

Published: 12/16/66

Identification

The User Checkpoint Dump Decision Module
S. H. Webber

Purpose

The user checkpoint dumping process is executed periodically to accumulate, on a distinct body of detachable storage, copies of frequently used segments. This accumulation allows the secondary storage reload process (section BH.3.02), when reloading secondary storage following a catastrophe, to locate and load many missing segments while scanning minimal amounts of off-line storage. This section describes the user checkpoint dump decision module and, when read in context with section BH.2.00, gives a complete description of the user checkpoint dumping process.

Introduction

The only purpose of the user checkpoint dump is to insure that the secondary storage reload process, operating during the second phase of the reload, finds many missing segments quickly and without scanning inordinate amounts of detachable storage. Thus, the user checkpoint dump consists of copies of many popular segments and normally only the most recent checkpoint dumps are preserved. To create a checkpoint dump, one complete hierarchy scan is carried out and all segments which have been used since a set date and time are dumped by the process. This pre-set cut-off time is determined by subtracting a fixed time span from the current time. The time span is set by a system administrator and will probably have a value of 1 - 3 weeks.

User Checkpoint Decision Module

The user checkpoint decision module is called by the hierarchy scan module to consider each branch the latter finds in the file system hierarchy. The decision module first compares the date/time-last-used (dtu) of the entry with the cut-off time. If the entry has not been used since the cut-off time, the decision module returns to its caller leaving no instructions for the dump module. The dump module is not called since no work is to be done with this entry. If, however, the entry has been used since the cut-off time, additional decisions must be made.

If the entry is a branch, the type of branch is determined:

- a. If the entry is a non-directory branch, the segment-dump-switch associated with the entry in the position segment (BH.2.00) is turned ON, and a normal return is made. This action results in a call to the dump module by the scan module and the eventual dumping of the segment. In addition the call-dump-module switch is turned ON.
- b. If the entry is a directory branch, then the scan deeper return is made which results in the subsequent searching of that inferior directory.

The decision logic is diagrammed in Figure 1.

The user checkpoint dump decision module is called by the following:

```
call user_checkpoint (dsw, psp, deeper);
```

dsw is the call-dump-module switch. The decision module sets this ON only if the segment defined by the current entry is to be dumped. The scan module will call the dump module for this entry only if dsw is ON on return of control from the decision module.


psp is a pointer to the base of the position segment.

deeper is an alternate return to be used if a directory branch is encountered, i.e. the immediately inferior directory should be scanned next by the process.

These arguments are declared by the following PL/I declaration.

```
dcl dsw bit (1),  
     psp pointer,  
     deeper label;
```

dtu = data/time-last-used
pct = present cut-off time
fsw = file dump switch
dsw = call-dump-module
 switch



User Checkpoint Dump Decision Module

Figure 1