

Published: 12/23/66

Identification

The Secondary Storage Reload Process
S. H. Webber

Purpose

After the reconstruction process has finished its work (section BH.3.01) the file system hierarchy is in the following state: 1) all directory and some non-directory segments are present in their latest backed-up version, 2) all missing user segments are flagged by the retrieval trap switch being ON in their branches. The function of the reload process is to find and load as many of these missing segments as is desirable for system performance.

Due to the nature of the data dumped by the system and the user checkpoint dumping processes, most user owned segments are loaded by the secondary storage reload process while the hierarchy reconstruction processes accounts for most segments required for Multics operation. Since the reload process can execute while Multics is operating normally with users, some users may find that some of their personal segments are not yet present and will have to wait for their arrival.

Introduction

The directory entries present in the hierarchy defining missing segments are characterized by the following:

1. The date/time last modified and the unique identification present in the entry define exactly which version of which segment is missing (no directories are missing since they have all been loaded by the preceding execution of the hierarchy reconstruction process).
2. The retrieval trap switch is set ON signifying that the associated segment is indeed missing.
3. Information is present telling exactly where in backup storage the correct version of the missing segment may be found.

The reload process scans certain portions of backup storage loading any missing segments that it finds. If this search were exhaustive all missing segments would eventually be found. However, searching such a large body of data is burdensome and very uneconomical. Therefore, only certain areas which are known to contain a large number

of missing segments are processed. Specifically, 1) all detachable storage created by the incremental dumper since the beginning of the last user checkpoint dump and before the beginning of the latest system checkpoint dump and 2) the latest user checkpoint dump are scanned.

The Reload Algorithm

All segments found in these areas of detachable storage are loaded unless one or more of the following conditions exists:

1. The segment was previously removed by the multilevel system and not since replaced on secondary storage.
2. The segment is a directory.
3. A directory entry defining the segment does not now exist within the hierarchy. (i.e. it was deleted since this segment version was dumped.)
4. The retrieval trap switch in the branch defining the segment is OFF in the hierarchy signifying that the segment is not really missing.
5. The unique identification and date/time-last-modified referring to the segment on backup storage are not identical to the same parameters found in the corresponding directory entry in the hierarchy. This indicates that the segment is not the proper version of the data that should be loaded.

Whenever a segment is loaded the retrieval trap switch is turned OFF, thus releasing the segment to users. All logical records read from backup storage which include directory segments are ignored. The above reload algorithm is described in Figure 1.

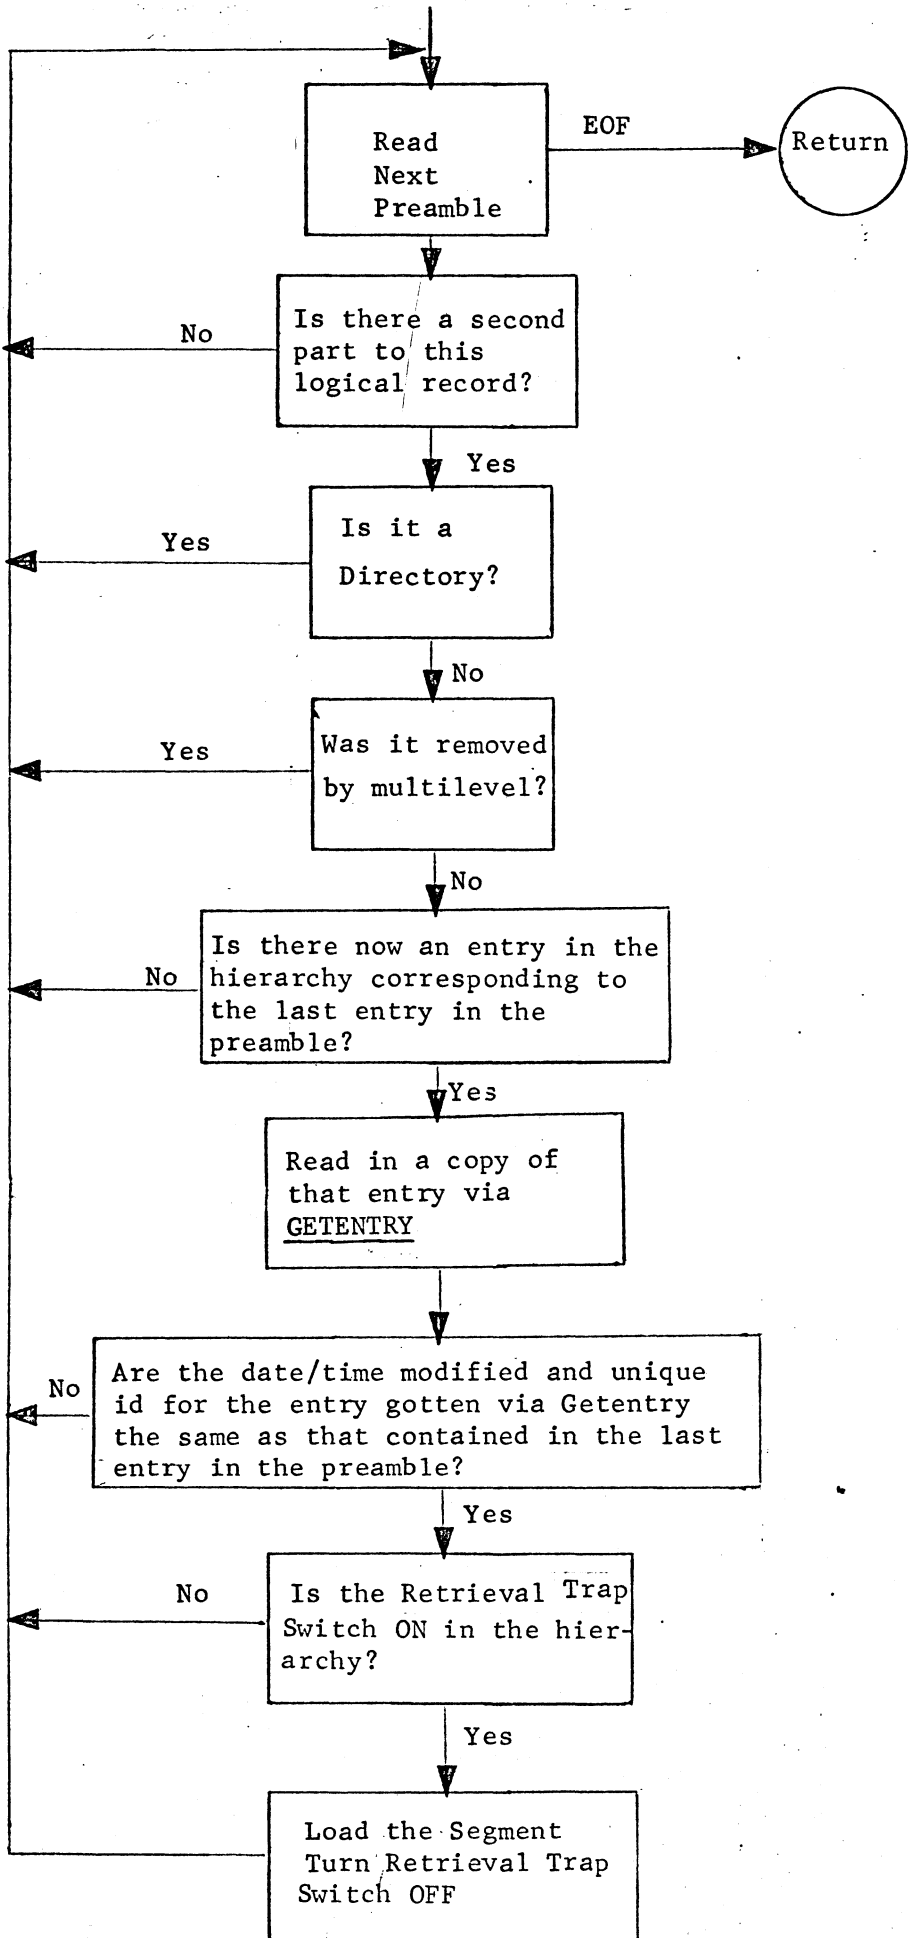


Figure 1.

The Reload Process