NSW File Types

Dave Maynard
Jon Postel
10 March 75 1

Introduction 2

This is a preliminary note on the file types to be supported by the file packages for the initial NSW implementation, comments are encouraged and should be directed to Jon Postel [POSTEL at BBNB, or (415) 326-6200 x3718]. 2a

NSW Physical File Types 3

This document specifies the currently defined physical file types within the NSW, and specifies the PCP encodings used to communicate the files among various PCP processes. The actual PCP format , i.e. PCPB36 , PCPB8 or PCPTXT, used on the connection must be agreed upon between the two ends of a physical channel but is irrelevent to this discussion. 3a

Physical File Attributes: 3b

The Physical file type is specified by three attributes: 3b1

DATA TYPE: 3b1a

This attribute has the value CHARACTER or BINARY and specifies whether the file is comprised of character strings or bit strings. Since it is clearly possible to encode any file as either data type this attribute is not an absolute constraint on the contents of the file but rather an indication of the most advantageous encoding to use. 3b1a1

RECORD TYPE: 3b1b

Record Type indicates the record structure of the file in the originating process. It has the following legal values. 3b1b1

FIXED: The file consists of records of fixed size. 3b1b1a

VARIABLE: The file consists of records of variable size. 3b1b1b

STRUCTURE TYPE: 3b1c

This attribute specifies whether the file is a simple

sequence of records or whether there is a more complex record structure. The legal values are: 3b1c1

SEQUENTIAL: The file is transmitted as a sequence of records. 3b1c1a

SPARSE: Each record carries a record number along with it. The list of pairs (record number,record data) are simply ordered on record number. That is the record number of each record is greater than that of its predecessor. The record number of the first record cannot be less than zero. 3b1c1b

RANDOM: Each record carries a record number along with it. The constraints on record numbers are that they are unique, and that record numbers are non-negative. 3b1c1c

PCP encodings of files 3c

CHARACTER (FIXED/VARIABLE) SEQUENTIAL 3c1

LIST(%datarecord% CHARSTR, . . .) 3c1a

CHARACTER (FIXED/VARIABLE) (SPARSE/RANDOM) 3c2

LIST (LIST(%recordnumber% INTEGER, %datarecord% CHARSTR), . . .) 3c2a

BINARY (FIXED/VARIABLE) SEQUENTIAL 3c3

LIST(%datarecord% BITSTR, . . .) 3c3a

BINARY (FIXED/VARIABLE) (SPARSE/RANDOM) 3c4

LIST (LIST(%recordnumber% INTEGER, %datarecord% BITSTR), . . .) 3c4a

Use Types: 3d

In addition to Physical File Type each NSW file also has an attribute called use type which is assigned at creation time by the creating tool. This attribute is used to give an indication of the semantic content of the file. It is our intention that the WM will store a matrix whose entries are of the form (process name, package name, procedure name) and that is indexed by (source file physical type, source file use type, destination file physical file type, destination file use type). The procedure thus indexed has as parameters (source file name, destination file name) and will either return TRUE

indicating the destination file has been successfully created and entered into the NSW file system or FALSE indicating failure of the conversion procedure. 3d1

In the initial phases of the NSW it is expected that this conversion matrix will be extremely sparse. Indeed many of the elements of this matrix will never be implemented, for example the task of converting a 360 cobol object file into a fortran source file seems well beyond expected advances in the state of the art. However some of the entries in this conversion matrix will be supplied by the file packages of each TBH (NSW Tool Bearing Host). In addition tool purveyers may find it in their interest to supply elements of the matrix corresponding to the use types most commonly created or requested by their tool. This allows a potential user to integrate the use of this tool more easily with other tools he uses. 3d2

In the case where the use type of a file is undefined or where the element of the conversion matrix needed is empty it seems advantages to supply default conversions based soely on physical file type. In fact the set of conversions based solely upon physical file type should form the minimum set of conversions provided by the file package of each TBH. 3d3

The following is a first cut at defining the conversions based solely on physical file type. 3d4

Physical File Type conversions: 3d5

The following conversions are defined separaely for each Physical file attribute. Conversion between physical file types is accomplished by performing each of the three possible translations (one for each attribute concurrently. 3d5a

Some of the following conversions take arguements which specify conversion parameters. I am as yet unclear exactly who specifies these , how and when. It seems that the requestor and supplier of the file must negotiate the proper values for these parameters. In the case where the requestor is a user this is fairly straight forward, however the case in which the requestor is a tool which in turn might want to consult the user is less clear. 3d5b

Attribute conversion primitives: 3d5c

CHARACTER => BINARY 3d5c1

Each Character is simply converted to an eight bit

byte containing the ASCII character code in the low order 7 bits. 3d5c1a

BINARY => CHARACTER 3d5c2

Treat each 8 bit bite as containing one ASCII character in the low order 7 bits. 3d5c2a

FIXED => VARIABLE 3d5c3

Preceed each record by the character/bit count for the record. 3d5c3a

VARIABLE => FIXED 3d5c4

PARAMETERS (fixedrecordlength, fillcharacter, break, append) 3d5c4a

fixedrecordlength = INTEGER 3d5c4a1

fillcharacter = CHARSTR 3d5c4a2

break = BOOLEAN 3d5c4a3

append = BOOLEAN 3d5c4a4

If the input record is shorter than the requested fixedrecordlength and append is FALSE the record is padded with the fill character/bit , IF however append is TRUE a new input record is fetched and is inserted in the current output record beginning with the next unused character/bit position. This continues until the current output record is full at which point a new record is begun if break is true, otherwise the unused portion of the input buffer is discarded. 3d5c4a5

If the input record is longer than the fixedrecordlength and break is FALSE the record is truncated with the truncated portion being lost. If However break is true a next record is begun. This is repeated until the entire input record has been processed. If append is FALSE then the last fixed record is padded, otherwise the next input record continues filling this current fixed length record and this process is continued until the last input record is processed at which point the last output record is padded if necessary. 3d5c4a6

SEQUENTIAL -> SPARSE and SEQUENTIAL -> RANDOM 3d5c5

Parameters (initrecnum, recinc) 3d5c5a

initrecnum - INTEGER 3d5c5a1

recinc - INTEGER 3d5c5a2

Each input record is assigned a record number beginning with initrecnum and incrementing by recinc. 3d5c5a3

SPARSE -> SEQUENTIAL 3d5c6

Record numbers are simply discarded. 3d5c6a

RANDOM -> SEQUENTIAL 3d5c7

The receiving process collects all the input records, sorts them by record number if necessary and then discards the record numbers. 3d5c7a

SPARSE -> RANDOM 3d5c8

No conversion needed, SPARSE is a proper subset of RANDOM. 3d5c8a

RANDOM -> SPARSE 3d5c9

The receiving process collects all the input records and sorts them by record number if necessary. 3d5c9a

NSW File Types

(J25538) 10-MAR-75 00:57;;; Title: Author(s): Jonathan B. Postel/JBP; Distribution: /DSM([INFO-ONLY]) JEW([INFO-ONLY]) NPG([INFO-ONLY]) ; Sub-Collections: SRI-ARC NPG; Clerk: JBP; Origin: < POSTEL, FILETYPES.NLS;4, >, 10-MAR-75 00:33 JBP ;;;;####;

KIRK 10-MAR-75 02:54 25539

Viewspec v feedback bug

Now that we have displays that are clear to read if slow to re-creaty (anD fXll ofQfunny chWracters) it is important that an old bug be fixed with the viewspec window. Viewspec v does not register in the window after it returns from being "highlighted". Viewspec u appears when in fact v is in effect.

KIRK 10-MAR-75 02:54 25539

Viewspec v feedback bug

(J25539) 10-MAR-75 02:54;;;; Title: Author(s): Kirk E. Kelley/KIRK;
Distribution: /FEED([ACTION]) NPG([INFO-ONLY]) ;
Sub-Collections: SRI-ARC NPG; Clerk: KIRK;

RWW 10-MAR-75 11:16 25540

part of a reply hit CR by mistake stay tuned for next message

Duane, because of all the hassles with ELF, BBN and the network (we seem to be running about 600 baud) I'm about a week behind on the proposal, About wed there should be a very crude draft in (watson,nswprop at bbnb, that lists what I am proposing. The one area Bill and I have not settled on is with extensions to NLS. At very leaxst I think there needs to be one or two guys to work on things users want , KWAC made a big list, Finney at pentagon just send a good list etc. I think some measurement and tuning now that we have it split is also needed. 1

RWW 10-MAR-75 11:16 25540

part of a reply hit CR by mistake stay tuned for next message

(J25540) 10-MAR-75 11:16;;;; Title: Author(s): Richard W. Watson/RWW; Distribution: /DLS([INFO-ONLY]) ; Sub-Collections: SRI-ARC; Clerk: RWW;

RWW 10-MAR-75 11:24 25541

mess cont

Final report hanging in there still some work but essentially all written, push on Dirk. Could come to RADC if nexessary on next trip east let me know. Putting your computer where your mouth is and working in the great computer network market place in the sky gives me nostalgia for the past. ELF is crashing about 12 times aday and disconnecting from the net another 6. BBN went up and down like a yo yo acouple days last week. Characters move through the net like mollasses, there's no where to go but up. DCE while I'm in San Diego this week is goin g to organize something to work with Walker. We are working with Retz to stabilize ELF but we are first major group to use it and as with our early TEnex use and net use it will take time. NSW is not likely to be very stable in July as hassles above will show up there also and we essentially have our productivity cut in half untill we can get a decnet working environment again. Its all for the best I am sure and all net users will benefit from our pain(thats what I keep teelling myself) See yah Dick 1

mess cont

(J25541) 10-MAR-75 11:24;;;; Title: Author(s): Richard W. Watson/RWW; Distribution: /DLS([INFO-ONLY]) ; Sub-Collections: SRI-ARC; Clerk: RWW;

POOH 10-MAR-75 11:48 25542

@ seems ok

Thanks, it seems @ for statement names is ok, i'll let you know if it doesn't work again. 1

POOH 10-MAR-75 11:48 25542

@ seems ok

(J25542) 10-MAR-75 11:48;;;; Title: Author(s): Ann Weinberg/POOH;
Distribution: /FEED([INFO-ONLY]) ; Sub-Collections: SRI-ARC; Clerk:
POOH;

DVN 10-MAR-75 13:01 25543

Journal Indices Missing from BBN

The system says that the files listed in the links to files in this journal item are not online,

DVN 10-MAR-75 13:01 25543

Journal Indices Missing from BBN

From JUL 72 Thru JUL 73
<:ebtz> 1

AUTHOR INDEX
<catalog, arcjaincn1, 0:wD> 1a

NUMBER INDEX
<catalog, arcjnincn1, 0:wD> 1b

A-F TITLEWORD INDEX
<catalog, arcjtafincn1, 0:wD> 1c

G-O TITLEWORD INDEX
<catalog, arcjtgoincn1, 0:wD> 1d

P-Z TITLEWORD INDEX
<catalog, arcjtpzincn1, 0:wD> 1e

From JUL 73 Thru DEC 73
<:ebtz> 2

AUTHOR INDEX
<catalog, barcjaincn1, 0:wD> 2a

NUMBER INDEX
<catalog, barcjnincn1, 0:wD> 2b

A-F TITLEWORD INDEX
<catalog, barcjtafincn1, 0:wD> 2c

G-O TITLEWORD INDEX
<catalog, barcjtgoincn1, 0:wD> 2d

P-Z TITLEWORD INDEX
<catalog, barcjtpzincn1, 0:wD> 2e

From JAN 74 Thru JUN 74
<:ebtz> 3

AUTHOR INDEX
<catalog, carcjaincn1, 0:wD> 3a

NUMBER INDEX
<catalog, carcjnincn1, 0:wD> 3b

A-F TITLEWORD INDEX
<catalog, carcjtafincn1, 0:wD> 3c

Journal Indices Missing from BBN

G-O TITLEWORD INDEX
<catalog, carcjtgoincnl, 0:wD> 3d

P-Z TITLEWORD INDEX
<catalog, carcjtpzincnl, 0:wD> 3e

From JUL 74 Thru DEC 74
<:ebtz> 4

AUTHOR INDEX
<catalog, darcjaincnl, 0:wD> 4a

NUMBER INDEX
<catalog, darcjnincnl, 0:wD> 4b

A-F TITLEWORD INDEX
<catalog, darcjtafincnl, 0:wD> 4c

G-O TITLEWORD INDEX
<catalog, darcjtgoincnl, 0:wD> 4d

P-Z TITLEWORD INDEX
<catalog, darcjtpzincnl, 0:wD> 4e

DVN 10-MAR-75 13:01 25543

Journal Indices Missing from BBN

(J25543) 10-MAR-75 13:01;;; Title: Author(s): Dirk H. Van Nouhuys/DVN; Distribution: /FEED([ACTION]) JCP([INFO-ONLY]) KIRK([INFO-ONLY]) ; Sub-Collections: SRI-ARC; Clerk: DVN;

lunch

this primer demonstrates the commands used for writing a memo, editing it, and distributing it to other people. this process is explaimed for tnls which is the typewriter version of nls. 1

SGR 10-MAR-75 15:22 25544

lunch

(J25544) 10-MAR-75 15:22;;;; Title: Author(s): Susan Gail Roetter/SGR; Distribution: /KS([ACTION]) REM([INFO-ONLY]) ; Sub-Collections: SRI-ARC; Clerk: SGR;

DIA 10-MAR-75 15:36 25545

Statistics on OFFICE-1: Before and After New Drum

This is a brief summary of the differences in the OFFICE-1 system before and after the second drum was added, as shown by superwatch statistics. 1

Source: 2

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

Drum capacity: 3

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

Drum usage: 4

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

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

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

Number of users: 5

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

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

Response: 6

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

Efficiency: 7

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

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

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

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

Summary: 8

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

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

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

DIA 10-MAR-75 15:36 25545

Statistics on OFFICE-1: Before and After New Drum

(J25545) 10-MAR-75 15:36;;; Title: Author(s): Don I. Andrews/DIA;
Distribution: /DCE([ACTION]) JCN([ACTION]) JHB([ACTION])
RLL([ACTION]) RWW([INFO-ONLY]) NPG([INFO-ONLY]) ;
Sub-Collections: SRI-ARC NPG; Clerk: DIA; Origin: < ANDREWS,
O-1-NEWDRUM,NLS;3, >, 10-MAR-75 15:33 DIA ;;;;####;

DIA 10-MAR-75 19:08 25546

OFFICE-1 Disk Performance

I was about to suggest that we need a person whose only responsibility was to find out whose eating all the cookies between here an BBN, what mods we need for the pie-slice scheduler and how it hurts us, why it takes so long to do a jump link in NSW (when its up), etc. 1

Now I have an additional point. 2

When OFFICE-1 came up we had to make mods to superwatch to get statistics. When they were done, I took a look at the statistics and also noticed that the info on the disk queue was missing. Everybody was very busy and fixing that got a very low priority. It finaly was fixed some time ago, probably just after OFFICE-1 went to 256K. 3

Evidently NOBODY has been regularly giving the statistics from OFFICE-1 a thoughtful eye. 4

When I looked at them to find out what changes ocurred by adding another 600 pages to the drum, I noticed that the disk takes about 120 to 200 ms. per page transfer!!!!Multiply that by a queue length of about 7 and you have user frustration and lots of I/O wait time. Note: users are not actually waiting 7*200 ms, because reads are given priority over writes. 5

The old SRI-ARC system was getting a page transfer in about 35-45 ms (but with two disk controllers). The difference is so great that my first thought is that the statistics code is screwed up. But bugs in that code would result in wildwild numbers, not just extra large ones. So I suspect the following: 6

(1) The disk driver is not overlapping seeks at all. 6a

(2) The hardware is malfunctioning and seeks take a very long time or are done several times. 6b

(3) Those CALPOMP disks are just bad news. 6c

(4) Perhaps all of the above. 6d

Obviously, we shall persue the matter. If the statistics are bad you have my apologies ahead of time. 7

DIA 10-MAR-75 19:08 25546

OFFICE-1 Disk Performance

(J25546) 10-MAR-75 19:08;;; Title: Author(s): Don I. Andrews/DIA;
Distribution: /DCE([ACTION]) JCN([ACTION]) JHB([ACTION])
RLL([ACTION]) RWW([INFO-ONLY]) NPG([INFO-ONLY]) ;
Sub-Collections: SRI-ARC NPG; Clerk: DIA; Origin: < ANDREWS,
STATS-0-1-DISK,NLS;3, >, 10-MAR-75 16:29 DIA ;;;;####;

KIRK 11-MAR-75 01:39 25547

New version of NLS-8 at BBNB

Commands branches, Move File or Message and Help should work now. Insert Date has been made a first level command. The Useroptions Include command now remembers the directory as well as the program file name. 1

KIRK 11-MAR-75 01:39 25547

New version of NLS-8 at BBNB

(J25547) 11-MAR-75 01:39;;;; Title: Author(s): Kirk E. Kelley/KIRK;
Distribution: /SRI-ARC([INFO-ONLY]) ; Sub-Collections: SRI-ARC;
Clerk: KIRK;

JBP 11-MAR-75 03:43 25548

confusing useroptions

i cant figure out how to set my control characters so that escape and control b do different things when i am in tnls as it is esc is mapped into control b and i dont see how to undo the mapping. 1

JBP 11-MAR-75 03:43 25548

confusing useroptions

(J25548) 11-MAR-75 03:43;;;; Title: Author(s): Jonathan B. Postel/JBP; Distribution: /FEED([INFO-ONLY]) ; Sub-Collections: SRI-ARC; Clerk: JBP;

JBP 11-MAR-75 03:46 25549

terminal type confusion

it would be nice if the tenex trmstat and terminal type is commands both used the same identifiers for terminals, as it is trmstat uses numbers while terminal type is uses names and nowhere can a user correlate these. 1

JBP 11-MAR-75 03:46 25549

terminal type confusion

(J25549) 11-MAR-75 03:46;;;; Title: Author(s): Jonathan B. Postel/JBP; Distribution: /FEED([INFO-ONLY]) ; Sub-Collections: SRI-ARC; Clerk: JBP;

DVN 11-MAR-75 15:13 25550

please tell archive about <arcdocumentation>

the archive system has not learned that the directory that used to be called <documentation> is <arcdocumentation> at bbn and so we can't interogate . there are files in there i'd like. 1

DVN 11-MAR-75 15:13 25550

please tell archive about <arcdocumentation>

(J25550) 11-MAR-75 15:13;;;; Title: Author(s): Dirk H. Van Nouhuys/DVN; Distribution: /FEED([ACTION]) KIRK([INFO-ONLY]) POOH([INFO-ONLY]) JCP([INFO-ONLY]) ; Sub-Collections: SRI-ARC; Clerk: DVN;

HGL 11-MAR-75 18:25 25551

Smoking

I agree with Charles' message on smoking both in meeting areas and the console room. 1

HGL 11-MAR-75 18:25 25551

Smoking

(J25551) 11-MAR-75 18:25;;;; Title: Author(s): Harvey G. Lehtman/HGL; Distribution: /SRI-ARC([INFO-ONLY]) ; Sub-Collections: SRI-ARC; Clerk: HGL;

CHI 11-MAR-75 19:23 25552

NSW Frontend Procedures

sent via SNDMSG to Balzer and Crocker.

NSW Frontend Procedures

Introductory notes: 1

The following document (Journal # 25552) is a description of the PCP externally callable procedures in the NSW Frontend for special interactions with the user and for managing display terminals. As always, comments are invited and welcome. This description will be modified if comments and reactions seems to indicate that there are holes in the design. However, until it is updated, this should serve as the current description of the Frontend (Command Language Interpreter) callabe procedures and data structures. I expect we will come up with some more compact encoding of the data structures used herein, but this should serve as a logical description in any case. 1a

The FE will make an External data structure named WINDOWS (actually code inside the FE) available to each tool that will contain the window-id of its default and created windows, and all attributes for those windows. 1b

This data structure will be of the form 1b1

LIST (terminal-class, default-text-window, printer-window, graphiics-window, tool-created-windows) 1b1a

where the third and fourth elements may be EMPTY. 1b1b

terminal-class: INTEGER [0 %line-at-a-time typewriter, 1 %full-duplex typewriter, 2 %full duplex display with pointer%] 1b1b1

default-text-window: LIST (window-id, type, diag-coords, window-att); 1b1b2

See definitions below. 1b1b2a

printer-window: EMPTY / default-text-window; 1b1b3

graphics-window, graphics-window: EMPTY / default-text-window; 1b1b4

tool-created-windows: (each) default-text-window; 1b1b5

The Frontend will allow certain of the display primitives to be used for graphics display manipulation. This usage will be detected on the basis of the window-id that is used in the call. This is also true of printers and tape cassette units that are part of the terminal configuration (used in conjuction with Line Processors for July 1975). For the graphics terminal, the CLI will be able to mark selections made by the user on the graphics

display in two ways: 1) by drawing a dot at the specified location and 2) by redrawing characters that are being selected. Any other marking will have to be done by the tool using the write-literal procedure. 1c

To support this, the Line Processor should be slightly changed to specify the device it will support on its Copy Printer port as part of its response to the Interrogate command. This in turn will tell the CLI which functions the tool can reasonably perform on this terminal. 1d

PROFILE UPDATING PROCEDURES: 2

new-profile (user-id, profile); 2a

This allows the USER-OPTIONS tool to update the user's interaction profile in the FE when the user runs this tool to change the way the FE behaves towards him. The profile is assumed to be a simple bitstr that is already properly setup for use by the FE. This will minimize the amount of time needed to load the user profile. 2a1

user-id: INTEGER 2a2

profile: BITSTR 2a3

new-toollist (user-id, toollist); 2b

This allows the WM to update the list of tools this user is allowed to run whenever conditions warrent. This list is used to give help and tool name recognition to the user. It in no way "grants the user actual access to the tool. For July 75 the tool list will be defines as described below. This may change as we learn more about how we can help the user in this domain. 2b1

toollist: LIST (%toollist% LIST (LIST (%user's tool name% CHARSTR, %WM's tool name% CHARSTR), ...), %entry tool% (INTEGER %index to toollist% / EMPTY)) 2b2

PRESENTING INFORMATION TO THE USER (WOULD NORMALLY BE DONE THROUGH CML STATEMENTS): 3

show (message, confirmflag); 3a

The message is presented to the user. If confirmflag is TRUE, the user may not continue until he confirms that he saw the message. In this case, the SHOW procedure does not return to the caller until the user has confirmed the message. 3a1

message: CHARSTR %may contain formatting chars such as CR/LF% 3a2

confirmflag: BOOLEAN 3a3

IF TRUE, user must confirm that he saw message before he may continue. 3a3a

show-error (message, confirmflag); 3b

This procedure is used to present error or warning messages toothe user, Confirm flag is used as in SHOW. 3b1

GETTING CLARIFYING INFORMATION FROM THE USER (AMBIGUOUS FILES NAMES, ETC) 4

get-info (message, string => string-2); 4a

GET-INFO presents message to user, presents string to user as starting value. User may edit this or start over. In either event a new string is returned as string-2. This allows WM to interact with user to clarify ambiguous file names, etc. 4a1

message, string, string-2: CHARSTR. 4a2

MANIPULATING DISPLAY TERMINALS (NOT NECESSARY FOR MOST TOOLS) 5

MANY CHANGES TO THE SCREEN AT ONCE: 5a

batch-display-commands (display-commands => ids); 5a1

This routine is used to effect many changes to the screen at once. display-commands is a data structure that is decoded and in general results in calls on the rest of the display primitives described below. The form of display-commands is as follows: 5a1a

LIST(LIST (opcode, params), ...) 5a1a1

opcode: CHARSTR %procedure name -- could make this an integer??% / EMPTY %default to last procedure name used% 5a1a2

params: as appropriate for given procedure but EMPTY can be used to default things like window-id and string-id. 5a1a3

A list of string and line-segment identifiers is returned. 5a1b

ids: LIST (LIST (%string-id% INTEGER, LIST(%line-seg-id% INTEGER, ...)), ...) / EMPTY 5a1b1

WINDOW MANIPULATION: 5b

create-window (old-window-id, type, diag-coords, new-window-att -> window-id); 5b1

Used to create new text windows. Windows must be created with respect to an old window. The coordinate system is relative to the window and is in termms of character positions 5b1a

new-window-att: EMPTY %use old window values% / window-att; 5b1b

window-att: LIST (%window-visible% BOOLEAN, %window-priority% INTEGER, string-att) 5b1c

%window priority is an integer from 1 to 10 (1 being highest priority). Whenever two windows overlap, the text of the higher priority window will dominate (note this only effects the overlapped area, not necessarily the whole window). The lowest priority of the current tool's window is higher than those of active, but not current tools. The NSW EXEC also has a window, which is high priority when the user is talking to it. A high priority window that has nothing displayed in it will not effect the display of lower priority windows. The CLI's command feedback window is always higher priority than tool or NSW EXEC windows. The window that is created by the CLI when it starts a new tool will have priority 1 for that tool. The tool may change this if desirable.% 5b1c1

%A window (as well as individual line segments and strings within it) may be set invisible, in which case the user sees no image from that window.% 5b1c2

The coordinates in string-att is ignored. 5b1c3

old-window-id, window-id: INTEGER 5b1d

diag-coords: LIST (%upper-left x% INTEGER/EMPTY, %upper-left y% INTEGER/EMPTY, %lower-right x% INTEGER/EMPTY, %lower-right y% INTEGER/EMPTY); %not necessarily in character coordinates.% %EMPTY implies use of old value in window identified by old-window-id% 5b1e

type: INTEGER [0 = random, 1 = sequential, 2 = graphics] 5b1f

Special note: For display terminals with graphics capabilities, this primitive can be used to create graphics windows also. 5b1g

delete-window (window-id); 5b2

Delete the specified window. The window-id will no longer be valid, nor, of course, will any string-i's or line-seg-id's that belonged in that window. Any image on the user's terminal of this window is cleared. 5b2a

May be used for graphics windows as well as text windows. Can only delete windows that belong to you. Windows created for the tool by the CLI at run-tool time do not belong to the process and thus cannot be deleted. They can be manipulated in other ways, however. 5b2b

clear-window (window-id); 5b3

Deletes the contents of the window. Frees all string-id and line-seg-id's. Removes the image from the user's display. 5b3a

May not be used for graphics or printer windows. 5b3b

scroll-window(window-id, change, (EMPTY/string-id-1), (EMPTY/string-id-2)); 5b4

change: INTEGER % >0 => scroll up that many lines, < 0 => down% 5b4a

This scrolls a sequential window CHANGE lines (up if CHANGE is positive, down if negative). If the window is a random window, the whole window or a group of strings may be scrolled together. In this case CHANGE has the same effect as with a sequential window. 5b4b

set-window-attributes (window-id, window-att); 5b5

Set the specified attributes for the window. See definition of window-att. 5b5a

STRING MANIPULATION: 5c

write-string (window-id, string-att, string => string-id, LIST(%line segment ids% INTEGER, ...)); 5c1

Write the specified string at the specified location with the specified attributes. The string consists of individually attributed and positioned line segments (which do not cross line boundaries). An identifier for the string as a whole and for the individual line segments are returned. 5c1a

```
string-att: line-seg-att;                                          5c1b

   If EMPTY use window defaults as string defaults.                5c1b1

   MUST supply cords for origin of string.                         5c1b2

string:  LIST( linesegment, ...) / string-addr                     5c1c

   The coordinates in line segments are relative to
   origin-cords, which are relative to window cords.               5c1c1

linesegment: LIST ( line-seg-att, (CHARSTR /
line-seg-addr));                                                   5c1d

line-seg-att: EMPTY/ LIST ( cords/EMPTY %highlight% INTEGER
[0 %default it%, 1 %highlight it%, 2 %dont highlight%],
%visible% INTEGER [0 %default it%, 1 %visible%, 2
%invisible%], %selectable% BOOLEAN/EMPTY, %selector code%
INTEGER/EMPTY);                                                    5c1e

   %highlight: make this line segment stand out from the
   rest of the text on the display (in a manner that is
   appropriate for the device).%                                   5c1e1

   % visible: A line segment or string's image can be made
   visible or invisible to the user.%                              5c1e2

   % selectability: a line segment can be made selectable
   (in DSEL's and SSEL's only) or not selectable.  If
   selectable, then a selector code can be stored with it to
   be used by the selection processor.  This allows line
   segments to be selectable in some contexts and not in
   others.%                                                        5c1e3

   IF EMPTY use string defaults.                                   5c1e4

cords: LIST( %x -- in character positions (0 is leftmost
position)% INTEGER, %y -- in line positions (0 is top line)%
INTEGER)                                                           5c1f

string-addr: string-id / LIST(window-id, string-id);               5c1g

NOTES:                                                             5c1h

   Could also be used to write on copy-printer if have
   special window-id for that.                                     5c1h1

   Also serves copy-string function.                               5c1h2
```

May be used to write text into graphics windows, also. This implies FE know how to write text on graphics device. 5c1h3

replace-string (window-id, string-id, string-att, string); 5c2

Replaces the specified string with a new string or with a copy of a string already in a window belongiing to this process. Note that the old string-id now applies to the new string. Note also that the strings position within the window can be changed during the replace. 5c2a

move-string (window-id, string-att, string-addr => string-id, LIST(%line segment ids% INTEGER, ...)); 5c3

Note set-string-attributes can be used to change the pposition of a string within a window. Move-string can be used to move a string from one window to another. It is equivalent to using write-string to copy a string and then delete-string to delete the old copy. 5c3a

delete-string (window-id, string-id); 5c4

Delete the specified string and free the string-id. All line-segments that are part of the string are deleted also, of course. 5c4a

set-string-attribute (window-id, string-id, string-att); 5c5

Set the specified attributes for the specified string. Note that the position of the string within the window can be changed with this primitive. See definition of string-att (defaults will be the old values, in this case). 5c5a

reposition-string (window-id, string-id, cords); 5c6

This is just a special case of set-string-attributes, but the frequency with which it is done warrents a separate, more efficient call. 5c6a

LINE SEGMENT MANIPULATION; 5d

Write-line-segement(window-id, string-id, linsegment => line-seg-id); 5d1

Append a new line segment to the specified string. The identifier for the new line segment is returned. 5d1a

Note that his serves the copy function also. 5d1b

replace-line-segment (window-id, string-id, line-seg-id, linesegment); 5d2

Replace a specified line segment with a new or a copy of an old line segment. Note that attributes of the string can be changed during the replace. 5d2a

line-seg-addr: line-seg-id / LIST (window-id, string-id, line-seg-id) / LIST (string-id, line-seg-id); 5d3

move-line-segment (window-id, string-id, linesegment); 5d4

Note, this is equivalent to using write-line-segment to copy a line segment and then delete-line-segment to delete the old copy. 5d4a

set-line-segment-attributes (window-id, string-id, line-seg-id, line-seg-att); 5d5

Set the specified attributes for the specified line segment. 5d5a

reposition-line-segment (window-id, string-id, line-seg-id, cords); 5d6

This is just a special case of set-line-segment-attributes, but its frequency warrents a special, more efficient call. 5d6a

SECONDARY DEVICE MANIPULATION: 6

write-literal(window-id, literal-string); 6a

literal-string: CHARSTR 6a1

This is treated in a device dependent manner. It is expected to be used to drive secondary devices (such as a graphics display) attached to Line Processors but will also be used for other things. The window-id will tell the FE how to get the liiteral string to the correct device. The literal string will be passed to the device unchanged. 6a2

CHI 11-MAR-75 19:23 25552

NSW Frontend Procedures

(J25552) 11-MAR-75 19:23;;; Title: Author(s): Charles H. Irby/CHI;
Distribution: /NPG([INFO-ONLY]) RWW([INFO-ONLY]) NSW([
INFO-ONLY]) ; Sub-Collections: SRI-ARC NPG NSW; Clerk: CHI;
Origin: < NSW-SOURCES, FE-PRIMITIVES.NLS;9, >, 11-MAR-75 19:19 CHI
;;;;####;


MOUTON

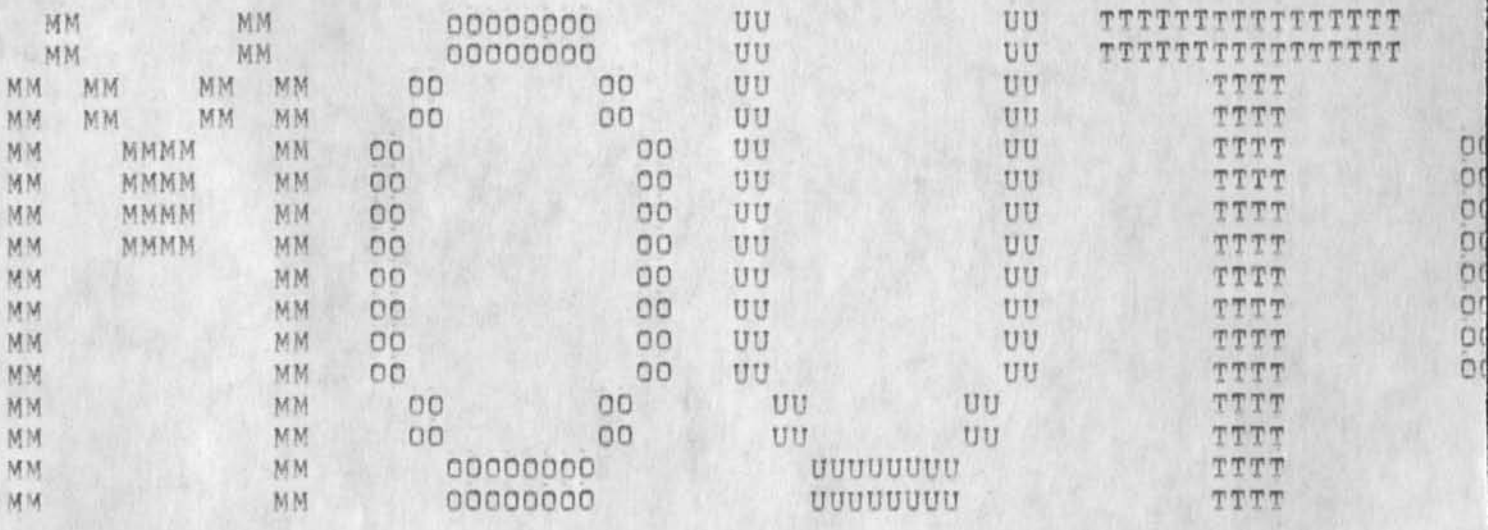

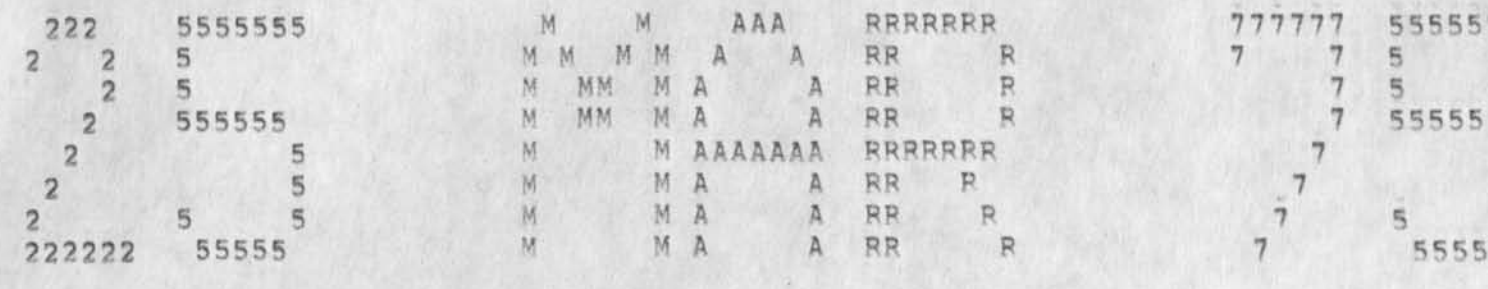

DVN 11-MAR-75 22:47 25553

Matrix Feeds Sandy

Sandy/feed/bugs
Sandy/bugs/feed
Feed/sandy/bugs
Feed/bugs/sandy
Bugs/feed/sandy
Bugs/sandy/feed.

1

DVN 11-MAR-75 22:47 25553

Matrix Feeds Sandy

(J25553) 11-MAR-75 22:47;;;; Title: Author(s): Dirk H. Van Nouhuys/DVN; Distribution: /SLJ([ACTION]) ; Sub-Collections: SRI-ARC; Clerk: DVN;

DVN 11-MAR-75 23:17 25554

Command Branches do and do not Work at Office-BB&N

Well, they work in the sense that the commands are carried out in the cases I tried, but the goings on on the screen and among the lineprocessor error lights would have to be seen to be believed. That's OK for running most branches, but if you are writing a branch and want to watch the action for debuging purposes, you could not. I have watched branches run smoothly at office-1, and display smoothly, so it is possible.

1

DVN 11-MAR-75 23:17 25554

Command Branches do and do not Work at Office-BB&N

(J25554) 11-MAR-75 23:17;;;; Title: Author(s): Dirk H. Van Nouhuys/DVN; Distribution: /FEED([ACTION]) KIRK([INFO-ONLY]) JMB([INFO-ONLY]) ; Sub-Collections: SRI-ARC; Clerk: DVN;

DVN 12-MAR-75 11:58 25555

Template and Example of Standard SRI Resume including Output Processor Directives

Mike Plako made this.

DVN 12-MAR-75 11:58 25555

Template and Example of Standard SRI Resume including Output Processor Directives

name, title
dept
div

1

Specialized professional competence 1a

. 1a1

Representative research assignments at SRI (since 19yy) 1b

. 1b1

Other professional experience 1c

. 1c1

Academic background 1d

. 1d1

Publications and patents 1e

. 1e1

Professional associations and honors 1f

. 1f1

date 1g

DVN 12-MAR-75 11:58 25555

Template and Example of Standard SRI Resume including Output Processor Directives

CAROLA ELLIOTT, RESEARCH ANALYST
TECHNOLOGY APPLICATIONS DEPARTMENT
ENGINEERING SYSTEMS DIVISION

2

Specialized professional competence 2a

- Computer programming; programming documentation; computer simulation models; applied statistics; data reduction 2a1

Representative research assignments at SRI (since 1971) 2b

- Statistical programming and application of statistical analysis programs (SPSS) for the income maintenance and tax study of the Urban and Social Systems Division 2b1
- Computer simulations of air defense weapons systems 2b2
- Modification of the Mixed Air Battle Simulation (MABS) model to enlarge its command and control capabilities (written in FORTRAN) 2b3
- Modification of the SRI Countersurveillance Reconnaissance Effectiveness Evaluation (SCREEN)/Ground Computer Program (written in FORTRAN) 2b4
- Assisted in computerized data reduction efforts 2b5

Other professional experience 2c

- Data processing documentation specialist, County of Santa Clara; analyzed and indexed source material and programs (written in COBOL) of the various county administration systems; established and maintained a program and documentation library 2c1
- Supervisory library technician, U.S. Naval Postgraduate School; maintained and updated the data base of the library's computerized information storage and retrieval system; supervised data processing; bibliographic literature searches 2c2

DVN 12-MAR-75 11:58 25555

Template and Example of Standard SRI Resume including Output Processor Directives

Academic background 2d

. B.A. in mathematics (with distinction, 1970), San Jose State University; graduate courses in applied statistics, linear programming, and projective geometry and analysis, San Jose State University; programming course in GPSS 2d1

December 1974 2e

DVN 12-MAR-75 11:58 25555

Template and Example of Standard SRI Resume including Output Processor Directives

(J25555) 12-MAR-75 11:58;;; Title: Author(s): Dirk H. Van Nouhuys/DVN; Distribution: /DIRT([INFO-ONLY]) DPCS([INFO-ONLY]) MAP2([INFO-ONLY]) PWO([INFO-ONLY]) PMK([INFO-ONLY]) DLS([INFO-ONLY]) ; Sub-Collections: SRI-ARC DIRT DPCS; Clerk: DVN; Origin: < PLACKO, SRIBIO,NLS;3, >, 3-MAR-75 09:02 MAP2 ;;;; ####;

JEW JBP 12-MAR-75 12:05 25556

Reply to Section II of BBN PCP Critique

This memo is in reply to SECTION II of Rick Schantz's and Bob Thomas' 4-MAR-75 critique of PCP and related NSW protocol designs; a forthcoming memo will address the higher-level design issues raised in SECTION I. In the present document, each of Rick's questions/comments/complaints has been reproduced, and our comments inserted beneath it in square brackets. 1

We greatly appreciate the time and effort that Rick and Bob have invested in reviewing our designs. We solicited their comments on Version I (ARC # 23904-6) published 2-SEP-74 and they provided very helpful and detailed written remarks, like those replied to here. Jon and I then spent a day at BBN in late October or early November, discussing Rick's and Bob's comments with them, and then spent the next month incorporating their suggestions (almost without exception) into what became the Version II design documents. These more recent comments will in like manner be reflected as far as possible in Version III. 2

Document: The Procedure Call Protocol (PCP.TXT) 3

4c3c3: PCP should stand by itself, without regard to higher level packages. Alternatively, if they are necessary to PCP itself, the things these packages do should become "kernel" implementations, not higher level packages. Package handles for example, seem to be a higher level construct and really may have no place in the PCP specification. The motivation for this comment is your statement "when PCP is used outside of the higher level framework provided by these packages, PH and PKH should be set to zero." Either these concepts are fundamental to PCP itself, or they are not. If they are not then you apparently are making low level provisions for an arbitrary set of high level functions. If so I am not convinced that other high level functions will not require other adaptations of the "kernel" PCP, which if true means we have a weak foundation. 3a

[I believe you are right. We can kernalize the concept of packages by defining two new PCP messages -- OPNPKS and CLSPKS, moving the OPNPKS and CLSPKS procedures from PSP (whose procedures are usually remote with respect to their caller) to PIP (whose procedures are local with respect to their caller). By so doing, we eliminate the need for PSP to be considered always open, and convert it from a miscellany of procedures to a "data store manipulation package" (DSMP) containing CRTDATA, DELDATA, RDDATA, WRDATA, LCKDATA, and UNLCKDATA.] 3a1

4c6b: What if all higher levels are not interested in the note? Must notes always propagate to the root node? 3b

[Yes they must. I know of no way that the system can know

which procedures along the thread of control may be interested in any particular NOTE and which will not. If the procedure which issues the NOTE intends it only for his immediate caller, he can and should use a coroutine return rather than a NOTE. It has been pointed out to us that there is no need that the issuer of a NOTE be required to wait while it propagates up the control thread, nor that it propagate back down at all, since by definition it can return no information to its issuer. I plan to make this change provided I can think of no problems involved with it.] 3b1

4c7: I don't understand this at all. It appears that simply by placing an exclamation point in the procedure name somehow makes this a privileged call. 3c

[I'm sorry; the documentation is misleading here. The exclamation point is purely a documentation construct, used to inform the reader that the procedure referred to (either in the definition of that procedure, or in an example of its use) is in effect a remote, system, internal subroutine, callable only by the system. The concept enters the protocol in the form of a BOOLEAN parameter in the CALPRO procedure, which must have the value TRUE to successfully cause a priviledge remote procedure to be invoked (see PCP -- 24459,5a2a1).] 3c1

4d1: Does the interrupt take place immediately (on receipt of the message) or at the next convenient time (i.e. after exiting a possible critical section)? 3d

[The latter. In the Tenex implementation, the CF first sets a lock shared by the CF and the PF to be interrupted. A procedure executing in the PF can therefore avoid interruption during a critical block of code by setting the lock upon entry to it and clearing it upon exit.] 3d1

5a2a3a: What is the state of the result arguments when a procedure returns with "failure"? 3e

[As in the case of a SUCCESS return, but unlike an ABORTED return (where the results are required to be an INTEGER error code and a optional CHARSTR error message), the system places no constraints upon the number or form of the results returned in a FAILURE return. The caller of course, must know what to expect from any particular procedure it calls, and that information must therefore be part of the procedure's documentation.] 3e1

5a2a4: The call handle is a locally generated item returned to the user immediately if he requests an out of line CALPRO, or is

returned when a procedure does a temporary return for an in line CALPRO. In any case the CH is a local identification tag. The CH parameter of the CALPRO message data structure, on the other hand, only has meaning between the two PCP implementations (i.e. "systems"). It is a convenience only, to use the same number for both purposes, and is made possible by requiring the callee to echo the caller's ch on all returns. It should be pointed out that this is an implementation efficiency consideration, not a system requirement. 3f

[The CH returned by the CALPRO procedure is always generated by the local (to CALPRO) system. The fact that the same handle is both returned to the user and sent to the remote process in the CALPRO message is indeed an implementation convenicnce, but it is not simply to make such a convenicnce possible that the callee's system is required to echo the CH in each RTNPRO message. Since the caller's system will in general have several CALPRO messages to the remote process outstanding simultaneously, some tag is required on each incoming RTNPRO message so that the required association between call and return can be made.] 3f1

5a2a5i: Is the integer [reslist=0] just a place holder i.e. an empty position in the list? If so, why use 0 when you mean empty? If a data structure could exist by the name "0" then a potential source of confusion is introduced. 3g

[First, I assume you mean "arglist", rather than "reslist". Second, although an EMPTY would have done just as well since the INTEGER can have but a single value, I've defined ARGLMSK as I have to maintain symmetry with RESLMSK, where the INTEGER can have two values. Third, if a data store named "0" were being referred to, the name would be encoded as a CHARSTR (all data store names are CHARSTRs) and embedded in "dselector", which is a LIST.] 3g1

5a2b1: whose ph is this? If it is the ph of the caller, how did the callee learn of it, when it is a process other than the one at the opposite end of the channel? 3h

[This point is confusing in the documentation. (PCP -- 24459,5a1d) defines this parameter, which, contrary to my documentation, is not really a process handle but rather routing information by which the receiver of a message selects the next process along a logical channel. A logical channel results from a call to ITDPRCS and provides a message path which follows the process tree structure, rather than short-circuiting it as a physical channel does. The logical channel consists of the two processes introduced, one or more

processes between them, and the physical channels which interconnect these several processes. A message is sent via the logical channel by passing it from one process to another until the ultimate destination process is reached. The routing parameter which appears in the message when it reaches process "i" identifies process "i+1". Before process "i" passes the message along, it stores the next routing code in the message.] 3h1

5a2c1: What happens to the old arguments if new ones are presented in say an interrupt continue RSMPRO? Do we start over again, or is it dependent on the particular procedure? Also, how can we tell the difference between the case where we want to respecify only 1 of many parameters (i.e. leave the others empty on resume), and the case where we want to respecify the others also with value empty (i.e. with no parameter)? 3i

[First, new arguments may only be supplied via RSMPRO after the callee makes a temporary return of subtype GENERALCOROUTINE or HELP. In particular, new arguments cannot be supplied following an interruption of the callee via INTPRO. But more fundamentally, RSMPRO does not supply new values for the arguments supplied in the original call, but rather communicates an independent and, as far as the system is concerned, arbitrary set of parameters to the callee. Perhaps the use of the word "arguments", or the use of the RSMPRO procedure to resume a procedure following both an interrupt by the caller and a coroutine return by the callee is confusing? I welcome any suggestions you may have.] 3i1

5a2c3: On co-routine temporary return you can not respecify the result list mask- why not? 3j

[I guess because I'm not convinced that argument- and result-list masks will be sufficiently useful to justify their existence -- anywhere. I agree, however, that they should be available consistently, i.e. always or not at all]. 3j1

5a2d2: This should read "temporary return of subtype interrupted". 3k

[Yes.] 3k1

5a2f3: What does "should probably break off communication" mean? 3l

[The implication of receiving a PCPERR message is that one or both connected processes are violating the (procedure call) protocol, either because they are running different versions of PCP or because one implementation has a bug. "Break off communication" was intended to sugest that this kind of error

be treated as catastrophic . Error recovery procedures at every level will be looked at more carefully in second-year NSW and are admittedly ill-understood at present.] 3l1

5b2b1: What is an inter-job channel? As yet unspecified? 3m

[For a third party to effect the establishment of a physical channel between two processes not currently connected by one, the system must determine each's physical location so that an appropriate form of inter-process communication can be selected. The assumption implicit in the current definition of "chntypsel" is that host address, job number within the host, and process number within the job, are the three parameters which determine the kind(s) of physical channels that can be used to reach it. For example, if the two processes' host addresses differ, ARPANET Host-Host Protocol must be employed. If the two processes reside on the same host but different jobs, an inter-job IPC facility may be available. And so forth. 3m1

As I believe the question was meant to point out, the three-level host-job-process address is presumptuous; the concept of a job or process may not exist in some systems, or the entity may not be addressable (as is true of processes in Tenex). We have already generalized PRCADDRs (in PMP's CRTPRC procedure) from "<host> <filename>" to "<host> <intrahostaddr>". A similar generalization feels right here.] 3m2

5b2b3: What is a process number in TENEX? 3n

[A job-global fork handle, should Tenex ever provide it. In the absence of a clean and efficient inter-process communication facility within Tenex, we will probably use the NCP, i.e. ignore the fact that the processes reside on the same host, in the hope that it will become one (e.g. by bypassing the IMP).] 3n1

5b3a4c: I don't understand what "prcnam" is and what it is used for. Some examples might help. 3o

["Prcname" is probably a poor term, and in any case, the concept is poorly developed in the documentation. PRCNAM is simply a character string which can be used to identify the process to a human user/programmer. "Works Manager", "Frontend", and "Editor" are NSW examples. We forsee the use of such names by (in particular) debuggers in presenting information to the programmer, e.g.: 3o1

Works Manager ADDR="OFFICE-1 WM,SAV" PH=3] 3o1a

5b3b3: What happens to any inferiors or other connections which the deleted process may have created? 3p

[Assuming that the process' user-code termination routine fails to delete/break them, they will be deleted (via DELPRC) and broken (via SEPPRCS) by the system.] 3p1

5b3d3: If the IPC implementation can support it (e.g. network connection) why can't we send a second message before notification of acceptance of the first? Also, is there an acknowledgement message sent by the receiver of a CALPRO (for example) message, other than the RTRNPRO which may not come for some time? 3q

[We can. I've found that restriction unnecessary and removed it. There is currently no "I got your request" acknowledgment. Very easily added, such an acknowledgment sounds like a good idea, especially if we introduce the notion of formally queuing some procedure call requests. How about making an intermediate acknowledgment optional, sent only by the callee's system when the callee indicates (either at compile- or run-time) that its execution requires significant real time, or when the request is queued by the system?] 3q1

5b3e1 and 5b3g1: Why separate the process specification from the port specification in order to create a channel? 3r

[Creation of a physical channel between two processes by a third is a multi-step process involving first selection of a form of IPC (intra-host or inter-host, 36-bit or 8-bit, ASCII or binary, etc.) on the basis of the two processes' relative locations, allocation in each process of a port consistent with the selected IPC form, and finally the cooperation of both processes in actually establishing the channel. 3r1

Each process provides as a data store, a list of the forms of IPC it supports. The third party obtains this list from process 1 via RDDATA, presents it to process 2 via ALOPOR, and ALOPOR selects and returns to the third party, an appropriate IPC form by noting process 1's location and effectively ANDing 1's list with its own. The third party then calls ALOPOR in process 1, giving it only one form of IPC from which to choose, namely that selected by process 2. At this point an IPC port has been allocated in each process, and the third party invokes the CRTCHNEND procedure in parallel in each, thereby establishing the connection. 3r2

Since the other process' location is information required in the IPC form selection and port allocation processes, it must be conveyed as an argument to ALOPOR. The port allocated by

the other process, on the other hand, cannot be conveyed via ALOPOR, since that same procedure invoked in the other process is what is employed to obtain it.] 3r3

5b3e1: What is the relationship between CRTPRC CRTCHNEND? Does create process use the create channel end? Can create channel end be used without new process logical connection creation? 3s

[CRTPRC and CRTCHNEND are independent; in particular, the former does not call the latter. CRTCHNEND, when invoked in parallel in two separate processes, establishes a physical channel between them, in a way similar to third-party file transfers in the current ARPANET FTP.] 3s1

5b3g2: Is it the case that the implementation of ALOPOR uses the remloc specification to select one of the elements of the channel type list as the best type channel to implement in these circumstances? If so, why specify the list at all with respect to whether it is inter-host, inter-job, etc. Let the "system" select and create the best type in this context- the user shouldn't be concerned with the type of channel implementation. 3t

[Yes, that is precisely the case, as I hope the discussion two questions above has demonstrated. But regarding the second part of your question, please carefully note the following, because a gross misunderstanding is evident (no doubt the fault of the documentation): the system DOES select the channel type, and all the details ARE completely hidden from the user code. The scenario two questions back is a description of how the CRTPHYCHN procedure is implemented by the system! The PMP procedures which correspond one-for-one with the IPC CRTCHNEND, DELCHNEND, ALOPOR, and DELPOR procedures are effectively internal (albeit remote) subroutines of higher-level PMP procedures which are called by the user. It's the necessity for this class of procedures which leads to the notion of priviledged procedures, and it's that notion which in turn prevents the invokation of such remote system procedures by user code. 3t1

Since the Version 2 documentation was produced, I have isolated these four procedures along with similar procedures used in creating logical channels in a separate, new package called PMXP (the Process Management Extension Package).] 3t2

Document: The Procedure Interface Package (PIP.TXT) 4

2b4: Is the signalling of event completion for out of line calls an interrupt, or must a data structure be polled to determine procedure completion without having to block? 4a

[In Tenex, it's being implemented as a pseudo interrupt.] 4a1

2b4: What do you mean by "in either case the caller must resume (or abort) the callee via RSMPRO to obtain its results? I thought the return would have already generated the results- (unless you mean a RSMPRO which doesn't actually get sent because the procedure has already completed, but rather all it does is get the arguments which have already been returned. 4b

[The RTNPRO message whose arrival from the callee's process signals the event associated with the call does indeed have contained the results of the procedure. However, the results are queued within the caller's system, and the caller must invoke RSMPRO to obtain them. The signalled event communicates to the user the fact of the return, while his invokation of RSMPRO communicates to him the details of the return.] 4b1

2c1: CH can not get released when using RSMPRO. 4c

[It can if the return is a permanent one.] 4c1

3b1: Comment: RSMPRO by not allowing for a ph parameter, implies that the "system" maintains a ch to ph mapping. True? 4d

[True.] 4d1

3c2: Does a user get control back immediately after executing INTPRO, or only after the procedure has returned in interrupted state? 4e

[Currently, the latter.] 4e1

Document: The PCP Support Package (PSP.TXT) 5

2b1: What does the package construction buy you? If certain packages are important to all processes, why not make these automatically opened (i.e. directly available). For others, I don't see the distinction between opening another package and creating another process which implements that package (only?). Perhaps it is an efficiency argument, i.e. it is too expensive to create a new process. 5a

[The openning of a package provides a convenient point at which to apply access controls to a family of procedures. One could of course include access checks in every procedure to which access had to be restricted, but that represents in general, considerable overhead. It's the same argument that leads to the concept of opening a file before elements of the file are

read or written; the necessary access checks are done once, rather than with every read/write. 5a1

The access check is performed by a routine supplied by the user-code, and no constraints are placed on the algorithm employed, which may consider the login parameters of the process' creator (now an argument to CRTPRC), the identity of the requesting process, and so forth. 5a2

The concept of packages also reduces process overhead, since initialization for families of procedures (i.e. packages) which aren't used during a particular creation of a process is bypassed (because package initialization is performed when the process is opened). 5a3

Processes also provide a convenient verbal handle for talking about procedure families.] 5a4

3c3a: In order to call RDDATA, we must specify a ph to CALPRO. A RDDATA parameter is dselector*, which also has a ph specification. Since ph's are relative to a given process, this leads to a potential conflict if the ph specified in dselector is not the one for the local (to the data structure) process. If it is not the ph of the local process, it will be in terms of the calling process, not the one executing the RDDATA procedure. Therefore, perhaps there is redundant specification of the ph field, if it is in both places. Or is it the case that you can specify data structures in an indirect manner? 5b

[The dselector is the general, system-wide means for addressing a data store, or any of its elements. When used as an argument to RDDATA/WRDATA, which by definition manipulate local data stores, the PH within the dselector must have the value SELF. In other contexts, e.g. within an argument- or result-list mask, the PH may have any value and is necessarily evaluated with respect to the callee's process (if I knew how to implement system-global process handles in a network environment, I surely would). 5b1

Document: The Process Management Package (PMP.TXT) 6

2b1: I disagree with this view. I consider the "servers" listening for initial requests on each constituent host to be part of the initial configuration. Thus we have at best, a number of independent trees which get interconnected by having these servers "introduce" the created process to the requesting process (on anther host, perhaps). These servers (tree roots) are known by virtue of a system wide fixed name i.e. the host and network contact socket. The user may not see the server processes and thus

can view his multi-machine creation as a single tree, but the MPSS is in fact an interconnection of somewhat independent trees. [Noted in passing: the "server" listening for connection attempts may have to manage the socket space alloted to the MPSS, so it is probably not totally void of function]. 6a

[We too have found it necessary (in addition to being able to create a new process instance) to be able to interconnect (or "splice", to use our terminology) previously existent and independent process trees (see the PCPVCHANGES document). In the NSW, the FE, the WM, and some tools will be the root processes of separate trees which will be spliced together.] 6a1

2c1a1: Why must the file (process) be run at HOST? The host parameter simply serves as part of the full syntax of the file name in order to distinguish/locate it. One can conceivably execute a remote CRTPRC procedure in an inferior and supply a filename on my host, but desire to have it run on the host of the inferior process. In your definition, you are using the HOST parameter to mean two potentially different things. Admittedly separating the two parameters (location of existence and location of execution) may present certain difficulties, but I think the idea is worth exploring. 6b

[Burying a file transfer within the CRTPRC procedure may well prove to be something one wants to be able to do. As mentioned elsewhere, I am generalizing PRCADDR to specify a host address and an intra-host address; in the latter, we can easily include the address of the host from which a copy of the process to be run is to be obtained.] 6b1

2d1: I don't understand what you are trying to accomplish with these generic process names. Perhaps some examples would help. 6c

[Answered elsewhere.] 6c1

2e1: If A and B have handles on each other, and A releases its handle on B (assume no hierarchichical relationship between them), does B get actively notified that his handle for A is no longer any good? [I don't mean is a table entry somewhere deleted, I mean is any executing code notified of the change in environment]. Can process B do anything about the release of the communication path? Can an inferior at any time just detach from its superior by releasing its handle on the superior? 6d

[A cannot release its handle for B. The third party which allocated each's handle for the other via ITDPRCS must release them both via SEPPRCS. Whenever A knows B, B also knows A. When the two processes are separated, a user-code routine in

each process IS invoked by the system, notifying the user code that its handle for the other process is about to be released and giving it the opportunity to take whatever action it wishes. Each process' user code is also notified, by a similar mechanism, when the processes are first introduced, allowing the user code to refuse he introduction or to perform whatever action may be required.] 6d1

At present, a created process cannot initiate its own deletion, although we may want to provide such an ability in the future. A process that was spliced to, however, MAY initiate disconnection from the process that spliced to it.] 6d2

2f1b: Why would a process want to be included in the path between two of its introducees, except to monitor or alter their communication? That is, won't most introductions also ask for direct process connection, thereby indicating that this should be a default? 6e

[I believe that most connections resulting from ITDPRCS will indeed use physical channels. The logical channel exists for two reasons. The first is to permit two processes to know each other, either for a very brief interval or for very infrequent interchanges, without creation (in the first case) or maintenance (in the second) of a physical channel. The second reason is that I suspect that there might be situations in which one process has to be able to ascertain its relationship to another with respect to the process tree (e.g. so that the system can verify the legality of some operation). 6e1

I don't understand the "default" suggestion. What could be easier than giving a BOOLEAN argument the value TRUE to obtain a physical channel, or FALSE to do otherwise?] 6e2

2f1d: Can more than one logical channel exist between two processes? If so, what does this mean? 6f

[If one were to apply ITDPRCS to the same process pair twice in succession, two logical (and physical) channels would indeed be created (and two pairs of process handles allocated. Although I see no particular utility to such a strategy, neither do I see any reason to specifically prohibit it, although that would certainly be easy enough to do.] 6f1

2f1e: Must a process ("system") know about all of the segments of all of its inferior's logical paths? How does the recipient of a request know which process along the line generated it? Since handles are relative, must an intermediary know all the names for a particular inferior, to forward its messages? 6g

[I hope that the answer to a previous question has clarified this. The system code with in a processkeeps in its tables information about only those logical channels which pass through that process, and for each such channel, only one piece of routing information in each direction. Only the terminal processes of each logical channel (i.e. the introduced processes) can transmit procedure call request via it; there is therefore no problem involved in determining the source of the request.] 6g1

3a1b: Can a superior use CALPRO to call the CRTPRC procedure in his inferior, thus directly creating a "grandson"? Do we then have a direct handle on a non-direct inferior? 6h

[CRTPRC can be called remotely only if the user code in the process that contains it permits its PMP package to be opened remotely. However, I believe that PMP procedures are not useful as remote procedures and will not be used as such. But if CRTPRC were called remotely, the effect of the call would be the same as if it had been called locally. CRTPRC always creates a process directly inferior to the process which contains CRTPRC itself, and the process handle that results is relative to the new process' direct superior.] 6h1

3a2b: What happens to other processes, channels, etc. of the deleted process? 6i

[Answered elsewhere.] 6i1

3a3c: With respect to "subsequent exchanges"- aren't these exchanges specified? Just having the handle passed as an argument to a subsequent procedure seems inadequate, since the "system" should somehow validate its use by the introduced process. 6j

[I don't fully understand this question. Use of both handles HAS been validated by the system code in each introduced process, once the ITDPRCS procedure has returned. "Subsequent exchanges" between user code in the third party and each of the newly-introduced processes are in general required to give the newly-available process handle meaning to the process' user code. In the NSW, the WM creates the tool as a direct inferior and then introduces it to the FE. The WM then calls BG$_{NN}$SW in the tool, passing it as an argument its process handle for the FE.] 6j1

3a4b: What is the effect on any on-going communication between the separated processes? 6k

[There are a number of design options here, and I'm not yet

sure which is best. One could "drain" the channel before deleting it, i.e. the system could mark the channel unavailable for use, refuse any subsequent attempts to use it, wait for any pending calls to complete, and THEN delete the channel. Alternatively, the system could force an ABRPRQ for each outstanding call and then delete the channel immediately. There are probably other options. Any suggestions?] 6k1

3a5b: What if we can't allocate COUNT more processors? It seems that we have no return argument which says how many we are willing to allocate. Can we only return failure, which means we allocate no more? [Aside: I have my doubts about this whole notion of reserving a number of "processors" within a process. Isn't an individual process a better model for these potential parallel executions, with the number of processors only instrumental in determining how many events can be (seem to be) in actual execution simultaneously. Is this construct just another "efficiency" motivated feature to avoid the overhead of process creation, or is there more to it? 61

[I have already redefined ALOPCRS to allocate as many processors as it can, with the number specified by the caller treated as an upper bound, and to return the number actually allocated as a result. The concept of processors arises exactly as you surmise -- from the need to allow within the model, processes which can execute two or more procedures simultaneously. Processor allocation is simply a mechanism by which a process A can prevent itself from being effectively "locked out" of another process B by one or more of the other processes which may have access to B. Like you, I am uncertain of the utility of the whole concept of processor allocation, and we'll eliminate it if it doesn't provi useful.] 611

3b1c: (related to 3a3c) How does this channel get passed to processes ph1 and ph2 for their use? 6m

[Answer related to 3a3c too. The user code in the third party must,] pass the newly-assigned port handles to their respective processes 6m1

3c1b: When do we return from this procedure, when the channel is created? If so, this obviously has to be done in parallel with the other end. If this is like RFC, is there the concept of one side "listening" (i.e. a passive participant) to help with synchronization? 6n

[Yes. Yes. Yes; we've added a BOOLEAN argument with value ACTIVE=TRUE / PASSIVE=FALSE to ALOPHYPOR.] 6n1

3c3b: Why must REMLOC be specified here and not with REMPORT during the create? Is it because the site determines the type of channel to be selected? If so, that whole function should probably be consolidated and also made transparent to he user i.e. by default, select the best type of channel for the two items to be connected. 6o

[Already answered.] 6o1

3d4b: What if the other half is not deleted in parallel? I think I'm confused as to what deleted in one direction means. Is it like deletion pending? 6p

[Since it's the third party in the "middle" of the logical channel that deletes it (via SEPPRCs), the system must run the logical channel in both directions to flush the routing information from each process along the channel. It does this by calling DELLOGCHNHLF in each of the two adjacent processes, which in turn propagate the request until it reaches the terminal processes. These two calls to DELLOGCHNHLF are made by the third party ('s system code) in parallel for efficiency; since they ARE called by the system, we can assure that they WILL be called in parallel.] 6p1

3d5b: I don't understand what a duplicate logical channel is for? I thought logical channels were paths for passing control, and physical channels were for implementing logical channels and also for auxiliary connections, if needed. In such a model, what does a duplicate logical channel do? 6q

[When one of the processes to be introduced to another via ITDPRCS is itself known to the third party via an introduction, creation of the new logical channel is accomplished by duplicating and extending the logical channel that supports the existing introduction. The word "duplicate" describes the manner in which certain logical channels are constructed from others, not the kind of channel that results. As the documentation obviously failed miserably to make clear, DUPLOGCHNHLF is one of several PMP remote procedures used strictly as internal, system subroutines and never called directly by user code.] 6q1

Document: PCP Data Structure Formats (PCPFMT.TXT) 7

3b1a: The BNF definition of key is probably incorrect in that it is not a general data structure; rather it is a data structure that can not have a key (i.e. not optional). 7a

[True. Laziness in documentation on my part.] 7a1

3b1d: Why does the definition of BITSTR have a length indicator at its head but INTEGER does not? Can't we have an arbitrarily long sequence of digits to form an integer? 7b

[Because they're inconsistently defined; sorry. The PCPIXT specification of the integer's value can indeed consist of an arbitrary number of digits ("$" means "arbitrary number", though I nowhere define the notation). One difference: the length of the BITSTR becomes a property of the data structure when assembled in the destination machine, while the length of the INTEGER does not. Integers must be represenable in a fixed number of bits internally (probably 16).] 7b1

3b1g: The "length" field is the count of the number of items in the data structure, I presume, and not the length of the entire structure. 7c

[True.] 7c1

4b1b1c: Why is value field of SIXBIT string limited to a single word? 7d

[No good reason. But I'm not convinced of the utility of these alternate CHARSTR encodings, anyway, and may eliminate them entirely.] 7d1

4b1b6a2 and 4b1b6c2: I assume that for header type list, each element in the value field has its own header to indicate its type and delimit its key and value. With repeated header, how can you know where one element ends and the next one begins if you don't have individual headers indicating this? 7e

[The lengths of the value and key fields which, by definition in this case, are the same for every element, are specified in the common header. This format might be used, for example, to transmit a LIST of key-less INTEGERS, in which case only a single word of overhead (i.e. one header) would be required, and the values of the integers would be contiguous.] 7e1

4b1b6b1: Shouldn't element count really be "repeat count" with element count equal to one understood? 7f

[Yes.] 7f1

Document: PCP ARPANET Interhost IPC Implementation (PCPHST.TXT) 8

3a2: What has to happen if the PRCADDR is specified as $F HOST FILENAME? 8a

[Briefly and simply, CRTPRC establishes a full-duplex Network connection via ICP to some fixed "PCP" contact socket at the indicated host, and transmits the filename to the connected process, which assumes it to be the name of a file in some fixed local directory.] 8a1

3e2: What if ICP is needed to multiplex a set of sockets? Can create channel be used independently of create process? 8b

[ALOPOR allocates a "port", which in the case of PCPNET IPC, is a local socket pair. CRTCHNEND establishes a channel (i.e. a pair of Network connections) using the allocated socket numbers. The socket allocation is therefore accomplished not by an ICP, but by the same machinery that user and server use to implement ICP.] 8b1

Document: PCP Interfork IPC Implementation (PCPFRK.TXT) 9

2a1b: You will probably be waiting a long time for job global TENEX fork handles. 9a

[Then we'll use a less efficient intra-host IPC implementation (i.e. the NCP) for a long time.] 9a1

2c1a1: No size of window (i.e. number of pages) in descriptor? 9b

[No. That information is conveyed as "channel width" in CHNTYPSEL in ALOPOR (see PCP -- 24459,5b2b1), although not explicitly stated in the documentation.] 9b1

2c2: How does receiving a message from the other fork lock the window? Is it the case that the window was locked by the sender and never unlocked? 9c

[The term "lock" is defined in the paragraph you cite, and in particular it does not mean, "increment a word and find the result zero." But by your definition, yes, the sender never unlocks the window (i.e. never decrements the lock word).] 9c1

2c3: How does returning a message unlock the window? 9d

[By convention.] 9d1

2c3: With respect to seeing window enqueued and therefore sending NOP, it appears that if a process receives a message and discovers the window lock indicates enqueued, it can only be this recipient process that is enqueued (assuming a channel connects only two processes). In such a case, isn't sending a NOP "waking up" the

wrong process? If not, does the original sender have to send another NOP back to the enqueued process? 9e

[NO, No, No. If a process receives a message and discovers the lock enqueued upon, it means that the sender wishes to send a second message.] 9e1

2c4: I assume that two inferiors can not share a single window (channel) to a common superior. True? 9f

[True.] 9f1

2e1: How does a multi-window message sender know when it can place the second (and all subsequent) "packet" in the window? 9g

[In precisely the same manner as it knows when it can send the next message: by waiting for the NOP from the other fork.] 9g1

2e1: Does the length field of a multi-window message indicate the total message length or the length of the particular "packet"? Can a multi-packet message be sent in any segmet lengths? 9h

[The latter. Yes. The scheme for transmitting multi-packet messages has been changed somewhat during implementation.] 9h1

3e2: What are you trying to avoid by specifying that remote fork's handle must be larger than the local fork's? 9i

[This was a first-pass kludge for effectively designating one fork as active (and responsible for doing the pMAPs (or for sending the RFC) and the other as passive. We have since added a boolean argument to ALOPOR to perform this function.] 9i1

Document: The File Package (FILE.TXT) 10

3b1d: Does the "use type" exist in the file descriptor? There is no mention of it in the file descriptor data structure. 10a

[Use types are a topic of current concern, we are trying to come to grips with our exact needs here. The use type was intended to be part of the file descriptor (see 3b3a1a).] 10a1

3b2a: What about APPEND access? EXECUTE? 10b

[This is a general question really -- how much file system should the file package provide? only what is needed for NSW, or enough to be generally useful, or a complete files system with elaborate access controls (al la Multics)? We chose the middle path but perhaps this was unwise. In any case we are

open for discussion on this topic, We do need to decide quickly what has to be done for NSW,] 10b1

3b2a: What does it mean, for example, to have public read but not creator read? The rules for access checking are not fully specified i.e, order, override, etc. 10c

[As indicated above we see this as a fluid area, We do need to have the access rules more carefully spelled out, In general the access controls are modeled on the Tenex controls so the rules should apply (where are they documented?),] 10c1

3b3a: What about file size, byte size/count, record format in the file descriptor? 10d

[The format or structural description of a file is an attribute of the file like (or part of) use type and may be kept as part of the file descriptor, We have a recent note on structure types, The issue of file size needs to be addressed,] 10d1

3c1: What does "the files within a process" mean? Do you mean the files accessible to a user? 10e

[This phrase was meant to indicate the files under control of the file package process,] 10e1

3c1a1: Can a user OPEN another users login directory? If so, what does MEM access mean if the other user has the password? Is this password the login password of the owner, in which case giving out the password means that someone can pose as OWNER, If not, is there a password which other people use, but which the owner can/does not use? 10f

[Good question, what we intended was something like the Tenex group concept, but we didnt get it right,] 10f1

3c2a: A thought- we may want the ability to APPEND new files to a directory, but not have the access to alter existing directory items as would be the case if we were granted write access, 10g

[Sounds useful to me, again we have to question the scope of the file system design to be implemented in the file package,] 10g1

3c3a1a: Don't you mean "directory creator"? 10h

[Yes,] 10h1

3c3a1b: Where is the directory password kept? Associated with the directory descriptor I assume, 10i

[Yes, it should be right after <crtor> -- <password> %password% CHARSTR.] 10i1

3ei: A thought- why isn't a directory just a particular file type, with only member access (i.e. with password)? 10j

[Two views on this, one is that it is an implementation choice; or two is that it is an extension of the file system to include graphs of directories(again -- what is the scope of this file system?).] 10j1

4a1b: Counter point- In order to use remote files, the user must actively first create a remote process, then execute remotely the OPNDIR procedure. Contrast this with an approach where the user executes OPNDIR locally and specifies a remote directory as a parameter. 10k

[There is a remote process in either case, and the user must supply the same type of identifing information in either case and the remote system wants to have accounting control in either case.] 10k1

4a1b: Is "DID" unique to a single process? Assuming it is, if I want to have an inferior use files in a directory I have opened, must I execute all calls for him (or alternatively, have him access the files using my procedure calls)? 10l

[Yes, but there could be a generalization of the OPNDIR call to allow another argument -- the handle of the process for which the directory is opened (could be a list too). That is the superior creates a file package process (FPP) and some other inferior process (OIP), the superior then introduces the OIP to the FPP and thus knows the handle by which the OIP is known to the FPP, the superior can then call OPNDIR in the FPP with the OIP handle as a an argument.] 10l1

4a3b3: When we specify a NETC: return, we are telling the file package to do RFCs and send the result of the operation on these network connections. Implied in this is the parallel execution of a procedure doing the matching RFCs (listening?). Therefore, the failure of one is related to the failure of the other, which may make for more complicated user programs, than would a single local command e.g. "get file", which has a single success/failure return. 10m

[True. If the caller is on the receive end of the connection it then has more information on any connection failure. But in general to move a file between two arbitrary file packages

controlled from a third place this type of rendezvous is necessary, though i would suggest using physical channels.] 10m1

4b1b and 4c1c1: Why must you use the full TEXT file name, after it is created, for subsequent file operations? Why is there no concept of OPEN file and use of a handle in place of the text file name for manipulating the file? 10n

[We primarily thought of this package for manipulating whole files, or major portions of files, but as you suggest it might be well to have file handles and open file and close file calls.] 10n1

4b1c4: Must all the files in the file list have the same access descriptor? By your definitions it appears so. 10o

[Yes, the list of files all have the same access. This could be easily changed to have the argument be a list of filename, access pairs.] 10o1

4b3b: Do we take failure return if the list lengths don't match? 10p

[Yes.] 10p1

4b4b: Why are the get unique file name and create file separate operations? Why not support an option on create new file which says generate a new name. (It seems that generated file names are useful mostly (only?) for temporary files, since users can not generally remember names of the form ..[s> etc. If this is in fact true, why use names at all? Why not have a CREATE temporary file primitive (procedure) which returns a handle on the file. When done, it can be released and it goes away, or it can be given permanent status by giving it a name.) 10q

[The Works Manager wants new names all the time and it wants unique names, it has another "virtual" file system on top of all this , so users never see the file package file names. But it is a good idea to have the create and unique operations combined (it prevents a race condition), but we may still need each as separate operations. Your suggestion for a temporary file type is interesting but i dont think it would be needed in NSW, also there are some problems with such things in terms of crash recovery.] 10q1

4b5b: I don't feel that these type of routines are file system operations. Rather, they are utility routines which operate on files, and as such belong elsewhere. 10r

[We agree, but this is the only utility routine and we didnt

want to create a new package (and document), but now i think it would have been better to do so. Any suggestions for a name and other procedures for this package ?] 10r1

4c3c3: Where is the file structure obtained from? It is not mentioned in the file descriptor. 10s

[The structure could be obtained by walking the file.] 10s1

Document: NSW Requirements on Tool Bearing Hosts 11

Most of this document (except for the interface to system calls section) has been obsoleted by our discussion of February 5. Correct? 11a

[Generally right, the change is really that the idea of a "TBH supervisor" has evaporated. There is a "dispatcher" listening on the contact socket that creates (when called) PCP processes that contain the PCP support stuff as listed in 1b2 and 1b3.] 11a1

Document: TENEX PCP Process Internal Structure (PCPTNXINT.TXT) 12

3a1: Although no ambiguity may result, it is sometimes very confusing to use the same name for different purposes, and using the same name without modifiers when describing things at the various levels of detail. For example, it is sometimes unclear when you mention "process" whether you are referring to an NSW process (i.e. multiple fork conglomeration) or a TENEX process (i.e. a single fork). 12a

[Yes; sorry. I will eliminate this source of confusion in the Version 3 documentation.] 12a1

4a1a1: Why does the processor fork contain the controlling fork environment and vice versa? Is it just to make it easy to create forks using the same map and to achieve parallel execution? 12b

[Yes, it WAS just a convenience. However, we have changed our strategy here, and are in fact NOT running the same SAV file in both CF and PFs.] 12b1

4b1a1: In this context, I assume "local process' superior" is meant in the TENEX sense, not in the NSW tree sense. 12c

[Yes.] 12c1

5b2: This seems like a terrible implementation strategy, and in some sense doesn't even fit into your model. 12d

[Yes, it's a bit of a kludge conceptually (although very clean in implementation) to get around the PF's inability to create a new inferior of the CF. I will probably take this approach only when the new inferior is to be interfaced to via inter-fork IPC, since Tenex does not similarly restrict creation of new Network connections (i.e. they can be owned by the job, not just by the fork which happens to create them). 12d1

With what model is this strategy inconsistent?] 12d2

6b2b1: Assume the object procedure we are calling is on a non-TENEX (i.e. not a 36 bit machine). We then use SYCALL and specify the object procedure arguments as a PCPB36 data structure. Will the SYCALL implementation transform these parameters to another format (e.g. PCPTXT) if that is what the channel to the object procedure's process calls for? 12e

[Yes. Actually the conversion is even lower down, in the IPC primitives.] 12e1

6b3b1: This defines only a very primitive signal routine i.e. one with no memory. I have often found it useful to define a more powerful signal routine that can handle a number of "signals" before any wait occurs. 12f

[Thanks for the suggestion.] 12f1

6b4b1 and 6b4b2: Can SYLOCK succeed twice in succession? If so, it is not really a lock. In that vein, what do you mean by "remove the most recent lock", and something being locked in a way that DOESN'T preclude its being locked immediately by the new caller? 12g

[SYLOCK can succeed twice in succession if called both times by the same fork. SYLOCK pushes each lock request onto a stack; SYUNLK is understood to apply to the topmost lock request on the stack. The locking primitives are being generalized to optionally return a "lock handle" by which the request can be subsequently referenced.] 12g1

6c1a1: Shouldn't this be a full word for an ASCIZ string pointer? 12h

[Yes.] 12h1

6c2a1: Processes, then packages, then bundles-- what next? Are bundles relevant to the TENEX implementation only? I assume not. 12i

[Bundles ARE a construct of the Tenex implementation only. Although the concept has been eliminated during implementation,

a bundle is simply a set of packages coded in the same language, and thus able to share certain run-time facilities.] 12i1

6c2a3: I guess you forgot the half word pointer to BUPRSM. 12j

[Yes.] 12j1

6c2a3: Is it implied that a single package must all be in a single bundle (i.e. in a single implementation language)? 12k

[Yes.] 12k1

6c3a1: Half word Asciz string pointer? 12l

[Sorry again.] 12l1

Global comment: If I want to write a package, what do I have to do? I assume that I must write the bundle initializer and create a bundle data structure, but how do you tell the "system" about the new package, bundle? 12m

[By appending (at compile-time, usually) the bundle descriptor address to the user descriptor (see 6c1a5).] 12m1

Document: Network Virtual Terminal Package (NSW-NVT.TXT) 13

3a4a and 3a4b: Is it the case that all of these naming constructs exist: $F filename, $F host filename,$N host socket, $N host socket filename, and that different situations require different flavors of name. 13a

[These names (prcaddrs) are confused, we keep changing our minds on these and since documents are written on different days some inconsistency creeps in. The latest thinking is: "$F<host><sp><intrahost name>" where for Tenex the intra host name is a file name of save a file that contain the pcp support code along with theuser code. The information about the subsystem or user program to run under an NVTP will be passed via CRTPRC new argument "startup-info".] 13a1

3b2: Is it the intent that this also destroys the "tool" fork that was created by OPNNVT, or merely detaches it? 13b

[Yes, the subsystem or user program is to be destroyed.] 13b1

3c2 and 5c3: On the other hand, care must now be taken at the NVTP site to ensure that two consecutive out-of-line calls of SNDNVT (each containing character information) do not cause interspersing of character input to the tool. 13c

[Correct, this is a danger.] 13c1

Document: NSW Tool Package (NTP.TXT) 14

3a1: Does every NSW PCP process have to have its own BGNNSW procedure? 14a

[No, only those which constitute tools.] 14a1

3a1: Our discussions at the Feb. meeting have led to an additional parameter indicating whether or not to trap file references to the works manager. How will this work with new tools, which will have their own BGNNSW procedure executed, not one in the encapsulator? 14b

[I can't see why the WM need inform the tool whether it requires the services of the encapsulator. It seems more reasonable to me to build that knowledge into the tool. The NVTP will enable the encapsulator for every program run through it. Every other tool can either enable the encapsulator, or interact with the WM directly, whichever is appropriate. In the Tenex implementation, we should define a bit in the USRDSC which signifies the need for the encapsulator (I currently believe that the encapsulator should reside in the CF).] 14b1

Document: NSW Process Structure (NSWSTRUC.TXT) 15

3a1: As mentioned previously, there are the network server processes waiting to receive a PCP CRTPRC data structure on the ICP negotiated network sockets. This also constitutes an independent tree. 15a

[True. See 3c3b on the next page.] 15a1

3b: Is it the case that the FE preallocates all of its processors to implement its multi-processor nature (which means for example, recompiling if it is necessary to handle more terminals), or does it do this on demand, in which case it would seem to be better modeled by a CRTPRC instead of a create another PF (provision for which I haven't seen anyway). 15b

[Good point. The FE may indeed want to allocate PFs dynamically, and we should define a new SSS by which it can do so.] 15b1

3c3a: In the TENEX implementation, for example, how does the tool builder communicate to the controlling fork its desire to be a multi-processor process (i.e. have many PFs)? 15c

[By including the number of processors desired in the USRDSC, See (PCPTNXINT -- 24792,6c1a3).] 15c1

3c3b: How does PMP's CRTPRC know whether to create a new tool or splice to an existing one? 15d

[By means of the process address, which in one case will designate the standard PCP contact socket for the host on which the "logger" (to use an old Network term) will be listening, and in the other a tool-assigned contact socket on which the tool itself will be listening. In either case, however, CRTPRC performs precisely the same functions.] 15d1

3d1a: How does the FE address the WM in the CRTPRC splicing procedure- by a special network contact socket just for this function? Also, it seems to me to be a poor idea to use CRTPRC to splice trees. This CRTPRC has none of the creation overtones and privileges that the normal create has. 15e

[Yes; see above. You could be right, although accomplishing both operations with the same primitive has the advantage that the "creator" needn't know whether it's addressing a pre-existent process or not. Furthermore, superior processes have at present very few (if any) "priviledges" with respect to their direct inferiors, except of course the right to delete them, a right which the process still has in the case of a splice (where it means "disconnect", rather than "destroy").] 15e1

3da1: Instead of having the problems associated with multiplexing requests from the FE to the WM, another approach would have a separate WM process to talk to a separate FE process by just having their ports connected by a channel. This approach seems to be cleaner and less perilous with respect to interference from other users. 15f

[I'm not sure what "problems" would be eliminated. It would eliminate the need for a USERID as an argument to each WM procedure. It would also, however, require a pair of Network connections for each user.] 15f1

3e2: How does "a list of WM process addresses" fit into the tree structure and PCP itself? This also doesn't account for the most responsive WM, the least loaded WM, etc, for splicing. 15g

[I assume that the WMs will be spliced together, and that the FE will attempt to splice to successive WMs in the list until it finds one up and thereby gains access to the distributed WM. PCP as currently being implemented, permits this kind of complicated inter-connection of processes. Since this is not a

first-year NSW issue, no thought that I know of has gone into WM selection algorithms.] 15g1

3e3: How simple is terminal linking in a multi-WM implementation if two users get different WMs? 15h

[Only one more level of indirection than required by the scenario in 3e1 is required: one WM must first introduce its FE to the other WM. 15h1

3f2: But what does a works manager do if there is no FE (i.e. user)? In the sense implied, the tool (in the form of the tool dispatcher and running tools e.g. mailer) can also exist without the WM or FE. 15i

[Tend to so-called "batch" jobs. 15i1

Yes, as discussed elsewhere.] 15i2

3f3: Involuntarily severed can have several subclasses: temporarily, permanently, after timeout, etc. Some of these provide for continuing (after resynchronization) rather than restarting. 15j

[This entire area will be dealt with in second-year NSW.] 15j1

JEW JBP 12-MAR-75 12:05 25556

Reply to Section II of BBN PCP Critique

(J25556) 12-MAR-75 12:05;;; Title: Author(s): James E. (Jim) White, Jonathan B. Postel/JEW JBP; Distribution: /NSW([INFO-ONLY]) SRI-ARC([INFO-ONLY]) ; Sub-Collections: SRI-ARC NSW; Clerk: JEW;
Origin: < JWHITE, SEC2REPLY.NLS;3, >, 12-MAR-75 12:01 JEW ;;;;####;

DVN 12-MAR-75 18:48 25557

More on Mysterious DEX Document

Responds to <ijournal,32013,>

DVN 12-MAR-75 18:48 25557

More on Mysterious DEX Document

In 1973 I think it was we wrote the original to DEX User Guide and Primer. The Primer has been updated from time to time to reflect changes in the system and new terminals, but the User Guide has not because of lack of resources. I don't know how the architects came to get it. I suggest you throw it away. 1

DVN 12-MAR-75 18:48 25557

More on Mysterious DEX Document

(J25557) 12-MAR-75 18:48;;;; Title: Author(s): Dirk H. Van Nouhuys/DVN; Distribution: /JOAN([ACTION] dpcs and dirt notebooks please) JMB([ACTION]) JML([INFO-ONLY]) SGR([INFO-ONLY]) POOH([INFO-ONLY]) JR([INFO-ONLY]) KIRK([INFO-ONLY]) JMB([INFO-ONLY]) HGL([INFO-ONLY]) ; Sub-Collections: DPCS DIRT SRI-ARC; Clerk: DVN;

DCE 13-MAR-75 00:51 25558

Announcement -- all-ARC meeting next Monday

There will be an all-ARC meeting next Monday morning, 17 Mar, at 1000. Dick and Jim will fill us in on the ARPA-IPTO contractors' meeting this week (Wed-Fri), and there is much other food for thought that is ready to put on the table. Doug 1

DCE 13-MAR-75 00:51 25558

Announcement -- all-ARC meeting next Monday

(J25558) 13-MAR-75 00:51;;;; Title: Author(s): Douglas C. Engelbart/DCE; Distribution: /SRI-ARC([ACTION]) ; Sub-Collections: SRI-ARC; Clerk: DCE;

RLL 13-MAR-75 00:56 25559

bug: at bbnb, using message,subsys twice

tried to use message subsystem twice. on the second try move message command failed because [message],work file was not openable. ttried quiting to tenex and resetting. same results the first time i tried using message. ttried expunge, then reset it worked. soonds as if the work file is not expuunged after quiting message. alternattely, maybe the file should not be deleted until one leaves nls insted of when one leaves mesage. in any case the current situation at bbnb is a bug. (feed; don't tell me you received message, tell me when it is fixed thanks).

RLL 13-MAR-75 00:56 25559

bug: at bbnb, using message,subsys twice

(J25559) 13-MAR-75 00:56;;;; Title: Author(s): Robert N. Lieberman/RLL; Distribution: /FEED([ACTION]) ; Sub-Collections: SRI-ARC; Clerk: RLL;

RLL 13-MAR-75 01:01 25560

bug: undelivered journal mail

i sent the group ident arc-log a mail item, namely
<ijournal,25524,nls,> i received the author copy but arc-log did not receive the item. it is not important so don't fix it. this is to let you know that it happened (about 6-mar-75). repeat don't fix it. you got more impootant things to do. rob 1

RLL 13-MAR-75 01:01 25560

bug: undelivered journal mail

(J25560) 13-MAR-75 01:01;;;; Title: Author(s): Robert N. Lieberman/RLL; Distribution: /FEED([ACTION]) JDH([ACTION]) ; Sub-Collections: SRI-ARC; Clerk: RLL;

LPD 13-MAR-75 11:34 25561

Thoughts from Peter Deutsch about the Elf and Big Characters

Forwarded from a sendmessage by DVN.

LPD 13-MAR-75 11:34 25561

Thoughts from Peter Deutsch about the Elf and Big Characters

622 ;
Distribution: VANNOUHUYS
Received at: 13-MAR-75 02:45:09-EDT 1

Mail from PARC-MAXC rcvd at 13-MAR-75 0245-EDT
Date: 12 MAR 1975 2345-PDT
From: DEUTSCH at PARC-MAXC
Subject: Elf buffering
To: vannouhuys at BBNB 1a

Did you ever check out whether the Elf was doing anything sensible about packeting up big characters? Even if BBN is giving you plenty of message allocation so that the stuff doesn't get held
up by Host protocol, there are still all those damn RFNMs which mean
that you'd still be experiencing the net round-trip delay for each and every character (or at least more than once depending on the relative speeds of the network and the line processor). 1b

LPD 13-MAR-75 11:34 25561

Thoughts from Peter Deutsch about the Elf and Big Characters

(J25561) 13-MAR-75 11:34;;; Title: Author(s): L. Peter Deutsch/LPD;
Distribution: /CHI([INFO-ONLY]) JBP([INFO-ONLY]) JLE([
INFO-ONLY]) RWW([INFO-ONLY]) KEV([INFO-ONLY]) POOH([
INFO-ONLY]) ; Sub-Collections: NIC; Clerk: DVN;

DVN 13-MAR-75 14:28 25562

Does Steve Warshall Have and Ident?

Susan says Warshall of MCA should have an ident, but the system says it knows no Warshall, Can you help me? 1

DVN 13-MAR-75 14:28 25562

Does Steve Warshall Have and Ident?

(J25562) 13-MAR-75 14:28;;;; Title: Author(s): Dirk H. Van Nouhuys/DVN; Distribution: /MLK([ACTION]) SGR([INFO-ONLY]) ; Sub-Collections: SRI-ARC; Clerk: DVN;

JMB 13-MAR-75 18:00 25563

The Modify subsystem

Can I have the Modify Subsystem as one of my Included subsystems at o-1 [Useroptions command Include Subsystem Modify] so that I don't have to do Goto Programs, and Load Program Modify, before I can Goto Modify?--like I have done for Message subsystem (I can just Goto Message whenever I want). If so, how do I do it? 1

JMB 13-MAR-75 18:00 25563

The Modify subsystem

(J25563) 13-MAR-75 18:00;;;; Title: Author(s): Jeanne M. Beck/JMB;
Distribution: /FEED([ACTION]) ; Sub-Collections: SRI-ARC; Clerk:
JMB;

JMB 13-MAR-75 18:01 25564

Training

Susan, I will continue to send out formal (sort of) reports of classes & scheduled tutorials until the users file is ready and we catch up with all the info for past courses (including mine). Let me know when. 1

JMB 13-MAR-75 18:01 25564

Training

(J25564) 13-MAR-75 18:01;;;; Title: Author(s): Jeanne M. Beck/JMB;
Distribution: /SGR([ACTION]) ; Sub-Collections: SRI-ARC; Clerk:
JMB;

JBP 13-MAR-75 19:36 25565

Performance Measurements

Introduction 1

This is a collection of rough ideas on how to determine the performance bottlenecks in the combined "computer and communication system" of ARPANET Hosts and IMPS. 1a

From the users point of view there is very little difference in poor performance due to the serving Host or due to the network of IMPS, or even due to her terminal. 1b

To be able to improve the situation the operators of the combined system must be able to get their hands on measures of performance at several points in the combined system. 1c

The user must also have a means of checking to see if the system is lightly of heavily loaded, if the system is running or limping. That is the user must be provided with measures that she can correlate with actual performance (e.g. the Tenex load average). 1d

User Statistics 2

Network Load Average 2a

There could be available in one or two hosts or via RSEXEC a network load average statistics updated say every 5 minutes. This could represent the number of packets in the network normalized by dividing by the number of imps (number of lines), or perhaps the number of packets on imp output queues (again appropriately normalized). i expect that the NCC and NMC would be able to suggest an approritate measure. 2a1

Point to Point Load Average 2b

There could be a similar number for the shortest route between each pair of imps in the network. 2b1

Process Network Output Delay 2c

A measure of the time a process spends on output to the network system calls. For example in Tenex the time from the beginning of a bout or sout to the network and the return form such a call. 2c1

NCP Output Delay 2d

A measure of the elapsed time between the process handing the message to the NCP and the time it actually goes out the door. 2d1

Internal Measures 3

Performance Measurements

There must be a capability to measure every step of the path of data flow between a user and her service program. 3a

The things we need to know are: 3a1

delays 3a1a

throughput 3a1b

Max possible 3a1b1

measured 3a1b2

buffer sizes 3a1c

message sizes 3a1d

The steps in the path that need to be measured are: 3a2

users terminal 3a2a

line processor 3a2b

connecting cable 3a2c

TIP or ELF or ? 3a2d

connecting cable and interfaces 3a2e

source IMP 3a2f

intermediate IMPs and connecting circuits 3a2g

destination IMP 3a2h

connecting cable and interfaces 3a2i

NCP 3a2j

scheduler 3a2k

system call interface 3a2l

Process 3a2m

Design Philosophy 4

We need to understand the policies used to implement the data

sending and allocation sending in both TENEX and ELF, ant the role that buffers play in these strategies. 4a

JBP 13-MAR-75 19:36 25565

Performance Measurements

(J25565) 13-MAR-75 19:36;;; Title: Author(s): Jonathan B. Postel/JBP; Distribution: /NPG([INFO-ONLY]) ; Sub-Collections: SRI-ARC NPG; Clerk: JBP; Origin: < POSTEL, PERFORMANCE-MEASUREMENTS.NLS;3, >, 13-MAR-75 19:30 JBP ;;;;####;

DVN 14-MAR-75 14:39 25566

Commands Branches Doing Better at BB&N

Following 25554 I am glad to report that commands branches have run without undue confusion on the screen and lineprocessor panel most of the time yesterday and today. 1

DVN 14-MAR-75 14:39 25566

Commands Branches Doing Better at BB&N

(J25566) 14-MAR-75 14:39;;;; Title: Author(s): Dirk H. Van Nouhuys/DVN; Distribution: /FEED([ACTION]) JMB([INFO-ONLY]) ; Sub-Collections: SRI-ARC; Clerk: DVN;

DVN 14-MAR-75 15:09 25567

Still No Journal Indeces at BBN

We've go lots of pages, what's the problem. Grumble Mumble. 1

DVN 14-MAR-75 15:09 25567

Still No Journal Indeces at BBN

(J25567) 14-MAR-75 15:09;;;; Title: Author(s): Dirk H. Van Nouhuys/DVN; Distribution: /FEED([ACTION]) JCN([INFO-ONLY]) ; Sub-Collections: SRI-ARC; Clerk: DVN;

DVN 14-MAR-75 20:16 25568

Fourth NSW Work Breakdown and Schedule

Here is a new cut at allocating and scheduling work on NSW documentation following <24848,> <25305,> and <25395,>. (25395,2) gives more information about the divsion we are assuming among tools and the file structure. We face the unfortunate necessity of documenting tools before they are frozen. Under this circustance "DUE" means a piece of docuemtnation that is accurate, complete, and comprehensilbe to its appropriate audience for that tool or feature at that date. As discussed in (25305,) there is more work than the perople alocated can do (all of POOH, half of KIRK, and 3/4 of DVN). We are counting on help from MCA in certain areas (<25084,>) and from new person(s) whom we hope to conclude hiring soon. 1

Subject PERSON RESPONSIBLE <link to data base file> 2

DCPS DVN<arcdocumentation,dpcshelp,> 2a

DRAFT Scenario formatting a file for COM DUE:April 11 2a1

Scenario on formatting a file for COM DUE:April 25 2a2

DRAFT Data Base File DUE:March 28 2a3

Data Base File DUE: April 28 2a4

DRAFT Introduction DUE: April 4 2a5

Introduction DUE: April 18 2a6

Letter KIRK 2b

Scenario on writing, formatting, and sending a letter to go by U.S. mail. 2b1

DRAFT DUE April 1 2b1a

DUE April 15 2b1b

Sendmail KIRK<arcdocumentation,sendmail,> 2c

DRAFT Scenario on sending mail to a person in/not in the IDENT file DUE:April 11 2c1

Scenario on sending mail to a person in/not in the IDENT file DUE:April 25 2c2

DRAFT Data Base File DUE:March 28 2c3

Data Base File DUE: April 28 2c4

DRAFT Introduction DUE: April 4 2c5

Introduction DUE: April 18 2c6

Programs KIRK<arcdocumentation,programs,> 2d

DRAFT Scenario on for [writing and] compiling Cobol programs DUE:April 18 2d1

Scenario on for [writing and] compiling Cobol programs DUE:April 25 2d2

DRAFT Scenario on for [writing and] compiling L-10 programs DUE:April 25 2d3

scenario on for [writing and] compiling L-10 programs DUE:May 2 2d4

DRAFT Data Base File DUE:April 11 2d5

Data Base File DUE: April 28 2d6

DRAFT Introduction to Cobol Programming in NSW (let's try tp get some one who has writen a Cobol program involved in this as a co-author) DUE: April 4 2d7

Introductionto Cobol Programming in NSW DUE: April 28 2d8

useroptions POOH<arcdocumentation,useroptions,> 2e

DRAFT Scenario [Maybe a Cobol programmer at Gunter AFB setting himself up to work,] DUE:April 25 2e1

Scenario [Maybe a Cobol programmer at Gunter AFB setting himself up to work,] DUE:May 9 2e2

DRAFT Data Base File DUE:April May 2 2e3

Data Base File DUE: May 26 2e4

DRAFT Introduction to NSW Useroptions DUE: April 28 2e5

Introductiono NSW useroptions DUE: May 11 2e6

Caluculator POOH<arcdocumentation,calculator,> 2f

DRAFT Scenario [Maybe on doing part of your income tax,] DUE:April 15 2f1

Scenario [Maybe on doing part of your income tax,] DUE:April 25 2f2

DRAFT Data Base File DUE:March 28 2f3

Data Base File DUE: April 28 2f4

DRAFT Introduction DUE: April 4 2f5

Introduction DUE: April 18 2f6

Graphics POOH<arcdocumentation,graphics-terms,> 2g

DRAFT Scenario on Making a Flow Chart DUE:April 25 2g1

Scenario on Making a Flow Chart DUE: May 2 2g2

DRAFT Data Base File DUE:April 4 2g3

Data Base File DUE: April 11 2g4

DRAFT Introduction DUE: April 18 2g5

Introduction DUE: April 25 2g6

NLS Editor <arcdocumentation,Editor,> 2h

DRAFT Scenario showing basic editing and integrated with the tutorials being written by Applications. DUE:May 2 2h1

Scenario on showing basic editing and integrated with the tutorials being written by Applications DUE: May 16 2h2

DRAFT Data Base File DUE:April 11 2h3

Data Base File DUE: April 18 2h4

DRAFT Introductionfor NSW audience DUE: April 25 2h5

Introduction DUE for NSW audience: May 2 2h6

NLS Environment File POOH KIRK <arcdocumentation,nlsum,> 2i

This file contains information relevant to all tools that are based on NLS. When the qurey fails to find an item in the data base for a tool based on NLS, by default it goes to this file next It does not contain links to nodes in the data bases attached to tools.. 2i1

DRAFT Data Base File DUE:May 2 2i2

Data Base File DUE: 18pril 18 2i3

Front End/Works Manager DvN<arcdocumentation,fewm,> 2j

A general introduction to NSW which will take the intelligent user to a point where she/he can continue to learn on her/his own. 2j1

DRAFT DUE April 11 2j1a

Final DUE: April25 2j1b

Scenario loging into NSW and getting a tool. 2j2

Darft DUE April 18 2j2a

Final DUE May 2 2j2b

DRAFT Data Base File DUE: April 7 2j3

Data Base File DUE: April 21 2j4

Interface to Secretarial Functions 2k

DVN, an introduction written for people who will be performing secretarial functions. This document will reach accross several tools and include those elements that seem appropriate for the tasks assigned to the user. 2k1

DRAFT DUE *April 15 2k1a

DUE *April 25 2k1b

KIRK, a tutorial for secretarial tasks that will include inserting and modifying text, various means of printing text so it can be viewed, reading and sending mail. 2k2

DRAFT DUE *April 20 2k2a

DUE *April 29 2k2b

Sequential I/O 2l

POOH, a tutorial on how to use Sequential I/O. 2l1

DRAFT DUE *May 20 2l1a

Fourth NSW Work Breakdown and Schedule

DUE *May 29	211b
NSW Cue Card POOH	2m
Draft DUE May 15	2m1
DUE June 1	2m2

Fourth NSW Work Breakdown and Schedule

(J25568) 14-MAR-75 20:16;;; Title: Author(s): Dirk H. Van Nouhuys/DVN; Distribution: /DIRT([INFO-ONLY]) KS([INFO-ONLY]) LAC([INFO-ONLY]) WEC([INFO-ONLY]) ; Sub-Collections: DIRT NSW SRI-ARC; Clerk: DVN; Origin: < VANNOUHUYS, FOURTH.NLS;5, >, 14-MAR-75 20:08 DVN ;;;;####;

CHI 17-MAR-75 01:24 25569

Findings to date in Office-1 and BBNB performance problems investigation

Sent via SNDMSG to distribution list in <engelbart>slist.

CHI 17-MAR-75 01:24 25569

Findings to date in Office-1 and BBNB performance problems investigation

This memo (Journal # 25569) summarizes our findings to date with respect to Office-1 and BBNB performance problems as seen by NET users. 1

OPEN LINE ON TYMESHARE TIP 1a

While at the TYMESHARE TIP, Martin Hardy found that, indeed, the BRK light (on the INBUF lights on the MLC panal) was coming on for several minutes at a time and that when on the rate at which the line-number register lights changed was very drastically reduced. Martin, working with Charles Kelley and Dave Walden of BBN, found that BBN's earlier (28 FEB) patch to the TIP to prevent this from happening was not catching this case because it was coming from 208 modems (there are 2 connected to that TIP) which do not send continuous break characters, but rather repeated break characters (i.e., with a pause in between each break). The modems sent the break characters whenever the terminals connected to them were powered off. Until we find a better solution we have large signs on these teerminals instructing users not to turn power off, ever! 1a1

From Dave Walden's earlier memo, we conclude that this open line(s) will have a sever effect on TIP performance, but also suspect that it would degrade the IMP to some extent also. The ARPA and Mitre TIP's should be inspected to see if they are also suffering from this or similar problems. 1a2

OPEN LINE ON OFFICE-1 PDP-10 1b

Martin was also able to find an open line (another modem to Bell Canada) on the PDP-10's Data Line Scanner. Martin had earlier requested that TYMESHARE install inhibiters on all Data Line Scanner lines to prevent this sort of thing, but apparently this was not done to all lines. This problem has been eleviated and was undoubtedly placing a load on the PDP-10, although we don't know how much. 1b1

NCP RUNNING OUT OF INPUT BUFFERS FOR IMP 1c

Ken Victor patched some counters into the TENEX NCP at Office-1 to see if it was ever running out of buffers for IMP input and thus shutting the IMP input interface off. It was, as much as 50% of the time during some periods. Ken talked to BBN and found that they new this was a problem and had fixed it in release 1.33. To fix the problem in Office-1 without moving to 1.33 will take about two days of Ken Victor and Don Andrews time. My recommendation is that this be done but also that

Findings to date in Office-1 and BBNB performance problems investigation

Office-1 move to 1,33 as soon is as reasonable since it contains many improvements in network handling (Office-1 staff must evaluate other changes to 1,33 also, however, since not all may be desireable from their viewpoint). 1c1

Ken has also developed a patch to measure how much time it takes for the NCP to output messages to the IMP. However, since the 1 millisecond clocks were not working on Office-1 and BBNB, we could not get very accurate data here (since we were interested in the distribution of times not just a simple average time). 1c2

IMP-READY-FOR-NEXT-HOST-BIT light on IMP interface off much of the time 1d

Martin observed that this light flashed off for significant periods of time (up to about 1/2 second at times. This means the IMP has the output from the PDP-10 to the IMP turned off quite alot. In contrast, the associated light HOST-READY-FOR-NEXT-IMP-BIT was on most of the time. This may be explained by other of our findings outlined here but certainly warrents careful consideration by BBN NETWORK people. 1d1

MESSAGES VS. PACKETS ACROSS LONG NETWORK PATHS 1e

(I hope I am wrong on this one since the consequences are potentially very bad and will require low-level protocol and network topology changes.) When doing significant data movement through the net, it is desireable to send large (greater than 1000 bits) messages through the net. This is not only true of File Transfer but also for driving high speed terminals, printers, and the like. My understanding of what happens in the case of a big message from host HS with IMP IS being sent to host HR with IMP IR is as follows: 1e1

On reading the 1001'st bit from the host HS, IMP IS notices that it is a multi-packet message. 1e1a

If it already has an available 8-packet buffer reserved in IR, then it can continue to read bits from the HS host and send the message as several packets in parrallel. 1e1b

If it does not have such a buffer reserved, it blockes the HS host from transmitting any more bits to the IMP IS and sends an IMP-message to IR, asking it to reserve an 8-packet buffer. On ackanowledgement of this request the IMP IS unblocks the host-IMP interface and finishes processing the message. 1e1c

Please note that the host HS was unable to send anything to the IMP IS while this allocate process took place. Our recent measurements of the NET and Dave Waldens recently distributed measrements indicate that host-messages (packets) take about 50 ms to pass through each node in the path. Now, presumable IMP-messages traverse the net more quickly, lets say in half that time, or 25 ms per node (I hope they are faster). For long paths (like from Office-1 to ARPA TIP or BBNB to ARC ELF) this allocate can take a long time (650 ms for Office-1 ARPA TIP and 550 for BBNB ARC ELF)!!! 1e1c1

After sending this long message, if host HS does not send host HR another multi-packet message in 1/8 (right now it is patched to be 1/4) second, then IMP IS sends a de-allocate message to the IMP IR to free the buffer for other uses. 1e1d

Please note that for long paths like the ones discussed above, the delay to get a RFNM back to host HS so that it can send another (possibly) multi-packet message is on the order of 1/2 second!!! This means that there must be more than one process in HS sending multi-packet messages to HR to keep the buffer allocated. Now, I hope that my reasoning has some holes in it, because otherwise, it seems that file transfers between distant hosts or driving highspeed terminals or printers from distant hosts will cause serious network problems for host HS (since the IMP will be turned off alot) and for the user since there will be long pauses between messages from the host HS (because each message must wait for a round-trip IMP-IMP allocate request before it can be sent because the host HS will not send another message until the RFNM returns for the outstanding message which will take so long that the imps will de-allocate the multi-packet buffer in IR)!! 1e1d1

Note, when an IMP IS sends a single packet message to an IMP IR it must also reserve space for it by means of an allocate request. However, in this case the allocate and message are sent togeter so that a round-trip delay for the allocate is avoided. 1e1e

Given that the above is an accurate picture of the network behavior, it seems some changes are in order! For example: 1e1f

The host HS should be able to tell IMP IS that it wants to send a long message to host HR (perhaps it should always tell the IMP that it wants to send a message of

length L to another host before doing it). The IMP IS should respond to HS when it can send the message. This allows the IMP to get things set up without blocking the host-IMP interface while it is doing so. 1e1f1

The host HS should be able to have more than one outstanding message being sent to host HR on any given connection. There are now four bits avialable in the message id that could be used as sequence numbers so that messages could be retransmitted if necessary and so that HR would keep them in the right order. 1e1f2

Perhaps the length of the timeout period for releasing 8-packet buffer reservations in IR should be lengthened, although this does tie up lots of buffer space and is not very desireable in an overall sense. 1e1f3

The net topology should be changed so that there is a reasonable number (like max = 5) imps between any two hosts! From looking at a topo map of the net, it appears that even a few five mile long phone lines would help quite a bit. 1e1f4

We are now trying to imperically determine what the best short-term strategy is for us to use in driving display terminals. We are varying the number of packets our ELF will allow being sent to its terminals and measuring the time TTYTST takes to fill the screen. So far, we seem to be clearly winning to send single packets as opposed to sending 2-packet messages. We will try three, four and so on next week, although our current understanding leads us to believe that sending single-packet messages and waiting for the RFNM between messages is the best we can do given the current situation. More on this later. 1e2

1 MILLISECOND CLOCKS BROKEN AT OFFICE-1 AND BBNB PDP-10'S 1f

While Dan Lynch from sri-ai, who was glad to help us because his people have been having so much trouble using the TYMESHARE TIP, was running his program (see his recent memo) he discovered that it would run on all tenex's tried except Office-1 and BBNB. upon further investigation, it was learned that the 1 ms clock was not running in these two machines. Ray Tomlinson at BBBN does not seem to think this would have much effect on system performance. However, we find it a little hard to beleive that there are not some scheduler events, page-management, etc code that need finer than 50 ms resolution in timing things (the time is updated every 3 ticks of the 60

cycle 16 2/3 ms clock). We have requested that these clocks be fixed in both systems. In the TYMESHARE system, the clock was apparently never installed due to some communication breakdowns. We will try to see if there is any effect on performance next week. In any event we need finer clock resolution to do measurements of NCP performance. 1f1

ERRATIC PAGE TRANSFER TIMES FOR OFFICE-1 DISK 1g

Don Andrews found that disk transfers at the Office-1 machine varied from 40 ms to 600 ms. We are running diagnostics on the disks and will try to determine if disk driver is not overlapping seeks with transfers or something like that. The 1 ms clock being off should not have screwed up Don's measurements much but could have some effect on page transfers. 1g1

We checked the way Office-1's interrupts are set up with BBN. They said Office-1's interrupts were essentially the same as their TENEX's and did not think there was a problem there. We also recalled that the mag tape drivers used to hog our system until we patched it to low-queue itself. We must check that this patch is still in Office-1, since its absence might mean that using tapes during the day may be interfering with the disk (same priority level) or the IMP (by somehow blocking the I/O buss). 1g2

TIMING OF CHARACTER ECHOING 1h

Using an electronic millisecond timer, we timed how long various computers took to echo characters. Details of this measurement will be forthcoming in a separate memo. What was so interesting about this test was that at times when Office-1 was echoing characters to terminals accessing it through a TIP or ELF in 1 to 4 SECONDS it was echoing to local terminals (not using the net) in 200 ms! Most of the tests involved use of the EXEC's ; Comment Command which echos immediate and deferred. We compared this with typing a literal to TNLS and DNLS. TNLS was identical to the EXEC's ; (which was expected since it also uses immediate and deferred echoing with process-activation on punctuation and control characters). DNLS was somewhat slower because it activates on every character to watch for wrap-around on the display. We could be more risky in DNLS and use TNLS stategy until we get close to the end of the line and then wake up on every character. This could cause some problems in the rare case where the user types a line without spaces or punctuation. 1h1

TENEX PROGRAM TO IMP TO TENEX PROGRAM TEST 1i

Findings to date in Office-1 and BBNB performance problems investigation

Dan Lynch recently sent out a memo on this test. Although the data resolution was different for Office-1 and BBNB, it is interesting to note that it seems to take considerably longer to send a byte of data to the local IMP and back on Office-1 than on any other TENEX measured. Why?? Note that Dan's tests were run in the middle of the night. 1i1

NLS OUTPUT 1j

Just to make sure NLS was not just being slow to generate its output, we patched in code to see how often it got blocked because its 180 character output buffer was full. For a typical screen repaint, it got blocked four times. We are not yet clear on the wake-up condition for a process that is so blocked. We hope that it is an intellegent condition (not just that there is now room for a character or two in the buffer or there is only a small number of characters left in the buffer), but know that it is not based on the terminal's output speed. We will check this further. 1j1

CHI 17-MAR-75 01:24 25569

Findings to date in Office-1 and BBNB performance problems investigation

(J25569) 17-MAR-75 01:24;;; Title: Author(s): Charles H. Irby/CHI; Distribution: /SRI-ARC([INFO-ONLY]) WEC([INFO-ONLY]) DLS([INFO-ONLY]) ; Sub-Collections: SRI-ARC; Clerk: CHI; Origin: < IRBY, FINDINGS.NLS;3, >, 17-MAR-75 01:20 CHI ;;;;####;

Journal Test

This is a sentence. This is another sentence to test automatic editing. this is not another sentence. Too many spaces are here. 1

A blue dog ratted when the sun flew over the hedgerow. Tiny Fungi Eat Gas. 2

Putative Martian organisms which flourished in the aqeuous epoch may now be in cryptobiotic repose. 3

test 3a

Following Matsubara, temperature plays the role of an imagenary time. 4

When anarchy triumphs and I am king, I will outlaw hierarchy. 5

6

Journal Test

(J25570) 14-MAR-75 19:23;;; Title: Author(s): Dirk H. Van Nouhuys/DVN; Distribution: /SLJ([INFO-ONLY]) ; Sub-Collections: SRI-ARC; Clerk: DVN; Origin: < VANNOUHUYS, ONETEST,NLS;12, >, 14-MAR-75 19:21 DVN ;;;;(; EXTERNAL LINKS: <vannouhuys,dvn,>####;

DVN 14-MAR-75 19:35 25571

Invitation To DIRT Distribution

A journal distribution group and subcollection called DIRT (Documentation Instigation and Review Team) has existed for a long time at ARC. These days it consists of the people in Development writing user documentation and other people who are interested in keeping up on what's happening in this area. To it are addressed weekly reports e. g. (25403,) planning documents e. g. (25395,) and sometimes discourses and polemics e. g. (22714,). At present the documentation work in the Development group is of interest to NSW with a few minor exceptions. Therefore I suggest to the following people might want to join the DIRT distribution: WEC LAC KS SW. We will add them to the list only if they tell us positively they want to. We welcome suggestions for others invovled in NSW who should receive this information.

DVN 14-MAR-75 19:35 25571

Invitation To DIRT Distribution

(J25571) 14-MAR-75 19:35;;;; Title: Author(s): Dirk H. Van Nouhuys/DVN; Distribution: /WEC([ACTION]) LAC([ACTION]) KS([ACTION]) JOAN([ACTION] DIRT notebook please) DIRT([INFO-ONLY]) MCG([INFO-ONLY]) LEG([INFO-ONLY]) ; Sub-Collections: SRI-ARC DIRT; Clerk: DVN;

JLE KIRK 14-MAR-75 23:45 25572

New incantation for copying files to ARC ELF printer

Instead of NET:,063-10 ... type
NET:,202-10, <CR> ASC <CR><CR>
<CR>=Carriage RETURN

1

JLE KIRK 14-MAR-75 23:45 25572

New incantation for copying files to ARC ELF printer

(J25572) 14-MAR-75 23:45;;;; Title: Author(s): Joseph L. Ehardt, Kirk E. Kelley/JLE KIRK; Distribution: /SRI-ARC([INFO-ONLY]) ; Sub-Collections: NIC SRI-ARC; Clerk: KIRK;

nsw directories

11-MAR-75 14:48:54-EDT,2800;000000000000
Mail from BBN-TENEXA rcvd at 11-MAR-75 1448-EDT
Date: 11 MAR 1975 1448-EDT
From: SCHANTZ at BBN-TENEXA
Subject: NSW Directories
To: JWhite at BBNB, Postel at BBNB, Warshall at BBNB,
To: Millstein at BBNB, Carlson at ISI
cc: Schantz

This note is concerned with the question of directories on the tool bearing host (in particular TENEX) in which to run user tools. At our last meeting we tossed around the notion that each tool invocation would be run in a separate directory, selected from among some number of directories allocated to the NSW.
After speaking with Jim White yesterday, it appears that not only is this the selected means of avoiding tool conflict, but also that the works manager has agreed to manage these directories. That is, the WM will pass directory name, password, account strings along with the process creation request for creating the tool.
All this is fine with me as long as we understand the consequences of these decisions.
1) Having a fixed number of single user directories means that any "simultaneous" user tool requests above this number must be rejected or delayed until a directory frees up.
[On the other hand, any other approach will also have some upper bound on the number of tools it can manage at once, but this number may be greater than the number of directories the NSW wishes to maintain. Along these lines, having more than one TENEX system participate as a TBH means that we can accommmdate more simultaneous tool invocations, even with the separate directory approach.] Having separate directories per tool use guarantees the non-interference of tools, and makes the TBH tool monitoring code simpler, since it does not have to duplicate the services already done by the system. 2

2) In order for the works manager to utilize these directories (and select their use) it must supply proper passwords. Since these directories are not associated with specific users, the passwords must be available to or generated by the WM code, without benefit of user interaction to supply the password or an encription key to decode it. I think this may present a security loophole. 3

3) After a tool is finished, who is responsible for cleaning up any files which may be left in the selected directory. [I don't mean permanent NSW files which are presumably moved elsewhere. Rather temporary files, or files which were moved into the local directory so they could be read but were not changed.]
I assume that the directory will have to be initialized before it is reused by another user tool. 4

nsw directories

Please advise me as whether these decisions have been made and how (or if) the things I mentioned are a problem, and possible solutions. The TBH dispatcher requires resolution of these issues.

--Rick

JBP 15-MAR-75 03:03 25573

nsw directories

(J25573) 15-MAR-75 03:03;;; Title: Author(s): Jonathan B. Postel/JBP; Distribution: /JEW([INFO-ONLY]) DSM([INFO-ONLY]) CHI([INFO-ONLY]) ; Sub-Collections: SRI-ARC; Clerk: JBP;

JBP 16-MAR-75 00:45 25574

Problem in Copy Sequential

by the way the copy works in spite of the error!

JBP 16-MAR-75 00:45 25574

Problem in Copy Sequential

Everytime i do a Copy Sequential i get the error message "NLS system error: illegal file # passed to FLNTADR" this is in TNLS at BBNB,
--jon. 1

JBP 16-MAR-75 00:45 25574

Problem in Copy Sequential

(J25574) 16-MAR-75 00:45;;;; Title: Author(s): Jonathan B. Postel/JBP; Distribution: /FEED([ACTION]) BUGS([ACTION]) ; Sub-Collections: SRI-ARC BUGS; Clerk: JBP;

JBP 16-MAR-75 00:50 25575

Network Performance Effort by Dave Walden

15-MAR-75 07:49:18-EDT,2684;000000000000
Mail from BBN-TENEX rcvd at 15-MAR-75 0749-EDT
Date: 15 MAR 1975 0743-EDT
From: WALDEN at BBN-TENEX
Subject: ANOTHER DISCOVERY IN THE OFFICE-1/NET EFFORT
To: network performance group:, engelbart at OFFICE-1,
To: norton at OFFICE-1, watson at OFFICE-1, stone at OFFICE-1, To:
heart, walden, mckenzie, sutherland, burchfiel, strollo, To:
mclindon at ISI, carlson at ISI, blue at ISI, stubbs at ISI, To:
licklider at ISI, russell at ISI, walker at ISI 1

MARTIN HARDY HAS PROMISED TO WRITE THIS UP IN GREATER DETAIL, BUT I WON'T WAIT FOR HIM, BEFORE I GIVE YOU ALL A HINT WHAT HAS BEEN DISCOVERED: 2

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

INCIDENTALLY, WE HAD OUR FIELD ENGINEER AT TYMSHARE LOOKING AT THE TIP LIGHTS WHILE THE SRI GUYS WERE RUNNING THEIR EXPERIMENTS AND THIS APPEARS TO HAVE BEEN HELPFUL, ALTHOUGH ANY OTHER PAIR OF VISUAL INTEGRATORS (I.E., EYES) COULD HAVE DONE THE JOB AND I BELIEVE MARTIN HIMSELF WENT DOWN TO CUPERTINO THE DAY BEFORE YESTERDAY TO STUDY THE TIP LIGHTS. THE LESSON TO REMEMBER IS THAT IT IS OFTEN TIMES USEFUL IF SOMEONE CAN JUST LOOK AT THE IMP AND TIP LIGHTS WHEN THERE IS TROUBLE. CLUES ARE SOMETIMES OBVIOUS FROM LOOKING AT THE LIGHTS. 2b

REGARDS, DAVE 3

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

FOREVER, AND IT IS ONLY INCREASED PRESSURE ON THE EXISTING MECHANISMS WHICH IS NOW DEMONSTRATING THEM TO BE SUB-OPTIMAL. -------- 4

JBP 16-MAR-75 00:50 25575

Network Performance Effort by Dave Walden

(J25575) 16-MAR-75 00:50;;; Title: Author(s): Jonathan B. Postel/JBP; Distribution: /NPG([INFO-ONLY]) ; Sub-Collections: SRI-ARC NPG; Clerk: JBP;

JBP 16-MAR-75 01:21 25576

DPS questions from Doug Wells

15-MAR-75 22:23:53-EDT,4075;000000000000 Mail from MIT-ML rcvd at 15-MAR-75 2223-EDT Date: 15 MAR 1975 2221-EDT From: DMW at MIT-ML To: dmw at MIT-MULTICS, postel at BBN-TENEXB 1

Jon, We have now received and partially digested yours and Jim's responses to Bob's and Rick's critique, and once again I have some points to be clarified for me: 2

1) you say the 'schar' construct means "an arbitrary number"; does this mean "1 or more" or "0 or more", i.e., do you wish to allow the construct (for example in PCPTXT of a list of integers length 5 - all zero) "L5,IIIII" ?? 3

2) what about PCPHST PORT* being an INTEGER, presumably a networ socket ranges from 0 to 2**32-1 (not including host), This seems to conflict with the desire to have INTEGERs range only up to 2 ** 15 - 1, Now as it happens, Multics assigns small user socket spaces (old userid field) to processes originating from the net, but ... I also seem to have seen some pretty hefty socket numbers on the TENEX connections to Multics. 4

3) Perhaps host names should not be passed over the net, only decimal integers. The problem here is that we have seen no "official" documents which say in which case official host names exist, nor even that they should be recognized in either case. We on Multics who have distinguished upper from lower case for years are still not convinced that we should attempt to say that we should map all host names to upper case before matching them. Any thoughts in this area, or if anyone (including NSW documents) were to say that upper case should match lower case, we would be happy to do it for that system. 5

4) In the File Package, do the mentioned date and time's have a specified format, or are they "human-readable only"? 6

5) Once again I am unclear about the duties of the TBH dispatcher (as he is now known). We would still like to dedicate one process (a simple one - with few "free" resources) to accepting requests for processes. Multics will not (and here I speak from a philosophical view of security) create a process without the access-id of the created process being fully specified before the process is created, and will not change the access-id of an already created process. 7

The scenario I seem to desire (and here I use already specified packages only because I dont need to explain them) is as follows: 8

An entity desirous of a process on Multics (e.g., the WM) ICP's to socket 25 at multics. The process at the Multics end of the Network connection speaks PCP only. The desirous entity opens the Process Management Package (the only Package available in this

restricted environment), specifies a PRCADDR and USERID and a Multics process is created for the desirous entity. 8a

Notice that this would (in actual implementation) create a process tree on Multics (whose root is the "Initializer" process -- so that it corresponds to the actual Multics philosophy of a single process tree per Multics). Now while we could use a single access-id for all NSW users and run an interpretive access scheme (ala the TENEX file transfer server), we (Multics Network Group) have desires to actually use PCP (if possible) for inter-Multics communication. (That is, we have two Multices now, and expect to have two or three more within a year; and we will have to develop schemas for communicating betwee them. 9

Would you attempt to explain to me how my scenario above fits in with your current concept of the TBH dispatcher (and the $F <host> <intrahostaddr> concept), especially in regard to the timely specification of USERIDs. 10

By the way, we are currently starting actual implementation o PCP/PSP (and detailed design of PMP), and I am curious about how other implementations are progressing, especially regarding the possiblity of testing my interpretation of PCPTXT,TXT about encoding and decoding of data structures. 11

Thanks for listening, Doug Wells 12

JBP 16-MAR-75 01:21 25576

DPS questions from Doug Wells

(J25576) 16-MAR-75 01:21;;; Title: Author(s): Jonathan B. Postel/JBP; Distribution: /JEW([INFO-ONLY]) ; Sub-Collections: SRI-ARC; Clerk: JBP;

JMB 16-MAR-75 02:27 25577

System messages about critical events

In this time of users' sensitivity to orneriness of the Office-1 system, I suggest for the future that the operator be given a repetoire of pre-set system messages of explanation that can be sent out, as a matter of policy, within a half-hour of operator's noticing or doing each of the following serious situations: 1

Acknowledge that the system has recently crashed--Did it smash anyone's files or not? 1a

Explain when the file system has been reloaded back up to such-and-such-a-time, even if this happens at 5 am Calif time (for obvious reasons for east coast users). 1b

The downtime message upon login when system is scheduled down for an unusual time should include a reason that is understandable to most users (like, for this whole weekend) 1c

system messages about critical events

(J25577) 16-MAR-75 02:27;;; Title: Author(s): Jeanne M. Beck/JMB;
Distribution: /FEEDBACK([ACTION]) JCN([ACTION]) ;
Sub-Collections: SRI-ARC FEEDBACK; Clerk: JMB;

JMB 16-MAR-75 02:28 25578

DNLS response (on direct line to ARPA-TIP) Week of 10-MAR

Charles, pass this on to who ever is appropriate

JMB 16-MAR-75 02:28 25578

DNLS response (on direct line to ARPA-TIP) Week of 10-MAR

Starting Wednesday, DNLS response was noticably quicker at Office-1. Continued Thurs. This is my perception only; haven't heard anybody else say this. Still some stuttering in TENEX at Lineprocessor. And there is still inordinate delay in response to CTRL-T. 1

On Wed afternoon about 4 pm Eastern time Robert Lieberman tested DNLS response at both BBNB and O-1. BBNB was faster. 2

On Thurs about 6 p.m. Eastern time, I tried a comparison of DNLS while connected to 49 and then while connected to 43, and found not much difference in response--though the load on o-1 was very low at that hour, and I don't know the load at BBNB. 3

JMB 16-MAR-75 02:28 25578

DNLS response (on direct line to ARPA-TIP) Week of 10-MAR

(J25578) 16-MAR-75 02:28;;; Title: Author(s): Jeanne M. Beck/JMB; Distribution: /CHI([INFO-ONLY]) ; Sub-Collections: SRI-ARC; Clerk: JMB;

JMB 16-MAR-75 18:21 25579
item you and I could send out to push for dnls stuff at sri-wash

Rita, partly becuase Rob's demo and business discussions at ARPA Friday didn't work out too well (Connie felt imposed upon for several reasons), I would like to push for DNLS equipment at SRI-Wash. I thought of sending out something like this, what do you think (fill in the blanks or change whatever you wish and send back to me)?: 1

TITLE: Facilities for ARC people in Washington area
COMMENT:
AUTHOR(S):JMB RH
DISTRIBUTE FOR ACTION TO: MEH,JCN(is there an opportunity to get the DNLS equipment at SRI?),RLL(Connie told me later that she found the demo & discussions on Friday awkward--could we push for DNLS equip at SRI?)
DISTRIBUTE FOR INFO-ONLY TO: ARC-APP
MESSAGE:
BRANCH AT:
PLEX AT: <,d>

2

Jeanne Beck & Rita Hysmith have been assigned an office at SRI-Wash. It is room #..... Telephone: (703) 524-2053 Ext.398 & 247. This office can be used by all Org. 750 people. All mail & other materials for Applications people in Washington (for RH or JMB etc) should be mailed c/o Hysmith or Beck to the Wash office via the SRI interoffice mail pouch. 2a

Telephone messages for APP Menlo people traveling in Washington area should be .....
When Rita is in Wash, she can be reached at that office; but Jeanne is more often located at ARPA--telephone (202) 694-3510 or 3506. 2b

A collection of ARC documentation is available in the SRI office, and we usually have a TI there. We are not set up for DNLS, however, except a large monitor sits there and a leased telephone line .....
[what's the status of the line?]. Hopefully we could complete a DNLS workstation there in the future for demos by ARC-APP people. It would be much less awkward for demos for more than 2 visitors than ARPA is. At ARPA, the room in Connie McLindon's area containing the Lineprocessor station is the workplace for 3 people and a passageway for 2 others, and so most demos crowd ARPA's work. The office at SRI would be more private. 2c

JMB 16-MAR-75 18:21 25579

item you and I could send out to push for dnls stuff at sri-wash

(J25579) 16-MAR-75 18:21;;; Title: Author(s): Jeanne M. Beck/JMB;
Distribution: /RH([ACTION]) ; Sub-Collections: SRI-ARC; Clerk: JMB;