDAY TO STUDY THE TIP LIGHTS. THE LESSON TO REMEMBER IS THAT IT IS OFTEN TIMES USEFUL IF SOMEONE CAN JUST LOOK AT THE IMP AND TIP LIGHTS WHEN THERE IS TROUBLE. CLUES ARE SOMETIMES OBVIOUS FROM LOOKING AT THE LIGHTS.

65c

REGARDS, DAVE

65d

P.S., NOTICE THAT EVERYWHERE WE LOOK WE FIND PROBLEMS (MISSING CLOCKS, INTERFERENCE ON LINK 0, TERMINALS RUNNING OPEN, ETC.). WE ARE CONFIRMING MORE AND MORE EVERY DAY THAT THERE IS NO ONE BIG BUG WHICH SOMEBODY RECENTLY PUT IN WHICH ONLY HAS TO BE FOUND AND BE CORRECTED. WE ARE FINDING THAT EACH ONE OF THE INVOLVED SYSTEMS IS SOMEWHAT SUB-OPTIMAL IN ITS OPERATION BUT IT HAS BEEN THIS WAY FOREVER, AND IT IS ONLY INCREASED PRESSURE ON THE EXISTING MECHANISMS WHICH IS NOW DEMONSTRATING THEM TO BE SUB-OPTIMAL.

65e

14-MAR-75 1842=EDT ENGELBART at BBN-TENEXB: To Strollo re early fixing of 1-ms clock on BBNB TENEX Distribution: STROLLO at BBN, walker at isi, mclindon at isi, carlson at isi,, blue at isi, stubbs at isi, licklider at isi,, russell at isi, heart at bbn, walden at bbn,, mckenzie at bbn, sutherland at bbn, burchfiel at bbn,, stone at office-1, engelbart at office-1, norton at office-1,, watson at office-1, irby at bbnb, victor at bbnb,, hopper at bbnb, watson at

Network Performance Dialog, 31Jan-5Apr 75

bbnb, norton at bbnb,, engelbart at bbnb Received at: 16-MAR-75
14:41:31

66

Ted: Re your earlier message today: please install the 1-ms clock in BBNB TENEX as soon as you can. Our software guys want to learn everything they can, as soon as they can, toward getting improved service delivery to NLS users from both OFFICE-1 and BBNB. We'd appreciate learning what we can during the weekend, from the effect on BBNB, to be better armed for digging into OFFICE-1 when it gets re-powered Sunday night. Please communicate with (@BBNB) IRBY and VICTOR as to when to expect the change, Thanks, Doug Engelbart

66a

14-MAR-75 1557-EDT STROLLO at BBN-TENEXB: 1 MS CLOCK ON BBN SYSTEM B Distribution: WATSON, NORTON, engelbart at office-1, stone at office-1, heart, walden,, mckenzie, sutherland, burchfiel, mclindon at isi,, carlson at isi, blue at isi, stubbs at isi, licklider at isi,, russell at isi, walker at isi, chipman, depesa, clements,, tomlinson, allen Received at: 14-MAR-75 13:06:37

67

1) THE CLOCK IS INDEED NOT FUNCTIONING 2) I'VE TALKED WITH ALL THE TENEX WIZARDS HERE AND WE CANNOT IMAGINE HOW IT COULD HAVE A FIRST ORDER EFFECT ON THE RESPONSE PROBLEMS YOU GUYS ARE SEEING 3) WE WILL OF COURSE FIX IT, I SUGGEST WE WAIT UNTIL THE NORMAL PM ON TUESDAY AM TO MINIMIZE THE AMOUNT OF DOWN TIME BUT YOU GUYS CAN OVER RULE THAT IF YOU FEEL STRONGLY ENOUGH TO WANT TO INTRODUCE ADDITIONAL DOWN TIME, 4) IF SOME NOTICEABLE IMPROVEMENT IN PERFORMANCE IS BELIEVED TO HAVE OCCURRED AFTER FIXING THIS CLOCK, WE WILL BE HAPPY TO CONDUCT SOME EXPERIMENTS ENABLING AND DISABLING THE CLOCK TO SEE IF WE CAN FIGURE OUT WHAT IS GOING ON, TED

67a

14-MAR-75 1529-EDT MCKENZIE at BBN-TENEX: VDH Code for Stanford Medical Center (SUMEX) Distribution: NETWORK PERFORMANCE GROUP:, ENGELBART AT OFFICE-1,, NORTON AT OFFICE-1, WATSON AT OFFICE-1, STONE AT OFFICE-1,, HEART, WALDEN, MCKENZIE, SUTHERLAND, BURCHFIEL, STROLLO,, MCLINDON AT ISI, CARLSON AT ISI, BLUE AT ISI, STUBBS AT ISI,, LICKLIDER AT ISI, RUSSELL AT ISI, WALKER AT ISI Received at: 14-MAR-75 12:47:27

68

This Message is especially for STEVE WALKER -----

68a

The Stanford Medical Center Host is supposed to be connected to the Tymshare TIP as a VDH. Installing the VDH code in the Tymshare TIP "steals" a large fraction of the IMP buffers (i.e., the buffers used for store-and-forward, Host queues, reassembly; not the TIP terminal buffers. SUMEX requested that the VDH code be installed on the evening of 6 March; Steve Walker decided to postpone the installation of the VDH code during days (it is

Network Performance Dialog, 31Jan-5Apr 75

installed every night between 2000 PDT and 0330 PDT) until the morning of 18 March. I think that when that decision was made everyone hoped that the "office-1 problem" would be fixed by now. However, since the problem isn't yet fixed, or even diagnosed, and since Tymshare/Office-1 will be down ALL WEEKEND for site power changes, perhaps it would be reasonable for ARPA to further postpone the installation of the VDH code. Perhaps we do not want to introduce one more variable into the already complicated performance equations? Regards, Alex

68b

14-MAR-75 1507-EDT MCKENZIE at BBN-TENEX: Addendum to Walden's Message of 14 March 1108-EDT Distribution: NETWORK PERFORMANCE GROUP; ENGELBART AT OFFICE-1, NORTON AT OFFICE-1, WATSON AT OFFICE-1, STONE AT OFFICE-1, HEART, WALDEN, MCKENZIE, SUTHERLAND, BURCHFIELD, STROLLO, MCLINDON AT ISI, CARLSON AT ISI, BLUE AT ISI, STUBBS AT ISI, LICKLIDER AT ISI, RUSSELL AT ISI, WALKER AT ISI
Received at: 14-MAR-75 12:11:50

69

Walden's message mentioned that Dan Lynch discovered "the 1msec clock is not running at BBN or Office-1". Since "BBN" could refer to several different computers, it is crucial to point out that ONLY BBN TENEX B has a failing 1msec clock; Lynch found all other BBN TENEX systems to be running the 1msec clock correctly. Alex McKenzie

69a

14-MAR-75 1404-PDT LYNCH at SRI-AI: NCP timing data Distribution: LICKLIDER AT ISI, STUBBS AT ISI, BLUE AT ISI, CARLSON AT ISI, WALKER AT ISI, RUSSELL AT ISI, MCLINDON AT ISI, HEART AT BBN, WALDEN AT BBN, MCKENZIE AT BBN, SUTHERLAND AT BBN, BURCHFIELD AT BBN, ENGELBART AT OFFICE-1, NORTON AT OFFICE-1, WATSON AT OFFICE-1, STONE AT OFFICE-1, STROLLO AT BBN, TOMLINSON AT BBN, POSTEL AT BBN, LYNCH AT SRI-AI Received at: 14-MAR-75 14:03:31

70

In order to give everyone a look at the data that has been collected concerning the various times to do simple I/O to one's own IMP and back, I have collected the following representative data in the last 12 hours. The program that collects the data is at [sri-ai]<lynch>impsuk.sav and the source file is .mac. It is a very short program and will run on any TENEX and requires no special capabilities. The first column gives the time in milliseconds for the send (bout) to return from the monitor and the second column gives the time from that point to the time the character is received by the program (bin). The program will loop forever, so a 'c' is the way to stop it. bbnb bbnb
bbnc office-1 sri-ai isia isib

70a

4	22	0	150	3	70	0	200	5
21	5	93	6	52	4	25	0	50
97	0	650	8	61	3	114	3	58

Network Performance Dialog, 31Jan=5Apr 75

31	0	250	2	72	0	150	3	23
5	356	5	54 4	47	50	50	-1	59
0	350	10	20	2	252	6	55 4	25
0	1500	2	78	0	900	5	23	4
170	6	38 4	34	0	100	2	38	0
100	4	26	3	160	6	44 4	25	0
50	0	51	0	400	5	25	5	234
8	60 2	26	0	50	2	42	0	200
4	33	41	131	6	42 2	17	0	50
3	68	0	200	4	24	18	40	5
42	4	26	0	150	2	136	0	300
4	24	4	74	7	36			

70b

Note that it appears that bbnb and office=1 have broken milisecond clocks. Also, that bbnc has a peculiar bias in its clock? All of the above data was taken during periods of light to medium (1-3) load,

70c

14-MAR-75 1132-EDT WALDEN at BBN-TENEX: TERMINAL RUNNING OPEN SWAMPING THE TIP Distribution: NETWORK PERFORMANCE GROUP:, ENGELBART AT OFFICE-1,, NORTON AT OFFICE-1, WATSON AT OFFICE-1, STONE AT OFFICE-1,, HEART, WALDEN, MCKENZIE, SUTHERLAND, BURCHFIELD, STROLLO,, MCLINDON AT ISI, CARLSON AT ISI, BLUE AT ISI,, STUBBS AT ISI, LICKLIDER AT ISI, RUSSELL AT ISI,, WALKER AT ISI Received at: 14-MAR-75 09:02:43

71

THE HARDWARE FIX WHICH PREVENTS A TERMINAL WHICH IS RUNNING OPEN FROM SWAMPING THE TIP WAS PUT IN THE TYMSHARE TIP ONE ABOUT FEB, 28. THE 10X-TIP HAS HAD THE FIX FOR MONHS,

71a

14-MAR-75 1107-EDT WALDEN at BBN-TENEX: SOME ANSWERS TO SOME QUESTIONS AND A FEW OTHER THINGS Distribution: NETWORK PERFORMANCE GROUP:, ENGELBART AT OFFICE-1,, NORTON AT OFFICE-1, WATSON AT OFFICE-1, STONE AT OFFICE-1,, HEART, WALDEN, MCKENZIE, SUTHERLAND, BURCHFIELD, STROLLO,, MCLINDON AT ISI, CARLSON AT ISI, BLUE AT ISI,, STUBBS AT ISI, LICKLIDER AT ISI, RUSSELL AT ISI,, WALKER AT ISI Received at: 14-MAR-75 09:03:23

72

POSTEL AND I WERE WONDERIN IF THE LINE PROCESSORS AT OFFICE 1 COULD GENERATE MULTI-PACKET MESSAGES, I HAVE CHECKED AND PORTS 15 AND 17 (OCTAL) ON THE TYMSHARE TIP ARE CAPABLE OF HANDLING MULTI-PACKET MESSAGE ON BOTH INPUT AND OUTPUT,

72a

MCKENZIE GUESSES THAT THE TYMSHARE TIP MAY ALREADY HAVE THE HARDWARE FIX WHICH PREVENTS A TERMINAL RUNNING OPEN FROM OVER-BURDENING THE TIP. UN FORTUNATELY, EVERYONE WHO KNOWS ABOUT THIS IS OUT OF TOWN. I WILL CHECK WITH THE NEXT ONE TO CALL IN AND THEN SEND THE DEFINITIVE ANSWER,

72b

Network Performance Dialog, 31Jan-5Apr 75

REGARDING DOUG'S LIST OF POSSIBLE TROUBLES (DATED 10 MARCH):

72c

TIP 327 IS IN ALMOST EVERY WAY A BETTER SYSTEM THAN TIP 322, EXCEPTING THE ADDITIONAL SPACE IT TAKES. REGARDING SPACE AND TIP VERSION, TIP 322 HAS ABOUT 3000 WORDS FOR TERMINAL BUFFERING IN A 28K TIP; TIP 327 HAS ABOUT 2000 WORDS FOR TERMINAL BUFFERING IN A 28K TIP. OF COURSE, 32K TIPS HAVE 4K MORE WORDS FOR TERMINAL BUFFERING WITH EITHER SYSTEM. TYMSHARE IS A 28K TIP. HOWEVER, TYMSHARE, TO OUR KNOWLEDGE ONLY HAS 12 TERMINALS ATTACHED AND 2000 WORDS SHOULD BE SUFFICIENT FOR SO FEW TERMINALS. THUS, TYMSHARE IS RUNNING TIP 327 WHICH IS IN OTHER WAYS BETTER THAN TIP 322. OF COURSE, TYMSHARE WOULD BE BETTER OFF WITH 32K AND I STRONGLY ENCOURAGE THIS ADDITION.

72d

REGARDING DOUBLE VS. SINGLE BUFFERING, ON OUTPUT THE TIP DOES DOUBLE BUFFERING. ON INPUT IT USES A SINGLE BUFFER TO INPUT INTO AND TO SEND TO THE NETWORKFROM, WITH A FLOATING BOUNDARY BETWEEN THE INPUT PORTION AND THE SENDING PORTION. I DON'T REALLY UNDERSTAND DOUG'S QUESTION ABOUT SINGLE BUFFERING BEING MORE EFFICIENT THAN DOUBLE BUFFERING: PERFORMANCE WOULD BE TERRIBLE WITHOUT THE OVERLAPPING OF OPERATIONS PERMITTED THROUGH DOUBLE BUFFERING.

72e

DOUG MENTIONS A CHRONIC HARDWARE PROBLEM IN THE TYMSHARE IMP/TIP. COULD SOMEBODY ELABORATE ON THIS AND GIVE US ANY EVIDENCE YOU HAVE.

72f

DOUG MENTIONS THAT THE INITIAL TOPOLOGY LAYOUT OF THE NETWORK WAS CLAIMED TO HAVE NO MORE THAN 3 INTER-IMP HOPS. I THINK ACTUALLY, IF MY MEMORY SERVES, THE CLAIM WAS MADE THAT THE TOPOLOGY WOULD BE LAYED OUT SO THE ROUND TRIP DELAY FROM ANY HOST TO ANY OTHER WOULD BE UNDER .2 SECONDS AND TO DO THIS IS WAS NECESSARY TO LIMIT THE NUMBER OF INTER-IMP HOST TO 6. CURRENTLY THE MAX IN THE NET IS 14 INTER-NODE HOPS; ITTIS 13 HOPS FROM ARPA TO OFFICE-1; IT IS AS LIKELY TO BE TEN HOPS AS TO BE 3 HOPS; ETC. THE DISTANCE FROM ARPA TO OFFICE-1 HAS GONE UP BY TWO HOPS SINCE THE BEGINNING OF FEBRUARY I THINK. IF ANYBODY IS INTERESTED, I CAN PROVIDE THE COMPLETE HISTOGRAM OF NUMBER OF HOPS BETWEEN ALL PAIRS OF HOSTS. NOTICE, THE 13 HOPS NUMBER, ETC., GIVEN ABOVE ASSUMES NO BROKEN NODE OR LINE ON THE SHORTEST PATH.

72g

DOUG ASKS ABOUT NETWORK LOADING: GENERALLY SPEAKING THE NETWORK IS STILL NOT VERY HEAVILY LOADED IN TERMS OF THROUGHPUT AND THE PACKETS ARE MOSTLY SMALL. THE NETWORK MEASUREMENT CENTER IS CARRYING OUT AN EXPERIMENT TO DETERMINE THE AVERAGE DELAYS IN THE NETWORK AND THE AVERAGE QUEUEING AT EACH STORE-AND-FORWARD NODE. OF COURSE, NETWORK LOADS HAVE NEVER BEEN SO HIGH WITH A DAILY AVERAGE NOW OF 4.4M PACKETS (INTER-NODE) AND 1.6M PACKETS INTRA-NODE, BUT STILL THE NETWORK IS NOT NEAR CAPACITY.

72h

Network Performance Dialog, 31Jan-5Apr 75

I DON'T KNOW IF TENEX USERS CAN SET THEIR PRIORITY BITS IN MESSAGES. THE TIP USERS CANNOT. HOWEVER, THE TIPS DO SET THE PRIORITY BITS ON ALL TRAFFICE, HOST/HOST PROTOCOL COMMANDS ETC. FOR THE USERS. THIS IS THE WAY IT SHOULD BE I THINK BECAUSE THESE TIP USERS ARE TRYING TO GET LOW DELAY AND NEED THE PRIORITY. I DON'T KNOW WHAT TENEX DOES; I THINK IT TO SHOULD USE PRIORITY FOR MESSAGE TO TIPS. (THE MAG TAPE TIPS DO NOT USE PRIORITY FOR ANY TERMINALS BECAUSE THAT IS THE EASY WAY TO PREVENT PRIORITY FROM BEING USED ON THE BIG FILE TRANSFERS THE MAG TAPES ARE DOING.)

721

REGARDING THE IMP AND TIP FAIRLY INTERLEAVING TRAFFIC, THE TIP TRIES TO BE QUITE FAIR, BUT THE IMP HAS BUT FEW EXPLICIT MECHANISMS TO GUARANTEE FAIRNESS. PUTTING IN SUCH EXPLICIT FAIRNESS MECHANISMS WAS TO HAVE BEEN OUR NEXT IMP DEVELOPMENT TASK BEFORE WE GOT DIVERTED ONTO THE PRESENT PROBLEM. NOTE, WE DON'T THINK THAT FAIRNESS IS A SERIOUS IMP PROBLEM EVEN WITHOUT THE MECHANISMS AS THE PROGRAM GENERALLY DOES NATURALLY CYCLE THROUGH THE VARIOUS COMPETING TASKS. HOWEVER, AS I MENTIONED PREVIOUSLY, THERE APPEAR TO BE MANY PLACES WHERE THERE IS INTERFERENCE POSSIBLE BETWEEN COMPETING HOSTS AND USERS SO THEY WILL "SEE" EACH OTHER ALTHOUGHT PROBABLY NOT BLOCK EACH OTHER OUT FOR PROTRACTED PERIODS. WE ARE LOOKING INTO THIS ISSUE MORE.

72j

I THINK IT WAS NOT PREVIOUSLY MADE KNOW TO THE MEMBERS OF THE ADDRESS LIST THAT LYNCH DISCOVERED LAST NIGHT THAT THE 1MSEC CLOCK IS NOT RUNNING AT BBN OR OFFICE-1, ALTHOUGH IT IS RUNNING AT OTHER TENEX SITES. LYNCH SPECULATED THIS MIGHT CAUSE SOME UNNECESSARY DELAYS IN SCHEDULEING AND MORE SERIOUSLY HURT PAGE MANAGEMENT. TOMLINSON CONFIRMS THAT BBN 'S 1 MSEC CLOCK IS NOT WORKING BUT REFRAINS FROM SAYING THIS THIS WOULD MAKE MUCH DIFFERENCE IN PERFORMANCE, ALTHOUGH HE DOES AGREE THAT IT SHOULD BE FIXED TO SEE IF THINGS SUDDENLY GET BETTER.

72k

REGARDS, DVE

721

14-MAR-75 0857-P JCN: Office-1/Network Service Problems: What's Goin' On
Distribution: JCN KWAC FEED
Received at: 14-MAR-75 12:12
Location: (IJOURNAL, 32096, 1:w)

73

11-MAR-75 2315-PDT ENGELBART at OFFICE-1: To Walden on Net/Office-1 problems Distribution: WALKER AT ISI, MCLINDON AT ISI, CARLSON AT ISI, BLUE AT ISI,, STUBBS AT ISI, LICKLIDER AT ISI, RUSSELL AT ISI,, HEART AT BBN, WALDEN AT BBN, MCKENZIE AT BBN,, SUTHERLAND AT BBN, BURCHFIELD AT BBN, STONE AT OFFICE-1,, ENGELBART AT OFFICE-1, NORTON AT OFFICE-1, WATSON AT OFFICE-1, irby at bbn Received at:
12-MAR-75 07:00:18

74

Network Performance Dialog, 31Jan-5Apr 75

Dave: I think that your Sunday's memo was a very nice job. Our technical specialists will be matching it's quality I hope, Don's study helps, and we're digging further. Will follow some of the points you bring out. We are all suffering from some sort of shock/disorientation: the overwhelming amount of trouble appearing suddenly in us in either OFFICE-1 or BBNB over the NET have blown our working productivity and our morale. Under user-load levels that worked well before there are now atrocious delays. Last week, at 1700 PDT, through a 4800-baud modem direct to TYMESHARE TIP, with O-1 Load Ave under 1, it was as bad as what I mentioned I experienced at BBNB in the PS of my Monday memo (there I happened to be going through ELF, in DNLS also). I didn't look at BBNB load average then. The times mentioned were for discrete steps in executing a command: for instance, hitting the "ic" key pair to set up for an Insert Character Command, hitting the CA key to bug a character, hitting a single character as the replacing LIT, hitting the CA to execute. If that seems anomalous from your looking at the log data for BBNB then, well, that's the story we've had at Office-1, too. More to follow. Best regards, Doug

74a

11-MAR-75 2230-PDT ENGELBART at OFFICE-1: Don Andrews' memo on effect of new drum at Office-1 Distribution: WALKER AT ISI, MCLINDON AT ISI, CARLSON AT ISI, BLUE AT ISI, STUBBS AT ISI, LICKLIDER AT ISI, RUSSELL AT ISI, HEART AT BBN, WALDEN AT BBN, MCKENZIE AT BBN, SUTHERLAND AT BBN, BURCHFIELD AT BBN, STONE AT OFFICE-1, ENGELBART AT OFFICE-1, NORTON AT OFFICE-1, WATSON AT OFFICE-1 Received at:
12-MAR-75 07:00:29

75

After operations stabilized with the new drum, the Superwatch statistics were analyzed by Don Andrews and summarized in (IJOURNAL, 25545,), a copy of which is appended below. Third drum has been on order, expected in a month or so. More data to follow. Doug Engelbart ----- 1 This is a brief summary of the differences in the OFFICE-1 system before and after the second drum was added, as shown by superwatch statistics.

75a

2 Source;

75b

2A The statistics were taken from averages from 8:00 to 14:00 on 2/19/75, 2/20/75 and 2/21/75 vs. 3/5/75, 3/6/75 and 3/7/75.

75c

3 Drum capacity:

75d

3A Keep in mind the following: The drum went from 600 pages capacity to 1200 pages capacity. The SRI-ARC system's Bryant drum had 3000 pages capacity. The OFFICE-1 drum has different operating characteristics than the Bryant drum did. Even though there are two drums now, transfers do not take place in parallel. The addition did nothing more than expand the capacity.

75e

Network Performance Dialog, 31Jan-5Apr 75

4 Drum Usage:

75f

4A With one drum, the drum was busy about 7% of real time. It was idle the rest of the time. An average of about 14 pages were read per second. [On SRI-ARC, about 50-60 pages were read per second.] The system spent about 2% of real time in I/O wait with the drum busy (i.e. wait on drum or drum and disk).

75g

4B With two drums, it was busy about 27% of the time. An average of about 25 pages were read per second. I/O wait with the drum busy was about 4% of real time. [I think the same figure for SRI-ARC was about 8-10%].

75h

4C The rest of the I/O wait time was spent waiting on the disk alone. That figure went from about 15% to about 10% of real time. Although idle times were different for the two test periods, this indicates that the drum was doing its job better. But 10% is still terrible. The same figure was about 1-2% on SRI-ARC.

75i

5 Number of users:

75j

5A During the three days in Feb., there were an average of 16-17 users using time. This is the #AU statistic which includes detached and system jobs.

75k

5B During the three days in Mar. there were 21-22 active users. This represents a 25% increase in jobs. I would say this accounts for the increase in time spent in the scheduler (from about 7.5 to 9.5 %). Percent used is difficult to compare since idle time was not the same but the percent of non-idle time spent running user jobs went up from about 60% to about 70%.

75l

6 Response:

75m

6A The length of time it takes to execute an NLS interaction divided by the CPU time actually used is a measure of responsiveness of the system. Before the drum was added, that figure was about 7.5, although there was an miserable day when it was 15! After the drum was expanded, that figure was 8.0. This represents a slight loss in responsiveness but I would expect much more than that considering the number of users added.

75n

7 Efficiency:

75o

7A I consider the statistic %SYS to be a reasonable index of system efficiency. It is the time spent in system mode. The idea is to minimize it (maximize the time spent in user mode -- running user programs).

75p

Network Performance Dialog, 31Jan-5Apr 75

7B Before the drum was expanded, %SYS was about 86%. After, it was 84.5%. This represents a user-mode increase from 14% to 15.5% which is an 11% increase.

75q

7C I/O wait times cannot be compared directly because the idle times are not comparable. However, the non-idle I/O wait times can be compared although only large differences would be significant.

75r

7D Before expansion, the non-idle I/O wait time was about 23%, after expansion, it was about 15%. That difference is significant.

75s

8 Summary:

75t

8A The extra drum capacity made a definite difference. The system is carrying more users at about the same level of service.

75u

8B I would guess from comparisons with SRI-ARC statistics that OFFICE-1 would benefit still more from a third drum. The 1200 page drum is only used 25% of real time and it should be more like 50% or more. As more users are added, need for drum space goes up also.

75v

8C From the statistics, it appears that the disk is not operating correctly. (More on this in another message). If there is a problem and it gets fixed, it may change our evaluation of the drum situation. I don't know which way tho.

75w

11-MAR-75 1630-EDT WALDEN at BBN-TENEX: poor tenex response over the net Distribution: NETWORK PERFORMANCE GROUP:, ENGELBART AT OFFICE-1,, NORTON AT OFFICE-1, WATSON AT OFFICE-1, STONE AT OFFICE-1,, HEART, WALDEN, MCKENZIE, SUTHERLAND, BURCHFIEL,, MCLINDON AT ISI, CARLSON AT ISI, BLUE AT ISI,, STUBBS AT ISI, LICKLIDER AT ISI, RUSSELL AT ISI,, WALKER AT ISI Received at: 11-MAR-75 13:27:03

76

In yesterday's message from Doug, he mentioned that while he was creating his message yesterday morning, using Tenex B at BBN over the network was very painful. I have taken a look at the Tenex B load averages during the period I assume Doug was using Tenex B and the load average appears to have been over 5 a lot of the time and all the way up to 12+ for a period. Thus, it is not too surprising Tenex was ugly to use.

76a

Acutally I have a lot of load averages for most of the Tenexs in the net since the beginning of March and quite a lot of the time the load averages at Office-1, Utah, ISIA, and BBNB are very high; e.g., 5, 6, 7, 8, 9, 10, 11, 12.

76b

Is it possible that there really is not enough Tenex time

Network Performance Dialog, 31Jan-5Apr 75

available in the network any more. Certainly, on my local Tenex response time is not good a lot of the time.

76c

I suggest the following experiments: 1) Only permit one user on some network Tenex and have that user come in from a TIP somewhere close by. Now artificially increase the load average on the Tenex (e.g., by throwing cycles away) and see how the TIP users sees Tenex as a function of load. This experiment checks load, not interference or scarcity of other resources than bandwidth.

76d

2) Again only allow one user on the Tenex (again from a TIP). Give this user all the machine so there can be no high load average. Now artificially load up the IMP/Host interface and see how response appears to the user as a function of interface use.

76e

I guess I have a third experiment. 3) Again, with no heavy load on the Tenex, but with a number of users at TIP terminals using the system so they conflict at the IMP/Host interface, test the response as a function of number of users (and variety of TIPs from which they are coming).

76f

Regards, Dave

76g

p.s., which Tenex is going to volunteer to do the above experiments,

76h

11-MAR-75 11:19-EDT MCKENZIE at BBN-TENEX: Preliminary Report: Walker's Troubles Reported on 6 March Distribution: WALKER AT OFFICE-1, network performance group; engelbart at office-1, norton at office-1, watson at office-1, stone at office-1, heart, walden, mckenzie, sutherland, burchfiel, mclindon at isi, carlson at isi, blue at isi, stubbs at isi, licklider at isi, russell at isi, walker at isi Received at: 11-MAR-75 08:13:40

77

Steve, In your message of 6 March you reported on 4 specific interactions between ARPA and Office-1 (of which only three were troublesome). I have gone over the Network Control Center log and checked on the network events occurring at the three troublesome times. Here is a preliminary report.

77a

Tues 4 March 2005 EDT - No known network problems; no abnormal events recorded at the NCC.

77b

Thurs 6 March 1200 EDT - The Tymshare IMP rejected a few incoming packets thinking there was no available buffering. Actually, since you could type far ahead, the performance problem HAD to be AFTER the packets were delivered to the Host. Thus I don't believe you even saw the packet rejections by the IMP. However, the fact that the IMP had to reject a few packets because it was

Network Performance Dialog, 31Jan=5Apr 75

so busy may provide a clue that the Office-1 Host was processing a great deal of network traffic (or there might only have been a great deal of Tymshare-TIP traffic), The VDH code was not in the IMP at that time,

77c

Thurs 6 March 2000 EDT - Again the IMP rejected a few packets due to shortage of IMP space. Further, at about 2018 the IMP had to "block the interface" from the TIP to the network in order to obtain IMP buffer resources to send multi-packet messages from TIP terminal(s) to the Office-1 Host (the TIP has 2 ports with 180 characters [1440 bits] of buffering in the terminal-to-network direction), and the interface remained blocked for 15 seconds without the necessary buffering becoming available. This should not have affected you directly, but again may provide a clue to the amount of network traffic the Office-1 Host was required to handle. The VDH code was not in the IMP at that time,

77d

Note 1: When I say "a few packets" I mean not more than about one a minute for a small number of minutes,

77e

Note 2: Over the weekend the IMP code was tuned up a bit so even the few packet rejections which occurred last week would probably not occur this week under THE SAME load,

77f

Regards, Alex McKenzie

77g

10-MAR-75 1712=EDT ENGELBART at BBN=TENEXB: From ARC re Remote-Service Problems at Office-1 Distribution: WALKER AT ISI, MCLINDON AT ISI, CARLSON AT ISI, BLUE AT ISI,, STUBBS AT ISI, LICKLIDER AT ISI, RUSSELL AT ISI, HEART AT BBN,, WALDEN AT BBN, MCKENZIE AT BBN, SUTHERLAND AT BBN,, BURCHFIELD AT BBN, STONE AT OFFICE-1, ENGELBART AT OFFICE-1,, NORTON AT OFFICE-1, WATSON AT OFFICE-1 Received at: 10-MAR-75 14:11:50

78

as requested, we are listing the potential problem areas in this pass; we will send observation and measurement data in subsequent passes. we are building on the sample problem structure in walker's message of 9 mar,

78a

tipImp performance problems:

78b

the differences between tip 327 and tip 322.

78c

buffer sizes -- found to be important in the past, we understand that these have been reduced for various reasons, and that some tips have reverted to 322 with some buffer re-allocations. also, vdh software in the imp-side of tymshare tip has reduced buffer sizes; some uncertainty here w.r. what and when this software is running there,

78d

Network Performance Dialog, 31Jan-5Apr 75

double buffering == for a line processor and its hi-speed tip connection, is it as efficient as single buffer for the amount of buffer space used? 78e

possible to have chronic hardware problem at tipImp? some of our symptoms at tymshare tip make us wonder, 78f

tenex problems 78g

1,31 vs 1,33, 78h

office=1 physical configuration (full complement of core and disk; just doubled drum capacity, could add more,) 78i

possible thrashing when imp interrupts go past a certain point in frequency, 78j

ncp buffer sizes 78k

job inputOutput buffer sizes 78l

possible glitch in office=1 tenex (have updated portions to 1.32 level and, in addition, are running a recent version of netser == so, non-standard 1,31) 78m

pager performance == it had fewer associative registers than at sri-arc, being increased now, 78n

nls problems: 78o

no significant changes in nls during recent months, but changeover from nls=7 to nls=8 could be contributing, 78p

memory utilization by nls, perhaps too much page faulting, 78q

generally, potential glitch ... 78r

network problems: 78s

topology == we used to hear planned that no connection would have more than three inter-imp hops, our average user has about 10, 78t

loading == number and size of packets, 78u

if a user program can set priority on its messages (question?), is it possible that this may be being practised in an unwarranted way? when it seems subjectively sometimes that the imp

Network Performance Dialog, 31Jan-5Apr 75

is paralyzed for many seconds, one could picture a very large file going by in priority packets.

78v

generally, when heavy traffic of packets going through your imp, could the algorithms by which the imp chooses your message over one that is passing through somehow penalize you and make you wait for a break in traffic?

78w

past a certain loading level, does the packet-switched net jam up? (like cars on a freeway, where apparently they jam if their number brings the average speed below 19 mph.)

78x

office-1 user loading

78y

total number of people sharing the available slots has swelled considerably, resulting in a heavier-density usage of each slot (ratio now exceeds any in our previous experience).

78z

older users are increasing in skill and sophistication, enabling them to consume more resources.

78aa

line processors becoming available -- display-nls provides its greater power at greater service costs.

78aa

process=commands feature is being used more, and when used incurs a relentlessly solid loading on its slot.

78ab

we are very pleased to participate in coordinate effort, and are determined to make our end work.

78ac

doug engelbart and jim norton

78ad

P.S. Relevant comment: I used BBNb "this morning" to transcribe and launch this message. Two weeks ago, on local SRI-ARC TENEX, I could have done it easily in well under half an hour; so far I've been at it for over two and a half miserable and very unsettling hours. Essentially the same as I would have expected at Office-1 at its worst. Currently from five to twenty-five seconds delay for each simple operation. It reminds me of trying to work from England in Sept 73 when their TIP was just coming up, when NORSAR TIP was shaky, when there was a 1200-baud line from London to Norway and I think 4800 baud across the Atlantic -- and, with strange European terminals. (I actually think this is worse.) DCE

78ae

10-MAR-75 1349-EDT WALDEN at BBN-TENEX: Net mail from site BBN-TENEX Distribution: WALKER AT ISI, network performance group; engelbart at office-1, norton at office-1, watson at office-1, heart at bbn,, walen at bbn, sutherland at bbn, burchfiel

Network Performance Dialog, 31Jan-5Apr 75

at bbn,, mclindon at isi, carlson at isi, blue at isi,, stubbs at
 isi, licklider at isi, russell at isi,, walker at isi, stone at
 office-1 Received at: 10-MAR-75 10:48:23

79

Steve,

79a

ON THE SUBJECT OF OFFICE-1

79b

1. I think that all the network complaints are contradictory and
 confusing; I, frankly, don't know what to suggest the problem
 might be. I must look when there is trouble. Otherwise there are
 too many things to guess.

79c

2. I think the IMP and the TIP systems have never worked so well
 as they are currently working. Except for a couple of specific
 cases at specific sites, TIP 327 is in every way better than TIP
 322 (except available space), and the latest IMP is the best IMP
 system there has ever been.

79d

3. I have no reason to believe most of the following list are
 happening, but they might be -- we must look when badness is
 happening.

79e

4. My model of the net/Office-1 problem has several components:
 the source TIP (e.g., at ARPA) the source IMP (also at ARPA)
 the store and forward network (the IMPs in between the
 source and the destination) the destination IMP (e.g., at
 Tymshare) the destination Host (Office-1) the TIP on the
 destination IMP (the Tymshare TIP in the Tymshare IMP)

79f

I don't believe any of the TIPs in any of the IMPs between the
 source and destination are relevant. Even the TIP at the
 destination is not relevant (it shares no!! space with the IMP)
 except in that it may somehow be putting a bandwidth burden on the
 destination IMP (this is independent of the TIP version -- to
 first order, TIP 327 is no slower than TIP 322 I believe --
 however, with either TIP 322 or TIP 327, it is possible to have a
 bunch of terminals which load down the TIP and hence the IMP --
 I'll come back to this point later).

79g

5. I think the problems fall into several areas: a) working
 perfectly; b) working as it should but not as fast as the user
 would like -- this results from overloading and interference;
 c) hiccups (momentary hangups) -- almost certainly due to Host
 glitches (e.g., temporary hangup in the scheduler) or network
 events (e.g., a destination Host goes down while another Host
 is trying to talk to it); d) very, very bad performance over a
 protracted period -- unknown why -- could be anything -- will
 undoubtedly be many different things many different times.

79h

Network Performance Dialog, 31Jan-5Apr 75

Now the list of possible problems:

791

1. The source TIP might be cycling slowly at times. Reasons for this might be a) a port or two running open and eating up a lot of bandwidth processing breaks -- it is simple to see if this is hapening by looking at the light on the MLC when the trouble is happening; b) there may be too many terminals attempting to gain service simultaneously from the TIP -- we do not know what its current performance capability is -- we want to study this this month; c) there might be a big terminal (e.g., a computer such as the XGP) eating up a lot of bandwidth when it is running; d) maybe there is a "line processor" on the source TIP and maybe that eats up a lot of bandwidth; e) maybe there is some way a terminal can get in a hung state (i.e., it is not taking characters from the TIP) and the burden on the TIP of continually trying to send characters to this terminal loads it down so badly it can't service the other terminals very well (TIP 322 was known to have problems of this sort I believe).

791

NOTE: the TIP has never been claimed to support computers and cassette input, etc, etc. the TIP's development (i.e., currently backward toward TIP 322) can't be driven by users needs of which we were never told and which the TIP was never constructed to handle.

79X

2. There is a demonstrable problem when too many terminals on a source TIP simultaneously try to talk to a particular destination Host. This problem has to do not with overload, but with interference. In the Host to Host protocol, link zero must be shared by all connections between the source and the destination and the reverse is also true. In the case of a single Host trying to send output (including echos) to many terminals on a single TIP, the TIP has to send allocates to the Host for all the terminals on the same link zero. Further, under the traditional rules of Host/Host protocol, only one message may be outstanding on link 0 at a time. The TIP currently does the best it can (as does TENEX) in packing many allocates into a single link zero message, but there is still demonstrable interference between the terminals. For example yesterday we did the following experiment: 1) connected three terminals from a BBN TIP to Office-1; 2) connected another terminal on the same TIP to another Host; 3) connected a terminal on another TIP to office 1; 4) with two terminals on the one TIP running TTYTST simultaneously from Office-1, there was noticable stuttering of the output; 5) with three terminals simultaneously running TTYTST, stuttering increased by 50% to all 3 BBN/Office-1 terminals and futhermore was synchronized to first order; however while this was going on a TTYTST from Office-1 to a terminal on another TIP did not

Network Performance Dialog, 31Jan-5Apr 75

stutter and output from another Host to a terminal on the same TIP did not stutter. Thus, neither Office-1 or the TIP was overloaded, but the terminals were interfering with each other a lot. We did the same test between the NCC-TIP and BBN TENEX-C. In this case the same effect was noticable but it took more like 8 terminals to demonstrate it.

791

This is due to the smaller distance from the NCC-TIP to the BBN TIP. NOTE: In this case it's fair to claim that the TIP is perfectly programmed -- it is doing exactly what it should and the best it can under Host/Host protocol -- further, it has been this way forever -- there is no change here. With a dozen terminals simultaneously going from the ARPA TIP to Office-1 (an even greater distance than in our experiment) the ARPA TIP might well appear unusable. It appears that there is a significant mod one could make to the TIP which, by violating protocol in a non-harmful way, would improve this situation a great deal-- during our experiment we patched in a partial fix and things got much better. This mod would take several days of careful thinking to work out and would require the release of a new TIP system everywhere and a little space (e.g., some small amount of code and 3 bits per port). We can try the fix with a patch, but there are possible detrimental side effects of the partial fix that is possible with a patch, namely the possibility of hung connections happening more often.

79m

3. At the source TIP, all of the overloading of the TIP bandwidth suggested in point 1 above could be seen as slow TIP response. Similar overloading of the destination TIP could possibly eat up a lot of the destination IMPs bandwidth -- but we have looked at this and it doesn't appear to be happening.

79n

4. Long delays across the store and forward network are a possibility. these could come in several forms: 1) routing loops -- we don't think these are happening a significant amount now; 2) some sort of blocking upon entry to the s-a-f net at the source -- we have to catch a case and look at it -- but we have traced the ARPA TIP a great deal and seen no signs of this; 3) long minimum delays across the net due to its topology -- it is certainly close to a second across the network now (13 hops from ARPA to Office-1) -- whether this is all necessary or not depends on the latency due to queuing seen in the average node -- this needs to be measured -- and we need to look when a problem is happening. Recall that two of these hops (Gunter and Eglin) were added to the net since 1 January 1975; 4) insufficient reassembly buffering in the IMPs -- this is probably not a problem at ARPA, but might be a problem at Office-1 (especially when the vdh code goes in -- which has not happened yet). All in all, while we may be picking up a fraction of

Network Performance Dialog, 31Jan-5Apr 75

a second unnecessary delay in this area, I don't think this is causing multi-second delays. 5. The destination Host is slow in taking traffic from the network. We have measured Office-1 a lot for this. 10% of the time Office-1 takes more than 1/5 second to take a message from the IMP, and 5% of the time Office-1 takes more than 1/2 second to take a message. The same measurements show the IMP to be a very much more minor offender in adding delay.

790

6. IMP/Host (and the reverse) interface blocking. Everytime a destination Host (e.g., BBN) to which Office-1 or the ARPA TIP is sending goes down, there is a possibility of a temporary (e.g. 15 seconds) hicup while the destination Host is timed out. I don't think this is probably happening too often, but when it does, it can interfere with other conversations to other Hosts. Further, this effect usually can only happen when multi-packet messages are sent and this doesn't happen too often. The Host/IMP interface can also be blocked if there is some kind of a network deadlock, and TENEX continually retransmits messages to such deadlocked IMPs over and over again keeping the interface blocked -- the symptom of this is getting through a few characters every 15 seconds. Another cause of interface blocking could be a slow Host to which one of the Hosts involved (e.g., ARPA TIP, Office-1) is sending (such a Host might be the PDP-15 at ARPA, although I presently have no reason to suspect it gets hung in slow states). This problem has been seen at BBN when the PDP-11 which has the line printer on it was slow in taking traffic from one of the TENEX systems and the TENEX systems in turn appeared very slow to their terminal users because of interface blocking.

790

7. The IMP to TIP pseudo Host interface can have the same sorts of blocking and some others.

790

NOTE: if one catches a bad case, it is easy to see if one of these blocking situations is occurring.

791

8. Does the "line processor" use multi-packet messages? If so, then we should look into its causing situations in which Host interface blocking happens or in which it uses a lot of buffering.

791

9. The echoing delay from Office-1 is noticeable slower than output. I don't understand why this is so. Perhaps it is a clue. The TENEX network output routines are hard to control to make them make good use of messages to get high performance. This is a fundamental problem with Host Host protocol which tries to suppress message boundaries rather than making use of them for good performance.

791

10. My experience with SRI-ARC was that when even a few users were

Network Performance Dialog, 31Jan-5Apr 75

running nls, TENEX was much slower than the load averages indicated. By the way, the load average numbers are not necessarily indicative of how good service TENEX is giving the network,

79u

11. I know many, many things have been done between TENEX 1.31 and 1.33 to make it handle the network better. I don't know all the ways in which it used to be bad, but I know it is better now. Perhaps somebody else can make a list of all the old deadlocks and hangups TENEX had with regard to the network. TENEX has been known to have troubles with buffer allocation so input blocked output or the reverse, with scheduling the NCP, and with full connection tables,

79v

My preliminary conclusions:

79w

A) Slowness is due to interference and size (ie big net), not overload,

79x

B) Glitches are for unknown reason but are probably due to the Hosts, not the net or TIPS,

79y

regards, Dave Walden

79z

p.s., i have skimped on host troubles and emphasised net possible troubles, because i know most about the net -- not because i think the hosts are free from trouble,

79ae

9-MAR-75 1346-PDT WALKER at USC-ISI: Office-1 Response Problems - Refinement Distribution: ENGELBART AT OFFICE-1, NORTON AT OFFICE-1,, WATSON AT OFFICE-1, HEART AT BBN, WALDEN AT BBN,, MCKENZIE AT BBN, SUTHERLAND AT BBN, BURCHFIELD AT BBN,, MCLINDON, CARLSON, BLUE, STUBBS, LICKLIDER, RUSSELL, WALKER,, STONE AT OFFICE-1
Received at: 9-MAR-75 13:46:36

80

My 6 March request for a list of the potential causes of the "Office-1 response problems" is not to be construed as a "major piece of research" but rather as getting everyone to commit to paper in a concise form the things that they have been so eager to state orally. There are a specific set of problems (some subject to quick solution, some not) which are causing these difficulties. We need to all understand the magnitude of the problems and what can be done to solve them. Until they are written down we can't even start,

80a

I am convinced that a list of "ten major problem areas" can be generated which will cover the entire problem. Examples of entries on the list might include:

80b

Network Performance Dialog, 31Jan-5Apr 75

1. TIP performance problems: the differences between TIP 327 and TIP 322 and the implications of those differences. 80c

2. Tenex Performance Problems: The differences between Tenex 1,31 and 1,33, and their implications. 80d

3. NLS Problems: additional response time restrictions that imposes on Tenex. 80e

4. Network topology restrictions and the part they play. 80f

... 80g

I don't expect everyone to be able to comment in depth on each area, but urge everyone to mention each area that they can think of so that all get covered. I do not wish to "reveal to the world at large a longer list of things to blame their problems on". 80h

May we get on with this! 80i

Steve 80j

6-MAR-75 1746-PDT WALKER: Office-1 Response Problems Distribution: ENGELBART, NORTON, WATSON, HEART AT BBN, WALDEN AT BBN,, MCKENZIE AT BBN, SUTHERLAND AT BBN, BURCHFIELD AT BBN,, MCLINDON AT ISI, CARLSON AT ISI, BLUE AT ISI, STUBBS AT ISI,, LICKLIDER AT ISI, RUSSELL AT ISI, WALKER AT ISI, STONE Received at: 6-MAR-75 17:46:02 81

Response problems continue to plague use of the Office-1 system over the network. We need a concerted effort to resolve the causes of this problem and get on with solutions to it. I request that all interested parties (at least: SRI, BBN=Network, BBN=Tenex) submit lists of potential causes for the troubles addressed to me with copies to each addressee of this message by 1400 EDT 10 March. Each party should then iterate on these lists with arguments why a particular cause is either invalid or not likely. Initial iteration replies should be in by Wed in order that the PIs Involved can discuss this with me next week at the PI meeting. 81a

I am asking ISI to test the access to Office at various times during the day. Duane, would you report your responses from Office during the next week? 81b

Performance that I have observed recently: 81c

Tues even 4 Mar 2005 EDT Response: 10-20 sec delays, severe input problems (ie bell after 3-5 chars) 81d

Network Performance Dialog, 31Jan-5Apr 75

Thurs 6 Mar 1100 EDT Office load: 5, response time: good (2 to 8 secs), input problem: none, accepted chars as fast as typed, 81e

Thur 11Mar 1200 EDT Office load: 8, Response time: bad (15 to 30 secs), input rate: good, accepted chars as fast as typed; I could get a full line of text ahead of the response with no bells, 81f

Thu 6 Mar 2000 EDT , Response time : bad (20 to 45 seconds), Office load 0,5, input problem: bad, 5 to 8 chars result in bell. 81g

What other info should be reported? 81h

This problem is very serious and has dragged on far too long. The cooperation of all parties is requested to get an understanding of the cause and to take steps to correct it, 81i

Thanks, 81j

Steve 81k

31-JAN-75 1501-PST FEEDBACK: system response Distribution: STONE, bergstrom, feedback Received at: 31-JAN-75 15:01:32 82

Thanks for passing on your problems of the past week. Not that it'll make you feel any better but you're not the only ones, 82a

In answer to your questions, Steve Walker at ARPA is the network resource manager (that may not be quite the right title but close). I think he's the one who manages buffers etc. Dave Russell is probably most knowledgeable about the future of the ARPANET - he's working on transferring responsibility to DCA, 82b

Dick Watson says the front end will alleviate problems with buffers and loads to some extent because it will allow much larger buffers, and because much of the feedback to the user will be local the overall load on the net will be reduced, 82c

We're trying to keep in close touch with the NCC to solve problems with disconnections and other specific problems - please let us know if any new developments arise out there, Susan/FEED 82d

DLS 7-APR-75 10:57 32263

Network Performance Dialog, 31Jan-5Apr 75

(J32263) 7-APR-75 10:57;;; Title: Author(s): Duane L. Stone/DLS;
Sub=Collections: RADC; Clerk; DLS;

conferencing and nls

This notte describes minimal features of teleconferencing and how teleconferencing would relate to an expanded Journal subsystem in N&S.

conferencing and nls

TELECONFERENCING AND NLS

1

ABSTRACT

2

During our lunch with Larry Day on April 3, you were asking what must be added to NLS to give it minimal teleconferencing capabilities. Because I am just finishing the computer conferencing chapter for Roger, I have also been interested in this question. This memo is a brief and incomplete note on my view of teleconferencing in NLS.

2a

INTRODUCTION

3

I see two basic issues. The first arises because all teleconferencing systems today are stand-alone programs. All communication functions for the group must be handled by the teleconferencing system. In an NLS environment, this might mean that a teleconferencing system could have less power. On the other hand, conferencing may provide the general paradigm for group communication in NLS, into which present tools would be imbedded.

3a

At the end of this note, I discuss some thoughts about generalizing the Journal system to provide functions which include teleconferencing. Those comments, however, are undeveloped. I have included them only to indicate possible directions for research.

3a1

The second issue is what characteristics do modern teleconferencing systems have. This paper focuses on a somewhat more limited question, what are the minimal characteristics for a teleconferencing system.

3b

CHARACTERISTICS REQUIRED OF A TELECONFERENCING SYSTEM

4

At a minimum, several people must be able to read and write on a conference record simultaneously. The conference record may simply be a linear list, with entries appended to the end of the list as they come in.

4a

Because entries are usually appended at the end of a list in the order of their submission, there is no danger of two participants trying to rewrite the same section simultaneously. Therefore there is no danger of more than one person having access to a partial copy of the record simultaneously.

4a1

In fact, it is essential that each participant's partial copy be updated dynamically as new entries come in. Although many teleconferencing interactions are asynchronous,

conferencing and nls

interactions among on-line participants are both frequent and necessary. It would be unworkable if a participant could only reply to comments after several hours' delay or even a full day's delay.

4a2

One desirable feature comes close to being a requirement. This is the ability to send private (off-the-record) messages to other participants. Again, immediate delivery is desirable.

4b

These two features - simultaneous reading and writing, and private message transmission - might form a minimal teleconferencing system, but in fact no systems have ever been built with just these two primitive capabilities. This is because three other considerations almost always require further elaboration in practical teleconferencing systems.

4c

First, message sending must be simple, because a typical session may find a participant sending many small messages to the record and to other participants. It would be inconvenient if a complex protocol were necessary each time a message were sent.

4c1

For example, in FORUM, the command "To Forum" means that the following message will be inserted at the end of the record. The author's name, the date and time are automatically added to the message, and an accession number is assigned automatically to the entry. In private messages, the command is "To Lastname." Again, the message formatting is automatic.

4c1a

In NLS, we would probably use something like the sendmail system, with a new command like Send Entry to send a formatted entry to the record. Titling, keywords, and so forth, would be optional.

4c1b

A second pragmatic concern is that the conference record may become very long, so it is necessary to have special viewing helps available to the user and suited to the conference environment.

4c2

At a minimum, all systems keep a pointer at the first entry added since the participant last logged out. This allows the participant to examine new information.

4c2a

However even new information lists can become very long, so some selective dissemination of information stacks are provided for programmed or semiautomatic viewing. During a session, a user marks an entry, an author, or a text string. When there is any entry by the author,

conferencing and nls

referencing the entry, or containing the text string, a citation is sent to the stack for that entry. When the participant does something like a Show Status (of) Lists command, the names of the lists are shown, and an asterisk or some other mark is placed by stacks that have received citations since they were last accessed. As an extension of this, whenever a person logs on, he or she is shown a list of conferences, and asterisks are placed by conferences that have had activity since the participant last entered them,

4c2b

Notification can also be active. Whenever a new participant enters the conference, other participants are notified automatically. Alternatively, a teletype or other device may be left on in a filtered mode, with only messages by an author or with some other characteristic being printed out as they are entered by another author.

4c2c

Third, to assist further in easing the viewing burden, nontrivial conferences are usually broken into subconferences and these may be subdivided further. Subconferences may overlap in time, just as a committee in business may have permanent subcommittees and episodic subcommittees.

4c3

FEATURES THAT COULD BE ADDED TO BASIC NLS STRUCTURES

5

If NLS incorporates a state-of-the art teleconferencing system, rather than a minimal system, it could offer special features.

5a

At the broadest level, teleconferencing resembles the creation of "sets" in the Journal system. These sets are not simple subcollections, however. They must be partially modifiable by a conference leader, and they must provide the types of aids to sending and viewing that we have discussed above, probably including information retrieval functions within the set or subset. An expanded Journal system with expediting, alerting, editing, and structuring could form the core of a conferencing system.

5a1

Another germane feature of NLS is its file structure. A node now contains text only, but it will soon be able to have a comments branch at each node, to hold links to comment entries. This would remove a major problem in current computer conferencing systems - how to focus comments. While there is some need for free-form multilogue, current systems are increasingly feeling a need to include focusing devices. Voting is a commonly-used device, but NLS would add new dimensions, allowing people to discuss position papers and the like,

5a2

conferencing and nls

NLS also offers desirable aids to users, including live training, linking, feedback, the help subsystem and online documentation. This will allow for easier learning of the system by users.

5a3

Overall, a conference is merely an episode in an organization's ongoing communication processes, and it seems to me that conferencing could properly be a subset of an expanded journal system designed to serve an organization's total communication needs. It also seems to me that teleconferencing provides some insight into the complexities that must be introduced into the journal subsystem if it is to become a broad communication tool. The journal is not merely a mail system; it must eventually include facilities for submitting, organizing, reorganizing and reading information.

5b

RA3Y 7-APR-75 12:00 32264

conferencing and nls

(J32264) 7-APR-75 12:00;;; Title: Author(s): Raymond R. Panko/RA3Y;
Distribution: /RAH([ACTION]); Sub-Collections: SRI-ARC; Clerk:
RA3Y; Origin: < PANKO, DOUG-TELECON,NLS;3, >, 7-APR-75 10:08
RA3Y ;;;;####;

Teleconferencing and NLS

This is a brief note on my views of teleconferencing in NLS. It discusses minimal features of teleconferencing systems and how the Journal system could grow to encompass teleconferencing features. The discussion is superficial, but you may find it useful.

Teleconferencing and NLS

TELECONFERENCING AND NLS

ABSTRACT

During our lunch with Larry Day on April 3, you were asking what must be added to NLS to give it minimal teleconferencing capabilities. Because I am just finishing the computer conferencing chapter for Roger, I have also been interested in this question. This memo is a brief and incomplete note on my view of teleconferencing in NLS.

INTRODUCTION

I see two basic issues. The first arises because all teleconferencing systems today are stand-alone programs. All communication functions for the group must be handled by the teleconferencing system. In an NLS environment, this might mean that a teleconferencing system could have less power. On the other hand, conferencing may provide the general paradigm for group communication in NLS, into which present tools would be imbedded.

At the end of this note, I discuss some thoughts about generalizing the Journal system to provide functions which include teleconferencing. Those comments, however, are undeveloped. I have included them only to indicate possible directions for research.

The second issue is what characteristics do modern teleconferencing systems have. This paper focuses on a somewhat more limited question, what are the minimal characteristics for a teleconferencing system.

CHARACTERISTICS REQUIRED OF A TELECONFERENCING SYSTEM

At a minimum, several people must be able to read and write on a conference record simultaneously. The conference record may simply be a linear list, with entries appended to the end of the list as they come in.

Because entries are usually appended at the end of a list in the order of their submission, there is no danger of two participants trying to rewrite the same section simultaneously. Therefore there is no danger of more than one person having access to a partial copy of the record simultaneously.

In fact, it is essential that each participant's partial copy be updated dynamically as new entries come in. Although many teleconferencing interactions are asynchronous,

Teleconferencing and NLS

interactions among on-line participants are both frequent and necessary. It would be unworkable if a participant could only reply to comments after several hours' delay or even a full day's delay.

4a2

One desirable feature comes close to being a requirement. This is the ability to send private (off-the-record) messages to other participants. Again, immediate delivery is desirable.

4b

These two features - simultaneous reading and writing, and private message transmission - might form a minimal teleconferencing system, but in fact no systems have ever been built with just these two primitive capabilities. This is because three other considerations almost always require further elaboration in practical teleconferencing systems.

4c

First, message sending must be simple, because a typical session may find a participant sending many small messages to the record and to other participants. It would be inconvenient if a complex protocol were necessary each time a message were sent.

4c1

For example, in FORUM, the command "To Forum" means that the following message will be inserted at the end of the record. The author's name, the date and time are automatically added to the message, and an accession number is assigned automatically to the entry. In private messages, the command is "To Lastname." Again, the message formatting is automatic.

4c1a

In NLS, we would probably use something like the sendmail system, with a new command like Send Entry to send a formatted entry to the record. Titling, keywords, and so forth, would be optional.

4c1b

A second pragmatic concern is that the conference record may become very long, so it is necessary to have special viewing helps available to the user and suited to the conference environment.

4c2

At a minimum, all systems keep a pointer at the first entry added since the participant last logged out. This allows the participant to examine new information.

4c2a

However even new information lists can become very long, so some selective dissemination of information stacks are provided for programmed or semiautomatic viewing. During a session, a user marks an entry, an author, or a text string. When there is any entry by the author,

Teleconferencing and NLS

referencing the entry, or containing the text string, a citation is sent to the stack for that entry. When the participant does something like a Show Status (of) Lists command, the names of the lists are shown, and an asterisk or some other mark is placed by stacks that have received citations since they were last accessed. As an extension of this, whenever a person logs on, he or she is shown a list of conferences, and asterisks are placed by conferences that have had activity since the participant last entered them,

4c2b

Notification can also be active. Whenever a new participant enters the conference, other participants are notified automatically. Alternatively, a teletype or other device may be left on in a filtered mode, with only messages by an author or with some other characteristic being printed out as they are entered by another author,

4c2c

Third, to assist further in easing the viewing burden, nontrivial conferences are usually broken into subconferences and these may be subdivided further. Subconferences may overlap in time, just as a committee in business may have permanent subcommittees and episodic subcommittees,

4c3

FEATURES THAT COULD BE ADDED TO BASIC NLS STRUCTURES

5

If NLS incorporates a state-of-the art teleconferencing system, rather than a minimal system, it could offer special features,

5a

At the broadest level, teleconferencing resembles the creation of "sets" in the Journal system. These sets are not simple subcollections, however. They must be partially modifiable by a conference leader, and they must provide the types of aids to sending and viewing that we have discussed above, probably including information retrieval functions within the set or subset. An expanded Journal system with expediting, alerting, editing, and structuring could form the core of a conferencing system,

5a1

Another germane feature of NLS is its file structure. A node now contains text only, but it will soon be able to have a comments branch at each node, to hold links to comment entries. This would remove a major problem in current computer conferencing systems - how to focus comments. While there is some need for free-form multilogue, current systems are increasingly feeling a need to include focusing devices. Voting is a commonly-used device, but NLS would add new dimensions, allowing people to discuss position papers and the like,

5a2

Teleconferencing and NLS

NLS also offers desirable aids to users, including live training, linking, feedback, the help subsystem and online documentation. This will allow for easier learning of the system by users.

5a3

Overall, a conference is merely an episode in an organization's ongoing communication processes, and it seems to me that conferencing could properly be a subset of an expanded journal system designed to serve an organization's total communication needs. It also seems to me that teleconferencing provides some insight into the complexities that must be introduced into the journal subsystem if it is to become a broad communication tool. The journal is not merely a mail system; it must eventually include facilities for submitting, organizing, reorganizing and reading information.

5b

RA3Y 7-APR-75 12:07 32265

Teleconferencing and NLS

(J32265) 7-APR-75 12:07;;; Title: Author(s): Raymond R. Panko/RA3Y;
Distribution: /DCE([INFO-ONLY]) ; Sub=Collections: SRI-ARC; Clerk:
RA3Y; Origin: < PANKO, DOUG-TELECON,NLS;3, >, 7-APR-75 10:08
RA3Y ;;;;###;

Output Processor Directives

Bob,

1

What I've done in situations where I've had to change OUTPUT PROCESSOR directives in a file I was unfamiliar with was to either:

1a

- use content patterns to get a filtered view showing only those directives I was interested in (solving the problem of not knowing exactly what they are)

1a1

- possibly use the PSHOW command if you can arrange it so that the appearance of the pages you're interested in doesn't affect the appearance of the pages following them (e.g., force a page break before & after the pages of interest). You could take the updated file, insert the PSHOW & other mods, do the OUTPUT PROC., do a Delete Modifications, and process the rest of the file (possibly forcing a page renumbering).

1a2

That about exhausts my store of suggestions. Hope they're of some help. I agree that a SAVE/RESTORE DIRECTIVES command would be useful.

1b

Regards,

2

Frank

3

Output Processor Directives

(J32267) 7-APR-75 16:28;;; Title: Author(s): Frank G. Brignoli/FGB;
Distribution: /RMS2([ACTION]) AID([INFO-ONLY]) ;
Sub-Collections: NIC AID; Clerk; FGB;

Another sample memo, with a somewhat different formatting technique

This is a supplement to <32261>, in which I requested
development of a new user-prog, "MEMO."

Another sample memo, with a somewhat different formatting technique

Memorandum for: MR. ELFORD
MR. FOREHAND
MR. MAJETIC
MR. MANNING
MR. MASLOW
MR. MESSICK

1

Subject: Project Operations

Date: 8 APR 75

From: David A. Potter

2

Having finally returned from a notably unproductive stint of jury duty, I find that it is impossible (given the status of your various calendars) to convene a meeting of the Internal Advisory Committee until April 28. In the meantime, I have a project to run, and would hate to see us lose the entire month, which after all represents one-third of the remainder of the fiscal year.

3

I would like to proceed immediately to make all the arrangements necessary to begin data collection in the city of New Haven. This will require me to spend a few days there in meetings with District staff and I assume, local union representatives, in order to plan the Spring's activities and recruit, inform, and involve participating teachers. It is unlikely that much more than this can be accomplished during April; if possible, however, I should also like to proceed to discuss with participating teachers and staff the development of initial lists of basic teaching tasks, and to negotiate with them the specific data collection activities in which they will be involved.

4

Activities during this period will be built on the project plan presented in the attached document, which you will find to be a considerably revised and expanded version of the paper we discussed at our last meeting.

5

I would very much appreciate two things:

6

1. I would like to hear from you by Thursday, April 10, if you have any objections to my proceeding with project activities as noted in New Haven.

6a

Another sample memo, with a somewhat different formatting technique

2. At any time (but the sooner the better -- what else?), I would welcome comments or suggestions regarding the attached paper. Send them to me, or give me a call and I'll be glad to stop by and discuss it with you.

6b

Another sample memo, with a somewhat different formatting technique

(J32269) 8=APR-75 07:31;;; Title: Author(s): David A. Potter/DAP;
 Distribution: /NDM([ACTION]) ; Sub-Collections: NIC; Clerk: DAP;
 Origin: < POTTER, MEMO,NLS;1, >, 7=APR-75 10:47 DAP ;;;;
 ####;

Towards Benchmarks

There now exists a new facility (to be brought up soon, if not up by the time you see this) in NLS to record a session for use in the future by RUNFIL, to use the facility, issue the start record session command and, before giving the final confirmation, issue the "RUNFIL" format command. I suggest that the final NLS command given be a "quit" command. The file generated can then be edited to place the necessary Exec commands (term type line processor, no raise, NLS, etc.) at the front of the file,

1

With respect to measuring NLS, I offer the following possible strategy:

2

Rather than putting in calls to measurement procedures as commands at procedure entry and exit points, modify LIO to generate these calls automatically at procedure entry and return statements. LIO can have a data base file that would list procedures which we wish to measure, or specify some syntactical rule for procedure names that we wished to measure (e.g., all procedures whose first letter of the name starts with an "X").

3

Since measurement under the former procedure required a recompilation, this approach offers no more overhead and, in fact, less overhead since the measurement calls do not have to be edited in or out.

4

Towards Benchmarks

(J32270) 8-APR-75 17:05;;; Title: Author(s): Kenneth E. (Ken)
Victor/KEV; Distribution: /NPG([ACTION]) RWW([ACTION]) DCE([ACTION]) JML([ACTION]) ; Sub=Collections: SRI-ARC NPG; Clerk:
JML; Origin: < LEAVITT, KEN,NLS;2, >, 8-APR-75 17:02 JML
;??;####;

Reply to RLL's 25689

RLL's idea to substitute ARC-APP and ARC-DEV for individual membership in the ident SRI-ARC creates more problems than it's worth and goes against the grain of the way the identfile and journal are set up. Granted, the 2 groups are useful as groups, but substituting them for individual membership leaves a lot of room for error. First of all, whenever a new ARC person is added to the identfile, his or her ident is automatically added to the membership of SRI-ARC. If RLL's idea was implemented, I would then have to by hand remove the new ident from the SRI-ARC membership and by hand add it to ARC-DEV or ARC-APP. It seems silly to make it this complicated. Let's just leave the two groups as groups, and keep individual membership as is in SRI-ARC.

1

MLK 8-APR-75 19:52 32271

Reply to RLL's 25689

(J32271) 8-APR-75 19:52; Title: Author(s): Marcia Lynn Keeney/MLK;
Distribution: /JCN RLL MLK JDH; Sub-Collections: SRI-ARC; Clerk: MLK;

JHB JCP 8-APR-75 20:33 32272

Summary of the Changes in the New Version of Tenex == 1.33

Advance notice for Architects

Summary of the Changes in the New Version of Tenex == 1.33

Summary of the Changes in the New Version of Tenex == 1.33

This is a brief overview of changes that users will experience at Office-1 with the arrival of the 1.33 Tenex. Most of the important changes happen at much lower levels in the system than the non-programmer sees, but a few will be noticable to all users,

EXEC -

A few things that might be expected to change will not. The NETUSER, COMUSER, and NICUSER restrictions will remain on EXEC commands and NLS use just as they have in the past,

These refer to the restrictions placed on those who access Office-1 via the NET or commercial lines, and those under the auspices of the NIC respectively,

Commands which relate to the group allocation system and the "MESSAGE" command will not change,

Users will continue to be able to log in without typing "log"....

The differences are minor. When a user logs in, he'll be given a short statement about the date and time of his last login,

The message, "you have new journal mail" will not appear since the EXEC will no longer be involved in establishing the user's ident at login. In addition to not having idents to check for mail, it will not ask for the idents of those using multiple user directories (NLS still will),

A new command that a user can use without logging in is "NETLOAD". This is like a systat for systems that are up on the NET, producing a load average and number of users for each system,

PIE-SLICE-SCHEDULER -

The biggest change under 1.33 will be the advent of the pie-slice scheduler. In addition to regulating the number of users that can log in under an allocation group, the scheduler will split up the allocation of CPU time among users,

Initially, the scheduler will treat all users with one CPU slice and we will depend upon the existing slot allocation mechanism. When we have gained enough experience with the new system, we can allocate CPU as well as connect time to each purchased slot,

Summary of the Changes in the New Version of Tenex == 1.33

The command "GRPSTIS" will augment the GROUPSTAT and SYSTAT commands in the EXEC. It will list the number of users that are logged in under each of the system's pie-slice groups, how much of a share that group is allocated, and how the allocation compares with the amount of CPU time people logged in under that group are actually getting. Of course this will not be meaningful until the scheduler is fully implemented.

3a2

OTHER SUBSYSTEMS -

4

Some small changes will be visible in other subsystems, such as SNDMSG and READMAIL since new versions of these will be implemented under 1.33. Notably, users will have to type SNDM<esc> to get SNDMSG instead of the current alternatives. In general, 1.33 will contain the latest software produced by BBN and an improved interface to the Network, which is the primary reason for bringing the new Tenex up at this time.

4a

A trial run will take place beginning Friday afternoon this week, and if successful, 1.33 should be operating Monday morning. A log on message will announce the exact time.

5

Summary of the Changes in the New Version of Tenex == 1.33

(J32272) 8-APR-75 20:33;;; Title: Author(s): James H. Bair, Jeffrey
C. Peters/JHB JCP; Distribution: /SGR([ACTION] do you have an ident
for user services?) KWAC([INFO-ONLY]) PKA([INFO-ONLY]) SGR([
INFO-ONLY]) WEC([INFO-ONLY]) LAC([INFO-ONLY]) BEV([INFO-ONLY
]) JMB([INFO-ONLY]) RA3Y([INFO-ONLY]) RH([INFO-ONLY]) ;
Sub-Collections: SRI-ARC KWAC; Clerk: JHB; Origin: < BAIR,
MONITOR,NLS;3, >, 8-APR-75 20:15 JHB ;;;####;

User Productivity Report for April 1 and 2

DEX input and the need for TI stations....

User Productivity Report for April 1 and 2

User Productivity Report for April 1 and 2

Current Status

It has been found that TTYs should be made more available in the ARC work area (numerous requests have been made),

Three to four TI stations should be available in the general work area, and should not be removed, but retained as an alternative to DNLS for all users.

More terminals are on order as previously reported, and when they are available some terminals can be free for office/home usage.

Currently there is one direct dial line to Office-1. Tymshare has informed us that 9 additional 300 baud ports are available. We are exploring the cost of the 113 data-sets required for this access.

It is recommended, based on ARCers input, that second phones be installed in the offices of those who can justify the retention of a terminal in their office.

One DEX station is running with full time support. The temporary typist is anxious for work. DEX requests may be submitted at room J2028 (there is an inbasket).

A quick sample of late evening response from BBN showed 14 secs full screen recreate time through ELF and 11 secs through the direct line to the Tymshare TIP. If more basic problems (such as keeping the respective systems up) can be solved, samplings such as this might provide some useful information.

This may be compared to a recreate time of 4.5 sec from Office-1 under optimum conditions-- direct line, load of .3.

Ongoing Activities

Interviews and the integration of the responses into the data base is continuing. The recommendations and reports from meetings and other activities are behind somewhat due to current productivity levels and other commitments.

UPG meetings are being held most afternoons at 4PM. Any input to myself before hand or during the meeting with DCE is welcome.

User Productivity Report for April 1 and 2

(J32273) 8-APR-75 20:58;;; Title: Author(s): James H. Bair/JHB;
Distribution: /SRI-ARC([INFO-ONLY]) ; Sub-Collections: SRI-ARC;
Clerk: JHB; Origin: < BAIR, APR2UPG-REPORT,NLS;1, >, 4-APR-75
08:56 JHB ;;;;####;

IMM 9-APR-75 16:06 32276

test of file with idents

hang in there wiht me while i experiment

1

IMM 9-APR-75 16:06 32276

(J32276) 9-APR-75 16:06;;; Title: Author(s): Inez M. Mattiuz/IMM;
Sub-Collections: BELL-CANADA; Clerk: IMM;

Friday afternoon - Picnic time

We've got so many new people at ARC - it's been suggested that we have a picnic in Burgess for lunch tomorrow (12:30 to 2:00) to get to know everyone. Bring a brown bag - with food if you want to eat lunch and petty cash will supply beer and wine.

1

SGR 10-APR-75 10:20 32277

Friday afternoon - Picnic time

(J32277) 10-APR-75 10:20;;; Title: Author(s): Susan Gail
Roetter/SGR; Distribution: /SRI=ARC([ACTION]) ; Sub=Collections:
SRI=ARC; Clerk: SGR;

Copy of MIKE report to KSH re Progress To Date (mid-February) in NLS.

I haven't looked at this since the time I wrote it; I think parts of it are relevant, and you may want to copy them.

Copy of MIKE report to KSH re Progress To Date (mid-february) in NLS.

H.Q. Planning
Bell Canada

PROGRESS REPORT ON H.Q. PLANNING
EXPERIENCE WITH OFFICE-1

This is a DRAFT version

The term Computer-Augmented Management Systems has been applied to the use of computer power by managers to solve (or at least reduce) their problems in dealing with the large volumes of information flowing through their offices each day. With the aid of appropriate software routines, terminal equipment, and interaction procedures, the manager working through a Computer-Augmented Management System can expect to reduce the time required to prepare memos and reports, facilitate subsequent editing of these documents, speed the distribution of these documents to managers working in similarly computer-augmented environments (and at the same time keep track of the distribution of such documents), as well as facilitate the normal paper-bound activities that routinely clutter his desk each day, such as reminders, calendars, phone numbers, etc.

2

For the past year, the H.Q. Planning Department of Bell Canada has been heavily involved in the CAMS process. By far the most intensive effort in this area has been the involvement with the Augmentation Research Center at Stanford Research Institute. ARC developed an extensive computer software system geared toward assisting the knowledge worker in his day-to-day functions. This program is known broadly as NLS, or On-line System. This paper is a report on the progress we have made in the use of this system to date. It details our experiences in the following areas: standard applications of NLS, novel applications, use patterns and particular applications of H.Q. Planning staff, current problems with the augmentation process, improvements expected early in 1975, and an outline of our evaluation proposal.

3

(A second major aspect of a management augmentation environment is a computer conferencing facility. Bell Canada's interest in and experience with computer conferencing has been documented in "Computer Augmented Conferencing - Its Impact on Bell Canada). An update to this report is currently being drafted, and will be available by the end of December. The update will review H.Q. Planning's involvement with Turoff, CMI, conferencing based on mini-computers, and the conferencing aspects of NLS.)

3a

Standard Applications

4

NLS is used in the group for a number of standardized, fairly routine functions, which have been described in the Augmentation Research Center literature. Memos to others in the group (soft-copy output) and to other users on the system are one of the

H.Q. Planning
Bell Canada

PROGRESS REPORT ON H.Q. PLANNING
EXPERIENCE WITH OFFICE-1

This is a DRAFT version

most widely used features. In sending a memo to one person or to a group, the senders normally include a copy to themselves, thus maintaining an up-to-date file of their correspondence.

4a

Writing letters (hard-copy output) on the system is also quite popular with the H.Q. Planning users. Final drafts are either typed directly by the computer onto departmental letterhead, or a draft copy may be typed into final form by one of the clerical staff, depending on the user's preferences.

4b

The preparation of formal H.Q. Planning documents such as Business Planning Papers, presentations of research results, and submissions to technical journals is an ideal application of NLS. While a few members of the Business Planning Group still rely on pencil and paper for rough drafts (which may or may not be entered into the system for further editing), the majority of the group compose their thoughts directly at the terminal keyboard. The use of the system for this type of generation and documentation of original thought is developed to a high degree in most of the regular users of the system (where regular use would constitute at least seven to ten hours connect time per week.)

4c

The ability to monitor and comment on the work of other members of the group is increasing as a result of the group's use of NLS. By far the most common form of browsing through files is found in the Staff-Supervisor's taking regular looks at individuals' directories to see what shape a research program is taking on, what conferences or meetings have been attended, and what resulted from them, who the newer research contacts of the group are, and so on. At the peer level, a lesser degree of browsing takes place; what does occur is usually directed towards searching for a particular type of information.

4d

The idea of a developing community of knowledge workers working together and interacting heavily as a result of their large common information space has started to be realized for H.Q. Planning. It was certainly the case that this department belonged to a large invisible college of future researchers, but the gradual spread of NLS capability to diverse interest groups has opened up new avenues for development of common interests. Still in its infancy, the community of knowledge workers which is developing around H.Q. Planning include the following:

4e

H.Q. Planning
Bell Canada

PROGRESS REPORT ON H.Q. PLANNING
EXPERIENCE WITH OFFICE-1

This is a DRAFT version

Rome Air Development Center in Rome, New York is cooperating with us in identifying improvements to be made in the existing NLS framework, and forwarding these to the development group at ARC.

4e1

The ARPA group monitoring the performance and contribution of ARC is in regular communication with us to exchange information and ideas on the implementation of different aspects of NLS such as the Output Processor (sophisticated formatting of reports, including Computer-Output Microfilm routines, if desired), the Content Analyzer (sophisticated file searching routine), and DEX (software routine for off-line preparation of input).

4e2

Hudson Institute has been awarded an ARPA contract to look at methods of automating their system of Chartbooks (used as idea pieces for formal or informal presentations of research findings). Hudson chose to join the family of researchers using NLS for similar type functions, and they are currently very interested in H.Q. Planning's work in NLS-based information retrieval software. As soon as Hudson develops a little better feeling for some of the possibilities of NLS, we will be in regular communications with them, perhaps to the extent of modifying our programs to suit the particular applications they have in mind.

4e3

Novel Applications (not anticipated by ARC or us)

5

The work that we are currently doing in developing an NLS-based information retrieval package was not anticipated when we first decided to adopt an augmented management capability. As such, it is representative of a number of other applications which we have since found for NLS and its subsystems - applications which were not foreseen by us nor by any of the ARC staff.

5a

The information retrieval package that we are currently designing is based on an existing package that has been available to H.Q. Planning for some time now. This is the system that runs on the I.P. Sharpe IBM computer. It is a little cumbersome to operate, and its use requires the retriever of the information to use the IBM 2741 terminal, not his local terminal. The new NLS-based system will permit users to use their local terminals to access a custom NLS program (designed by Gwen Edwards and Penny Napke) and

H.G. Planning
Bell Canada

PROGRESS REPORT ON H.G. PLANNING
EXPERIENCE WITH OFFICE-1

This is a DRAFT version

retrieve forecasts from the Futures Information System data base, abstracts of most of the relevant futures literature contained in the department's local library, and abstracts of documents contained in national data bases such as NTIS, ERIC, and MATRIX (the on-line version of the ORBA clipping service), and Predicasts' F&S Index. This latter capability will be possible once we have developed the routine for accessing these data bases (on a subscription basis with the West Coast cataloguers), and selectively searching them and loading relevant abstracts into our own H.G. Planning futures data base. The individuals at ARC are very excited about our research in this area. Nothing like this has been attempted with NLS before, and they see real payoffs for themselves and for other users.

5b

The L-10 programming language was developed by the ARC staff in order to provide them with an appropriate high level language with which to compose the NLS operating routines; NLS is written in L-10. These programs were written by the ARC staff for the benefit of the non-technical NLS users (that is, 99% of the users), and they permit these users to perform minor editing routines, reformat statements into preselected formats, interact with the system in order to compose letter formats, etc. It was not expected that any of the users would develop sufficient interest in the NLS capabilities that they would start to write their own L-10 user-programs, but that is what has happened in H.G. Planning. We now have the capability of writing custom user-programs for reformatting statements into our own preselected formats, as well as writing larger programs to handle interactive systems such as the information retrieval package mentioned elsewhere. The insights, that we have acquired and will develop as our experience with L-10 increases, should be of great help to both ourselves and ARC in evaluating the potential contributions of systems such as this to the office environment of the future.

5c

Another unexpected development in our implementation of the concept of the augmented knowledge workshop has been our reliance on off-line preparation of data for input at a later time (with the interim data stored on cassettes of magnetic tape). The DEX routines were developed several years ago by the ARC staff, in an effort to headoff what they believed would be severe opposition to the type of physical office environment they were proposing. With the aid of DEX, it is possible for researchers to take advantage of the powerful editing and formatting capabilities of NLS as well

H.Q. Planning
Bell Canada

MIKE 10-APR-75 11:09 32278
11 APR 75

PROGRESS REPORT ON H.Q. PLANNING
EXPERIENCE WITH OFFICE-1

This is a DRAFT version

as the distribution capabilities, without the necessity of investing thousands of dollars in high-speed terminal and communications equipment. It was not until Bell Canada joined the OFFICE-1 system that this capability was fully recognized and appreciated. H.Q. Planning has found that the off-line preparation of material has other benefits as well, including the ability of NLS users to prepare material for the system even if the system is down, without the necessity of writing a draft in long-hand, and having it subsequently transcribed into the system by the clerical staff.

5d

A further benefit is the ability to compose long jobs with a number of various short inputs to them in an efficient manner. If one were keeping track of incoming bills and receipts, or perhaps documenting the responses of a number of different groups of people to a common question, it might be awkward to expect a secretary to sign onto the system and load a particular file in order to make one or two quick insertions, and then log off. With a mag, tape cassette on which to store the incoming material, the secretary would have only to flip the tape into the cassette recorder, turn the terminal to a local mode, and input the entries, leaving the tape in position to accept new entries as soon as they became available.

5e

Finally, the fact that secretaries and clerical staff members are using the system at all is surprising to most of the staff at ARC, and although H.Q. Planning certainly hoped to take advantage of the opportunity to train clerical staff on the system, there were no expectations that the venture would turn out so profitably. Where it was initially hoped that clerical staff could be interested in preparing cassette tapes for input, and perhaps eventually learning the input routines themselves, it is now the case that clerical workers are signing onto the system, inserting material from some other source, editing that material, formatting that material if necessary, updating the file, and sending the initiator of the work a message indicating that the job is now completed, and where it can be located. That this is possible at all is due in part to the care which has gone into the system and in part to the enthusiasm and aptitudes displayed by the clerical staff involved in the project.

5f

Use Pattern and Particular Applications within H.Q. Planning

6

H.G. Planning
Bell Canada

PROGRESS REPORT ON H.G. PLANNING
EXPERIENCE WITH OFFICE-1

This is a DRAFT version

Mike Bedford

6a

Mike is the most experienced NLS user in the group, and has made an effort to learn a little bit about each aspect of the NLS environment, whether it appears to be immediately productive or not. In this sense, he considers NLS as a subject for his research activities, as well as a tool for conducting those activities. He has a fairly highly developed knowledge of the Content Analyser (for making sophisticated searches through files for particular occurrences of strings or data structures), the Output Processor (permitting sophisticated formatting, including Computer-Output Microfilm routines), DEX (the routines associated with the off-line preparation of data and conversion of the data to an NLS file, and L-10, the NLS programming language. The ARC staff considers Mike to be the most experienced and advanced of the non-ARC users of NLS.

6a1

Inez Mattiuz

6b

Inez is the System Architect for the Augmented Knowledge workshop, and such, she is responsible for being familiar with all the aspects of work in NLS, despite the fact that she has been with the group for less than a year, and that she had no previous exposure to or experience with computer systems, she has met this responsibility. She is familiar with the Output Processor, with DEX, and with DNLS, the version of NLS employing the high speed terminal and the mouse-addressable cursor and five-finger keyset. (This was the original version of NLS. The version that can be addressed from a standard teletype-compatible terminal came as an afterthought to the ARC staff.) Inez should probably be spending more time with NLS, and conducting more of her work so as to take advantages of some of the NLS capabilities, but the pressures resulting from the other, non-architect, functions imposed by her job preclude this.

6b1

Larry Day
Phil Weintraub
Phil Feldman

6c

These users are all quite familiar with the basic text entry and editing capabilities of NLS. Also, they make regular use of the SENDMESSAGE capability, permitting them to send messages to

H.Q. Planning
Bell Canada

PROGRESS REPORT ON H.Q. PLANNING
EXPERIENCE WITH OFFICE-1

This is a DRAFT version

other members in the group (or to several members, if desired). This is used to keep others appraised of progress in a particular area of interest to the entire group, to ask for specific information from the group. Both Larry Day and Phil Weintraub make extensive use of the flexibility of "office hours" that the system permits, Larry having a permanent work station in his home, and Phil making heavy demands on the portable terminal. Larry has had the most experience with using the system to prepare custom reports on short notice by collecting relevant sections of previously written documents and composing them into a new document. In the future, these three individuals will receive exposure to and training in some of the additional features of NLS that, while not required to enter text and edit for later use, will significantly increase the power of the system for them. Specifically, they will learn about the journal system (as a means for storing documents in an on-line mailbox, and distributing the mailbox address to interested individuals, rather than the text of the entire document), "linking" as a means of addressing frequently cited references in a short-hand notation, and increased reliance on the structure of a document or file to make it more readily valuable to others interested in the subject material but not familiar with the logic implicit in the author's writing style.

6c1

Ken Hoyle

6d

Ken Hoyle does not use the system in a hands-on fashion, but nevertheless, his experience with NLS will provide us with a great deal of insight into the use of CAMS systems by senior managers. He uses the system to prepare (through Mary Vu) memos, notes, letters, and other correspondence. In fact, almost all his correspondence is prepared and stored on the system. He also keeps his personal calendar on-line. Since his secretary is signing on to the system several times a day, other members can use the system to check his calendar, to check the status of proposals submitted to him, and to send him messages about work in the group.

6d1

Mary Vu

6e

Mary Vu is the group's most experienced user of the Output Processor directives. These are the directives which permit formatting letters and memos to a particular format, and then

H.Q. Planning
Bell Canada

MIKE 10-APR-75 11:09 32278
11 APR 75

PROGRESS REPORT ON H.Q. PLANNING
EXPERIENCE WITH OFFICE-1

This is a DRAFT version

copying that format to other documents so that they will be printed with the same layout. (Common directives include specifications for skipping to a new page, leaving space for diagrams, indentation for different paragraphs, header statements and page number directions, and in COM, directives for choosing type fonts and type size, double column printing, etc.).

6e1

In the future, Mary will receive training in some of the user-programs, permitting her to structure the information in her files for maximum efficiency. These will include programs which will go through a file of memos she has typed, and sort them according to who they were sent, in chronological order. Also a subroutine might be written that would automatically alert her to call up specific memos at a specific date for further action or followup. Where there are several subjects that are topics of much correspondence, she might want to use a sort routine that would group them according to whether or not a particular word or phrase appeared in the text of the memo, letter, or report. These are basically simple routines that she could learn in a matter of minutes, yet which would increase the system's value to her considerably.

6e2

Don Atkinson

6f

Don has developed the ability to prepare his papers both on-line and off-line. He does not do the editing of the text on the system itself, preferring to work from a hard copy of the text and having the clerical staff amend the computer-stored version. He is familiar with the SENDMESSAGE capability, and uses it regularly to exchange messages with his counterparts at ARC. To increase the benefits received from the system, Don may find it necessary to take advantage of the group's portable terminal facility. He prefers working with the benefit of a hard copy of his material before him, and also prefers to do much of his preparation during non-"office hours". The portable terminal fills this need admirably.

6f1

Penny Napke

6g

Penny Napke is familiar with and capable in all the text entry and editing routines, although she does not use the system for preparing documents or memos as such. Most of her work in NLS

H.Q. Planning
Bell Canada

PROGRESS REPORT ON H.Q. PLANNING
EXPERIENCE WITH OFFICE-1

This is a DRAFT version

is dedicated to the Content Analyzer programs and the other types of programs that are executable (that is, they ask the user of the program for information, and then process the text or file according to that input). She is in regular contact with the programming staff of ARC, and with their encouragement is developing the NLS-based information retrieval package for the use of the rest of the group. Penny's main handicap now is that her other responsibilities in H.Q. planning do not permit her to spend enough time with the L-10 programming to make a great deal of headway. Complicating this is the fact that when she does make time for the L-10 work, the staff at ARC or the staff here in H.Q. Planning is often too busy to help her define the programming task. With the concentration of the responsibility for this project in Gwen Edwards's hands, this problem may be significantly alleviated. (During the past year, Mike Bedford, Phil Weintraub, Phil Marquette, and Inez Mattiuz have all shared responsibility for various aspects of this project.)

6g1

Gwen Edwards

6h

Gwen Edwards was exposed to this system very briefly while still at S.R.I., and she expects to be taking full advantage of it once her basic training or introductory period are completed here. It is too early to tell what aspects she will find most useful.

6h1

Jim Kollen
Mike Katsoulis

6i

Both Jim Kollen and Mike Katsoulis are reluctant to use the system for direct entry of their thoughts into a computer file. Both prefer to sit down with pencil and paper to rough out their thoughts and ideas and then turning to NLS and the clerical staff to edit the draft and format it for printing. Conventional wisdom leads one to believe that both Jim and Mike would be more productive (in a document preparation sense) if they were to make fuller use of the system's capabilities, but one of the advantages of this type of environment that has been stressed all along is that it is not necessary to convert an entire office staff in order to achieve efficiencies, nor is it necessary for any one individual to make a 100% conversion to full-blown augmentation. The system and its different

H.Q. Planning
Bell Canada

PROGRESS REPORT ON H.Q. PLANNING
EXPERIENCE WITH OFFICE-1

This is a DRAFT version

components are sufficiently flexible to adapt to individual work styles, rather than imposing particular work styles on individuals.

611

Current Problems with the Augmentation Process

7

The version of NLS that we are currently using is the eighth release of the Augmentation Research Center - NLS-8. It has been used by the ARC staff for the past year, but was made available to outside users only in November, 1974. It has a number of bugs which make it somewhat awkward to do anything out of the ordinary routine of NLS text entry and editing, and we believe that it was released too early for use by non-ARC staff. The conversion of H.Q. Planning to NLS-8 will be a gradual process, and there is some question whether the improvements we will take advantage of in NLS-8, will make up for the inconvenience caused by having to learn many new routines and commands. (This is not to say that NLS-8 was not worth the effort to develop. New users coming onto the system will find its increased prompting and greater flexibility quite useful, once the minor bugs are ironed out. Also, we suspect it to be the case that most of the members of H.Q. Planning may not have reached the stage of proficiency in NLS-7 where they would have noticed its limitations, and thus they may not fully appreciate the improvements found in NLS-8.)

7a

Each of the H.Q. Planning users of NLS has received training on the various aspects of the system needed for text entry and editing and for routing messages to different individuals, but there has been no training directly relating to the philosophy of NLS and its inherent ability to increase the information contained in a text string by taking advantage of structuring of the tree-type. Consequently, each user has developed his own relationship with the system and there is no commonality between users. If this commonality were to increase, each of the users would find that the information prepared by the other members was more valuable to him, since he would be able to access it more intuitively and manipulate it more freely. The users have not yet reached the point where we can cut and paste together information from each other's computer files, since those files have evolved independently, with no thought to the value they may hold for other users. (The question here is not so much one of maintaining similar file structures, but of adopting a similar philosophy

H.Q. Planning
Bell Canada

PROGRESS REPORT ON H.Q. PLANNING
EXPERIENCE WITH OFFICE-1

This is a DRAFT version

concerning how the system will be used by the group as a whole and keeping this philosophy in mind when composing documents.)

7b

This is a good example of how a lack of training may have lessened the full impact the system may have had on our operations, but considering that the main body of NLS experience and expertise is on the West Coast, and that our System Architect must devote a great deal of her time to non-NLS matters, we have progressed remarkably far in the past year.

7c

Improvements Expected in Early 1975

8

Since we first started to operate in the NLS environment, there has been a need for a high-speed printing capability. This was unavailable for some time for a number of reasons (indecision about which would best suit our needs, identifying interface problems with the high-speed CRT, acquiring circuit diagrams for the line processor, which makes operation in DNLS possible and which also drives the printer). The printer is now on our premises, and within a month we will have the Centronix printer producing our reports for us at a rate of 120 characters per second. This high speed will permit us to generate custom-generated reports in a matter of minutes, and have them typed out in only a few minutes more. This capability will be valuable if a member of the department faces an important deadline such as a flight departure, or if a visitor to the group wishes to take with him some of our information on a particular topic, or the minutes of the meeting he was attending (prepared while the meeting was in progress, and edited and printed out while he was putting on his coat). Before spring, we will have increased the speed of one of the multiplexor channels from 1200 baud to 4800 baud. This will permit us to operate in DNLS at a speed closer to that for which it was designed, and will further increase the speed with which we can generate printed reports.

8a

Use of DNLS will increase among the users in the department. DNLS, with its mouse and five-finger keyset provides NLS users with a very powerful editing capability; it is considerably faster for editing documents than its teletype counterpart, TNLS. To provide this increased capability will required an additional training commitment on our part, but we feel the returns justify the required resources. With the acquisition of a second portable terminal, we expect that there will be increased productivity in

H.G. Planning
Bell Canada

PROGRESS REPORT ON H.G. PLANNING
EXPERIENCE WITH OFFICE-1

This is a DRAFT version

the group, since the use characteristics of these machines in a remote work environment indicate that in most cases they supplement the amount of work done in the office, or done at home during office hours. That is, instead of merely permitting a knowledge worker to do his knowledge in a remote location, the portable terminals permit him to do more of this type of work, yet without imposing on him to any great extent.

8b

The use of portable terminals in meetings and seminars will permit Group members to have minutes recorded on the system in real-time, edited or formatted with a pre-prepared user-program, and either printed on the high-speed terminal for immediate distribution, or mailed to the participants later that same day if the meeting is away from the printer location.

8c

Evaluation of H.G. Planning Use of NLS

9

This report constitutes a progress report on H.G. Planning's use of NLS in an Augmented Knowledge Workshop environment. It is intended to provide a picture of how the system is being used at this time, and not an evaluation of the specific costs and benefits associated with its use. This evaluation is currently underway, but it is too early to accurately assess the benefits associated with the expenses incurred to date. There is a steep learning curve associated with the adoption of a new work style, and to take a snapshot view at this time would reveal this steep slope, but it would not accurately convey our impression that the slope is decreasing with time and with our additional expenditures. We believe we are just now beginning to experience during the past two to three months, the returns and dividends attached to our use of this system have begun to outweigh the costs incurred in getting to this point; not until now have the marginal benefits we've received surpassed the costs we've incurred in terms of dollars and manpower resources. During 1975, we expect these returns to accelerate. A full evaluation of our first year's experience with the system should indicate this trend. This evaluation will be completed by June, 1975.

9a

Copy of MIKE report to KSH re Progress To Date (mid-February) in NLS.

(J32278) 10-APR-75 11:09;;; Title: Author(s): Michael T.
Bedford/MIKE; Distribution: /GCE([INFO-ONLY]) ; Sub=Collections;
NIC; Clerk: MIKE;

JHB 9-APR-75 20:58 32279

User Productivity Group Status Report for April 3 through 8

DEX stuff needed and terminals for use....

User Productivity Group Status Report for April 3 through 8

User Productivity Group Status Report for April 3 through 8

Current Status

TI stations have been set up in the general work area, and the Terminette is available for use in the DEX room, J2028. The terminette was moved due to the noise problem cited by several ARCers. It now has a data set and is ready to use in relative isolation. The second Terminette did not have the lease renewed by MEH.

More TI terminals are needed for general use, particularly with the new hires. If anyone knows where a spare can be found, please let us know.

The Office landscape survey should detect the usefulness of the present terminal additions and indicate any changes. It is a periodic notation of each workstation's usage and current system conditions during the day.

Offline Support:

DEX is up and running, although it is dependant upon the direct dial line to G-1 for input. JML has worked through the numerous pitfalls and glitches while training a typist. One recommendation will surely be that DEX is smoothed out.

We plan to time the Cassette read program to determine the effective baud rate -- it appears to be much lower than capacity (partly due to the Net constraints the program now takes into account).

We checked on the TIP interface to card reader ("RJE") that was referred to once in Net lore, and found that there once was a plan to interface to a card-at-a-time device (using control characters for on and off control by the Net) but that it never got off the ground. We will also consider the Line Processor as a possible interface for the Terminette and the Net.

A DEX User Guide was found to be in need (currently, the Primers are up to date, however).

Equipment that can help in the short run is being costed: typewriters, phones for TNLS in offices, and tables. The old keyboard tables are being considered for conversion (\$200 for 6 tables).

Planned Activities

User Productivity Group Status Report for April 3 through 8

DNLS appears to be attractive to ARCers even though the reliability through ELF and AI is low and the response is consistently slow.

1b1

A controlled comparison of task completion time on DNLS vs. TNLS is in order to get a better picture of the real differences.

1b1a

The use of scope TNLS has some possibilities, and a CRT without an LP is being considered.

1b1b

The ways in which NLS resources can be used more effectively is growing.

1b1c

A TNLS Clinic is to be set up by SGR for those that are interested. Anyone who wishes to use DEX themselves should contact SGR, JML or JHB for training assistance.

1b2

Superwatch is being considered to take a look at the relative use of CPU/Connect time under various conditions. Perhaps with a large enough sample size, the ratio would be an indicator of productiveness -- the more CPU per unit connect time, the more effective use of resources.

1b3

A fundamental question has been raised concerning the optimum ratio between computer power and people. The status quo appears to be more people and less computer.

1b4

User Productivity Group Status Report for April 3 through 8

(J32279) 9-APR-75 20:58;;; Title: Author(s): James H. Bair/JHB;
Distribution: /SRI-ARC([INFO-ONLY]) ; Sub-Collections: SRI-ARC;
Clerk: JHB; Origin: < BAIR, APR8UPG-REPORT,NLS;3, >, 9-APR-75
20:53 JHB ;;;;###;