

DRAFT:

Date?

Hjg 11/25/66

Identification

Multilevel Storage Support

L. J. Lambert, K. J. Martin

Discussion

Multilevel Storage Management (BH.4) determines where in secondary storage a file in the Multics file hierarchy shall reside. The primary goal of Multilevel is to arrange file residence for the most efficient overall operation of Multics. This is accomplished by placing frequently accessed files on fast storage devices and less frequently accessed files on slower devices. In addition, overall system efficiency is better served by placing very large files on slower devices, leaving the space on faster devices for smaller files.

The owner of a file (that is, the user who pays for its residence on secondary storage) knows the most about his needs regarding his file. In nearly all cases the owner of a file will need no special handling for his files. The files which he accesses frequently serve him best when residing on fast access devices. His infrequently accessed files need not be on fast access devices and are more economical when residing on slow-access devices. Occasionally, although rarely, a user will find that Multilevel's treatment of his files does not serve his needs. For instance, he may know that although a file is large and used infrequently, when it is needed, the access must be prompt. The criteria normally used by Multilevel causes such a file to be placed on a device where access time would be slow.

As mentioned previously, it is desirable for the owner of a file to be able to influence the handling of his file by Multilevel Storage Management. The Transactor in support of this function provides the owner of the file with the ability to set limit stops concerning his file residence. These limit stops are meant to reflect upper and lower bounds of file residence expressed in terms of device type. This implies the ranking of devices by some criteria. The criteria chosen is speed of access. The limit stops are true limit stops in that a file will never be stored on a device class of high rank than specified by the upper limit and the lower limit prevents the file from drifting down to a device class of lower ranking.

Usage:

The Transactor must have at least two items of information. In some instances when both an upper and lower limit are to be established three items are necessary:

1. file name

2. device class of upper limit stop
3. device class of lower limit stop

The file name must always be specified. If the owner wishes to set only an upper limit he specifies only the upper limit stop. This is interpreted by Multilevel Storage Management to mean that this file will never be put on a device class of higher ranking.

If the owner specifies a lower limit stop only, it is interpreted to mean that the file will never drift to a storage class ranked lower than this limit. If both limits are imposed the file will be restricted to residence within the range specified by the limits.

It is important to note two things; first, the setting of one limit has no effect on the other limit. Secondly the default setting for all files is an upper limit equal to that of the highest ranked device class and a lower limit equal to that of the lowest ranked device class. These settings give Multilevel Storage Management the widest possible latitude when selecting file residence for this file.

Implementation:

The multilevel storage support facility of the Transactor is implemented by command. The command specifies the file name, an upper and/or lower limit by device class.

The procedure invoked by this command uses the items supplied to set the priority parameter (s) in the branch entry for the file. This is done by determining a unique device class number for each named device class and setting this number in the appropriate priority parameter item. Priority parameter 1 represents the upper device class limit and priority parameter 2 the lower device class limit. A value of zero for either parameter indicates a default condition.

In addition to setting the parameter (s), this procedure must check the authority of this user to make such a request. This authority, if present, is found in the "user profile" of the requesting user.

It goes without saying that the user must have some access to the file and write access to the directory in order to set the priority parameters. In addition to those access rights the account number to which the file residence is charged must be one to which the user is allowed to charge. This combination of criteria is enough to ~~at least partially~~ establish ownership of the file.

Each request is also checked for consistency to avoid ambiguities such as an upper limit being set which is of lower rank than the lower limit.

This involves not only checking of the items of request itself but also cross-checking against previous settings of the priority parameters.

The appropriate parameter values are determined from a data base which describes the ranking of device classes. This data base is described and maintained by Multilevel Storage Management (BH.4).