

DRAFT:

8/14/66

Identification

Transactor: Overview

L. J. Lambert, K. J. Martin

Introduction: The Transactor Overview is meant to provide the reader with a general picture of those functions performed by the Transactor. The detailed information is to be found in following BT. sections. In addition, some definitions are presented which are referred to frequently throughout the BT. sections.

Purpose: The Transactor's primary objective is to provide the user with a mechanism for exercising certain controls over processes operating within the Multics System. In fulfilling this objective the Transactor maintains records of all negotiations attempted by a user.

A user may negotiate with Multics in support of his efforts in the following areas:

- resource reservation
- response control
- multilevel storage management support
- transfer of resources
- system complaints

Definitions: Before proceeding any further it is necessary to define terms. Some of these definitions have been extracted from other sections and are restated here for consistency and ease of reference.

A process is a program in execution. It is defined by a set of machine conditions and an associated address space (a core image).

As used in the context of the BT sections, a job is an identifiable task which is bound by a login and logout.

An audit trail is a file of selected information used for verification or recovery purposes.

A user is a person entitled to use the system and is defined in a user list, a data base of the system skeleton (see BD. user list). A user, when he logs in, becomes identified with a group of processes called a user-process-group.

Two classes of user are defined and are distinguished by only their method of using the system. An absentee user makes arrangements for or initiates a group of processes. These processes do not require any further intervention to successfully complete. An absentee user

can be thought of as making arrangements via remote terminal or by simply carrying an appropriately prefaced card deck to the computer room. The absentee user next has contact with the process when the results are available. The present GECOS user falls into this category. An interactive user on the other hand desires to interact (via remote terminal) with his process as it is in execution. This class of user may or may not make prior arrangements for the execution of the process. A user can readily change status from interactive to absentee or from absentee to interactive. It is important to note that within the system there is no distinction between absentee and interactive status.

The term resource will be understood to mean any component of the system or logical part thereof. It is also reasonable, in some instances, to consider funds which can be exchanged for component usage as a resource. The following are typical examples of resources: processor time, tape drives, system software, disk tracks.

Resources can have two classifications determined by their ability to be shared. A shareable resource is one that can be used by many users at once without danger of conflict. A non-shareable or dedicated resource is one that cannot reasonably be used by more than one user at a time. An example of a dedicated resource is a tape drive.

A user's allocation of a resource is a statement of the amount of that resource which he is allowed to use in a particular period of time.

In addition to defining terms there are certain users who have a distinct function for the system.

An operator is a user with a special designation for Multics. He provides the manual operations required by Multics. These operations include such things as tape mounting, printer and card equipment set up. An operator is usually associated with one or more consoles. In most cases there are several operators.

The notion of user authority is introduced to protect against the undisciplined use of system features such as those offered by the Transactor. Each user of Multics will be administratively either given or not given the authority to perform certain operations. The various authorities are quite separate; a user could possess any combination of authorities.

A system administrator is a person(s) who is responsible for the efficient, secure and orderly operation of a Multics installation. By virtue of their function, system administrators are users possessing many authorities. One authority a system administrator should possess is the authority to alter all or some other users' authorities.

Components:

At each Multics installation some portion of the total resources will

be turned over to the Transactor for management. That is, the Transactor makes reservations for users from this portion of the available resources.

The resource reservation function of the Transactor allows the user to obtain a commitment from the system concerning reservable system resources. This commitment is in response to a request submitted by the user which defines: 1) the resources desired, 2) for how long, and 3) during what period. Specification of the period can be as flexible or as constraining as the user wishes. A higher degree of flexibility increases the chances of a request being filled.

Resources are reserved as shared or dedicated resources. The difference lies in the billing practices employed for each. A user is billed for only that amount of the shared resource that he actually uses. A user of a dedicated resource is charged for the amount of time requested regardless of the usage during this time period. In the case of early release of a dedicated resource, the user will be charged for the remainder of the reservation if it cannot be utilized by some other user.

Provision is made for the job which runs periodically with fixed frequency and fixed resource requirements. A reservation associated with this type of job is called a planned reservation. The weekly payroll is an example of such a job.

Another area considered by the Transactor is that of system response. This applies primarily to the interactive user and to response experienced at the console. Any user with the appropriate authority and funds is allowed to alter his response time if the system is capable of doing it. In return for improved service, the user pays more. Likewise, a user receives a reduced rate for slower service. The Transactor provides the mechanism by which a user requests changes to his response time.

Response time is not solely a function of more frequent processor time. Other things must be considered such as the present load on the system, I/O to dedicated resources, and the File System response to this process's requests (such as missing page fault).

The Transactor's primary objective in the area of multilevel storage support is to provide the user some influence over the charges incurred by residence of files which he owns on an auxiliary storage device. There are circumstances when an owner of a file desires to have a file always maintained on or above a certain level of access device to guarantee some desired access time. Similarly, an owner may be willing to accept slower access to a file in return for a reduced rate. It may also be true that the owner of a file is not willing to be charged for the residence of his file above a certain level. The Transactor provides the mechanism for the owner to express his desires to the Multilevel Storage Management Module.

Since it is not always possible for a user to anticipate needs accurately when resources are budgeted, a Transactor function designed for the transfer of resources between accounts is provided. The resource transfers involve any resources which could be allotted to a user.

The Transactor makes it possible for a user to register complaints about the system. The complaints are registered by the Transactor and referred to a system administrator for adjustment. The Transactor also provides the system administrator with information to which it has access. It does not attempt to make any adjustment.

All transactions must be carefully monitored, and records kept which are forwarded to the responsible administrator. These records are used by the administrator for reference if a conflict arises in any transaction area. These records would be kept for a reasonable period determined individually by each Multics installation.

Implementation:

The Transactor is implemented as commands, system processes and procedures. For instance, the resource reservation function of the Transactor involves all three. Normal reservations are made via command which in turn arrange for the awakening of processes. Other system processes call upon Transactor procedures to accomplish resource reservation. The specific implementations are detailed in the appropriate MSPM - BT sections.