CHI WLB 24-JAN-72 15:45 8642

The meeting called for in (Journal, 8614, 1) has been changed to an all-ARC meeting to discuss and launch mechanisms for dealing with questions of organizational and personal development. Everyone is encouraged to come prepared for active participation in this important meeting. 1

CHI WLB 24-JAN-72 15:45 8642

(J8642) 24-JAN-72 15:45; Author(s): Charles H. Irby, Walter L. Bass/CHI WLB; Distribution: Paul Rech, Stephen W. Miller, Michael D. Kudlick, George J Eilers, Donald R. Cone, Bonnar Cox, David R. Brown, Don Limuti, William R. Ferguson, Priscilla Lister, Robert L. Dendy, Linda L. Lane, Marilyn F. Auerbach, Walter L. Bass, Mary S. Church, William S. Duvall, Douglas C. Engelbart, Beauregard A. Hardeman, Martin E. Hardy, J. D. Hopper, Charles H. Irby, Mil Jernigan, Harvey G. Lehtman, John T. Melvin, Jeanne B. North, James C. Norton, Cindy Page, Bruce L. Parsley, William H. Paxton, Jeffrey C. Peters, Jake Ratliff, Barbara E. Row, Ed K. Van De Riet, Dirk H. van Nouhuys, Kenneth E. Victor, Don C. Wallace, Richard W. Watson, Don I. Andrews/SRI-ARC; Sub-Collections: SRI-ARC; Clerk: WLB;

BLP 25-JAN-72 2:01 8643

Cleaning NLS -- reply to WSD's 8637

This is a reply to (8637,) by WSD, which was in turn a reply to (8573,) by BLP. The subject is suggestions and conventions for cleaning NLS. 1

All of 8637 is contained herein. I have added comments and reissued it. To mark my comments, they all start with a UID ':. The UID's and there meanings: 2

AGREE: The suggestion will be added to the "working list". Of course other people are free to object. 2a

DISAGREE: The suggestion will not be added to the list unless I am outvoted or some new argument is persuasive. 2b

MOOT: The suggestion will be added as something to be resolved later. I don't have any strong feelings one way or the other. 2c

DON'T UNDERSTAND: I didn't understand the suggestion. 2d

MY MISTAKE: I said something wrong in 8573. It will be corrected. 2e

BLP 25-JAN-72 2:01 8643

Cleaning NLS -- reply to WSD's 8637

Errors 3

I agree philosophically with the idea of everything signalling, but I think that the restriction of having only one SIGNAL statement active at a time in a procedure, and not having anyway of explicitly activating others is sufficiently severe so that the exclusive use of SIGNAL for errors is not feasible. 3a

DISAGREE: The one active signal can check for which error code. 3a1

How about a global error string in which any routine signalling may place an error message. This provides an alternative to the passing of error strings. 3b

COMMENT: That's what sysmsg is. 3b1

Warning/progress messages 4

These should not be issued directly from core NLS routines, or at least should be issued in a way such that they are suppressed if the job is running as a library process or detatched. 4a

AGREE (mostly): Core NLS routines can issue them. If no one is armed for it, then it is merely ignored (library processes are thus cool). 4a1

Local Strings 5

The current need for global strings stems from the inability to return a local string as the result of a procedure. How about the dynamic string system we have talked about as a solution?? 5a

DISAGREE: The calling routine can pass the address of a string local to it and the called routine fills it. 5a1

Standard Format 6

I think that the description of what a procedure does should be the first sub-branch, followed (or including) a description of te arguments. 6a

MOOT: 6a1

Cleaning NLS -- reply to WSD's 8637

I think that the format for the description of the parameters is important...i.e. they should not be imbedded in text. 6b

DON'T UNDERSTAND: 6b1

I think that a like branch should describe possible returns and results. 6c

AGREE: 6c1

Re- standard local and formal parameter names....anything REFed in one place should always be REFed. 6d

I use and find very convenient z1, z2, z3, z4...for names of local text pointers. 6d1

MOOT: How about tp1... or t1.... 6d1a

While this is not as descriptive as mnemonic names, the FIND construct is sufficiently compact and complex that 6d1b

1: Comments should be liberally sprinkled throughout the statement saying what is being searched for, where pointers are pointing, etc. 6d1b1

AGREE: 6d1b1a

2: Under no circumstances should SWORK and SCNDIR be set up in one procedure, and assumed to be right in another. If another procedure wishes to use the values of swork and scndir, they should be passed as parameters. 6d1b2

AGREE: 6d1b2a

LOCAL declarations should always be in the order of LOCALS, TEXT POINTERS, STRINGS. 6d2

This makes perusing the stack much simpler. 6d2a

AGREE: 6d2b

Pretty-up format 7

Each L10 statement should correspond to one NLS statement except in the case of extended IF a AND b OR c....type of statements. 7a

Cleaning NLS -- reply to WSD's 8637

AGREE: 7a1

Specifically, the procedure actual parameter list should appear in the same statement as the procedure name in a procedure call. This makes it much easier for TNLS users. 7a2

AGREE: 7a2a

There should NOT be a space between the procedure name and the actual parameter list in a call. 7b

DISAGREE: 7b1

Operators and expressions separated by spaces is pushing things, I think. 7c

AGREE: 7c1

GOTO State 8

Should be done away with entirely, and replaced by the SIGNAL mechanism. 8a

AGREE: 8a1

See (nls,idtctl,) for an example of how this might work. 8b

Core-routines...CLIST 9

MY MISTAKE: X routines are the only things that should fiddle CLIST. 9a

I do not think that core routines should have anything to do with CLISTS. 9b

Except for markers. 9c

And if they must (because we cannot figure a way around it), it should be entirely and easily optional. 9d

The difference between RECDEF and CONST is not made entirely clear. 10

Couldn't RECDEF simply be a branch in CONST? 10a

MOOT: 10a1

JSYS Routines 11

An alternate approach to te JSYS stuff would be to have one routine which accepted as parameters register values, and a JSYS identification, and executed that JSYS in a particular way. 11a

DISAGREE: The JSYS could probably all be executed that way, but I don't think a general return mechanism could be worked out for all the JSYS's. 11a1

In any event, the JSYS routine/identification name MUST ALWAYS be the same as the name of the JSYS, i.e. no expanding such as changing lgtjfn to lgetjfn. 11b

AGREE: 11b1

If we use indivicual procedures, let us standardise the types of calls (i.e. Is a Destination Designator a String address, a JFN, a file number, a String Pointer ????? 11c

AGREE: At least to the extent possible. 11c1

If individual procedures are used, a directory similar to the JSYS directory should be published, which has the calling arguments and results. 11d

MOOT: They're subroutines like any other. 11d1

The JSYS routines should not SIGNAL. 11d2

DISAGREE: These subroutines ought to do error handeling the same way other subroutines do. 11d2a

Universal File Handle 12

AGREE!!! A very nice idea. 12a

A suggestion: 12a1

1 <= NLS File Number <= 100B 12a1a

JFN < 0 (take the negative of it) 12a1b

100B < string address <= 777777B 12a1c

We Need some type of NLS Universal file handle, which may be used anywhere a file is to be identified. 12b

The Universal File Handle should be able to be either: 12b1

(1) An NLS File Number 12b1a

(2) A JFN 12b1b

(3) The address of a string containing a file name. 12b1c

Appropriate library routines should be available for turning each of these things into each of the other things, wihout knowing what the thing you have is. 12b2

For example, If a routine is passed a UFH blap, and needs the JFN of the file for whatever he wishes to do, it calls the routine FINDJFN(UFH) which either SIGNALS, or returns the JFN. 12b2a

BLP 25-JAN-72 2:01 8643

Cleaning NLS -- reply to WSD's 8637

(J8643) 25-JAN-72 2:01; Title: Author(s): Bruce L. Parsley/BLP; Distribution: Walter L. Bass, William S. Duvall, Mary S. Church, J. D. Hopper, Charles H. Irby, Harvey G. Lehtman, John T. Melvin, Bruce L. Parsley, William H. Paxton/NPG; Sub-Collections: SRI-ARC NPG; Clerk: BLP;
Origin: <PARSLEY>REPWSD.NLS;2, 25-JAN-72 1:58 BLP ;

AAM 25-JAN-72 6:56 8644
TENEX EXECUTIVE-LEVEL MESSAGE FOR BBN-NET

JEANNE,
WHEN WE LOG IN, WE GET A MESSAGE WHICH SAYS "YOU HAVE A MESSAGE"
HOW CAN WE RECEIVE THE MESSAGE WHICH WE "HAVE"? (HINT: I
THINK WE MUST BE SRI PERSONNEL TO GET MESSAGES - IS THAT TRUE?) 1

AAM 25-JAN-72 6:56 8644

TENEX EXECUTIVE-LEVEL MESSAGE FOR BBN-NET

(J8644) 25-JAN-72 6:56; Title: Author(s): Alex A. McKenzie/AAM;
Distribution: Jeanne B. North/JBN; Sub-Collections: NIC; Clerk: AAM;

AAM 25-JAN-72 7:02 8645

RE: NFG MEETING

PEGGY,
DUE TO THE EXTREME RELUCTANCE MY CHILD IS EXHIBITING FOR LEAVING THE COMFORT AND SECURITY OF ITS CURRENT ABODE FOR THE COLD, CRUEL WORLD I MAY NOT BE ABLE TO GET TO THE NFG MEETING THIS FRIDAY. I'LL LET YOU KNOW MORE DEFINITELY BY TELEPHONE THIS THURSDAY.
REGARDS,
ALEX MCKENZIE

1

AAM 25-JAN-72 7:02 8645

RE: NFG MEETING

(J8645) 25-JAN-72 7:02; Title: Author(s): Alex A. McKenzie/AAM; Distribution: Peggy M. Karp/PMK; Sub-Collections: NIC; Clerk: AAM;

AAM 25-JAN-72 7:05 8646

LET ME KNOW IF YOU GET THIS MESSAGE 1

AAM 25-JAN-72 7:05 8646

(J8646) 25-JAN-72 7:05; Title: .HED="TEST MESSAGE TO "BBN-NET"";
Author(s): Alex A. McKenzie/AAM; Distribution: Julie B. Moore, Ellen
Westheimer, David C. Walden, Bernie P. Cosell, Joel B. Levin, Alex A.
McKenzie, Frank E. Heart, Will R. Crowther/BBN-NET; Sub-Collections: NIC
BBN-NET; Clerk: AAM;

SUPERWATCH OPERATION 1

Abstract 1a

Superwatch is an information gathering and formatting-printing program designed to help find out what is going on within our TENEX. 1a1

It is designed to put a very small load on the system while collecting information from it, so that it will not alter the operation of the system significantly. Measurements show that the load incurred can be kept as low as 0.1% (of real CPU time). 1a2

Superwatch also contains miscellaneous commands for changing system parameters and turning on and off various measurement code within the system. 1a3

Monitor Functions 1b

The TENEX monitor, as we get it from BB&N contains some information collection, in the form of meters, concerning what is happening within the system. 1b1

By meters I mean counters that continually increase in value. The difference between two readings, and the time interval between the readings, can be used to compute an average rate over the interval, for example. 1b1a

Nearly all of our information collecting has been in this form. 1b1b

We have added a considerable amount of metering at ARC, notably: 1b2

Some sampling is done at a 50ms rate. The results of this sampling are again meters, which register the sum of the samples. The meter differences and number of samples give us the average sampled value. 1b2a

In this way we sample: 1b2a1

Drum and disk queue length and per cent busy 1b2a1a

Number of jobs in the balance set, number of pages in balance set, number of free pages, number of reserved pages 1b2a1b

(under an of-off flag) PC sampling (the unmapped PC at the time of the 60Hz clock interrupt) 1b2a1c

At a slower rate, controlled by a process clock, we can sample the use of user memory pages. 1b2b

This sampling is done at a 500ms rate and is under the control of an on-off flag. 1b2b1

The user pages are categorized as private-modified, private-unmodified, shared-0, shared-1, shared-2. The types of shared pages refer to the number of processes referencing those pages (i.e. process bits in the process use field for that page). Shared-2 represents pages referenced by two or more processes. 1b2b2

From the metering of the number of pages of each type, we get a profile of the kinds of pages taking up user memory. 1b2b3

Meter type information is also collected from NLS via a JSYS, concering 1b2c

The number of tasks completed. 1b2c1

The elapsed execution time for a task. 1b2c2

The elapsed real time for the execution of a task. 1b2c3

The elapsed real time from one task to the next. 1b2c4

At this time, a task is defined an interaction on the character level. 1b2c5

We have also added several other meters to record such things as 1b2d

actual scheduler overhead, process clock overhead, IO wait 1b2d1

activity of garbage collector 1b2d2

page faults by type of fault and page 1b2d3

The monitor contains JSYS's to return the meter readings to a user program -- in the registers so that differences can be quickly taken. 1b3

The JSYS's have been written with speed in mind, so that the user program collecting information disrupts the system load as little as possible. 1b3a

More meters can be easily added. 1b3b

A user program - superwatch - collects the information from the monitor. 1b4

The program runs a a specified interval. At each interval it collects the meters from the system, takes the differences, and spills them onto a file (stat file). 1b4a

There are two modes of information collection: detail and no detail. One puts a significantly heavier load on the system and must be run at longer intervals, but collects more information. 1b4b

A separate command starts the program running on a 10 minute interval with detail turned on. This cound run as an autojob to provide a profile of the day. 1b4c

Another superwatch command will read the stat file, format and print it. 1b5

The compact file of (mostly) meter differences, can be read and converted to textual form after the collection is done. 1b5a

The superwatch program has a table of computed values, or measurements, that can be produced from the stat file. Some are simple meter readings, others are combinations of many readings. Other measurements can be easily added, if they are computed from the information in the stat file. Other measurements may require new meters within the system. 1b5b

It is expected that the user will wish to see a subset of all the possible measurements. Several subsets are provided as built in "print options." Also, a print option may be specified, item by item, with the option Set command. It is then referred to as print option zero, (normally an empty option). The measurements in the print options are specfied below. Of course, a stat file may be printed several times with different print options. 1b5c

PC sampling results are printed with a special command which requests a "GET" file with a DDT symbol table in it, as well as a stat file. 1b5d

Averages over the entire test may be printed, and/or values for each collection interval may be printed, either for the entire test or for a specified interval of the test, given two times of day. 1b5e

If a printed value overflows its field, asterisks are printed in the entire field. 1b5f

The collecting command has a special "print while running" mode that does the processing and formatting to any output file (e.g. TTY:) at collection time (as well as producing the stat file). 1b5g

A display command, which is like the "print while running" mode with the display tube as the output device. This displays a set of horizontal histograms in real time. 1b5h

COMMAND STRUCTURE 2

(Collect statistics) 2a

Specify: 2a1

The output file 2a1a

The program sampling interval in seconds. 2a1b

This determines the rate at which the program collects information from the system. 2a1b1

The total run time for the test in minutes. 2a1c

The program can be stopped at any time and the file can still be printed, however. 2a1c1

Do a ↑C and REENTER, then print the file. 2a1c2

Detail (reply Y or N). 2a1d

The primary difference is that Detail includes information by job and subsystem, memory status information, and program counter sampling results. 2a1d1

Print while running (reply Y or N) 2a1e

Allows one to see the results while the program is running. 2a1e1

The interval should be longer than the time needed to print the results. A shorter interval is essentially not acknowledged - it run at the minimum interval. 2a1e2

An output file and print option will be requested. 2a1e3

Be advised that the superwatch program will be a measureable load on the system if run in this mode -- depending on the interval of course. 2a1e4

Action: 2a2

The program samples information from the system at the specified interval, and writes a compact form on the output file. Averages are maintained as well. 2a2a

The file is kept in good form so the program may be terminated by ↑C without messing up the file. 2a2b

Averages are maintained on a 10 minute interval so that ↑C may make the averages somewhat incorrect. 2a2b1

(Display statistics) 2b

Specify: 2b1

Same as for Collect statistics except for a few specifications which are preset. 2b1a

The print option determines the contents of the display. 2b1b

If the print option has too many items to fit on the display, the last of them are ignored. 2b1c

Action: 2b2

The program produces a file as in Collect statistics, and also creates a display of the information specified in the print option. 2b2a

The display is updated at each interval, in the form of histograms and numbers. 2b2b

(Print) 2c

(Averages) 2c1

Action: 2c1a

Prints averages for the entire test. 2c1a1

(Statistics) 2c2

Action: 2c2a

Prints averages and statistics from each individual sample made when the collection was being done. 2c2a1

(Time interval) 2c3

Specify: 2c3a

the time interval as two daytimes (hrs and min). 2c3a1

Action: 2c3b

Prints individual samples that were made in the specified time period. 2c3b1

For all of the above: 2c4

Specify: 2c4a

The statistics file generated by a Collect or Dispplay command. 2c4a1

The output file. 2c4a2

The print option number. 2c4a3

(Interpret PC sample) 2d

Specify: 2d1

The statistics input file. 2d1a

Whether you want monitor symbols, or 2d1b

The GET file containing the relevant DDT symbol table. 2d1c

The output file. 2d1d

Action: 2d2

The PC sample results are written on the output file in both symbolic and numeric form. 2d2a

Global and local symbols are taken from the DDT symbol table in the get file. 2d2b

Number and per centage of executions in each "bucket" are printed. (See Turn ON PC sample). 2d2c

If monitor symbols are requested, the symbols for the monitor being used at the time of the collection are used. 2d2d

(Scheduler parameter set) 2e

Specify: 2e1

Your initials. 2e1a

(restricted) 2e1a1

Scheduling mode (Normal or Compile time). 2e1b

Action: 2e2

The scheduler parameters are set as specified. If the system rejected the parameters (which should not happen) a message to that effect is written on the terminal. 2e2a

The system logs the fact that the parameters were changed and specifies the user and initials. 2e2b

(Turn) 2f

(ON) 2f1

(Pc sample) 2f1a

Specify: 2f1a1

User bit 2f1a1a

=0 sample system PC 2f1a1a1

=1 sample user PC 2f1a1a2

anything else turns off the PC sampler 2f1a1a3

subsystem name (only if you set user bit = 1) 2f1a1b

a null subsystem name means: sample all user PC's 2f1a1b1

lower bound address (in octal) 2f1a1c

number of words per bucket (in octal) 2f1a1d

Action: 2f1a2

The system's program counter sampling code is invoked (on a 50 ms interval). 2f1a2a

If the PC sampling code is already on, a message to that effect is printed. 2f1a2b

Either system mode or user mode PC is sampled. IF user mode, either a specified subsystem or all user programs are sampled. 2f1a2c

The information is collected as counts of samples within specified ranges, (or "buckets") with 13 ranges and one count for out-of-range in each direction for a total of 15 counters. 2f1a2d

The ranges are specified as a lower bound and word count per range (or bucket). The word count is rounded down to a power to two, but the lower bound is not. For example, if the word count were 4 and the lower bound 123, the buckets would be: 2f1a2e

0-122 2f1a2e1

123-126 2f1a2e2

127-132 2f1a2e3

133-136 etc. 2f1a2e4

(Memory sampler) 2f1b

Specify: ca 2f1b1

Action: 2f1b2

Turns on sampling of core status tables. 2f1b2a

This must be on to collect memory usage information. 2f1b2b

(OFF) 2f2

(Pc sampler) 2f2a

Specify: ca 2f2a1

(Memory sampler) 2f2b

Specify: ca 2f2b1

(Response cutoff set) 2g

Specify: 2g1

Your id. 2g1a

(restricted) 2g1a1

The upper cutoff in ms. 2g1b

The lower cutoff in ms. 2g1c

Action: 2g2

The systems response cutoff parameters are set as specified. 2g2a

When the response index goes above the upper cutoff, a flag is set for the EXEC to prevent new logins, until the response index goes below the lower cutof, at which time the flag is reset. 2g2b

(Option set) 2h

Specify: 2h1

A string of option names (as they appear in printouts) delimited by CR, ending with a -2. 2h1a

A -1 item in the option string inserts a CR in the output at the corresponding point. 2h1b

All names contain all upper case leters. 2h1c

Action: 2h2

This option string may be referenced as print option zero. 2h2a

(High queue on dismiss) 2i

Specify: ca (requires Wheel status) 2i1

Action: 2i2

Does not appear to be necessary to do this to get satisfactory service. 2i2a

(Quit) returns to the exec 2j

LIST OF MEASUREMENTS THAT CAN BE PRINTED 3

Each item will be described in the following format: 3a

(name) 3a1

meaning: 3a1a

units: 3a1b

normal values: 3a1c

indications: 3a1d

System times, etc. 3b

(t ms) 3b1

meaning: time in collection interval actually experienced by collection program 3b1a

units: ms. 3b1b

normal values: several ms off of the specified interval, but should average pretty close to it. 3b1c

indications: way off indicates heavy load and bad response. 3b1d

(IDL) 3b2

meaning: Idle time, no jobs to run at all 3b2a

units: ten times per cent real CPU time 3b2b

normal values: zero during day 3b2c

(ST1) 3b3

meaning: BBN's I/O wait time, includes scheduler, process clocks and I/O wait 3b3a

units: ten times per cent real CPU time 3b3b

normal values: 200 to 300 (20% to 30%) 3b3c

indications: too large means excessive time in scheduling, clocks, or I/O wait 3b3d

(ST2) 3b4

meaning: BBN's page mang. time plus garbage collection time 3b4a

units: ten times per cent real CPU time 3b4b

normal values: 100 to 150 (10% to 15%) 3b4c

indications: too much time garbage collecting or too many page faults 3b4d

(%SK) 3b5

meaning: time spent scheduling -- with jobs to run 3b5a

units: ten times per cent real CPU time 3b5b

normal values: 150 to 250 (15% to 25%) 3b5c

indications: excessive scheduling 3b5d

(%CK) 3b6

meaning: time spent in process clocks 3b6a

temporarly set to the IMPCHK clock only 3b6a1

units: ten times per cent real CPU time 3b6b

normal values: 20 to 80 (2% to 8%) 3b6c

indications: excessive process clock work being done 3b6d

(%DW) 3b7

meaning: that portion of I/O wait time spent with drum not busy 3b7a

units: ten times per cent real CPU time 3b7b

normal values: zero 3b7c

indications: the system is waiting on the disk 3b7d

(IOW) 3b8

meaning: time spent looping in the scheduler with runnable jobs all waiting on the drum or disk 3b8a

units: ten times per cent real PCU time 3b8b

normal values: 80 to 120 (8% to 12%) 3b8c

indications: large indicates few jobs runnable or system parameters not correct for amount of core. 3b8d

(%U) 3b9

meaning: time spent running user programs 3b9a

units: ten times per cent real CPU time 3b9b

normal values: 500 to 700 (50% to 70%) 3b9c

indications: small indicates low system load to excessive time in scheduler, etc. 3b9d

(PFO) 3b10

meaning: page fault overhead 3b10a

units: 100 times per cent real CPU time per page fault or 100 times ms. per page fault 3b10b

normal values: 200 to 400 (2 to 4 ms per page fault) 3b10c

indications: large indicates excessive garbage collection of working sets 3b10d

(SKO) 3b11

meaning: scheduling overhead 3b11a

units: 100 times per cent real CPU time per scheduling or 100 times ms. per scheduling 3b11b

normal values: 100 (1 ms per scheduling) 3b11c

indications: efficiency of scheduler 3b11d

(SKD) 3b12

meaning: number of times through scheduler (SKDJOB) 3b12a

units: times per second 3b12b

normal values: 180 3b12c

indications: excessive rescheduling 3b12d

(HR) 3b13

meaning: hours reading from HP clock at time of sample 3b13a

(MIN) 3b14

meaning: minutes reading from HP clock 3b14a

(SC) 3b15

meaning: seconds reading from HP clock 3b15a

Drum and Disk activity, and related things 3c

Note: When both drums are being used, statistics are collected as though they were one device, and some measurements are somewhat strange. Both drums 100% busy results in %DB=200.0% busy. The drum queue length is actually about half of DQL for each drum, assuming equal page allocation. drum page time and total time are correct. 3c1

(%DB) 3c2

meaning: drum busy 3c2a

units: ten times per cent of real time 3c2b

normal values: 600 to 700 (60% to 70%) 3c2c

indications: drum usage 3c2d

(DQL) 3c3

meaning: average drum queue length when drum is busy 3c3a

units: ten times number of entries in queue including page in transit 3c3b

normal values: 30 to 40 (3.0 to 4.0) 3c3c

indications: 3c3d

(DPT) 3c4

meaning: drum page time - average time to transfer one page 3c4a

units: ms. times ten 3c4b

normal values: 60 to 100 (6.0 to 10.0) 3c4c

indications: large times indicate malfunction of drum 3c4d

Note: values taken over short interval (5 sec or less) are not very accurate. 3c4e

(DT) 3c5

meaning: drum time - total time a user program waits for a page on the average 3c5a

units: ms. times ten 3c5b

normal values: 300 (30.0 ms.) or better 3c5c

indications: drum response 3c5d

(DRR) 3c6

meaning: drum reads 3c6a

units: pages per second 3c6b

normal values: 50 to 60 pages per second 3c6c

indications: number of page faults and swapins 3c6d

(DRW) 3c7

meaning: drum writes 3c7a

units: pages per second 3c7b

normal values: 25 to 35 pages per second 3c7c

indications: pages returning to drum via garbage collection 3c7d

(%KB) 3c8

meaning: disk busy 3c8a

units: ten times per cent real time 3c8b

normal values: 500 to 700 (50% to 70%) 3c8c

indications: disk usage 3c8d

(KQL) 3c9

meaning: disk queue length 3c9a

units: ten times number of entries in queue including page in transit 3c9b

normal values: 25 to 50 (2.5 to 5.0) 3c9c

indications: 3c9d

(KPT) 3c10

meaning: disk page time - average time to transfer one page 3c10a

units: ms. per page 3c10b

normal values: 140 to 190 ms. 3c10c

indications: large indicates improper functioning of disk 3c10d

(KT) 3c11

meaning: disk time - total time a user program waits for a disk page on average 3c11a

units: ms. per page 3c11b

normal values: large values indicate disk swampped - see KPT 3c11c

indications: 3c11d

(DKR) 3c12

meaning: disk reads 3c12a

units: pages per second times ten 3c12b

normal values: 15 to 20 (1.5 to 2.0) 3c12c

indications: 3c12d

(DKW) 3c13

meaning: drum writes 3c13a

units: pages per second times ten 3c13b

normal values: slightly less than DKR 3c13c

indications: 3c13d

(XG) 3c14

meaning: number of calls to XGC working set cruncher 3c14a

units: calls per second times ten 3c14b

normal values: ? 3c14c

indications: too large indicates systems working set size parameter is too small - see PFO 3c14d

(DD) 3c15

meaning: number of pages written by DDMP (job zero - return to disk from drum guy) 3c15a

units: pages per second times ten 3c15b

indications: smaller than this indicates system not fully loaded 3d1d

(GJ) 3d2

meaning: number of runnable (go) jobs in system 3d2a

units: number of jobs times ten 3d2b

normal values: GJ >= BJ when fully loaded 3d2c

indications: if GJ is significantly greater than BJ, system is overloaded 3d2d

(#BR) 3d3

meaning: number of runnable jobs in the balance set 3d3a

units: number of jobs times ten 3d3b

normal values: should be better than half of BJ 3d3c

indications: too small indicates that jobs are hung up on drum or disk, see IOW, %DW 3d3d

(#BE) 3d4

meaning: number of balance set entries 3d4a

units: new jobs per second times ten 3d4b

normal values: 100 to 140 (10.0 to 14.0) 3d4c

indications: shows type of computing being done (interactiveness) and load on system 3d4d

(#BP) 3d5

meaning: number of balance set pages (sum of BJ's working sets) 3d5a

units: pages 3d5b

indications: 3d5c

(#UP) 3d6

meaning: number of user pages avaiable in the system 3d6a

units: pages 3d6b

normal values: about 145 to 160 depending on debug switch 3d6c

indications: 3d6d

(%B0) to (%B4) 3d7

meaning: per cent of jobs blocking for TTY input on queue (0 to 4) 3d7a

units: percent times ten 3d7b

normal values: %B0 should be greater than 900 (90.0%) - indicates proper queue lengths 3d7c

indications: 3d7d

(TG0) to (TG4) 3d8

meaning: time on go list for jobs on queue (0 to 4) 3d8a

units: ms. 3d8b

normal values: 10 to 20 3d8c

indications: system response: response index is an exponential average of TG0. 3d8d

(#0) to (#4) 3d9

meaning: number of jobs removed from go list (or entering balance set) on queue (0 to 4) 3d9a

units: number of jobs times ten 3d9b

normal values: 3d9c

indications: 3d9d

(#0) to (#4) 3d10

meaning: number of jobs removed from go list (or entering balance set) on queue (0 to 4) 3d10a

units: number of jobs times ten 3d10b

normal values: 3d10c

indications: 3d10d

Memory information 3e

(#RP) 3e1

meaning: number of reserved pages 3e1a

units: pages 3e1b

normal values: 100 to 120 3e1c

indications: 3e1d

(#RQ) 3e2

meaning: number of pages on replacable queue (free pages) 3e2a

units: pages 3e2b

normal values: 15 to 25 3e2c

indications: 3e2d

normal values: 80 to 120 depending on user core available (#UP) 3e3

(%PU) (%PM) (%S0) (%S1) %S2) 3e4

meaning: percentage of user pages in memory containing 3e4a

%PM: private modified 3e4a1

%PU: private unmodified 3e4a2

%S0: shared but unreferenced 3e4a3

%S1; shared, but referenced by only one process 3e4a4

%S2: shared, referenced by more than one process 3e4a5

units: per cent times ten 3e4b

normal values: 3e4c

indications: 3e4d

(#PM) (#PU) (#SO) (#S1) (#S2) 3e5

meaning: average number of user pages of the specified type 3e5a

units: pages 3e5b

normal values: 3e5c

indications: 3e5d

Information by job 3f

(JOBn) One value for each job: 3f1

meaning: The job number 3f1a

(TTYn) 3f2

meaning: TTY number (in decimal) or DET for detached job 3f2a

(SUBS) 3f3

meaning: name of subsystem that job is running as recorded by SETNM JSYS 3f3a

(USER) 3f4

meaning: the user name for that job 3f4a

(%JU) 3f5

meaning: per cent of real CPU time used by that job 3f5a

units: per centage times ten 3f5b

normal values: 0 to 900 (90%) 3f5c

indications: how the scheduler is dividing up the CPU among runnable jobs 3f5d

Information by subsystem 3g

(SUBn) One value for each subsystem in systems list of subsystems 3g1

meaning: the name of the subsystem (first 6 characters) 3g1a

(%SU) 3g2

meaning: per cent of real CPU time spent running that subsystem 3g2a

units: per centage times ten 3g2b

normal values: 0 to 900 (90%) 3g2c

indications: how much each of the subsystems are used. the name is set by SETNM JSYS and may be messed up by user program. 3g2d

(Tav) 3g3

meaning: average time between page faults (requiring page swap) 3g3a

units: ms. times ten 3g3b

normal values: 0 to very large numbers depending of compute-boundedness and system load. NLS is generally about 10 to 20 ms. 3g3c

indications: small numbers indicate that working set is not large enough or program poorly organized 3g3d

(%SPF) 3g4

meaning: shared page faults 3g4a

units: per cent of all swap-causing faults that are on shared pages 3g4b

normal values: ? 3g4c

indications: 3g4d

NLS statistics 3h

All measurements end with D, T or N for DNLS, TNLS and NTNLS. A task is defined as a character interaction - it includes feedback as well as command execution. 3h1

(AXD) (AXT) (AXN) 3h2

meaning: average execution time per task 3h2a

units: ms. 3h2b

normal values: 3h2c

indications: 3h2d

(SVD) (SVN) (SVT) 3h3

meaning: service given NLS by the system when executing - 50% would indicate that 100ms of execution time was obtained in 200ms. 3h3a

units: per cent real execution time 3h3b

normal values: 3h3c

indications: 3h3d

(ACD) (ACT) (AVN) 3h4

meaning: average connect time - total real time between completed character interactions 3h4a

units: ms. 3h4b

normal values: 3h4c

indications: 3h4d

(%UD) (%UT) (%UN) 3h5

meaning:per cent of real CPU time used per NLS user, average 3h5a

units: per cent real CPU time times ten 3h5b

normal values: 3h5c

indications: 3h5d

(TME) (TMT) (TMN) 3h6

meaning: number of task completed by NLS 3h6a

units: tasks per minute 3h6b

(%B1) (SKO) (PFO) (TMD) (DRR) (#BE) (SKD) 4a4d

(4) print option four 4a5

(5) print option five 4a6

requires detail - memory sampler on 4a6a

(IDL) (%SK) (%CK) (IOW) (ST2) (%PU) (%PM) (%SO) (%S1) (%S2) 4a6b

(#PU) (#PM) (#SO) (#S1) (#S2) (#RP) (#RQ) (BJ) (GJ) 4a6c

(6) print option six 4a7

requires detail - good for printing "big brother" collections 4a7a

(HR) (MIN) (SC) (IDL) (%SK) (%CK) (IOW) (ST2) (%U) (SKO) (PFO) 4a7b

(BJ) (GJ) (#RP) (#RQ) (TGO) (#O) (%BO) (%DB) 4a7c

(DQL) (DPT) (DT) (%KB) (KQL) (KPT) (KT) (%UD) 4a7d

(%UT) (%UN) (SUBN) (%SU) (TAV) (%SPF) 4a7e

normal values: 3h6c

indications: 3h6d

(%D) (%T) (%N) 3h7

meaning: per cent of NLS tasks completed by (DNLS, TNLS, NTNLS) 3h7a

units: 3h7b

normal values: 3h7c

indications: 3h7d

PRINT OPTIONS 4

The following print options are predefined. The measurements included in them are indicated. 4a

(0) print option zero 4a1

empty. May be filled with the option set command. 4a1a

(1) print option one 4a2

requires detail 4a2a

(%U) (IDL) (%SK) (CK) (IOW) (ST2) (BJ) (GJ) (#BR) (TGO) 4a2b

(#BE) (JOBN) (TTYN) (SUBS) (USER) (%JU) (SUBN) (%SU) (TAV) (%SPF) 4a2c

(2) print option two 4a3

requires detail - same as old watch but with subsystem name included. 4a3a

(JOBN) (TTYN) (SUBS) (%JU) 4a3b

(3) print option three 4a4

detail is not required 4a4a

(IDL) (%SK) (%CK) (IOW) (ST2) (%U) (%DB) (DQL) (DPT) 4a4b

(%KB) (KQL) (KPT) (BJ) (GJ) (#BR) (#RP) (TGO) (#O) 4a4c

DIA 25-JAN-72 13:50 8649

all about superwatch

(J8649) 25-JAN-72 13:50; Title: Author(s): Don I. Andrews/DIA; Distribution: Paul Rech, Stephen W. Miller, Michael D. Kudlick, George J Eilers, Donald R. Cone, Bonnar Cox, David R. Brown, Don Limuti, William R. Ferguson, Priscilla Lister, Robert L. Dendy, Linda L. Lane, Marilyn F. Auerbach, Walter L. Bass, Mary S. Church, William S. Duvall, Douglas C. Engelbart, Beauregard A. Hardeman, Martin E. Hardy, J. D. Hopper, Charles H. Irby, Mil Jernigan, Harvey G. Lehtman, John T. Melvin, Jeanne B. North, James C. Norton, Cindy Page, Bruce L. Parsley, William H. Paxton, Jeffrey C. Peters, Jake Ratliff, Barbara E. Row, Ed K. Van De Riet, Dirk H. van Nouhuys, Kenneth E. Victor, Don C. Wallace, Richard W. Watson, Don I. Andrews/SRI-ARC; Sub-Collections: SRI-ARC; Clerk: DIA; Origin: <ANDREWS>DOCSUPER.NLS;14, 24-JAN-72 20:57 DIA ;

WLB 25-JAN-72 16:36 8650

Peas Come to Your First POD Meeting

The first meeting of the Oak POD (Acorn?) will be held in Jim Norton's Office at 1:00 PM Wednesday 26 Jan 72. 1

WLB 25-JAN-72 16:36 8650

Peas Come to Your First POD Meeting

(J8650) 25-JAN-72 16:36; Title: Author(s): Walter L. Bass/WLB; Distribution: Beauregard A. Hardeman, J. D. Hopper, Don Limuti, Priscilla Lister, James C. Norton, William H. Paxton, Dirk H. van Nouhuys/BAH JDH DL PL JCN WHP DVN; Sub-Collections: SRI-ARC; Clerk: WLB;

DGE 25-JAN-72 17:00 8651

To Launch PODAC

INTRODUCTION 1

The purpose of the this memo is to launch a new activity within ARC -- our "Personal and Organizational Development ACtivity" (PODAC). 1a

It is to be a separate organizational setup from that for which we departmentalize our activities, hierarchically apportion our "pusher/contractor" responsibilities, etc. in the business of setting and pursuing our goals. Whereas the Executive Management Committee reports directly to me as a centralized body to coordinate ARC's executive activities, there will be a POD Committee (PODCOM) interacting with me in dealing with our PODAC matters. How PODCOM composition is established, and what PODAC is all about, is discussed below. 1a1

I feel that a conscious and continuous effort of this sort is so important to our goals that even if the trial approach being launched here doesn't work out, I expect to keep plugging away with PODAC experimentation and evolution toward meeting the basic needs outlined below. 1a2

This whole activity is aimed at serving two needs that exist within ARC: 1b

We who tell the world that we are learning how to show other teams how to achieve greater goal-pursuit effectiveness must constantly examine ourselves (the "example" that we are working with), both as an organizaton and as individuals, while making a conscious effort to understand how we are doing, and how we can improve. I am quite convinced that unless there is a strong, constant, and pervasive attitude among us that we want to keep developing ourselves, and unless we consciously keep trying to do so, then we are fooling ourselves about seriously pioneering this augmentation SYSTEM stuff. 1b1

To work on these things, we need a plentiul flow of information having to do with goals, attitudes, ambitions, feelings, etc. as they relate to the common pursuits, and purposeful discussion about strengths, weaknesses, and means for improvement, 1b1a

Then, assuming that such attitudes and commitments do exist, to carry out effective collaboration within a goal-oriented team requires a continuous seeking after information about the environment the group is in, and its

general state of affairs. Why this, why that, why not xxx, when ... etc.and giving of both as well as 1b2

DISCUSSION 2

PODAC's purpose is to facilitate communication and people interaction toward such ends as: 2a

Keeping everyone informed about goals, plans, decisions, etc 2a1

Keep everyone informed about the problems and opportunities facing ARC and its people and its goals. 2a2

Providing a multiway dialogue regarding personal development of every member of ARC. 2a3

Providing a forum for the expression of concerns, beliefs, ideas, feelings, dissension, etc. existing within any person or group in ARC about the way things are being done (or not being done), about our goals, etc. 2a4

Providing an organized mechanism for interactions among all parties toward affecting the understanding, beliefs and attitudes of each other, as a means of affecting the decisions and actions within ARC, toward what each thinks is the best goals, organization, products, behaviour, etc. 2a5

Better and more accurate statements of purposes, roles, modes of pursuing them, etc., are expected to be part of the product of the ongoing PODAC. 2b

In my initial design, PODAC and EMC are to have no direct affects upon one another. 2c

PODAC does not exist to vote on what ARC will do. PODAC will have no line-management responsibilities or authority. It is to be "orthogonal" to the management structure that commits resources, sets targets, hires, reviews, and is held accountable. 2c1

Similarly, EMC is but another set of people within the various sets of ARC's staff, and it holds no special role in PODAC beyond that. (Other sets are, for instance, CSO staff, PSO staff, women programmers, people over thirty, etc.) 2c2

SETUP PLANS 3

To Launch PODAC

Basic idea: 3a

Divide our whole activity into four groups 3a1

Aiming for balanced repesentation in age, sex, professional training, length of association with ARC, work roles, etc. 3a1a

As ARC grows, we would distribute new people around by some process that reasonably maintains uniformity in size and balance in membership, assuming that when we reached about 40 we would in some way split into five groups of eight each. 3a1b

Each group meets weekly for two hours. 3a2

I consider that attendance and participation will be part of the "employment conditions" for working within ARC, and thus expect to press strongly for complete and consistent attendance, 3a2a

The meetings are basically private to each group, and are to be devoted to dialogue about matters of personal and organizational development. I assume that there will be continuing interaction within PODAC about topics and modes for this dialogue. I also plan to interject my own ideas about bounds on topics and modes of dialogue, and the PODCOM and I will continuously review and set guidelines toward making PODAC more effective. 3a3

Each such group will appoint its own representative to the PODCOM. We will find some balsnce of time and duties that this PODCOM participation will carry with it, but we will seek a coordination and guiding function for it. 3a4

I expect to participate regularly, in PODCOM meetings, in periodic meetings with the groups during their private sessions -- in which the proceedings, involving me, will still be private. I aso would expect PODCOM and me to hold all-ARC metings as we feel would be useful. 3a5

We will run with this initial PODAC form for three months, and then evaluate it. We might chane it drastically, or not at all. I could kill it, conceivably, or it could get its change directions fromits own activity -- we'll wait and see. 3a6

Miscellaneous consideratons: 3b

Handling problem of size growth of lab: 3b1

Assume that we'd integrate new people into the existing n groups, until we'd have enough to form n+1 groups. 3b1a

Beginning at 4x8 (four groups of eight members each), we could grow to 4x10, split into 5x8, grow to be roughly 5x10, split into 6x8, etc. 5x5. 3b1b

Assume that questions of cost and value will become more clear, together with the dynamics, the problems, etc., and we'll count on a major evaluation -- in the meantime, we commit ourselves to the cost in order to explore this way. 3b2

DCE 25-JAN-72 17:00 8651

To Launch PODAC

Kickoff POD groupings:	3c
Cedar 9	3c1
Hardy	3c1a
Lane	3c1b
Lehtman	3c1c
North	3c1d
Parsley	3c1e
Peters	3c1f
Rech	3c1g
VanDeRiet	3c1h
Victor	3c1i
Fir 8	3c2
Church	3c2a
Ferguson	3c2b
Jernigan	3c2c
Melvin	3c2d
Page	3c2e
Ratliff	3c2f
Wallace	3c2g
Watson	3c2h
Oak 9	3c3
Bass	3c3a
Hardeman	3c3b
Hopper	3c3c
Kaye	3c3d
Limuti	3c3e
Lister	3c3f
Norton	3c3g
Paxton	3c3h
vanNouhuys	3c3i
Redwood 8	3c4
Andrews	3c4a
Auerbach	3c4b
Cone	3c4c
Dendy	3c4d
Duvall	3c4e
Irby	3c4f
Kudlick	3c4g
Row	3c4h

DCE 25-JAN-72 17:00 8651

To Launch PODAC

Full list of PODAC particpants:	3c5
Auerbach	3c5a
Bass	3c5b
Church	3c5c
Dendy	3c5d
Duvall	3c5e
Ferguson	3c5f
Hardeman	3c5g
Hardy	3c5h
Hopper	3c5i
Irby	3c5j
Jernigan	3c5k
Kudlick	3c5l
Lane	3c5m
Lehtman	3c5n
Limuti	3c5o
Lister	3c5p
Melvin	3c5q
North	3c5r
Norton	3c5s
Page	3c5t
Parsley	3c5u
Paxton	3c5v
Peters	3c5w
Ratliff	3c5x
Rech	3c5y
Row	3c5z
VanDeRiet	3c5a*
vanNouhuys	3c5aa
Victor	3c5ab
Wallace	3c5ac
Watson	3c5ad
Andrews	3c5ae
Cone	3c5af
Kaye	3c5ag

DCE 25-JAN-72 17:00 8651

To Launch PODAC

(J8651) 25-JAN-72 17:00; Title: Author(s): Douglas C. Engelbart/DCE;
Distribution: S.R.I. - Augmentation Research Center/&SRI-ARC;
Sub-Collections: SRI-ARC; Clerk: DCE;

HGL 25-JAN-72 17:10 8652

Cedars Arise and Form Your Roots

The first meeting of the Cedar pod will take place on Wednesday, 26 January 1972 at 1300. It will take place in the ARC Conference room. A fine time will be had by all. 1

HGL 25-JAN-72 17:10 8652

Cedars Arise and Form Your Roots

(J8652) 25-JAN-72 17:10; Title: Author(s): Harvey G. Lehtman/HGL; Distribution: Martin E. Hardy, Linda L. Lane, Harvey G. Lehtman, Jeanne B. North, Bruce L. Parsley, Jeffrey C. Peters, Paul Rech, Ed K. Van De Riet, Kenneth E. Victor/MEH LLL HGL JBN BLP JCP PR EKV KEV; Sub-Collections: SRI-ARC; Clerk: HGL;

MFA 26-JAN-72 0:30 8654

Your baseline stuff is ready 1

(J8654) 26-JAN-72 0:30; Author(s): Marilyn F. Auerbach/MFA;
Distribution: Don I. Andrews, William S. Duvall/DIA WSD;
Sub-Collections: SRI-ARC; Clerk: MFA;

EHF 26-JAN-72 6:25 8655

Request For Documents

I have had a request from Phil Messing for the following documents:

Twinkle-A Syntax Language For A Translator Writing System
NIC #7520

Design And Implementation Of A Table Driven Compiler System
NIC #3315

Could you send these to me as soon as possible?
thank you,
Ernie Forman 1

EHF 26-JAN-72 6:25 8655

Request For Documents

(J8655) 26-JAN-72 6:25; Title: Author(s): Ernest H Forman/EHF;
Distribution: Jeanne B. North/JBN; Sub-Collections: NIC; Clerk: EHF;

JCN 26-JAN-72 11:05 8656
Journal submission ideas for consideration

Journal submission ideas JCN 26 JAN 72 1

Bill, this is a request for consideration of the new Journal submitting option Doug has suggested and you and I have discussed recently. It appears to us that a change that would permit people to set up for entry all the data they intend to give the Journal system before actually entering the Journal system is needed. This would allow offline or off-journal-line entry WHEN THE USER HAS THE NEED and would allow him to perform almost all of his part of the submission work even when the journal (or the whole system for that matter) was down. 1a

One way to get the input information organized would be for the highest level statement to contain all data the user must, and wished to enter about the document. Such a branch might look something like: 1a1

(jenter) (username,filename,branch:zwn)
struc: branch
title: Here's a title
dist: DCE WSD
author: JCN
subcol: SRI-ARC DSS
etc. for any journal commands ? 1a2

Text of the message or document (unless the link specifies another location for the item being submitted?) 1a2a

More text of the message or document 1a2a1

More text of the message or document 1a2a2

Also, for those times when a user wishes to go directly to the Journal, the ability to submit directly from DNLS would also make it easier and, I think, really would encourage many more ideas to flow into the record. Having to go to TNLS -with its limitations- and losing your DNLS job state (return rings, compilations, window setups, etc.) is a real barrier to the use of the Journal. There are thoe magic moments when the idea is fresh where the entry process should be quick and easy.. system up or down, in or out of DNLS..etc. 1a3

I would be interested in any reactions you have to these ideas and would like to help any way I might... I personally would find the feature sextremely useful, since

Journal submission ideas for consideration

I am frequently out of phase with the Journal up or system up times. 1a4

for a preassigned number - use a name like: (jenter9034jcn) ? 1a5

JCN 26-JAN-72 11:05 8656

Journal submission ideas for consideration

(J8656) 26-JAN-72 11:05; Title: Author(s): James C. Norton/JCN; Distribution: Douglas C. Engelbart, William S. Duvall, Charles H. Irby, J. D. Hopper/DCE WSD CHI JDH; Sub-Collections: SRI-ARC; Clerk: JCN;

JBP 26-JAN-72 11:06 8660

ASCII on file ?

is the USASCII document on file at nic, does it have a nic number ?
i plan to reference the ASCII standard in some other documents and would like to hhave a nic number for it. 1

JBP 26-JAN-72 11:06 8660

ASCII on file ?

(J8660) 26-JAN-72 11:06; Title: Author(s): Jon B. Postel/JBP;
Distribution: Jeanne B. North/JBN; Sub-Collections: NIC; Clerk: JBP;

MSC 26-JAN-72 11:11 8661

New Versions of NLS

Plus ca change, plus c'est la meme merde - or - The future just ain't what it used to be. 1

Every time a new version of NLS gets brought up, new commands, new bugs, and the gods create general chaos for a while. To conserve energy, it seems reasonable to localize the confusion in one person, so I'm going to be co-ordinatoress of new NLS systems. 2

In this connexion, I present the following proposals for your delight and delectation. 3

NIC-NLS is strictly for fixing bugs. New features, command changes, etc, should not be implemented in NIC-NLS. 3a

There are two main reasons for this restriction: 3a1

It assures that NIC-NLS will (virtually) always be in a state such that it can be made the running system, so we can fix a crippling bug relatively quickly. 3a1a

We can make NIC-NLS the running system without worrying about changing a lot of documentation. 3a1b

We can bring up a new system (REL-NLS or NIC-NLS) three times a day - at 5:00 AM, at 1:00 PM, and any time after 6:00 PM. The person who brings up the new system should be available for the next two hours, so he can handle any problems that may arise. 3b

Making REL-NLS into the running system -- 3c

When you want to make REL-NLS into the running system, let me know 5 days ahead. I'll handle the testing described below, check with the rest of the programmers, and coordinate documentation changes with Marilyn. (If your changes require much new documentation, a week or 10 day notice would be appreciated.) 3c1

This doesn't mean that there will be five days while nobody can do anything, and I'm not asking for five days after you complete your changes, just an advance estimate. There may, however, be a day, while testing goes on, and while we copy the NLS/REL-NLS files to NIC-NLS, when we can't change the source files. I'm asking for such a large lead time so we'll have time to do the documentation, and so other programmers can plan their work in accord with bringing up the system. 3c1a

New Versions of NLS

Whenever REL-NLS is considered to be ready to be brought up, it should be used for at least five console hours by several people before it is made the running system. I'll make sure this is done. 3c2

I've made changes to NLS utility to make copying the NLS/REL-NLS files into NIC-NLS as painless as possible. Documentation of this procedure is forthcoming. Whover makes REL-NLS into the running system is responsible for seeing that the NIC-NLS files get updated. Obviously, this should be done as soon as possible after making REL-NLS into the running system. 3c3

I'll ask everyone not to change the NLS sources when we're ready to make a new subsystem from REL-NLS. Let me know when NIC-NLS has been updated, so I can tell everyone that NLS files can be modified. 3c4

Whenever you make a new <SUBSYS>NLS from NIC-NLS or from REL-NLS -- 3d

Let me know bfore -- I'll try to be the expert on the state of our multifarious versions of NLS. 3d1

A branch in <NLS>STATUS named NEWSUBSYS should be updated to include the following information: 3d2

the source (NICL-NLS) of the new system 3d2a

the ident of the person who brought it up, 3d2b

the time it was ssaved in SUBSYS, and 3d2c

the name and version of the backup system. 3d2d

I didn't make up these procedures just to be officious. Hopefully, they will centralize the information about the state of NLS and make the transition between systems less painful for users. I realize that all of this formalization is damned annoying and might delay changes to NLS, If you can think of less painful alternatives, I'm not committed to playing Madame Lafarge. 4

MSC 26-JAN-72 11:11 8661

New Versions of NLS

(J8661) 26-JAN-72 11:11; Title: Author(s): Mary S. Church/MSC; Distribution: Walter L. Bass, William S. Duvall, Mary S. Church, J. D. Hopper, Charles H. Irby, Harvey G. Lehtman, John T. Melvin, Bruce L. Parsley, William H. Paxton, Don C. Wallace, Ed K. Van De Riet, Richard W. Watson, William R. Ferguson, Marilyn F. Auerbach, John T. Melvin, Kenneth E. Victor, Robert L. Dendy/NPG DCW EKV RWW WRF MFA JTM KEV RLD; Sub-Collections: NIC NPG; Clerk: MSC;

JBP 26-JAN-72 11:18 8662

test

this test msg entered at 1111 on 26 jan 72.

1

JBP 26-JAN-72 11:18 8662

test

(J8662) 26-JAN-72 11:18; Title: Author(s): Jon B. Postel/JBP;
Distribution: Chuck S. Kline/CSK; Sub-Collections: NIC; Clerk: JBP;

JBP 26-JAN-72 11:33 8663

Complaint ?

Is it true that on line delivery of messages and documents to UCLA-NMC is not happening ? It appears as if there was some sort of mix up when you changed "ucla-7" to "ucla-mnc". 1

JBP 26-JAN-72 11:33 8663

Complaint ?

(J8663) 26-JAN-72 11:33; Title: Author(s): Jon B. Postel/JBP;
Distribution: John T. Melvin, Richard W. Watson/JTM RWW;
Sub-Collections: NIC; Clerk: JBP;

CSK 26-JAN-72 13:44 8666

Fix of UCLA-NMC

Dick,
I have fixed up most of the errors in
identfile using the identification submode.
Thanks for your reply. Things seem to work now.
Charley Kline

1

CSK 26-JAN-72 13:44 8666

Fix of UCLA-NMC

(J8666) 26-JAN-72 13:44; Title: Author(s): Chuck S. Kline/CSK; Distribution: Richard W. Watson, Jon B. Postel, Chuck S. Kline/RWW JBP CSK; Sub-Collections: NIC; Clerk: CSK;

OAK 26-JAN-72 17:02 8668

The Oak POD held its first meeting from 1400 - 1615 in JimNorton's office on 26 Jan 1972. Walt Bass and Dave Hopper were chosen to represent the Oak POD on PODCOM, and a regular meeting time of 3:00 PM Wednesday afternoons was agreed upon. 1

(J8668) 26-JAN-72 17:02; Author(s): Walter L. Bass, Beauregard A. Hardeman, J. D. Hopper, Don Limuti, Priscilla Lister, James C. Norton, William H. Paxton, Dirk H. van Nouhuys/OAK; Distribution: Don I. Andrews, Marilyn F. Auerbach, Donald R. Cone, Robert L. Dendy, William S. Duvall, Charles H. Irby, Michael D. Kudlick, Barbara E. Row, Martin E. Hardy, Linda L. Lane, Harvey G. Lehtman, Jeanne B. North, Bruce L. Parsley, Jeffrey C. Peters, Paul Rech, Ed K. Van De Riet, Kenneth E. Victor, Mary S. Church, William R. Ferguson, Mil Jernigan, John T. Melvin, Cindy Page, Jake Ratliff, Don C. Wallace, Richard W. Watson, Douglas C. Engelbart/RED CEDAR FIR DCE; Sub-Collections: NIC RED CEDAR FIR; Clerk: WLB;

Meeting between ARC and TYMSHARE

MINUTES OF THE MEETING BETWEEN ARC AND TYMSHARE PERSONNEL
24 JAN 1972 1

This meeting was called to discuss the possibility of using TYMSHARE for some of ARC's computing needs. ARC named several kinds of service requirements and TYMSHARE described its embryonic interest in the ARPA net and in TENEX. It was agreed to meet again in about a month when TYMSHARE might have a firmer commitment to the net and to TENEX. 2

Participants from SRI were Bob Dendy, Charles Irby, Smokey Wallace, Ed van de Riet, Mike Kudlick, Dick Watson, Jim Norton, and Mimi Church; from TYMSHARE - Max Beere, Bill Combs, Allen Gartner. 3

1. TYMSHARE currently has a computer network that its customers use. It has interfaces similar to the IMP and TIP; different kinds of machines are linked through the network. It uses slow speed lines - about 150 baud- for input, a maximum of 4800 baud for output. TYMSHARE is interested in joining the ARPA net, and has had some discussions with Larry Roberts. 3a

2. TYMSHARE has three PDP-10's now, running the DEC time-sharing monitor. They are interested in TENEX, but it's not clear how interested. 3b

3. Four possible arrangements between ARC and TYMSHARE were mentioned: 3c

a. TYMSHARE to provide the raw computing resource for the NIC. 3c1

b. TYMSHARE to maintain the ARC in-house machine. 3c2

c. TYMSHARE to provide raw computing support for both service and development needs. (ARC would not have an in-house large machine.) 3c3

d. TYMSHARE to provide backup spare drum insurance. 3c4

All of these arrangements (except perhaps the last) presume that TYMSHARE is a member of the ARPA net and that it provides TENEX service. Everyone agreed that if TYMSHARE were to provide service to ARC, it should begin with NIC support, then perhaps expand its service later. 3d

4. TYMSHARE would, hopefully, be able to provide more reliable and more economical service than ARC has now. The

economical question is open if we could get cheaper service from anther subsidized ARPA site. TYMESHARE, because of multiple PDP-10s, spare subsystems and parts, leverage with manufacturers, etc. could probably provide more reliability than we can. 3e

5. Several questions require further negotiation before any commitment can be made: 3f

a. Finances - this would have to be worked out with ARPA, especially the question of NIC service. 3f1

b. Special purpose hardware - how would our dislay system run with a remote computer (we would probably have a mini here) 3f2

c. Monitor modifications - how would ARC modify and debug a monitor that is being maintained by TYMSHARE (how would we get hands on the operator's console)? 3f3

6. Any plan would have to fit in with possible ARPA plans for a central TENEX facility, possibly not run by TYMSHARE. 3g

7. We would want to go slow in all this with any initial agreements (except possibly drum backup) starting mimimum 6 months to a year. Any agreement of this type would probably take that long for negotiations and all paaties being prepared. 3h

RWW MSC 26-JAN-72 17:04 8669

Meeting between ARC and TYMSHARE

(J8669) 26-JAN-72 17:04; Title: Author(s): Richard W. Watson, Mary S. Church/RWW MSC; Distribution: Paul Rech, Stephen W. Miller, Michael D. Kudlick, George J Eilers, Donald R. Cone, Bonnar Cox, David R. Brown, Don Limuti, William R. Ferguson, Priscilla Lister, Robert L. Dendy, Linda L. Lane, Marilyn F. Auerbach, Walter L. Bass, Mary S. Church, William S. Duvall, Douglas C. Engelbart, Beauregard A. Hardeman, Martin E. Hardy, J. D. Hopper, Charles H. Irby, Mil Jernigan, Harvey G. Lehtman, John T. Melvin, Jeanne B. North, James C. Norton, Cindy Page, Bruce L. Parsley, William H. Paxton, Jeffrey C. Peters, Jake Ratliff, Barbara E. Row, Ed K. Van De Riet, Dirk H. van Nouhuys, Kenneth E. Victor, Don C. Wallace, Richard W. Watson, Don I. Andrews, Steve D. Crocker, Lawrence G. Roberts/SRI-ARC SDC2 LGR; Sub-Collections: SRI-ARC; Clerk: MSC;
Origin: <CHURCH>MINUTES.NLS;2, 26-JAN-72 16:57 MSC ; ;

RWW 26-JAN-72 17:13 8670

We Fixed the ID File

Jon, Thanks for the message. We blew it and didn't get the user field changed. Should be ok now. 1

RWW 26-JAN-72 17:13 8670

We Fixed the ID File

(J8670) 26-JAN-72 17:13; Title: Author(s): Richard W. Watson/RWW; Distribution: Jon B. Postel, Jeanne B. North/JBP JBN; Sub-Collections: SRI-ARC; Clerk: RWW;

EHF 27-JAN-72 6:29 8671

I have had a request for BBN Report #2184, which we do not have on file. Can you provide us with a copy, or do I have to ask BBN directly?

Ernie Forman

P.S. The document title is MULTI-LINE-CONTROLLER 1

EHF 27-JAN-72 6:29 8671

(J8671) 27-JAN-72 6:29; Author(s): Ernest H Forman/EHF; Distribution: Jeanne B. North/JBN; Sub-Collections: NIC; Clerk: EHF;

JCN 27-JAN-72 8:13 8672

Proposed New Journal Submission Features - Pre-specified Catalog Data

Journal submission ideas JCN 26 JAN 72 1

Bill, this is a request for consideration of the new Journal submitting option Doug has suggested and you and I have discussed recently. It appears to us that a change that would permit people to set up for entry all the data they intend to give the Journal system before actually entering the Journal system is needed. This would allow offline or off-journal-line entry WHEN THE USER HAS THE NEED and would allow him to perform almost all of his part of the submission work even when the journal (or the whole system for that matter) was down. 1a

One way to get the input information organized would be for the highest level statement to contain all data the user must, and wished to enter about the document. Such a branch might look something like: 1a1

(jenter) (username,filename,branch:zwn)
struc: branch
title: Here's a title
dist: DCE WSD
author: JCN
subcol: SRI-ARC DSS
etc. for any journal commands ? 1a2

Text of the message or document (unless the link specifies another location for the item being submitted?) 1a2a

More text of the message or document 1a2a1

More text of the message or document 1a2a2

Also, for those times when a user wishes to go directly to the Journal, the ability to submit directly from DNLS would also make it easier and, I think, really would encourage many more ideas to flow into the record. Having to go to TNLS -with its limitations- and losing your DNLS job state (return rings, compilations, window setups, etc.) is a real barrier to the use of the Journal. There are thoe magic moments when the idea is fresh where the entry process should be quick and easy.. system up or down, in or out of DNLS..etc. 1a3

I would be interested in any reactions you have to these ideas and would like to help any way I might... I

personally would find the feature sextremely useful, since I am frequently out of phase with the Journal up or system up times. 1a4

for a preassigned number - use a name like: (jenter9034jcn) ? 1a5

JCN 27-JAN-72 8:13 8672

Proposed New Journal Submission Features - Pre-specified Catalog Data

(J8672) 27-JAN-72 8:13; Title: Author(s): James C. Norton/JCN; Distribution: William S. Duvall, Douglas C. Engelbart, Charles H. Irby/WSD DCE CHI; Sub-Collections: SRI-ARC; Clerk: JCN;

JBN 27-JAN-72 9:07 8673

How to Find the Message You Have

Alex - At your site, when you receive the message that you have a message, it means the site has a message, and you will find that the BBN-NET Directory will have a file MESSAGE.TXT. To get the message you copy it to tty thus: COPY MESSAGE.TXT alt TTY: CR 1

JBN 27-JAN-72 9:07 8673

How to Find the Message You Have

(J8673) 27-JAN-72 9:07; Title: Author(s): Jeanne B. North/JBN; Distribution: Alex A. McKenzie, Marilyn F. Auerbach, Cindy Page/AAM MFA CXP; Sub-Collections: SRI-ARC; Clerk: JBN;

DRC 27-JAN-72 9:45 8676

course test for course pso hostessees

now is the time for all good men to come to the aid of the pso girls 1

DRC 27-JAN-72 9:45 8676

course test for course pso hostessees

(J8676) 27-JAN-72 9:45; Title: Author(s): Donald R. Cone/DRC; Distribution: Donald R. Cone, Dirk H. van Nouhuys, James C. Norton, Richard W. Watson/DRC DVN JCN RWW; Sub-Collections: SRI-ARC; Clerk: DRC;

RJR 27-JAN-72 9:43 8677

course test message

how's your new football? 1

RJR 27-JAN-72 9:43 8677

course test message

(J8677) 27-JAN-72 9:43; Title: Author(s): Rilla J. Reynolds/RJR; Distribution: Donald R. Cone/DRC; Sub-Collections: NIC; Clerk: RJR;

SCB 27-JAN-72 9:45 8678

course test sample journal session

this is a sample message 1

SCB 27-JAN-72 9:45 8678

course test sample journal session

(J8678) 27-JAN-72 9:45; Title: Author(s): Stephen C. Butterfield/SCB; Distribution: Stephen C. Butterfield/SCB; Sub-Collections: NIC; Clerk: SCB;

AKS 27-JAN-72 9:48 8679

sample journal session

welcome to sunny california 1

AKS 27-JAN-72 9:48 8679

sample journal session

(J8679) 27-JAN-72 9:48; Title: Author(s): A. Kenneth Showalter/AKS;
Distribution: Dirk H. van Nouhuys/DVN; Sub-Collections: NIC; Clerk: AKS;

WLB 27-JAN-72 9:54 8680

Proposed Minor Changes to the Identification System

Group Address 1

In perusing the IDENTFILE I find that in every group entry the group address is identical to the group coordinator. This indicates to me that the group address field is superfluous and should be eliminated. 1a

This elimination would both (slightly) compact the identfile and (greatly) improve the efficiency of certain programs dealing with the processing of membership and distribution lists. 1a1

Clean-Up 2

Some fields in the identfile (e.g., memlists and affiliations) seem to come in several flavors -- involving whether there are 0 1 or 2 semi-colons preceeding the EOL. This results in some of the get-field routines returning garbage (semi-colons) along with the desired entry string. 2a

I propose that we standardize on having NO semi-colons preceeding any EOF, make sure that PSO personnel know this, and clean-up the existing entries (an A/F program could be used to do both this and the group address fix). 2a1

Directories 3

In producing the NIC id directory, it is desireable to suppress the listing of certain groups (e.g., NEWIDS) which are only of system interest. This is done by inserting into the Comments field for that group identfile entry the string "SPECIAL". 3a

I think that it will be increasingly desirable to have groups which are of interest only to certain other groups or affiliates -- such as our NLS team, buyer, interest, etc. groups. 3b

For these it will not only be necessary to suppress their listing from the NIC directory but to make it possible to prepare directories for individual groups or affiliates -- e.g., ARC. 3b1

To accomplish this I propose that we abandon the SPECIAL convention and in its place institute a new "Directories" field for groups: all groups which should appear in the NIC directory must have "NIC" in this field, the designators "SRI-ARC" and "NPG" will appear in NLS-related groups, etc. 3b2

I propose to implement (with Mimi's help) the above changes unless there are any objections. This will involve: 4

Changing GETIADD to get the group address from the coordinator field. 4a

Changing the IDENT system user-interaction programs to delete the seperate group address and to handle the "Directories" field. 4b

Cleaning up the IDENTFILE as outlined above. 4c

Setting the "Directories" fields in existing IDENTFILE entries (can be done with an A/F followed by hand editing of a few entries). 4d

Preparing a memo for PSO and documentation people outlining these changes and their implications for clerical processes. 4e

WLB 27-JAN-72 9:54 8680

Proposed Minor Changes to the Identification System

(J8680) 27-JAN-72 9:54; Title: Author(s): Walter L. Bass/WLB; Distribution: Douglas C. Engelbart, Charles H. Irby, James C. Norton, Jeanne B. North, Barbara E. Row, Richard W. Watson, William S. Duvall/DCE CHI JCN JBN BER RWW WSD; Sub-Collections: SRI-ARC; Clerk: WLB;
Origin: <BASS>MSS.NLS;3, 27-JAN-72 9:49 WLB ;

JBL 27-JAN-72 9:53 8681

Reply to test.

I got your message (8641?) of 15 January. 1

JBL 27-JAN-72 9:53 8681

Reply to test.

(J8681) 27-JAN-72 9:53; Title: Author(s): Joel B. Levin/JBL;
Distribution: Alex A. McKenzie/AAM; Sub-Collections: NIC; Clerk: JBL;

JBL 27-JAN-72 10:07 8682

End of series

Jeanne, thanx for all your help. It is all working quite well, as far as I can tell, and pretty buglessly and as advertised even. Sorry to be such a pest, but I'm pretty well done for now. 1

JBL 27-JAN-72 10:07 8682

End of series

(J8682) 27-JAN-72 10:07; Title: Author(s): Joel B. Levin/JBL;
Distribution: Jeanne B. North/JBN; Sub-Collections: NIC; Clerk: JBL;

Mary S. Church
Augmentation Research Center
Stanford Research Institute
Menlo Park, California 94025

To:
Engelbart Copy

8683

MSC 27-JAN-72 11:24 8683

Copying REL-NLS and NLS Files to NIC-NLS

NLS utility has been modified to copy source files to the NIC-NLS directory, and a RUNFIL command file created to copy the rel files. 1

Source file copies 2

NLS utilty (invoked through the TASKS file) now recognizes a "Copy" command. The file specified in "Copy filename" is assumed to be in the NLS directory unless a directory name is explicitly specified. This file is searched for a "FILE ...." statement; if one is found, then the statement is scanned for a phrase of the form "% .....<.....>....%". If this is found, then the text enclosed by the angle brackets is replaced by "NIC-NLS". (So rel files compiled from NIC-NLS sources have a better chance of winding up in a predictable place.) In any event (whether or not the "FILE..." statement is found), the file is output to <NIC-NLS>filename. 2a

A branch in TASKS, named COPY contains the appropriate commands to copy all of the current NLS source files into the NIC-NLS directory. This branch should be copied to the TODO branch, and NLS utilty run, to perform the copies. 2b

Rel file copies. 3

<NLS>COPREL contains commands for RUNFIL that 3a

1) delete all current versions of a rel file from the NIC-NLS directory, then 3a1

2) copies the rel file from REL-NLS into NIC-NLS directory. 3a2

Neither the source nor the rel file copy procedures has the audacity to expunge the NIC-NLS directory. This should be done after these processes have completed successfully. 4

MSC 27-JAN-72 11:24 8683

Copying REL-NLS and NLS Files to NIC-NLS

(J8683) 27-JAN-72 11:24; Title: Author(s): Mary S. Church/MSC; Distribution: Walter L. Bass, William S. Duvall, Mary S. Church, J. D. Hopper, Charles H. Irby, Harvey G. Lehtman, John T. Melvin, Bruce L. Parsley, William H. Paxton/NPG; Sub-Collections: NIC NPG; Clerk: MSC;

DIA 27-JAN-72 13:51 8685

concerning times for POD meetings

It is clear that, so far as our use of the system is concerned, no two POD meetings should occur at the same time. Perhaps we can put the POD meeting times on the activities chart... 1

DIA 27-JAN-72 13:51 8685

concerning times for POD meetings

(J8685) 27-JAN-72 13:51; Title: Author(s): Don I. Andrews/DIA; Distribution: Paul Rech, Stephen W. Miller, Michael D. Kudlick, George J Eilers, Donald R. Cone, Bonnar Cox, David R. Brown, Don Limuti, William R. Ferguson, Priscilla Lister, Robert L. Dendy, Linda L. Lane, Marilyn F. Auerbach, Walter L. Bass, Mary S. Church, William S. Duvall, Douglas C. Engelbart, Beauregard A. Hardeman, Martin E. Hardy, J. D. Hopper, Charles H. Irby, Mil Jernigan, Harvey G. Lehtman, John T. Melvin, Jeanne B. North, James C. Norton, Cindy Page, Bruce L. Parsley, William H. Paxton, Jeffrey C. Peters, Jake Ratliff, Barbara E. Row, Ed K. Van De Riet, Dirk H. van Nouhuys, Kenneth E. Victor, Don C. Wallace, Richard W. Watson, Don I. Andrews/SRI-ARC; Sub-Collections: SRI-ARC; Clerk: DIA;

FIR POD Meeting Notes

Notes on FIR POD meeting - Jan. 26, 1972

1. Several issues and problem areas were discussed, but the major focus was ARC working hours and ARC members' salaries.

A. The question of hours will continue to be discussed, and at the moment, the Firs have no uniform position on the matter.

B. However, concening salaries, some specific recommendations were made.

There seems to be a question as to whether ARC salaries are equitable, both within the group, and for the computer market of this area. To answer these two questions will require information on the salary scales of ARC, other groups within SRI, and neighboring and related industries.

The gathering of such information will require some time and effort. Toward this goal, it was suggested that a Salary Study Committee be set up. Smokey has volunteered to be a member, and if such a committee is actually set up, other volunteers will be needed.

2. Bill Ferguson was chosen as PODCOM member.

WRF 28-JAN-72 2:18 8692

FIR POD Meeting Notes

(J8692) 28-JAN-72 2:18; Title: Author(s): William R. Ferguson/WRF; Distribution: Paul Rech, Stephen W. Miller, Michael D. Kudlick, George J Eilers, Donald R. Cone, Bonnar Cox, David R. Brown, Don Limuti, William R. Ferguson, Priscilla Lister, Robert L. Dendy, Linda L. Lane, Marilyn F. Auerbach, Walter L. Bass, Mary S. Church, William S. Duvall, Douglas C. Engelbart, Beauregard A. Hardeman, Martin E. Hardy, J. D. Hopper, Charles H. Irby, Mil Jernigan, Harvey G. Lehtman, John T. Melvin, Jeanne B. North, James C. Norton, Cindy Page, Bruce L. Parsley, William H. Paxton, Jeffrey C. Peters, Jake Ratliff, Barbara E. Row, Ed K. Van De Riet, Dirk H. van Nouhuys, Kenneth E. Victor, Don C. Wallace, Richard W. Watson, Don I. Andrews/SRI-ARC; Sub-Collections: SRI-ARC; Clerk: WRF; Origin: <FERGUSON>FIR-NOTES.NLS;4, 28-JAN-72 2:14 WRF ; ;

KEV 28-JAN-72 3:03 8693

auto-logout of inactive jobs

We have a program that normally runs detached that will logout inactive jobs 1

any attached job that has been inactive for time t1 will be sent two control-c follwed by a warning message that says if they continue inactive for time t2 they will be automattically logged out 1a

if they wish to continue they merely type continue to the exec 1a1

KEV 28-JAN-72 3:03 8693

auto-logout of inactive jobs

(J8693) 28-JAN-72 3:03; Title: Author(s): Kenneth E. Victor/KEV; Distribution: Bob Van Tyul, Jeanne B. North, Robert L. Dendy, John T. Melvin, Kenneth E. Victor, John W. McConnell, Peggy M. Karp, Dan L. Murphy, Rod M. Fredrickson, Peter H. Lipman, Don C. Wallace, Carl M. Ellison, Ted R. Strollo/TUG; Sub-Collections: SRI-ARC TUG; Clerk: KEV; Origin: <VICTOR>AUTO-LOGOUT.NLS;1, 27-JAN-72 20:51 KEV ;

DVN AKS 28-JAN-72 10:27 8698

Course Test: A Statement

THESE COMMENTS ARE IN UPPER CASE

DVN AKS 28-JAN-72 10:27 8698

Course Test: A Statement

This is a demonstration of entering edited things in the journal. 1

DVN AKS 28-JAN-72 10:27 8698

Course Test: A Statement

(J8698) 28-JAN-72 10:27; Title: Author(s): Dirk H. van Nouhuys, A. Kenneth Snowalter/DVN AKS; Distribution: Richard W. Watson, Rilla J. Reynolds/RWW RJR; Keywords: edited; Sub-Collections: NIC; Clerk: BER;