

To: Distribution
From: T. H. Van Vleck
Date: February 11, 1972
Subject: Load Control Groups

This memo describes the initial implementation of a load control group mechanism for Multics. The implementation is split into an initial part, which is very easy and solves most of our current problems in the load control area, and a second part, which puts in several fancy features.

The first phase will provide the biggest change to the user interface: it introduces "CTSS-style" load control groups, and the notion of primary and secondary status for users. The load control groups can be allocated a quota of primary load units representing a guaranteed number of users; users in excess of the quota will be logged in as secondary, and will be subject to preemption if a user who receives primary status logs in; and preemption within a load control group across project boundaries will work.

This first phase will be accomplished by the installation of two procedures, and the creation of a single new data segment.

We will attempt to install the load control programs and tables in a way that will not change the appearance of the system to the average user, and then later perform the actual division of projects into groups by editing the various system tables.

Load Control Groups - 2/2/72

page 2

The second part of the party-group implementation could provide some or all of the following features:

- 1) Scheduled service
- 2) Edit-only service
- 3) Load Control group administrators
- 4) load control group control of some project attributes

All of these functions are easy to code but involve more programming than the basic load control group mechanism. Furthermore, these functions may not be necessary at all. We should live with the basic implementation for a month or two before dedicating the resources necessary to provide what may turn out to be nothing but "frills".

There are a few miscellaneous tasks which fall between part one and part two. These tasks include:

- 1) Modifications to "who" to show group and secondary status
- 2) Administrative reports by group
- 3) Minimum waiting time after user is preempted before he can login again.

Detailed description of part one

Only one module in the system is involved in the actual load-control decision. Replacing this program, "load_ctl_", with a version that uses a new strategy, is straightforward.

The new version of load_ctl_ will discover the group associated with each project by inspecting a slot in the SAT which is currently ignored. This slot holds an 8-character group identifier, which will be looked up in the "master group table" (MGT) to find the current and maximum number of primary load units for the group.

The other module that will be installed is a simple-minded editor which allows the system administrator to create and edit the MGT. Group administrators who wish to change parameters of a project in their group and the moving of projects from one group to another will be handled by the system administrators on a request basis. These changes should be infrequent enough so that not much work is involved.

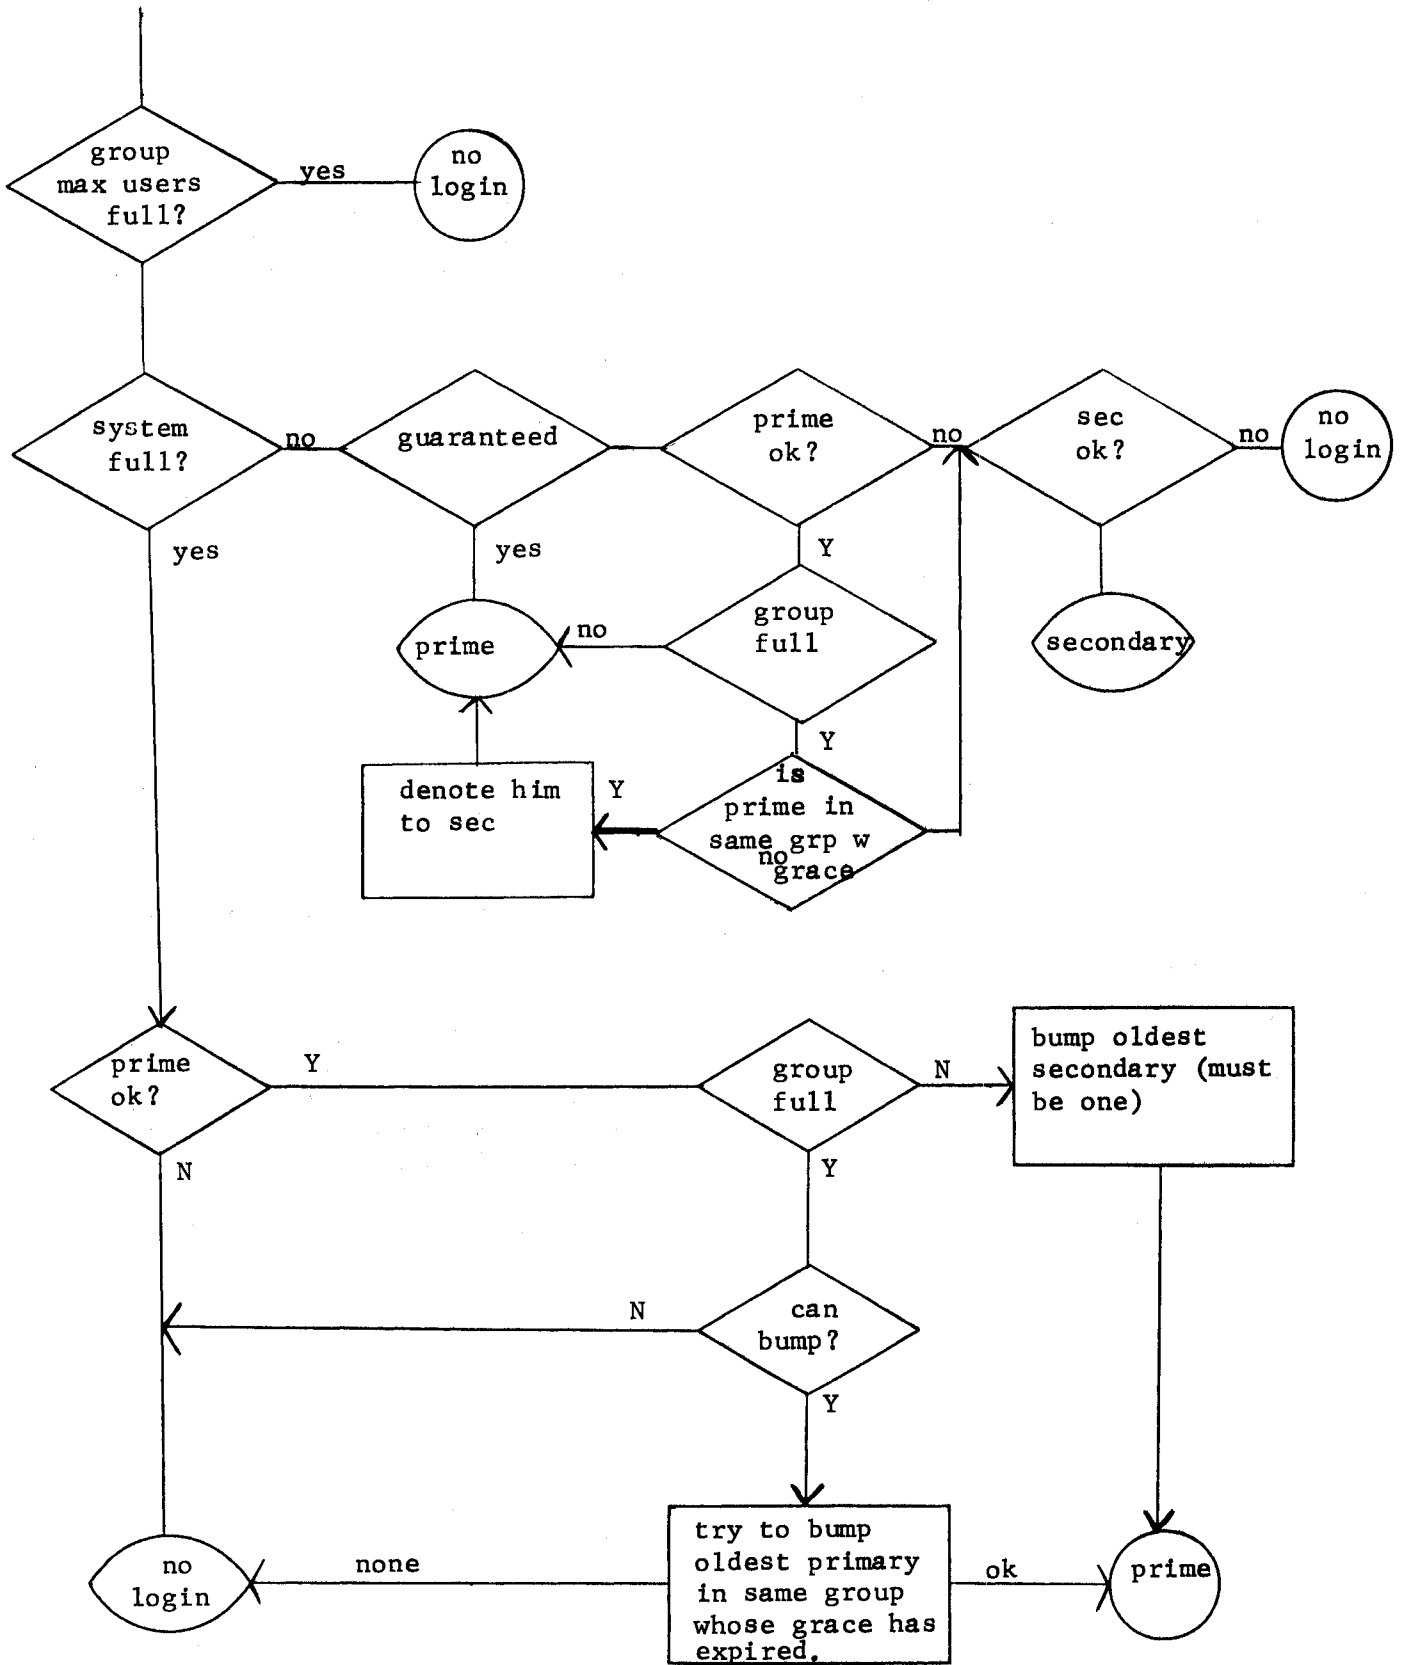
The MGT for the basic load control group mechanism will contain little more than the group identifier, the current and maximum number of primary units, and the current number of secondary units for the group. Later, we may want to add pointers to a schedule data base, attribute masks, administrator identifiers, and a count of edit-only users.

The flowchart provided as figure 1 shows the guts of the login decision. The algorithm references the following data items:

- 1) System maxunits (already in SAT)
- 2) System current units (already in SAT)
- 3) System max. users (already in SAT)
- 4) System current users (already in SAT)
- 5) Group max. primary units (new, in MGT)
- 6) Group current primary units (new, in MGT)
- 7) User "guaranteed_login" attribute (already in pdt)
- 8) User "prime_ok" attribute (new, in pdt)
- 9) User "secondary_ok" attribute (new, in pdt)

- 10) User primary flag (new, in answer table)
- 11) User "preempting" attribute (already in pdt as "bumping")
- 12) User "grace" (already in pdt)
- 13) User login time (already in answer table)
- 14) User group identifier (new, in answer table, from sat)
- 15) Absolute maximum number of users in group (new, in MGT)

load_ctl_ -- decision at login



Load Control Groups - 2/2/72

page 5

User Interface

A user who logs in and is awarded primary status will get one of two messages, depending on whether he has a grace which is meaninglessly large, or not. If the user will become subject to bumping within the next 24 hours, a message of the form

You are protected from preemption until 1527.

Will be typed just before the "Name, Proj logged in..." message.

If the user has a grace which is greater than 24 hours, he will just be told

You are protected from preemption.

Users who are given secondary status at login will be told

You are subject to preemption.

even though "preemption" does not mean quite the same thing here, since secondary users can be preempted by users from any group.

No message will be sent to the user if he is demoted from primary to secondary after his grace runs out, or if he is promoted to primary from secondary as a result of some primary user in his group logging out. This is mostly because the answering service has no clean communication channel to created processes, other than grabbing the typewriter by force. The "who" command will be modified to show primary or secondary status, and time bumpable, and the "user_info_" subroutine can be changed to return these parameters. Users may then write whatever programs they want to list who is bumpable and when, or to check their own status.

The answering service will still grab the typewriter away from the user process to give him a three-minute warning before he is to be preempted. The message will be of the form

From Multics: Secondary user preemption.

You will be logged out in 3 minutes.

If emergency preemption or group preemption takes place, the message will say "Group preemption" or "Emergency preemption" instead of "Secondary user preemption".

Users who get such a message should clean up and log out. Logging right back in again, in the hope of preempting some other user in the same group whose grace has also run out, is

Load Control Groups - 2/2/72

page 6

considered antisocial. The group supervisor will be able to detect this practice easily, by examining a filtered copy of the answering service message log which contains only those messages pertaining to his load control group.