Send all comments, criticisms and contributions to the editor
T. W. Lanzatella

IN REGARDS TO BINARY FILE QUEUE TYPES

The last issue of the Deadstart Systems Newsletter announced the installation of a
new queue type, BIFT or Binary File Queue type. Two unforeseen problems reared
their ugly heads when we attempted to introduce this change to the production
system. One, XEDIT stopped functioning; an easy remedy was to reassemble and SYSEDIT
the new version. Two, SYSTEM 2000 stopped functioning; again, an easy fix would have
been to reassemble and place a new version on CALLPRG - if we had had a source of the
program. Unfortunately, we do not have a source, hence, the only alternative was to not
run the new tape, which we didn't.

NOTICE OF CHANGES TO THE OPERATING SYSTEM

All changes announced in Vol. 1 No. 6 issue of the Deadstart Systems Newsletter
except the installation of BIFT will be receiving their production system debut on
this next tape.

Tim Salo contributed the following modifications:

1. A modification to COMCTID which repairs an origin type check.

2. A modification to DOCMENT which primarily removes the ident DOCINFO. The ident
   DOCINFO installed the UM parameter, an undocumented, unused and poorly implemented
   parameter to create old info file records. This modification also fixes the
   following bugs:

   a. Mode 1 if too many list cards.
   b. Delete extra page between copies.
   c. Change name to University of Minnesota from CDC on the foot of the page.
   d. Insert an RFL = entry point thus, changing DOCMENT's type from OVL to ABS.

3. Tim also modified CPMEM to dump the deferred contents of the B registers for time-
   sharing users.

Kevin Matthews modification to PDUMP to accommodate incremental permanent file
dumps was added with this tape.

The long awaited loader statistics modifications to LINK are finally being added to
the system. The following description is due to Rick Matthews.

The Loader Statistics, as requested by Mike Frisch, consist of package names and
usage counts for routines loaded off system libraries, i.e., SYSLIB, FORTRAN, SYSIO,
et al. The method for storing and updating these counts is as follows:

1. A Random Access, Direct Access, Permanent File by the name ZZZZZS1 is on the
   master device on user number LIBRARY. It is public catalog type, mode is write,
   and passworded. The file consists of $n-512_{10}$ word records, where n is the
   number of libraries and a directory. The first word of each record is the given
   library. Initially, the rest of the record is zeroed.

2. The loader will try to access this permanent file 10 times before abandoning statistics gathering. If found, it will read the directory into a 100_8 word managed table. After verifying the record type, it will compress the table to only directory entries. If not a directory record (OPLD), it abandons gathering.

3. When a system library has been found, the loader searches the directory for a matching library name. If found, it allocates 1000_8 words for the list and reads it into core. While satisfying references, including all ULIB references, it updates any entries found and will add new entries as long as space exists in the table. After the library has been completed, a random rewrite replaces the record, and at end of load, the file is returned.

This process adds about 1300_8 words to field length required to load; 1000_8 for table of package names, 100_8 words for directory, 200_8 words for code and managed table entries.

PROPOSED CHANGES TO THE OPERATING SYSTEM

The following proposals were received from John Strait:

1. UNPAGE

   In the past few months UNPAGE has been enjoying rather heavy usage. Of the CALLPRG programs only SPSS has exceeded UNPAGE, and it has been in the top 10 - 20% of all control cards. Therefore, I propose that it be placed on the Deadstart tape as a system control card.

2. PASCAL

   As the CICS department will be making use of PASCAL starting in fall 1975, and much of its use is from TELEX, a PASCAL subsystem should be added to TELEX. This will require changes to the PASCAL compiler, TELEX, and certain associated common decks and peripheral processor programs.

LETTERS TO THE EDITOR

Alan Johnston lodged the following complaints/opinions about Don Hamnes' proposal to change QFM.

1. This change should not be implemented until level 10.

2. The change will (expletive deleted) the users who call LFM.

3. The 1000B CM buffer is too big and perhaps useless.

4. The move results in inefficiency; against KRONOS philosophy.

5. Having 1CJ call QFM is stupid and inefficient.

6. The access idea is poor security and can be performed in an alternate easier way.

It would be nice to call an overlay that performs the proposed function, not calling another PP.

Kevin Matthews contributed the following comments concerning Jim Mundstock's ECS roll-out proposal.

A. It seems awkward to have several parts of the system make sure that ECS jobs have at least 4K of central. It might be easier to have 1RO make sure that 4K is available. If 1RO can't find 4K, it could do one of two things:

    a. Roll-out ECS inefficiently, a PRU at a time.

    b. Not roll-out the ECS, and in that case leave the ECS attached to the rolled job, as is done presently.

B. Checkpoint/Restart

    Checkpoint and Restart currently don't work with USER ECS. Can it be worked into this scheme easily? Will the user ECS in the middle of the central memory portion of the roll-out file cause problems here?

C. ECS Memory Scheduling

    Currently, there is no scheduling based on user ECS requirements. This is because a job which requests user ECS that is not available will roll-out for three minutes, then roll-in to try again. Jobs that have ECS have it even when they are rolled, so that ECS availability is not considered for roll-in. This will have to change when ECS is rolled in and out, and has apparently not been considered in the current proposal. 1SJ should not schedule a job for roll-in unless the ECS required is available.

D. The current plan is to have 1RI call MEM to request the user ECS needed. This is bad because 1RI simply wastes time waiting for MEM to finish, and because interlock problems may occur whenever one PP requests (and waits for) another PP to do something. I think a better solution would be to have MEM be called directly (either by the system or by 1RI) with a special call for this use only. Then MEM can call 1RI when the user ECS has been obtained.

The ever prolific Dennis Lienke submitted the following letter which comments on practically everything in sight.

Re: 29 July 75 Deadstart Systems Newsletter (Vol. 1 No. 6)

1. Hamnes' Proposal

    a. If the caveat in point 9 is true ("... Binary ... may not work), then what is the sense of point 1 (translation of calls - LFM calls QFM)?

    b. The precedent cited for SFM function 00 requiring a CM buffer is, in my opinion a bad precedent. In addition, the SFM function 0 only requires 100B words, not 1000B.

    c. The reason for the requirement of 1000B word buffer was "QFM is crowded." - that would seem to call for splitting up QFM into overlays - and there are numerous precedents for that.

2. Elliott's proposal

    (PERMIT) How many account numbers could be accommodated?

3. EJM ECS Roll-out

The question of recovery of files was not addressed - is that not a consideration?

4. Marisa Riviere - System Maintenance

Suppose a person has a semi-private file with attach 'accounting' information associated with it (I have such a file). Does the facility exist for updating (i.e., rewriting) the file without losing the information?

5. Bob Williams - Random note

The entire question of such files could well be made moot by the Record Manager. It contains facilities for word addressable files, and many types of 'indices'. I think that our time could be better spent on other areas, assuming:

a. MNF will ultimately use CRM.

b. Not "many" users have conversion problems (you could simply tell everyone to use FTN 4), or make every one use CRM facilities directly from whatever compiler they're using (except RUN).