JBL 12-JAN-73 12:09 13787

Thanx

By the way, does Bill Naylor have a NIC ident?
Pax

1

JBL 12-JAN-73 12:09 13787

Thanx

(J13787) 12-JAN-73 12:09; Title: Author(s): Levin, Joel B./JBL;
Distribution: Postel, Jonathan B./JBP; Sub-Collections: NIC; Clerk: JBL;

COMMAND USAGE ANALYSIS: Report on Phase I. 1

A command usage analysis package is now part of the experimental NLS system and will soon be transferred to the running system. This report is intended to describe its function and implementation, to give examples of it use, and to propose a number of improvements and extensions for this type of analysis. 2

I. Command frequency monitoring in DNLS. 3

The DNLS system, as it runs today, is driven by a language comprising over 150 main commands - a "command" being understood here as either a one-letter or a two-letter code initiating a specific NLS operation. (note that some of these commands, for instance ej for Execute Journal or gq for Goto Query, may trigger entire subsystems that have their own command sub-structure. The sub-structures are beyond the scope of this analysis.) 3a

For the purposes of display, these commands can be roughly divided into "editing commands" and "subsystems", although these names are not perfectly accurate. 3b

In the first group we find a matrix of 8 operations (copy, delete, move,...) that apply to 11 qualifiers: branch, character, group, etc. All combinations in this matrix are valid commands, including such obscure processes as "Xset Invisible". The table of editing commands therefore contains 88 elements. In addition we include in the first table such one-letter commands as Append, Break, Null file, Quit and Viewspec. 3b1

Table 1. 3b2

	brn	chr	grp	inv	lnk	num	plx	sta	tex	vis	wrd	tot
copy	-	-	-	-	-	-	-	-	-	-	-	-
delt	-	3	-	-	-	-	-	-	-	-	-	3
inst	-	2	-	-	-	-	-	2	1	-	1	6
move	-	-	-	-	-	-	-	-	-	-	-	-
repl	-	1	-	-	-	1	-	-	1	-	1	4
subs	-	-	-	-	-	-	-	-	-	-	-	-
trsp	-	-	-	-	-	-	-	-	-	-	-	-
xset	-	-	-	-	-	-	-	-	-	-	-	-
TOT	-	-	-	-	-	-	-	-	-	-	-	-

appd 0

```
brek 1
null 1
quit 0
vspc 1      pt= 45300     rt= 782422                              3b2a
```

The second group is also a matrix of seven operations that may apply to 19 qualifiers, but not all entries are valid commands. The operations are: Execute, Freeze, Goto, Jump, Load, Output and Update. 3b3

Table 2. 3b4

	Exec	Frez	Goto	Jump	Load	Outp	Updt
a	-	-	-	-	-	-	-
b	-	-	-	-	-	-	-
c	-	-	-	-	-	-	-
d	-	-	-	-	-	-	-
e	-	-	-	-	-	-	-
f	-	-	-	-	1	-	-
h	-	-	-	-	-	-	-
i	-	-	-	6	-	-	-
j	-	-	-	-	-	-	-
l	-	-	-	-	-	-	-
m	-	-	-	-	-	-	-
n	-	-	-	-	-	-	1
o	-	-	-	1	-	-	-
p	-	-	-	-	-	-	-
q	-	-	-	-	-	-	-
r	-	-	-	-	-	-	-
s	-	-	-	1	-	-	-
t	-	-	-	-	-	-	-
u	-	-	1	-	-	-	-

3b4a

In the implementation of an internal counting mechanism for these commands it is not feasible to restrict the analysis to the combinations shown above in Tables 1 and 2, because NLS is an evolving system where new commands appear and disappear: since the beginning of this phase several commands have been added and two have been taken out (Execute WSI measurements and Execute 940 file). The analysis package uses a 30x30 matrix to accumulate command counts, so that it becomes a simple matter to reflect such changes. 3c

In the following we describe how the package can be used as a statistical tool by the individual user (section II) through a set of simple commands, and how we intend to put it to work in a general analysis of DNLS usage (section III). In the last

part (section IV) recommendations are offered for the design of the next phase. 3d

II. Using the command analysis package on-line. 4

In the early design stage for this work there was some discussion of the existing statistical commands in NLS. The discussion is documented in the Journal (12793,) (13110,) (13143,) (13191,) and led to the decision to separate timing studies of the Superwatch type from command usage analysis. The "Execute WSI Measurements" command was taken out. 4a

The analysis package which is available for on-line use is part of the "Goto Use Measurements" sub-system and is called by typing the two letters: gu. 4b

There were two sub-commands at this point (Begin and End Measurements) that triggered a timing mechanism which is probably obsolete. (Question: is this code really obsolete and, if so, could we speed up NLS by taking the measurement flag test out of all core-NLS routines?). Pending some decision on this code, I have not touched the "Begin" and "End" subcommands. Remember, however, that they are NOT part of the command usage mechanism, which is ON automatically and requires no user action to initialize it. 4c

G[oto] U[se Measurements] F[requency count] 4c1

is the command that gives a user access to the current state of the matrix. This may be followed by one of the commands: 4c2

D[isplay]
T[ables]
R[eset]
S[ave] 4c3

The Display command is intended for debugging and maintenance purposes rather than for general use. It shows the state of the entire string of counters, in rows of ten numbers with a running index to the left. 4c4

The Table option overlays the current status of Table 1 over the user's display (which is preserved and can be restored with a Command Delete). Hitting a CA (command Accept) will call Table 2. 4c5

The Reset command will reinitialize the counting machinery. 4c6

The Save command will i) automatically create a file named QBVMXYZ.LAN in the user's directory if none exists, XYZ being the user's ident, ii) write Table 1 and Table 2 as NLS statements in that file in LIFO fashion, iii) update the file and iv) reset all counters. 4c7

The file that results from use of the Save option is under user control and can be edited, formatted and processed like any other NLS file. 4c8

This set of commands therefore provides a flexible mechanism for monitoring one's own use of DNLS, running special statistical experiments, etc. 4c9

III. Automatic statistics-gathering. 5

Whenever a DNLS session terminates normally, the system executes code that has an effect similar to that of the Save command described above. The statistics that have been accumulated during the session are written out as NLS statements in that user's analysis file. 5a

A user program named "Sweeper" that we intend to run every night will gather up these statements in a single master file and reset the user files so that directory space problems are minimized. 5b

The information gathered into the master file is the following: For each DNLS session, 5c

1) The user ident
2) The date and termination time of the session
3) rt, the real time duration of the session
4) pt, the CPU time used
5) The count statistics obtained for all commands. 5c1

All times are expressed in milliseconds. 5c2

IV. Proposal for next phase. 6

The question arises of processing the information in the master file in order to produce meaningful statistics on the utilization of DNLS commands. The following is a proposal along these lines. 6a

1. Daily, 6b

A printout of the master file will be produced with viewspecs y and K (statement signatures). 6b1

A consolidated table of command usage could be generated. 6b2

Time distributions reflecting session duration (both CPU and real time) could be printed out. 6b3

2. Weekly, 6c

A table showing the distribution of most common DNLS commands for the ten largest users would seem to be a useful result. It would give some indication of variability in user behavior (programmers vs. non-programmers, for instance). 6c1

A general table of commands with their overall usage frequency would also be useful, together with a separate table of commands that have not been used at all during that period. 6c2

A general user population profile would be a good way of graphically reflecting command usage. 6c3

Individual profiles could also be generated, using as a basis the command ranking derived from observation of the entire user population. (See example below from my own experiments). 6c4

3. Monthly, 6d

We could again produce consolidated tables and profiles. 6d1

Some tables could be selected for their usefulness in timing studies. For example, assuming we had ten tables giving CPU time and real time on the basis of ten frequently-used commands, we could solve the linear system to obtain the time coefficients for each command, thus producing a very accurate measure of command efficiency. 6d2

4. Example: 6e

This proposal can be illustrated by the consolidated tables

**and a profile obtained from the three DNLS sessions that were required to enter and edit the present report. 6f

Table 3. 6f1

Command usage analysis: Report on Phase I.

	brn	chr	grp	inv	lnk	num	plx	sta	tex	vis	wrd	
copy	-	-	-	-	-	-	-	-	-	-	-	
delt	-	7	-	-	-	-	-	-	-	-	1	
inst	-	10	-	-	-	-	-	30	1	-	4	
move	-	-	-	-	-	-	-	1	-	-	-	
repl	-	5	-	-	-	1	-	-	3	-	4	
subs	-	-	-	-	-	-	-	-	-	-	-	
trsp	-	1	-	-	-	-	-	-	-	-	-	
xset	-	-	-	-	-	-	-	-	-	-	-	6f2

appd 2
brek 1
null 0
quit 0
vspc 4 pt= 307 sec. rt= 5789 sec. 6f2a

The table above was obtained by adding together the elements of three matrices resulting from the three separate DNLS sessions. Similarly we can derive a table of usage for the main subsystems: 6f2b

Table 4. 6f3

	Exec	Frez	Goto	Jump	Load	Outp	Updt	
a	-	-	-	-	-	-	-	
b	-	-	-	-	-	-	-	
c	-	-	-	-	-	-	-	
d	-	-	-	-	-	-	-	
e	-	-	-	2	-	-	-	
f	-	-	-	-	3	-	-	
h	-	-	-	-	-	-	-	
i	-	-	-	11	-	-	-	
j	-	-	-	-	-	-	-	
l	-	-	-	-	-	-	-	
m	-	-	-	-	-	-	-	
n	-	-	-	-	-	-	4	
o	-	-	-	4	-	-	-	
p	-	-	-	-	-	-	-	
q	3	-	-	-	-	3	-	
r	-	-	-	-	-	-	-	
s	-	-	-	7	-	-	-	
t	-	-	-	-	-	-	-	
u	-	-	3	-	-	-	-	6f3b

These tables lead to the following ranking of commands: 6g

30	is	5	rc	4	iw	3	eq	1	rn	
11	ji	4	v	3	rt	2	a	1	it	
10	ic	4	rw	3	lf	2	je	1	ct	
7	dc	4	jo	3	oq	1	dw	1	ms	
7	js	4	un	3	gu	1	b			6g1

And they result in the following profile: 6h

```
is  ********************************
ji  ***********
ic  **********
dc  *******
js  *******
rc  *****
v   ****
rw  ****
jo  ****
un  ****
iw  ****
rt  ***
lf  ***
oq  ***
gu  ***
eq  ***
a   **
je  **
dw  *
b   *
rn  *
it  *
ct  *
ms  *
```
6h1

These measurements indicate that in the course of the three sessions, 24 commands appeared, of which 18 were used more than once, and six more than five times. The timing studies also lead to an estimate of the overall cost of producing such a document, which (at current BBN rates) would be of about $50 for computer cost alone. 6i

There are probably many other useful results that could be derived from the information contained in the master file as described above. In a future phase, the timing studies could be refined and they could be combined with an analysis of command sequences. 6j

---------- 6j1

JFV 12-JAN-73 14:05 13788

Command usage analysis: Report on Phase I.

(J13788) 12-JAN-73 14:05; Title: Author(s): Vallee, Jacques F./JFV; Distribution: Agent, Station, Hoffman, Carol B., Lee, Susan R., Michael, Elizabeth K., Dornbush, Charles F., ARC, Guest O., Feinler, Elizabeth J. (Jake), Handbook, Augmentation Research, Kelley, Kirk E., Meyer, N. Dean, Byrd, Kay F., Prather, Ralph, White, James E. (Jim), Vallee, Jacques F., Kaye, Diane S., Rech, Paul, Kudlick, Michael D., Ferguson, Ferg R., Lane, Linda L., Auerbach, Marilyn F., Bass, Walt, Engelbart, Douglas C., Hardeman, Beauregard A., Hardy, Martin E., Hopper, J. D., Irby, Charles H., Jernigan, Mil E., Lehtman, Harvey G., North, Jeanne B., Norton, James C., Page, Cindy, Paxton, William H., Peters, Jeffrey C., Ratliff, Jake, Row, Barbara E., Van De Riet, Edwin K. (Ed), Van Nouhuys, Dirk H., Victor, Kenneth E. (Ken), Wallace, Donald C. (Smokey), Watson, Richard W., Andrews, Don I./SRI-ARC; Sub-Collections: SRI-ARC; Clerk: JFV;
Origin: <VALLEE>REPORT.NLS;17, 12-JAN-73 10:09 JFV ;

JRP 12-JAN-73 11:09 13789

APL Question

Bob, Thanks for your message. We don't have APL at UCSB. Maybe you heard "UCSD". We have Culler-Fried OLS which is contemporary and similar to APL. We hav a group that is learning how to use APL at UCSD, Multics and UCLA-CCN. We'll be glad to share our findings with you as we come by them. While I'm composing, here's a question for you. Do you have information about the Distributed Operating System for TIPS that we can have? Ron Stoughton hasn't heard anything about the effort for some time. 1

JRP 12-JAN-73 11:09 13789

APL Question

(J13789) 12-JAN-73 11:09; Title: Author(s): Pickens, John R./JRP; Distribution: Bressler, Robert D. (Bob)/RDB2; Sub-Collections: NIC; Clerk: JRP;

REK2 11-JAN-73 20:57 13790

ARPA Network Information Center
Stanford Research Institute
Menlo Park, California 94025

NIC #13637
Packet Radio Temporary Note #9

Susan Lee (SRI-NIC)
20-DEC-72

Packet Radio Meeting 1

The next meeting of the Packet Radio Group has been scheduled for January 22 and 23 at Collins Radio Corporation in Dallas, Texas. The meeting will begin at 9:00 am. each morning and will likely continue through mid or late afternoon on Tuesday. 2

If you need assistance with local arrangements, please contact Francis Dickson at Collins, (214) 235-9511 ext. 2806. I look forward to seeing you there.

3

Robert E. Kahn
ARPA
11-DEC-72 4

Masinter, Larry M.
Stanford University - Heuristic Programming
Department of Computer Science
Serra House
Stanford, California 94305

To:
Masinter, Larry M.
Stanford University - Heuristic Programming
Department of Computer Science
Serra House
Stanford, California 94305

Author Copy

REK2 11-JAN-73 20:57 13790

ARPA Network Information Center
Stanford Research Institute
Menlo Park, California 94025

NIC #13637
Packet Radio Temporary Note #9

Susan Lee (SRI-NIC)
20-DEC-72

(J13790) 11-JAN-73 20:57; Title: Author(s): Kahn, Robert E./REK2 ; Distribution: Abramson, Norman, Dickson, Francis, Frank, Howard, Kahn, Robert E., Kleinrock, Leonard, Metcalfe, Robert M. (Bob), Shapiro, Elmer B./na fd hf rek2 lk rmm ebs ; Sub-Collections: NIC; Clerk: KIRK; Origin: <KELLEY>PRT9.NLS;2, 11-JAN-73 18:13 KIRK ; --a-

LMM 11-JAN-73 22:22 13791

sample message

Hello Les. I think this is a XGP. 1

LMM 11-JAN-73 22:22 13791

sample message

(J13791) 11-JAN-73 22:22; Title: Author(s): Masinter, Larry M./LMM;
Distribution: Earnest, Lester D./LDE; Sub-Collections: NIC; Clerk: LMM;

LMM 11-JAN-73 22:28 13792

funny business

Hello Ram. This is a message. 1

LMM 11-JAN-73 22:28 13792

funny business

(J13792) 11-JAN-73 22:28; Title: Author(s): Masinter, Larry M./LMM; Distribution: Nevatia, Ram/RN; Sub-Collections: NIC; Clerk: LMM;

NJN 12-JAN-73 13:51 13793

Nobody loves a 4k Imlac

Larry-- Sorry myresponse has taken so long, but I have been extremely busy, and the information I have uncovered so far has not been encouraging. In short, there is nothing for 4k no option Imlacs except TSE, the standard text and edit package from Imlac Corp. 1

Every site that has Imlac software takes advantage of every option and every word of memory that they have. 1a

This includes BBN-TENEX with 16k and all options, UTAH with 8k and some options, and DMCG with 8k, hardware modified Imlacs. 1a1

UCLA-NMC has some Imlacs that I haven't checked out yet but they are becoming user only and it is very difficult to get onto their system now. Hence they are useless. 1a2

TX-2 does alot of graphics but mostly storage tube; their system is obscure anyway. Likewise for UCSB. Harvard is the only site I know of with a lovely higher level language with graphics primitives, but they are for an ARDS only. DMCG's Muddle (a lisp-like language) has graphics primitives but for 8k Imlac. 1a3

If you decide you want the TSE package, and you can toggle or ddt a bootstrap tty loader into your Imlac, starting at location 40 (octal), then we can try to get a TSE paper tape (we have one at BBN) into a Tenex file which you can then load into Imlac using a loader either at BBN or UTAH. 1b

You have to be careful loading over the net though; make sure that the file is sent over as 8 bit image, so you don't lose pieces of it. 1b1

Let me know what you think and I can make the necessary arrangements and give you further, more detailed instructions.
--Nancy 2

NJN 12-JAN-73 13:51 13793

Nobody loves a 4k Imlac

(J13793) 12-JAN-73 13:51; Title: Author(s): Neigus, Nancy J./NJN;
Distribution: Masinter, Larry M./LMM; Sub-Collections: NIC; Clerk: NJN;

NJN 12-JAN-73 14:24 13794

double-check

Dave-- I added and changed some items in the format; for my version see (BBN-NET,NETREF,1:w). Note (Settings), (Documentation) and (Protocols). I don't think ANTS and especially TIP belong in this document. Certainly the TIP Users Guide is available to everyone who wants one, and I presume the Illinois people provide ANTS Manuals for all their users. I thought we were documentaing servers. No need to kill ourselves.
--Nancy

NJN 12-JAN-73 14:24 13794

double-check

(J13794) 12-JAN-73 14:24; Title: Author(s): Neigus, Nancy J./NJN; Distribution: Crocker, David H./DHC; Sub-Collections: NIC; Clerk: NJN;

CHI 14-JAN-73 12:59 13795

Response to (13776,), Making nls use the Net more Eficiently

Dave, Tenex treats network lines in the following manner. Periodically, it checks to see if there is anything waiting to go to the network lines. If so, it sends whatever has accumulated so far. If we increase the time period, the user will get larger messages less often and slowwer response towhat should be quick feedback. We will look into measureing this and improve it if we can. Charles Irby. 1

CHI 14-JAN-73 12:59 13795

Response to (13776,), Making nls use the Net more Eficiently

(J13795) 14-JAN-73 12:59; Title: Author(s): Irby, Charles H./CHI; Distribution: Crocker, David H./dhc ; Sub-Collections: SRI-ARC; Clerk: CHI;

CHI 14-JAN-73 13:04 13796

Response to (13775,), Print Journal fixes

Dave, Jim white has already corrected the Print Journal problems in the experimental NLS. These fixes will become available in the running system sometime in the next week or so. -- Charles. 1

CHI 14-JAN-73 13:04 13796

Response to (13775,), Print Journal fixes

(J13796) 14-JAN-73 13:04; Title: Author(s): Irby, Charles H./CHI;
Distribution: Crocker, David H., Hopper, J. D., Kaye, Diane S., Lehtman,
Harvey G., Irby, Charles H., Kelley, Kirk E./dhc bugs ;
Sub-Collections: SRI-ARC BUGS; Clerk: CHI;

CHI 14-JAN-73 13:15 13797

Measure the network Clock frequency

Paul, has any measurement been done which would help us determine an optimal periodisity for the Schedular's process clock which outputs to network lines? Is it as good as we can get it? If no measurement has been done to date, could you look into doing some? -- Charles. 1

CHI 14-JAN-73 13:15 13797

Measure the network Clock frequency

(J13797) 14-JAN-73 13:15; Title: Author(s): Irby, Charles H./CHI; Distribution: Rech, Paul, Andrews, Don I./pr dia (I realize that this question is really directed toward you -- I am simply playing the GAME) ; Sub-Collections: SRI-ARC; Clerk: CHI;

DHC 13-JAN-73 10:58 13798

Triple-check

Nancy, Nix on .protocols. That is part of the previous branch, .programs (if you prefer, we could call is .processes.) Protocols, even if not implemented quite that way, appear as processes. Why add an extra category?

Same problem with .documentation. That is part of the .info (your name = .contact.) branch. Check my (ucla-nmc,netref,) again. You have an earlier version.

I guess I agree with you about the TIP and ANTS, for the moment at least.

Glad you put in 'Settings'. I added 'case' to it. 1

DHC 13-JAN-73 10:58 13798

Triple-check

(J13798) 13-JAN-73 10:58; Title: Author(s): Crocker, David H./DHC; Distribution: Neigus, Nancy J./NJN; Sub-Collections: NIC; Clerk: DHC;

DHC 13-JAN-73 11:03 13799

QFD

It is very irritating not be able to use the 'qfd' command in tenex,
and therefore to have to print out a rather long directory list.

I can't believe it is computationally expensive to use that commands. 1

DHC 13-JAN-73 11:03 13799

QFD

(J13799) 13-JAN-73 11:03; Title: Author(s): Crocker, David H./DHC; Distribution: Hopper, J. D., Kaye, Diane S., Lehtman, Harvey G., Irby, Charles H., Kelley, Kirk E./BUGS; Sub-Collections: NIC BUGS; Clerk: DHC;

DHC 13-JAN-73 11:21 13800

more hassle with SENDPRNT

A few extra bits of information. I did a TENEX 'Type' of <ucla-nmc>planalc.seq, with the TENEX 'indicate formfeed' mode and it showed the ↑L's.P I did it in the 'formfeed' mode and the ↑L's were stripped off. (I was on a ti, so they may have actually been sent.

That makes it seem as it sendprnt is the rogue.

1

DHC 13-JAN-73 11:21 13800

more hassle with SENDPRNT

(J13800) 13-JAN-73 11:21; Title: Author(s): Crocker, David H./DHC; Distribution: Wallace, Donald C. (Smokey)/DCW; Sub-Collections: NIC; Clerk: DHC;

DVN 12-JAN-73 22:09 13801

No Easey Way To Make Windows, User Sequence Generators, reply to (,13777)

Within NLS a hunk of code called the Sequence Generator decides in what order statments are called from the random files to be printed (or displayed). 1

The default Sequence Generator goes by a canonical walk, down each branch and strating from the top of the next. 1a

But you can write your own sequence generator, to print the statements in your file in any selection or order you please. 2

Of course you have to know some L1100 to do it. 2a

User-writen Sequence Generators are described in (journal,8043). 2b

It might be a little more pratical to write content analyser patterns for each case. E.G. in the case of folklore you couldd exclude statements in all caps by means of: 3

[LL]; 3a

For more on how to use the content analyser, see (journal,9246,10b4) and (journal,12020) 3b

DVN 12-JAN-73 22:09 13801

No Easey Way To Make Windows, User Sequence Generators, reply to (,13777)

(J13801) 12-JAN-73 22:09; Title: Author(s): Van Nouhuys, Dirk H./DVN; Distribution: Crocker, David H., Neigus, Nancy J., Forman, Ernest H., Auerbach, Marilyn F./dhc njn ehf mfa ; Sub-Collections: SRI-ARC; Clerk: DVN;

IWC 15-JAN-73 12:15 13802

ECOM communication from Ira

Comments on Sig's first communication 1

1a I have read Sig's statement of the Pitt role on the ECOM project. This message is in response to that statement and his desire to begin the dialog over the network. 1a

1b I think that our first task is to arrive at a clear statement of the purpose of the measurement effort. By this I mean some indication of what we hope to gain from the measurements. 1b

1b1 Measurement without any purpose is simply a meaningless academic exercise. A measurement program with a goal in mind will make it easier to agree on the appropriate measurement technology. 1b1

1b2 Some examples of what I mean by "purpose" of measurement are the following: 1b2

1b2a Terminal operator comparisons 1b2a

1b2b Input device evaluations 1b2b

1b2c Evaluation of particular programmed functions (e.g., facilities for inputting and/or manipulating data) 1b2c

1b2d Determination of most commonly used functions (so that they may be optimized) 1b2d

1b2e Comparisons of alternative program packages for accomplishing the same task 1b2e

1b2f Measurement of latency of response of various program stages to determine ergonomic satisfaction 1b2f

1b2g Measurement of terminal operator satisfaction 1b2g

1c I think that the example of an experiment cited in Sig's comments is too general to be of any real value. It seems to me to focus on how well the MEDEA system is taught, rather than how well it is designed. 1c

IWC 15-JAN-73 12:15 13802

ECOM communication from Ira

(J13802) 15-JAN-73 12:15; Title: Author(s): Cotton, Ira W./IWC; Distribution: Pyke, Thomas N., Treu, Siegfried/TNP ST; Sub-Collections: NIC; Clerk: IWC;

RWW 15-JAN-73 10:34 13803

Seminars for the next three weeks

Seminars for the next three weeks:
Wed Jan 17 Harvey on documentation and other conventions for NLS.
Wed Jan 24 Smokey on Bsys, the new backup file system.
Wed Jan 31 Paul and Jacque on NLS command usage counts, command timing meeasurements and related subjects. 1

RWW 15-JAN-73 10:34 13803

Seminars for the next three weeks

(J13803) 15-JAN-73 10:34; Title: Author(s): Watson, Richard W./RWW ; Distribution: Agent, Station, Hoffman, Carol B., Lee, Susan R., Michael, Elizabeth K., Dornbush, Charles F., ARC, Guest O., Feinler, Elizabeth J. (Jake), Handbook, Augmentation Research, Kelley, Kirk E., Meyer, N. Dean, Byrd, Kay F., Prather, Ralph, White, James E. (Jim), Vallee, Jacques F., Kaye, Diane S., Rech, Paul, Kudlick, Michael D., Ferguson, Ferg R., Lane, Linda L., Auerbach, Marilyn F., Bass, Walt, Engelbart, Douglas C., Hardeman, Beauregard A., Hardy, Martin E., Hopper, J. D., Irby, Charles H., Jernigan, Mil E., Lehtman, Harvey G., North, Jeanne B., Norton, James C., Page, Cindy, Paxton, William H., Peters, Jeffrey C., Ratliff, Jake, Row, Barbara E., Van De Riet, Edwin K. (Ed), Van Nouhuys, Dirk H., Victor, Kenneth E. (Ken), Wallace, Donald C. (Smokey), Watson, Richard W., Andrews, Don I./sri-arc ;
Sub-Collections: SRI-ARC; Clerk: JBN;

bill naylors nic ident is WEN. 1

(J13804) 13-JAN-73 13:11; Title: Author(s): Postel, Jonathan B./JBP; Distribution: Levin, Joel B., Naylor, William E./JBL WEN; Sub-Collections: NIC; Clerk: JBP;

JBP 13-JAN-73 13:17 13805

ucla-nmc is not a service site and has no macro processot.
suggest you refer to the resources notebook maintained by the
network information center. i also suggest that it is inapproiate
to direct such requests to the princple investigators but rather
to the network technical liasons. 1

(J13805) 13-JAN-73 13:17; Author(s): Postel, Jonathan B./JBP; Distribution: Silberski, Robert, Poh, Susan S./RS2 SSP; Sub-Collections: NIC; Clerk: JBP;

JBP 13-JAN-73 13:24 13806

i am somewhat reluctant to specify that systat,netstat,date-time are official functions assigned to standard sockets without having a document specifing their characteristics. such a specification must be implementable on all hosts not just tenexs. could you look into this situation and perhaps come up with a specification for these functions? --jon 1

JBP 13-JAN-73 13:24 13806

(J13806) 13-JAN-73 13:24; Author(s): Postel, Jonathan B./JBP;
Distribution: Neigus, Nancy J./NJN; Sub-Collections: NIC; Clerk: JBP;

Response to LPD (Journal 13688)

Ref: Journal 13617

Peter:

Susan Poh's request for information concerning 'Macro - Processors' arose from our fulfilling an ARPA contract to assist potential ARPANET users.

Specifically, the inquiry came from the Urban Mass Transportation Administration of the Department of Transportation (DOT.) Their interest is relative to the development of an Urban Mass Transportation Planning system.

Please re-consider your response and assist us in answering the original question tendered.

Incidentally, we did find a macro processor which developes code for FORTRAN and GLYPNIR. The processor was written by Klaus H. and Gertrud Hain of the Naval Research Laboratory, Washington, D.C. 20390.

Thank you for your re-consideration.

Most sincerely,
Jean Iseli

1

JI 15-JAN-73 11:20 13807

(J13807) 15-JAN-73 11:20; Author(s): Iseli, John/JI; Distribution: Deutsch, L. Peter, Mitchell, James G., North, Jeanne B., Newell, Allen, McCarthy, John, Dines, Thomas R., Roberts, Lawrence G., Heart, Frank E., Glaser, Edward L., Marill, Thomas M., Cheatham, T. E., Stone, Duane L., Forgie, James W., Vorhaus, Alfred H., Pyke, Thomas N., Uncapher, Keith W., Weissman, Clark, Feigenbaum, Edward A., Kleinrock, Leonard, Harris, David O., Pratt, William K., Evans, David C., Engelbart, Douglas C., Raphael, Bertram, Slotnik, Daniel L./LPD PI; Sub-Collections: NIC PI; Clerk: JI;

AAM 16-JAN-73 10:56 13808

test

this is a test

test

a 1

AAM 16-JAN-73 10:56 13808

test

(J13808) 16-JAN-73 10:56; Title: Author(s): McKenzie, Alex A./AAM;
Distribution: McKenzie, Alex A./AAM; Sub-Collections: NIC; Clerk: AAM;

AAM 16-JAN-73 6:43 13809

Nancy, This is the note from Postel re:FTP

JBP 8-JAN-73 10:40 13728
objection to rfc 414
Message: I think it is a mistake to have line travel from the server to the user without a numeric code prefix even in status replies. this proposal strikes me as ligimatizing some programmers lazyness a puts an inconsistency into the protocol. i urge that every line from server to user have a numeric code prefixed. if necessary a code could be defined as continuation-of-previous-line.

1

AAM 16-JAN-73 6:43 13809

Nancy, This is the note from Postel re:FTP

(J13809) 16-JAN-73 6:43; Title: Author(s): McKenzie, Alex A./AAM; Distribution: Neigus, Nancy J./NJN; Sub-Collections: NIC; Clerk: AAM;

AAM 16-JAN-73 8:53 13810

OUTPUT DEVICE TELETYPE makes page numbering mistakes, maybe

The Output Device Teletype command seems to screw up page numbers a little bit. For example, when I printed out the document which described the "BUGS and NP" system (mjournal,13733) it took 2 Teletype pages; the first had no page number and the second had page number 1.

AAM 16-JAN-73 8:53 13810

OUTPUT DEVICE TELETYPE makes page numbering mistakes, maybe

(J13810) 16-JAN-73 8:53; Title: Author(s): McKenzie, Alex A./AAM; Distribution: Hopper, J. D., Kaye, Diane S., Lehtman, Harvey G., Irby, Charles H., Kelley, Kirk E./BUGS; Sub-Collections: NIC BUGS; Clerk: AAM;

AAM 16-JAN-73 8:57 13811

Use of QUOTATION MARKS in JOURNAL TITLES

It seems to be impossible, at least the last time I tried it, to put quotation marks in the "Title" of a Journal message. The Journal system seems to take the first quotation mark as some sort of delimiter. 1

AAM 16-JAN-73 8:57 13811

Use of QUOTATION MARKS in JOURNAL TITLES

(J13811) 16-JAN-73 8:57; Title: Author(s): McKenzie, Alex A./AAM; Distribution: Hopper, J. D., Kaye, Diane S., Lehtman, Harvey G., Irby, Charles H., Kelley, Kirk E./BUGS; Sub-Collections: NIC BUGS; Clerk: AAM;

BAH 16-JAN-73 10:08 13812

Dec 26-29, A Week In Review

You will note another column has been added to the analysis file, CON/CPU, which is the reverse of CPU/CON and in different terms. This is to give a yet clearer picture of what the figures in the ratios represent.

Dec 26-29, A Week In Review

1

WEEK: DEC 26-29, 1972 8:00-17:00 2

IDENT	CPU HRS	CON HRS	CPU/CON	% SYS	CON/CPU
(DIA)	.186	5.232	.036	.8	28:1
(MFA)	.761	12.033	.063	3.2	16:1
(WLB)	.101	1.253	.081	.4	12:1
(KFB)	0.000	0.000	0.000	0.000	00:0
(CFD)	.735	18.183	.040	3.1	25:1
(DCE)	.819	20.375	.039	3.4	25:1
(JAKE)	.247	5.667	.044	1.0	23:1
(WRF)	0.000	0.000	0.000	0.000	00:0
(BAH)	.243	5.780	.042	1.0	24:1
(MEH)	.435	13.773	.032	1.8	31:1
(JDH)	.037	3.133	.012	.2	83:1
(CHI)	.200	5.586	.036	.8	28:1
(MEJ)	.123	4.428	.028	.5	36:1
(DSK)	1.184	22.411	.053	5.0	19:1
(KIRK)	.532	15.023	.035	2.2	28:1
(MDK)	.232	4.376	.049	1.0	20:1
(LLL)	.020	.246	.081	.1	12:1
(SRL)	.004	.097	.041	0.0	24:1
(HGL)	.613	15.021	.041	2.6	24:1
(NDM)	1.242	18.795	.066	5.2	15:1

2a 2b 2c 2d 2e 2f 2g 2h 2i 2j 2k 2l 2m 2n 2o 2p 2q 2r 2s 2t 2u 2v

(EKM)	.183	10.291	.018	.8	56:1	2w
(JBN)	.255	15.756	.016	1.1	62:1	2x
(JCN)	0.000	0.000	0.000	0.000	00:0	2y
(WHP)	0.000	0.000	0.000	0.000	00:0	2z
(JCP)	1.117	17.493	.064	4.7	16:1	2a@
(JR)	0.000	0.000	0.000	0.000	00:0	2aa
(PR)	.112	4.589	.024	.5	42:1	2ab
(BER)	1.205	17.430	.069	5.1	14:1	2ac
(JFV)	.412	10.996	.037	1.7	27:1	2ad
(EKV)	0.000	0.000	0.000	0.000	00:0	2ae
(DVN)	.728	10.951	.066	3.1	15:1	2af
(KEV)	0.000	0.000	0.000	0.000	00:0	2ag
(DCW)	.768	22.164	.035	3.2	28:1	2ah
(RWW)	.329	9.695	.034	1.4	29:1	2ai
(JEW)	.525	6.748	.078	2.2	13:1	2aj

3

Highest CPU User: NDM 1.242 hrs Lowest CPU User: SRL .004 hrs 4

Highest CON User: DSK 22.411 hrs Lowest CON user: SRL .097 hrs 5

Highest CPU/CON: LLL .081 Highest CON/CPU: JDH 83:1 6

Average Ratio Of Group: .045 Total % Sys By Group: 56.1% 7

BAH 16-JAN-73 10:08 13812

Dec 26-29, A Week In Review

(J13812) 16-JAN-73 10:08; Title: Author(s): Hardeman, Beauregard A./BAH ; Distribution: Hoffman, Carol B., Lee, Susan R., Michael, Elizabeth K., Dornbush, Charles F., ARC, Guest O., Feinler, Elizabeth J. (Jake), Handbook, Augmentation Research, Kelley, Kirk E., Meyer, N. Dean, Byrd, Kay F., Prather, Ralph, White, James E. (Jim), Vallee, Jacques F., Kaye, Diane S., Rech, Paul, Kudlick, Michael D., Ferguson, Ferg R., Lane, Linda L., Auerbach, Marilyn F., Bass, Walt, Engelbart, Douglas C., Hardeman, Beauregard A., Hardy, Martin E., Hopper, J. D., Irby, Charles H., Jernigan, Mil E., Lehtman, Harvey G., North, Jeanne B., Norton, James C., Page, Cindy, Paxton, William H., Peters, Jeffrey C., Ratliff, Jake, Van De Riet, Edwin K. (Ed), Van Nouhuys, Dirk H., Victor, Kenneth E. (Ken), Wallace, Donald C. (Smokey), Watson, Richard W., Andrews, Don I./SRI-ARC ; Sub-Collections: SRI-ARC; Clerk: BAH ;

JAN 2-5, A Week In Review

BAH 16-JAN-73 13:40 13813

JAN 2-5, A Week In Review

1

WEEK: JAN 2-5, 1973 8:00-17:00 2

2a

IDENT	CPU HRS	CON HRS	CPU/CON	% SYS	CON/CPU
					2b
(DIA)	.174	6.369	.027	.9	37:1 2c
(MFA)	.488	8.984	.054	2.4	18:1 2d
(WLB)	0.000	0.000	0.000	0.000	00:0 2e
(KFB)	.084	4.297	.020	.4	50:1 2f
(CFD)	.506	12.656	.040	2.5	25:1 2g
(DCE)	.630	19.913	.032	3.2	31:1 2h
(JAKE)	.004	.056	.071	0	14:1 2i
(WRF)	.921	11.122	.083	4.6	12:1 2j
(BAH)	3.341	25.321	.132	16.8	8:1 2k
(MEH)	.284	8.334	.034	1.4	29:1 2l
(JDH)	.035	2.531	.014	.2	71:1 2m
(CHI)	1.547	22.256	.070	7.8	14:1 2n
(MEJ)	.161	10.660	.015	.8	67:1 2o
(DSK)	.278	4.974	.056	1.4	18:1 2p
(KIRK)	.419	11.130	.038	2.1	26:1 2q
(MDK)	.514	14.280	.036	2.6	28:1 2r
(LLL)	0.000	0.000	0.000	0.000	00:0 2s
(SRL)	.054	1.572	.034	.3	29:1 2t
(HGL)	.797	16.632	.048	4.0	21:1 2u
(NDM)	.516	11.063	.047	2.6	21:1 2v

(EKM)	.080	5.153	.016	.4	62:1	2w
(JBN)	.136	4.661	.029	.7	34:1	2x
(JCN)	.466	7.807	.060	2.3	17:1	2y
(WHP)	.001	.009	.111	0	9:1	2z
(JCP)	1.104	17.967	.061	5.5	16:1	2a@
(JR)	0.000	0.000	0.000	0.000	00:0	2aa
(PR)	.225	9.555	.024	1.1	42:1	2ab
(BER)	.290	8.554	.034	1.5	29:1	2ac
(JFV)	.478	11.990	.040	2.4	25:1	2ad
(EKV)	0.000	0.000	0.000	0.000	00:0	2ae
(DVN)	.177	5.463	.032	.9	31:1	2af
(KEV)	1.284	17.732	.072	6.4	14:1	2ag
(DCW)	.221	9.295	.024	1.1	42:1	2ah
(RWW)	.061	1.366	.045	.3	22:1	2ai
(JEW)	.377	8.005	.047	1.9	21:1	2aj

3

Highest CPU User: BAH 3.341 hrs Lowest CPU User: WHP .001 hrs 4

Highest CON User: BAH 25.321 hrs Lowest CON user: WHP .009 hrs 5

Highest CPU/CON: BAH .132 Highest CON/CPU: JDH 71:1 6

Average Ratio Of Group: .046 Total % Sys By Group: 78.5% 7

BAH 16-JAN-73 13:40 13813

JAN 2-5, A Week In Review

(J13813) 16-JAN-73 13:40; Title: Author(s): Hardeman, Beauregard A./BAH ; Distribution: Hoffman, Carol B., Lee, Susan R., Michael, Elizabeth K., Dornbush, Charles F., ARC, Guest O., Feinler, Elizabeth J. (Jake), Handbook, Augmentation Research, Kelley, Kirk E., Meyer, N. Dean, Byrd, Kay F., Prather, Ralph, White, James E. (Jim), Vallee, Jacques F., Kaye, Diane S., Rech, Paul, Kudlick, Michael D., Ferguson, Ferg R., Lane, Linda L., Auerbach, Marilyn F., Bass, Walt, Engelbart, Douglas C., Hardeman, Beauregard A., Hardy, Martin E., Hopper, J. D., Irby, Charles H., Jernigan, Mil E., Lehtman, Harvey G., North, Jeanne B., Norton, James C., Page, Cindy, Paxton, William H., Peters, Jeffrey C., Ratliff, Jake, Van De Riet, Edwin K. (Ed), Van Nouhuys, Dirk H., Victor, Kenneth E. (Ken), Wallace, Donald C. (Smokey), Watson, Richard W., Andrews, Don I./SRI-ARC ; Sub-Collections: SRI-ARC; Clerk: BAH ;

JLM 15-JAN-73 6:43 13814

iadpcommitee agenda

PURPOSE:To resolve the best vehicle to provide suuport tp admin and also to identify a piece of the admin job which was feasible to address. 1

ATTENDEES: 2

Col wilson ,col cavel(amin chief)major zara and lt.col o,keefe,(esd),capt ?(afsc/acd,f.allen,a marks,and ?(admin staff office) and j mac (radc). 2a

DISCUSSION:Major Zara opened the meet by describing his view of the problem.He stressed that it seemed to esd that teir was technology around but it had never been applied systematically to this area and worse yet because the area was not directly weapons systems orientation and with the austere budget it was tough to defend this kind of program.Col Wilson echoed agreement and stated that he was coonvinced that the need was tere but only a proram that talked of cost saving had a chance and that know pitch of improved communication or productiivity would sell at this particluar time. 3

JLM 15-JAN-73 6:43 13814

iadpcommitee agenda

(J13814) 15-JAN-73 6:43; Title: Author(s): McNamara, John L./JLM;
Sub-Collections: RADC; Clerk: JLM;
Origin: <MCNAMARA>FADMINTRIP.NLS;1, 11-JAN-73 10:05 JLM ;

JRP 15-JAN-73 10:28 13815

AMES RESOURCE EVALUATION

Bill, we have organized another group of students here to investigate ARPANET resources for actual user usage. At least one student wants to investigate the ILLIAC, TENEX, 360/67 complex at ames but I don't know if there is any "meat" for his project as of yet. We'd like to find and document network resources that might be useful to the local community here (and thus to the entire network also) . If there is anything really worth looking into at this point in time could you let me know?? Thanks. P.S. The kid is approaching 22 pounds. 1

JRP 15-JAN-73 10:28 13815

AMES RESOURCE EVALUATION

(J13815) 15-JAN-73 10:28; Title: Author(s): Pickens, John R./JRP;
Distribution: Jones, William P./WPJ; Sub-Collections: NIC; Clerk: JRP;

JRP 15-JAN-73 12:08 13816

Status of SCRL

Here are your answers. 1) initial limited operation is scheduled for 1 FEB 1973...full operation is scheduled for 1 MAR 1973. 2) they are running on a PDP-11 with a modified version of the ANTS operating system. 3)the person to get in touch with regarding more information is DAVE RETZ at either 805-965-3011(SCRL) or 805-961-3221(UCSB CSL). Hope that helps. 1

JRP 15-JAN-73 12:08 13816

Status of SCRL

(J13816) 15-JAN-73 12:08; Title: Author(s): Pickens, John R./JRP; Distribution: Neigus, Nancy J./NJN; Sub-Collections: NIC; Clerk: JRP;

Thanks for your reply (it was quick enough). I would like to be able 1

to load TSE from over the net, as all we have now is a home brew 2

thing that has not even the TSE graphics stuff. We do have the ROM 3

bootstrap loader, so all I need is the file on TENEX. There isn't 4

much urgency in this as of yet, as I am waiting until we get 5

a higher speed interface (right now we only have 10 cps). I agree 6

that it is a primative IMLAC, but after all, it is serial number 1. 7

Thanks again, 8

Larry Masinter 9

(J13817) 13-JAN-73 15:12; Author(s): Masinter, Larry M./LMM;
Distribution: Neigus, Nancy J./NJN; Sub-Collections: NIC; Clerk: LMM;
Origin: <SU-HP>REQ.NLS;1, 13-JAN-73 15:07 LMM ;

LMM RC BGB NSS 13-JAN-73 3:53 13818

Request for accounts

We would like to request the opening of three new accounts on the USC-ISI facility. These will be used to further our work on the Stanford Heuristic Programming Project. We expect to use not more than 200 disk pages and roughly 10 hrs/month of CPU time with each account.

The three accounts will be under the names (Bruce) Buchanan, (Ray) Carhart and ( ) Sridharan. We can use the password "DENDRAL."

We have communicated with John Melvin concerning these accounts, and he suggested that this request be checked with you.

Thank You

Dr. R. E. Carhart

LMM RC BGB NSS 13-JAN-73 3:53 13818

Request for accounts

(J13818) 13-JAN-73 3:53; Title: Author(s): Masinter, Larry M., Carhart, Ray, Buchanan, Bruce G., Sridharan, N. S./LMM RC BGB NSS; Distribution: Dolan, Bruce A./BAD; Sub-Collections: NIC; Clerk: LMM; Origin: <SU-HP>REQ.NLS;1, 13-JAN-73 3:47 LMM ;

MCK 16-JAN-73 12:23 13819

Jim--

I got your list of junk SMFS files (thanks). It's always nice to regain unused disk space. So you encode host address in filename, eh? And here I thought 'AC' stood for 'ArChive'. If I had done that pgm, I'd have used hex designation rather than counting letters of the alphabet: '02' for SRI-ARC, '1F' for CCA, etc. Each to his own, I guess. 1

I'm glad you like my SMFS mapdisk; it was a lot of fun to write. I didn't even need too much help from Steve McCampbell. I expected you to suggest that I print out passwords, but I also thought you would suggest that I list the files in alphabetical order by owner's userid, also by alphabetical order. If you had, I'd have promptly told you where to go, and then agreed with you. It's a litle too hard for a dinky little one-week project. Besides, I don't have enough room across the page for data set name, file name, and two passwords. I really don't need data set name, except when certain clowns leave junk files around - it's easier to scratch them with IEHPROGM by ds name than with SMFS by filename. 2

Yes, I save userid, account, and site in the file. I wanted to store it all in the DSCB so that my mapdisk program did not have to open the data set to retreive it, but I could find no way (in problem state) to store stuff there. So, I just tacked it on the end of the first header record behind the first password. That way, SMFS doesn't mind if its there or not on existing files. Due to the (sloppy) way this is done, it is very unlikely that an SMFS directory command will ever be written (sigh). 3

By the way, you might like to know that there is a bug in your SMFS interface. If one says q[uit], and then CONTINUE, someone, probably you, types CR-LF and the word 'none', CR-LF and your prompt character. If one then says '?', he gets a list of your commands, but the word 'quit' has the 'q' missing (no, I'm not kidding). At that point, a 'q' is refused, and a 'u' is recognized as 'userid'. Take that, Superwhite. Also, it looks to me and my log like you still allocate too small a file. 4

Anyway, happy programming

--Mark 5

(J13819) 16-JAN-73 12:23; Author(s): Krilanovich, Mark C./MCK;
Distribution: White, James E. (Jim)/JEW; Sub-Collections: NIC; Clerk:
MCK;
Origin: <UCSB>JL.NLS;1, 16-JAN-73 12:21 MCK ;

MDK 15-JAN-73 9:57 13820

Disk Space Allocation Study

This is in response to DIA's comments on the disk space allocation request that NIC made to Analysis.

MDK 15-JAN-73 9:57 13820

Disk Space Allocation Study

This is in response to DIA's comments [1] on the disk space allocation request that NIC made to Analysis [2]. I really want to thank Don for pointing out the need for more information, and for asking some specific questions. I hope this note clarifies the problem as we see it. This note is organized as follows: 1

Current Utilzation 2

Statistics 2a

What follows is a summary of actual disk allocation on one day during the week of January 1st. (Given the relatively static nature of file directories, I would assume that this one sample is fairly representative of what goes on day by day. However, more statistics are clearly needed.) 2a1

SR	ND	TP
0 - 125	77	3425
125 - 250	26	4213
250 - 500	17	6161
500 - 1000	3	2417
1000 - 2000	9	12422
2000 - up	2	4467
totals	134	33105

2a1a

2a1b

where: 2a2

SR = File Space Allocation Range (in Pages)
ND = Number of Directories in this Range
TP = Total Pages Used by Directories in this Range 2a2a

Observations 2b

1) Approximately 10% of the directories (14 of 134) use about 58% of the disk space (19306 of 33105 pages); 2b1

2) Approximately 25% of the directories (31 of 134) use about 77% of the space (25467 of 33105 pages); 2b2

3) Approximately 75% of the directories (103 of 134) use only 25% of the space (7638 of 33105 pages). This group includes all individual users except DCE, DIA, NDM, DCW, WSD, KEV, and JCN. 2b3

Further Information 2c

Of the 14 largest directories (those with 500 or more pages),

6 are program directories ... ... (total, 8614 pages)
TENEX, REL-NLS, NLS, MPS, SUBSYS, NIC-NLS

3 are journal directories ... ... (3538 pages)
JOURNAL, KJOURNAL, MJOURNAL

4 are NIC directories ... ... (6446 pages)
DOCUMENTATION, CATALOG, NIC-WORK, NIC

1 is a user directory ... ... (708 pages)
NORTON. 2c1

As mentioned above, the total space used by these 14 directories is 58% of all used disk space. 2c2

Current Practices 3

What Happens When Disk Space Shortage Becomes Acute 3a

When in a panic situation (no more disk space), the computer says, in effect, 3a1

"I'M out of space and I can't do anything about it. No user can continue to work until you all get rid of some files." 3a1a

What the system OUGHT to say is, 3a2

"YOU'RE out of space, Mr. User. You've exceeded your allocation and you'll have to get rid of some of your files. Meanwhile, the other users can continue to do work unless they also exceed their allocations." 3a2a

Note: This in fact sometimes happens in NLS, I know. It's happened to me. But usually it is the computer, and not a user, that runs out of space, and

this then penalizes ALL users, even those that haven't exceeded their allocations. 3a2a1

Since usually it is the computer, and not a user, that runs out of space, Operations must save the day. What they do is the following: 3a3

1) run DELD to expunge deleted files; 3a3a

2) urge users to delete some of their files; 3a3b

3) delete some redundant versions of files in certain large directories (NIC-NLS, DOCUMENTATION, NLS, for example) and put them back when the users request it, or when the space shortage is alleviated. 3a3c

4) run BSYS to archive files that have been inactive for N days (whatever N is) and to archive files that users have requested be archived; 3a3d

5) run TRIM in the daytime to delete all but the latest version of each file. 3a3e

NOTES: 3a3f

- Step 5) is a last resort; TRIM is normally not run in the daytime. 3a3f1

- Step 3) is where the NIC sometimes gets hurt, since what appear to be redundant versions may in fact not be, as far as the NIC staff member who is working with that file is concerned. 3a3f2

The reason that NIC sometimes gets hurt by Step 3) whereas other users don't, is simply that NIC has the largest files (other than programs) on-line, and so it is easiest for operations to recover larger amounts of space from the NIC directories than from any other directories. 3a3f2a

- Step 2) has rather limited payoff, in my opinion, since the users only account for 25% or less of the disk space used. 3a3f3

The Allocations As Meaningless Numbers 3b

The reason that the allocations are meaningless numbers is that since they are so large, they do not act as an

effective control on the use of disk space. That is, the computer runs out of space before any user does, which is not the way it should be. 3b1

In other words, why should I worry about my file space if I've been conditioned to know that (a) I can get all the space I want as long as I beat the others to what's available, and (b) chances are very good that the system will run out of space before I do, since everyone is getting all the space he wants. 3b1a

Note: I know this isn't quite true; everyone doesn't get all the space he wants. But the effect of excessively large allocations is essentially just that, in my opinion. 3b1a1

My contention is that if allocations were by groups of persons doing related work, and if the sum of all allocations were less than the physically available space, not only would the system probably never run out of space, but: 3b2

ONLY USERS IN A GROUP WOULD RUN OUT OF SPACE, AND IT WOULD BE UP TO THEM TO DELETE, ARCHIVE, OR WHATEVER, TO KEEP THEIR USAGE IN ACCORDANCE WITH THEIR BUDGET. 3b2a

Incidentally, I don't profess to know how this should really be implemented. I am merely trying to give a different viewpoint and am asking Analysis to study the problem thoroughly. 3b3

For example, couldn't we look at the whole problem of what CATEGORIES of disk space are needed, and set up a system that addresses this problem? 3b3a

Some categories might be : 3b3b

active program files
seldom-used (archivable) program files
NIC files used by NIC staff
NIC files used by Network persons
active user files
temporary user files
archivable user files
files that must remain on-line no matter how seldom they are used 3b3b1

Specific Comments on DIA's Response 4

I hope the above discussion gives a better framework than my previous journal item did for the type of study I think is necessary. To answer specifically the points raised in DIA's response: 4a

1) The Allocation Numbers. 4b

In my opinion, it isn't a question of whether everyone is going to be pushing his allocation limit at the same time. 4b1

If everyone does, then we know we either need more physical disk space, or we need to question seriously what files we consider "must" be on line. 4b1a

Rather, what is at question is whether the addition of one file (or statement) should usually cause ALL users to scurry and get rid of files. I don't think this should be the case. 4b2

Therefore I think a better solution than the present one of over allocation is one that gives space limitations which 4b2a

a) force users to think about what they have on line, and 4b2a1

b) force only those users that exceed their allocations to have to delete or archive files. 4b2a2

It is this type of thinking which led me to state that in the present method, the allocation numbers are meaningless as an instrument for controlling file space utilization. 4b2b

In my opinion, the sum of all space allocations should be less than the total available space, to provide some safety margin for temporary files and the like. 4b2b1

2) The NIC Files 4c

The situation on the NIC files is primarily due to the fact that they are large files in large directories which are used by several users. 4c1

As I have mentioned above, when Operations has to save the day, the NIC directories are one of the first places it

goes to delete files. This sometimes, though not always, causes serious delays for NIC staff members. 4c2

3) Archiving 4d

The present method of archiving files that haven't been referenced in x days is unsatisfactory because it doesn't fully take into account the statistics of the situation: Namely, 77% of file space is used by 25% of the directories, and most of these files are usually not deletable or archivable (programs, documentation, and the like). 4d1

Consequently, automatic archiving affects the small directories much more than the large directories, and can't really alleviate our space problem substantially. 4d2

I would prefer a method which puts the responsibility on the user, and provides him an incentive to keep only those file on line that MUST be on line. 4d3

References: 5

[1] (HJOURNAL, 13720, 1:w) DIA's response seeking more information. 5a

[2] (HJOURNAL, 13690, 1:w) NIC's initial request to Analysis 5b

MDK 15-JAN-73 9:57 13820

Disk Space Allocation Study

(J13820) 15-JAN-73 9:57; Title: Author(s): Kudlick, Michael D./MDK; Distribution: Hoffman, Carol B., Lee, Susan R., Michael, Elizabeth K., Dornbush, Charles F., ARC, Guest O., Feinler, Elizabeth J. (Jake), Handbook, Augmentation Research, Kelley, Kirk E., Meyer, N. Dean, Byrd, Kay F., Prather, Ralph, White, James E. (Jim), Vallee, Jacques F., Kaye, Diane S., Rech, Paul, Kudlick, Michael D., Ferguson, Ferg R., Lane, Linda L., Auerbach, Marilyn F., Bass, Walt, Engelbart, Douglas C., Hardeman, Beauregard A., Hardy, Martin E., Hopper, J. D., Irby, Charles H., Jernigan, Mil E., Lehtman, Harvey G., North, Jeanne B., Norton, James C., Page, Cindy, Paxton, William H., Peters, Jeffrey C., Ratliff, Jake, Van De Riet, Edwin K. (Ed), Van Nouhuys, Dirk H., Victor, Kenneth E. (Ken), Wallace, Donald C. (Smokey), Watson, Richard W., Andrews, Don I./sri-arc ; Sub-Collections: SRI-ARC; Clerk: MDK;
Origin: <KUDLICK>DISKSPACE.NLS;11, 15-JAN-73 8:50 MDK ;

NJN 15-JAN-73 10:40 13821

Socket Numbers

The socket number list in RFC 433 showed that there is a wide variation in socket assignments. In an effort to correct this we have suggested the following:

Socket	Assignment
0-63	Network Wide Standard Functions
1	Telnet
3	File TRansfer
5	Remote Job Entry
7	Echo
9	Discard
11	Systat
13	Date-Time Function
15	Netstat
19	Character Generator
64-127	Host Specific Functions
65	Speech Data Base at tx-2
67	Datacomputer at CCA
69	CPYNET

For Tenex sites this requires three changes. Echo moves to socket 7, CPYNET to 69, and SYSTAT to 11. (These numbers are in decimal.)

We would appreciate your making these changes promptly to comply with the rest of the sites on the net. Thank you, Nancy Neigus (BBN)

NJN 15-JAN-73 10:40 13821

Socket Numbers

(J13821) 15-JAN-73 10:40; Title: Author(s): Neigus, Nancy J./NJN;
Distribution: McConnell, John W./JWM; Sub-Collections: NIC; Clerk: NJN;

NJN 16-JAN-73 6:33 13822

new sockets

jon-- Bob Thomas called me yesterday and requested to socket assignments for the RSEXEC he is working on now. Since this is still of a somewhat experimental nature I gave him 247 and 249. I sent messages to all Tenex sites re socket changes; we will see who the cooperative ones are now. Am updating my copy of phonebook. Received your package; letter is on its way. --nancy 1

NJN 16-JAN-73 6:33 13822

new sockets

(J13822) 16-JAN-73 6:33; Title: Author(s): Neigus, Nancy J./NJN;
Distribution: Postel, Jonathan B./JBP; Sub-Collections: NIC; Clerk: NJN;

NJN 16-JAN-73 6:36 13823

BBN Report request

Susan-- BBN Report 2491 is going to the printers some time this week. It should be back at the beginning of next week at which time a copy would ordinarily be sent ot the NIC. If you want more than this one copy please let me know before then and we will send as many as you request. --Nancy 1

NJN 16-JAN-73 6:36 13823

BBN Report request

(J13823) 16-JAN-73 6:36; Title: Author(s): Neigus, Nancy J./NJN;
Distribution: Lee, Susan R./SRL; Sub-Collections: NIC; Clerk: NJN;

RADC 16-JAN-73 10:57 13824

this is a test from ray

how now brown cow

RADC 16-JAN-73 10:57 13824

this is a test from ray

(J13824) 16-JAN-73 10:57; Title: Author(s): Rome Air Development Center (ISIM)/RADC; Distribution: Stone, Duane L., Cavano, Joel P., McNamara, John L./DLS JPC JLM; Sub-Collections: RADC; Clerk: RADC;

RADC 16-JAN-73 11:01

13825

todays AHI lesson

this is my present to you

RADC 16-JAN-73 11:01

13825

todays AHI lesson

ghave a good day. why everybody 1

everybodygoodmorning all 2

RADC 16-JAN-73 11:01

13825

todays AHI lesson

GHAVE 3A KOOD DAY. 2a

3

RADC 16-JAN-73 11:01

13825

todays AHI lesson

why everybody 3

hello everybodygoodmorning 4

RADC 16-JAN-73 11:01

13825

todays AHI lesson

h pp .0ow are you my buddysolong for today. 5

5

Lee, Susan R.
Stanford Research Institute
Augmentation Research Center
333 Ravenswood Avenue
Menlo Park, California 94025

To:
Lee, Susan R.
Stanford Research Institute
Augmentation Research Center
333 Ravenswood Avenue
Menlo Park, California 94025

13040

Author Copy

RADC 16-JAN-73 11:01

13825

todays AHI lesson

(J13825) 16-JAN-73 11:01; Title: Author(s): Rome Air Development Center (ISIM)/RADC; Distribution: Cavano, Joel P./JPC; Sub-Collections: RADC; Clerk: RADC;
Origin: <RADC>RAY.NLS;2, 15-JAN-73 11:16 RADC ;-1

DCE 16-JAN-73 16:27 13826

Phone Log: 16 Jan 73, with Allan Kessler, MIT ADMINS group

DCE 16-JAN-73 16:27 13826

Phone Log: 16 Jan 73, with Allan Kessler, MIT ADMINS group

I tried to reach Stuart McIntosh; wanted to learn about his progress in (possibly) connecting his PDP-15 data-management system onto the ARPANET. Tom Marill had offered him a TIP port at CCA, when we talked at ICCC. 1

Stuart is at a conference in Hawaii, and will be on vacation for a few weeks after that. Alan answered the phone; says he doesn't know about the Network stuff. Will have Stuart communicate when he gets back. 2

(I am sure that someday we will want to make use of ADMINS tools from our Workshop.) 3

DCE 16-JAN-73 16:27 13826
Phone Log: 16 Jan 73, with Allan Kessler, MIT ADMINS group

(J13826) 16-JAN-73 16:27; Title: Author(s): Engelbart, Douglas C./DCE ; Sub-Collections: SRI-ARC; Clerk: DCE ;

KEV 16-JAN-73 14:04 13827

archive status of nls files

It would be nice if nls propagated archive status when it did an update file 1

KEV 16-JAN-73 14:04 13827

archive status of nls files

(J13827) 16-JAN-73 14:04; Title: Author(s): Victor, Kenneth E. (Ken)/KEV; Distribution: Victor, Kenneth E. (Ken), White, James E. (Jim), Dornbush, Charles F., Michael, Elizabeth K., Vallee, Jacques F., Mitchell, James G., Deutsch, L. Peter, Kaye, Diane S., Andrews, Don I., Bass, Walt, Hopper, J. D., Irby, Charles H., Lehtman, Harvey G., Wallace, Donald C. (Smokey), Ferguson, Ferg R., Norton, James C., Watson, Richard W./npg dcw wrf jcn rww ; Sub-Collections: SRI-ARC NPG; Clerk: KEV;

JCN 19 JAN 73 7:03PM 13828

Transmittal Letter to Craig Decker

Stanford Research Institute
Augmentation Research Center
333 Ravenswood Avenue
Menlo Park, California 94025

Craig Decker
303 N Charles Street
Apartment 2-N
Baltimore, MD 21218

Dear Craig:

Here are two key documents describing the activities of the Augmentation Research Center. 1

You'll note that the 1962 document contains the plans for a long term research project directed toward augmented human intellect. We are still working within that framework. 2

The June 1971 report should give some idea of fairly recent developments. Our May 1972 report has just recently been approved by RADC and is in the process of final printing. I'll send you a copy of that report when it is available. 3

I was pleased to meet you and hope this information will be useful. 4

Sincerely,

James C. Norton
Assistant Director
Augmentation Research Center

ber

JCN 19 JAN 73 7:03PM 13828

Transmittal Letter to Craig Decker

(J13828) 17-JAN-73 14:58; Title: Author(s): Norton, James C./JCN;
Sub-Collections: SRI-ARC; Clerk: BER;
Origin: <ROW>NORTON13828.NLS;1, 17-JAN-73 13:22 BER ;

AAM 17-JAN-73 5:53 13829

SENDMESSAGE as seen from a TIP

This is addressed to DCW, DCW3, and NP

AAM 17-JAN-73 5:53 13829

SENDMESSAGE as seen from a TIP

Smokey et al,

To review my comment of December to you; the SENDMESSAGE feature would be much nicer for a TIP user if , after the "message was sent, the program sent two extra messages (which could contain either the TELNET NOP or the ASCII NUL) before closing the connection. This would give the simple-minded TIP a chance to finish typing out the message before typing "R CLOSED", which currently appears in the middle of the typescript. 1

AAM 17-JAN-73 5:53 13829

SENDMESSAGE as seen from a TIP

(J13829) 17-JAN-73 5:53; Title: Author(s): McKenzie, Alex A./AAM; Distribution: Wallace, Donald C. (Smokey), Walden, David C., Irby, Charles H., Kelley, Kirk E./DCW DCW3 NP; Sub-Collections: NIC NP; Clerk: AAM;

CFD 17-JAN-73 10:29 13830

Some Measurements of NLS Startup Time

The exceedingly long startup time for NLS is due principally to the file operations done during startup. It does not appear that significant gains can be made by simply moving the constant initialization code to load time initialization. Verifying the user's ident is the single most expensive operation-- maybe we should consider some alternate method for accomplishing this.

Some Measurements of NLS Startup Time

STARTUP MEASUREMENTS FOR NLS

A set of 10 timing probes was patched into the running NLS system in the startup code in order to get a handle on where NLS is gobbling up so much time during initialization. The following table represents the average elapsed time in milliseconds between time probes. It clearly shows that validating a user's ident is the single most expensive operation.

A random variation of about 10% was noted in the timings, evidently due to fluctuations in machine load and paging considerations.

These figures are approximately half of the timings reported earlier by Charles Irby, and as yet no reason for the discrepancy is known. The load average was about 3.5-4.0 when these timings were taken.

TABLE - SELECTED TIMING PROBES

Address	elapsed	delta
startup	0	144
initch	144	107
intmeas	251	37
getdev	288	2
initch	290	105
filinit	395	27
getinit	422	1242
makemf	1664	306
initdis	1970	201
openid	2171	--

CFD 17-JAN-73 10:29 13830

Some Measurements of NLS Startup Time

(J13830) 17-JAN-73 10:29; Title: Author(s): Dornbush, Charles F./CFD ; Distribution: Irby, Charles H., Watson, Richard W., Hopper, J. D., Kaye, Diane S., Vallee, Jacques F., Michael, Elizabeth K., Lehtman, Harvey G., Norton, James C./CHI RWW JDH DSK JFV EKM HGL JCN ; Sub-Collections: SRI-ARC; Clerk: CFD ;

CFD 17-JAN-73 12:04 13831

Some Measurements of NLS Startup Time

The exceedingly long startup time for NLS is due principally to the file operations done during startup. It does not appear that significant gains can be made by simply moving the constant initialization code to load time initialization. Verifying the user's ident is the single most expensive operation-- maybe we should consider some alternate method for accomplishing this.

STARTUP MEASUREMENTS FOR NLS 1

A set of 10 timing probes was patched into the running NLS system in the startup code in order to get a handle on where NLS is gobbling up so much time during initialization. The following table represents the average elapsed time in milliseconds between time probes. It clearly shows that validating a user's ident is the single most expensive operation. 1a

A random variation of about 10% was noted in the timings, evidently due to fluctuations in machine load and paging considerations. 1b

These figures are approximately half of the timings reported earlier by Charles Irby, and as yet no reason for the discrepancy is known. The load average was about 3.5-4.0 when these timings were taken. 1c

1d

TABLE - SELECTED TIMING PROBES 1e

1f

Address	elapsed	delta	1g
-------	-----	-----	1h
startup	0	144	1i
initch	144	107	1j
intmeas	251	37	1k
getdev	288	2	1l
initch	290	105	1m
filinit	395	27	1n
getinit	422	1242	1o
makemf	1664	306	1p
initdis	1970	201	1q
openid	2171	--	1r

CFD 17-JAN-73 12:04 13831

Some Measurements of NLS Startup Time

(J13831) 17-JAN-73 12:04; Title: Author(s): Dornbush, Charles F./CFD ; Distribution: Irby, Charles H., Watson, Richard W., Hopper, J. D., Kaye, Diane S., Vallee, Jacques F., Michael, Elizabeth K., Lehtman, Harvey G., Norton, James C./CHI RWW JDH DSK JFV EKM HGL JCN ; Sub-Collections: SRI-ARC; Clerk: CFD ;