

TO: A. Bensoussan
C. Clingen
F. J. Corbató
J. Gintell
N. Morris
R. Roach
J. Saltzer ✓
M. Schroeder
S. Webber

FROM: V. Voydock
R. Feiertag

DATE: September 2, 1971

SUBJECT: Access Control Proposals

Enclosed is a document describing our proposals on access control, as modified by the meeting of August 13, and subsequent discussions. The proposal concerning gate lists has been tabled until Schroeder finishes documenting his plan to make it possible for a user to use more than one protected subsystem. There will be a meeting to discuss these proposals on Wednesday, September 8, 1971, at 1:00 p.m. in room 511 at 545 Technology Square.

This document describes a number of proposed changes to the access control mechanism that have evolved over the space of several months. These include:

- I. A new set of directory access modes, and a plan to make all supervisor modules enforce access the same way.
- II. Associating two ring numbers with a directory.
- III. Removing the append access mode on non-directory segments, and introducing the "maximum length" attribute.
- IV. Having one set of ring brackets per segment.
- V. Adding a "deleteable" attribute to segments.
- VI. Eliminating CACL's and replacing them with initial ACL's.
- VII. Adopting the idea of extended access and using it to implement access control on directories.
- VIII. Re-defining access control for system daemons.

This section describes the newly proposed set of directory access modes and presents a plan for making all supervisor modules enforce access in the same way. In addition, a proposal to associate two ring numbers with a directory is introduced.

First we must establish some terminology. An A-operation is an operation performed on the attributes of a segment. This includes modifying its names, its ACL, its "deleteable" attribute, and its maximum length, deleting it, and listing its attributes. It also includes adding a segment to a directory. A D-operation is an operation performed on the data of a segment. This includes initiating the segment, reading it, writing it, executing it, truncating it, setting its bit count, and getting its bit count.

The first major change proposed is to get rid of the "execute" directory access mode and to say that one's right to perform a D-operation on a segment is completely determined by the access information appearing on that segment's ACL. That is, to initiate a segment one needs non-null access to that segment, to get its bit count or read it, read access, to execute it, execute access, to write it, truncate it or set its bit count, write access. One's access to the directory containing the segment is not taken into account when performing a D-operation. Because a certain class of users have need of more security than this plan provides, we have had to modify this proposal to provide a sure means of denying access to an entire subtree of the hierarchy. This will be discussed in detail below.

The second major change is to discard the "append" directory access mode and say that modify access on a directory includes the right to add entries. The argument for this is that it simplifies the access control mechanism with no significant loss of flexibility since no one has found a use for append access without modify access.

D-operations on a segment are controlled by one's access to that segment (almost). A-operations are controlled by one's access to the directory containing that segment. Let us now describe the various directory access modes:

- I. use If a user does not have use access on a directory, he has no access to the contents of that directory or any of its subdirectories (i.e., to the subtree of the hierarchy whose root node is that directory). In no way can he obtain any information about anything in that subtree. Since it does not make sense to have any other access attribute without having use access, the file system will not allow this.
- II. status (formerly read) If a user has status access on a directory, he can list the contents of the directory and find out any and all information about the attributes of any entry in that directory, he cannot add entries or change the attributes of existing entries.
- III. modify (formerly write) If a user has modify access on a directory, he may add entries to that directory and change the attributes of existing entries. He cannot list the attributes of existing entries.

Use access has been added to provide another level of security for a certain class of users. For example, some project administrators would like to allow sharing of information among project members (including the right of one project member to give or deny access to other project members). At the same time they would like to be sure that no project member, either accidentally, or deliberately, can give access to anyone outside that project to any of this information. He can do this by giving project members use access in the project directory and everyone else null access there. One could argue that a project administrator has to trust his project members and that they could (for instance) print out a copy of confidential information and give it to someone who has no right to have it. This is true, but in the event of an information leak, use access narrows down the possible sources of the leak. The project administrator who has used use access to limit access to information knows that the leak must have occurred outside the system. That is, by someone making a physical copy of the information rather than by someone inadvertently or deliberately giving an unauthorized person access to the information online.

Now let us consider what may happen if a user tries to perform an operation on an entry whose pathname is >Dl>...>Dn>E. Seven distinct error conditions (related to access control) may occur.

- I. error_table_\$noentry ("Entry not found")

E does not exist.

- II. error_table_\$no_directory ("Some directory in path specified does not exist")

One of D1, ...Dn does not exist.

- III. error_table_\$incorrect_access ("Incorrect access to directory containing entry")

The user does not have correct access on Dn to perform the operation.

- IV. error_table_\$moderr ("Incorrect access on entry")

The user does not have the correct access on E to perform the operation.

- V. error_table_\$delete_off ("Attempt to delete segment which is not deleteable")

The user tried to delete E and the deleteable attribute of E is off.

- VI. error_table_\$null_access ("No access to subtree of hierarchy")

The user ^{does not have "null"} ~~has null~~ access on one of D1, ..., Dn.

- VII. error_table_\$no_info ("Insufficient access to return any information")

The user does not have enough access to be given any information.

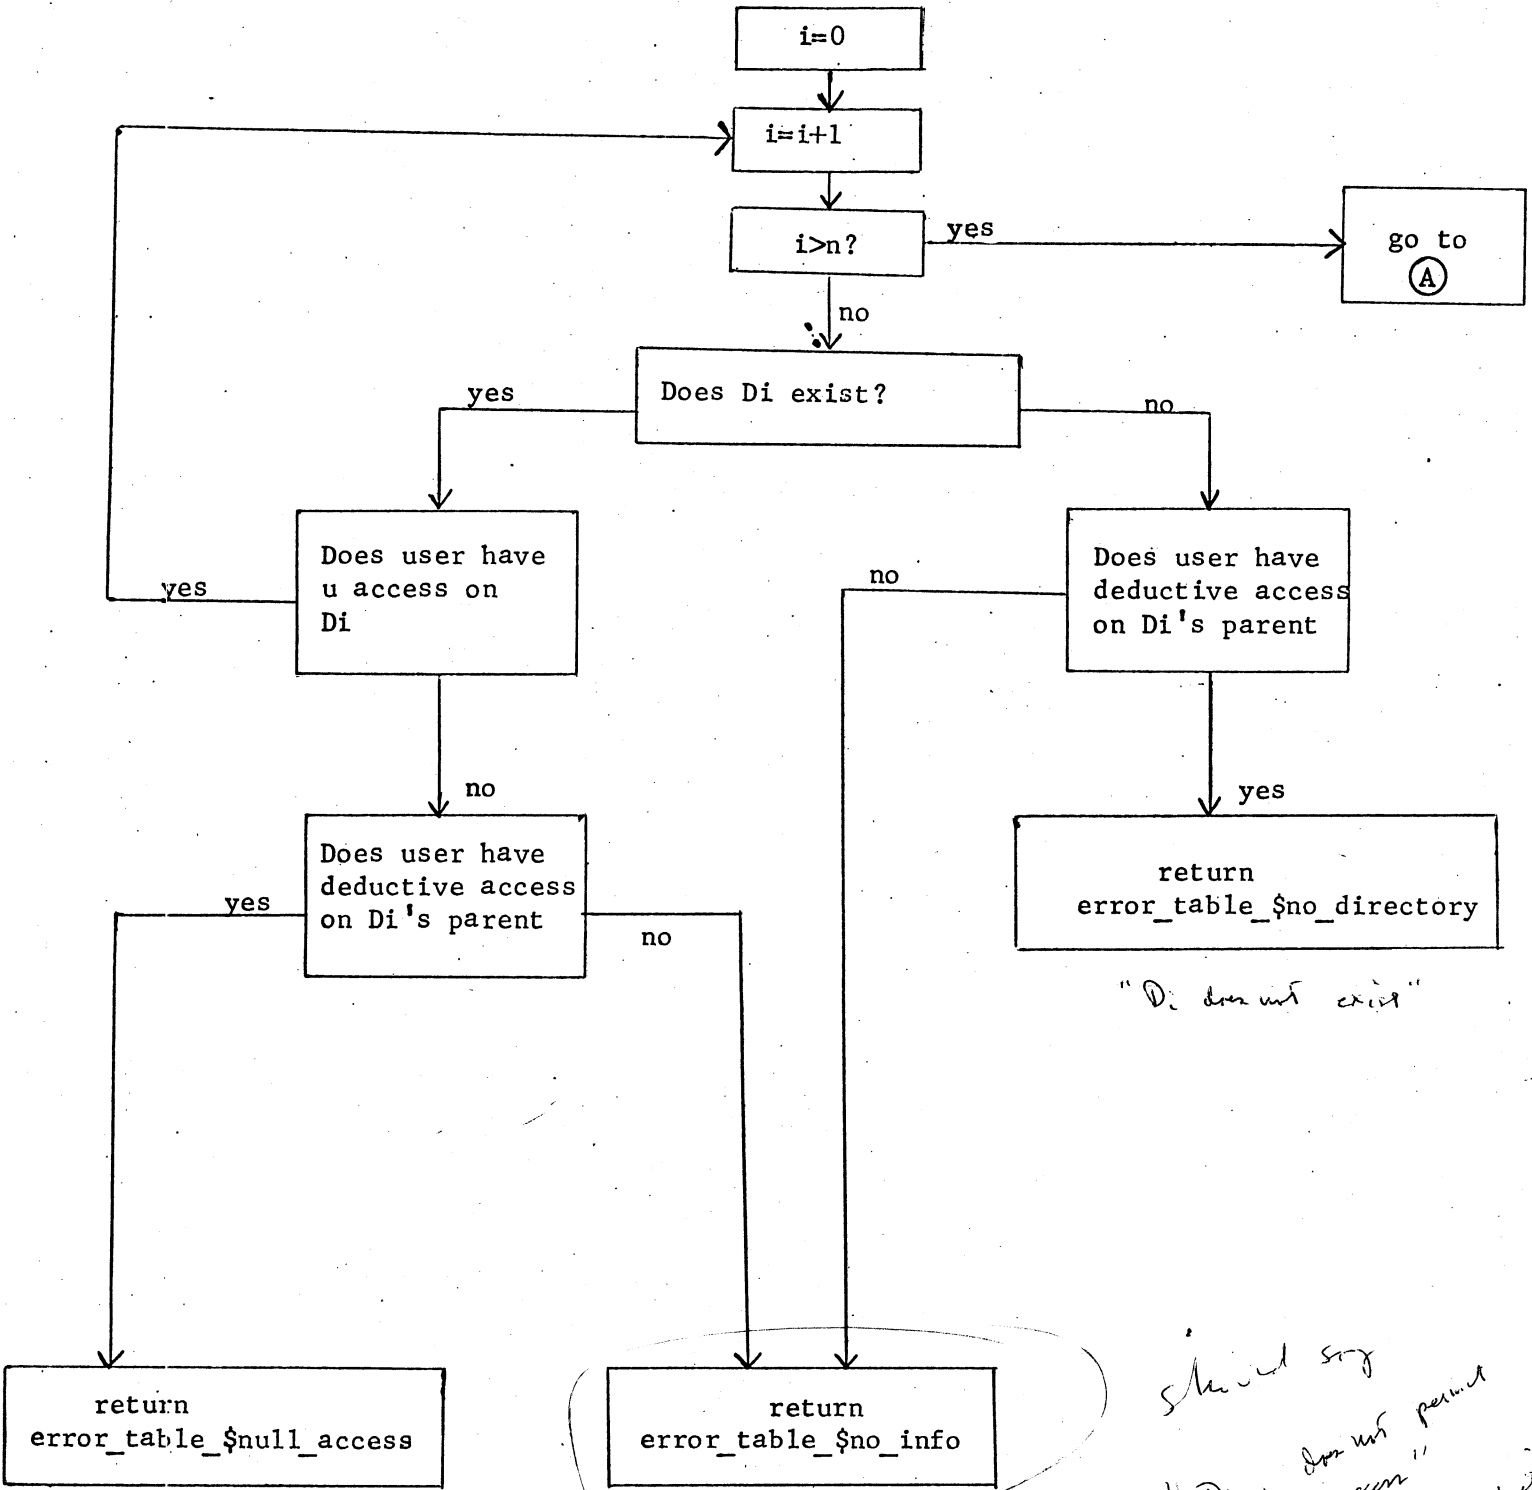
The following flow charts describe what access checks must be made by all modules of the supervisor that manipulate segments. These checks should be made when a segment fault, linkage fault, incorrect access fault and undefined access fault occurs as well as by the file system primitives.

The following principles are implicit in the flow charts. The motivation for them is that they simplify the access checking mechanism and that they give away little (if any) information that couldn't be determined by experimentation.

1. If one has non-null access on a segment (directory or non-directory) one has the right to know of its existence and one's effective access to it.
2. If one has deductive access on a directory, one has the right to know of the existence of particular entries in it and one's effective access to them, where deductive access is defined to be S, M, or some combination thereof.

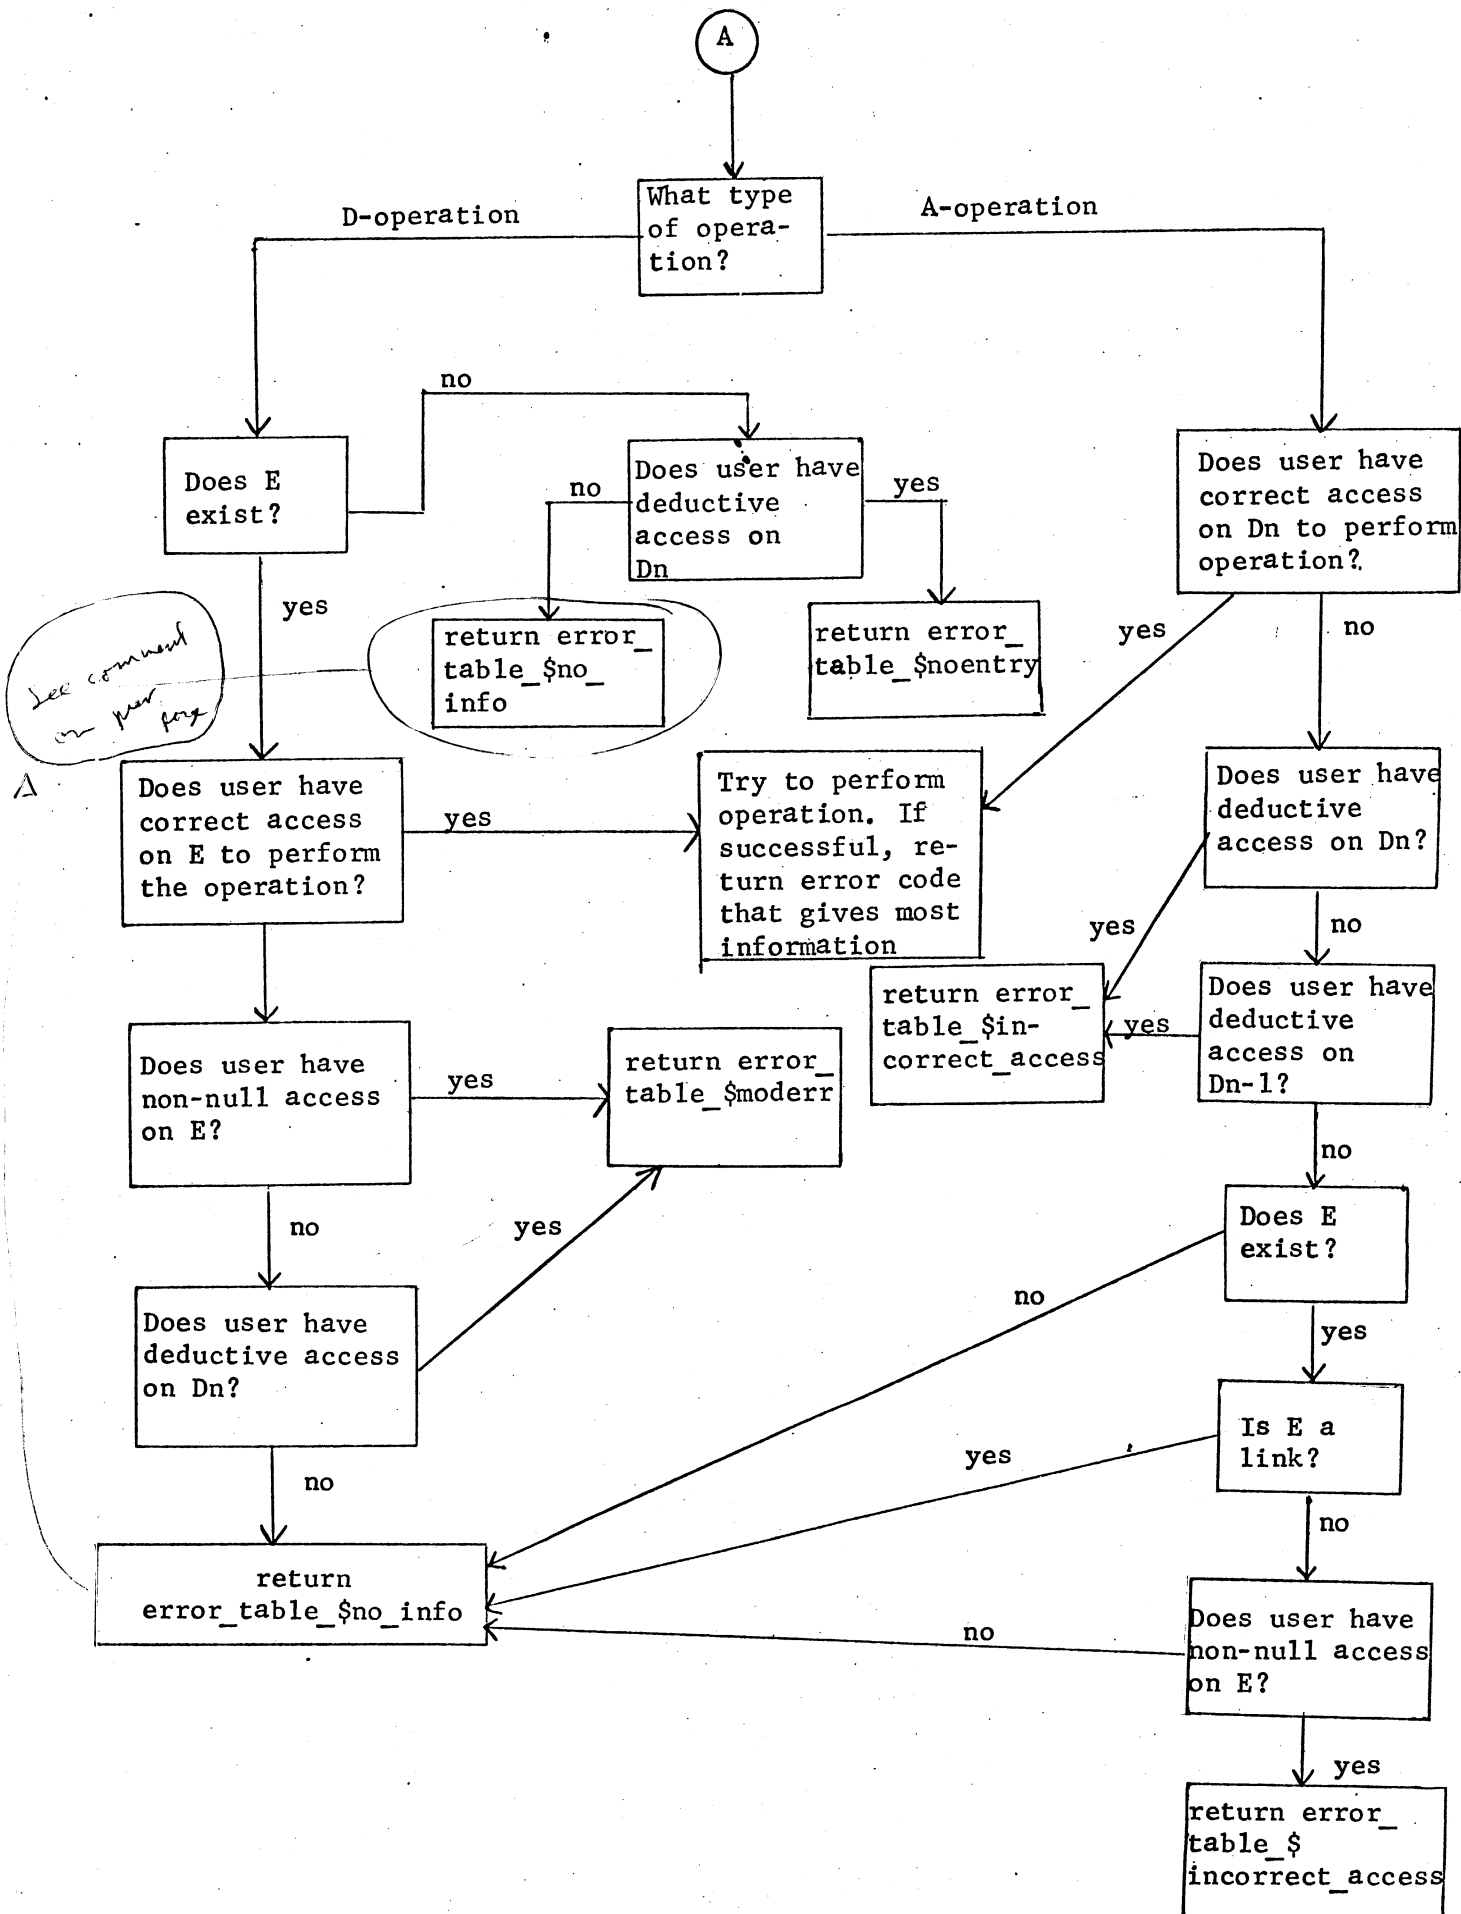
As an aside, note that these principles imply that if one has non-null access to a segment or deductive access on the directory containing the segment, the status primitive should admit the segment exists and return one's effective access to it -- even if one does not have status permission in the directory containing the segment.

Let us now consider the flowcharts:



"Di does not have u access"

Should say "Di does not permit access" where i is the first directory going back with the check in access.



Finally, we propose that every directory have two ring numbers, r1 and r2 associated with it. R1 is called the modify ring and is defined to be the highest ring in which the M attribute applies; r2 is called the status ring and is the highest ring in which the S and U attributes apply. We require that r1 be less than or equal to r2. It is possible that one would want a third ring number to control u access, but there is no obvious use for it. Therefore, for simplicity, we will stick to two ring numbers.

Consider now the access attributes of non-directory segments. They currently have the access modes read, execute, write, and append. The latter mode, append, was intended to allow a process to add data to the end of a segment but not allow modification of the data already in the segment. Unfortunately, we are not currently able to implement this mode. The append attribute is currently used to allow growing of the segment, i.e., add new pages to the end of the segment. The current use of the append mode is not well known or well used. It is primarily used to artificially set a maximum length on a segment, a feature that should be more properly implemented by adding a maximum length attribute to a segment. Since there is currently no proper use of the append attribute it should either be deleted from ACLs or it should have no interpretation, i.e., reserved for a later proper implementation.

Besides the access modes, segments also have sets of ring brackets. The current association of a set of ring brackets with a segment and a user has the disadvantage of being difficult to explain and visualize. With the current scheme a segment exists in different rings for different processes. A great deal of simplification is achieved by having only one set of ring brackets associated with a segment. This simplification causes no loss of functional capability because any accessing rights that can be granted by multiple sets of ring brackets on a segment can be achieved by having a procedure in a privileged ring simulate the access associated with the segment. This modification also solves the problem of what ring brackets are to be associated with a process not specified on the ACL. Clearly with one set of ring brackets, those are the only brackets that apply.

The current delete primitive requires both write permission on the segment and modify permission in the directory in order to delete a segment. This property has been used as a means of providing self-protection against accidental deletion of segments, i.e., if the segment does not have write permission, it cannot be deleted. This has the strange property of protecting object segments but not protecting data segments against deletion. It, therefore, seems more useful to provide an attribute which allows any segment to be protected. For this purpose the "deleteable" attribute is introduced. If the deleteable attribute is off, the segment cannot be deleted. This added protection eliminates the necessity for requiring write permission on a segment in order to delete it. Therefore, the delete primitive will require modify permission in the directory and the "deleteable" attribute being on in order to delete a segment.

The CACL is a means by which access to a group of segments can be controlled easily. Unfortunately the grouping used by the current CACL mechanism, i.e., all segments in a single directory, is not an appropriate one. It is usually not the case that all segments in a particular directory want similar access. Secondly, since the CACL is logically appended to the ACL of a segment the effect of changing a CACL upon the access to any particular segment is unclear. It depends on the contents of that segment's ACL. Thirdly, in a multiple ring situation, the rules concerning modification and use of CACLs become complex and unworkable and render the CACL useless. For these reasons the CACL is to be eliminated from Multics. The detailed arguments are given in the memo on CACLs dated June 7, 1971.

Some useful features of CACLs will be preserved. Access to large classes of segments can be modified by use of the star convention in ACL commands. Also default initial values for ACLs can be established by the use of the initial ACL.

The initial ACL is a means by which a user can specify the ACL to be added to a newly created segment in a specific directory. Each directory will contain two sets of initial ACLs, one for newly created directories and one for newly created non-directory segments. Each of these two sets will contain an initial ACL for each ring.

Each initial ACL will consist of a list of star names, i.e., file system entry names which may include the star ("*") character to indicate a class of names. With each star name will be associated a list of ACL entries. When a new segment is created via a call to append, the appropriate initial ACL will be found by using the type of the segment (directory or non-directory) and the current validation level. The list of star names is then searched for the first such star name that matches the name on the new segment. The list of ACL entries associated with this star name is then used to form the ACL of the new segment. The ACL entries specified in the call to append are then added to the new ACL.

New primitives and commands will be provided to manipulate initial ACLs. Separate commands will be provided to set entries (add or change), list entries, and delete entries for both initial ACLs applying to directories and non-directory segments. The validation level at the time of the operation will determine which ring's initial ACL is involved. When an entry is added to an initial ACL it is checked to make sure the specified user id is valid. This guarantees that the initial ACL can be validly added to a new segment with no possibility of error. Star names will be ordered on the initial ACL in a manner similar to the way process group id star names are ordered on ACLs. The most specific names will be listed first in a manner that favors specificity in leftmost components. Rather than state the precise algorithm the following listing will indicate the ordering.

a
x.y.z
a.b.c.d
a.*.b
a.**
*.a.b
..a
**

The file system currently supports a feature called extended access, the purpose of which is to allow subsystems a convenient way of specifying access attributes other than the standard attributes, on segments which those subsystems manage. Extended access is implemented simply as a set of bits in each ACL entry which the subsystem may set and interpret as it pleases; the file system does not interpret these bits in any way. Currently, extended access is used only by the message segment primitives as a means of specifying access to message segments.

Directory control is another logical candidate as an application of extended access. Currently the directory access attributes occupy the same bits as the standard segment access attributes and the directory ring brackets also occupy the same storage as the standard segment ring brackets. This duplicity of use has led to the unfortunate result of unwanted similarity between non-directory segment and directory segment access attributes. In order to allow full independence in the selection of directory access attributes, and make directory control and segment control more independent, the directory attributes should be handled separately from segment access attributes. The obvious separate mechanism is that of extended access. Directory segments would have standard access which would be rw for all users on the ACL with ring brackets of 0,0,0 on the segment and would also have some form of the sum directory attributes as extended access. As stated above this has the significant advantage of permitting the segment management facility to compute access in a uniform manner for all segments without having to special case directories as is presently done.

As currently implemented extended access provides extra bits for access attributes on each ACL entry. However, the proposed modification of making ring brackets a per segment rather than per ACL entry attribute means that some standard access attributes are no longer in the ACL entry. For consistency, therefore, an extended ring bracket field should be added to each branch as an extension of the standard ring bracket field in order that subsystems using extended access can treat ring brackets in a manner akin to the standard access.

Multics currently supports several service processes known as daemons. The daemon processes now includes backup, dumper, retriever, io, and translator. It is expected that sometime in the future this list will change to include backup, dumper, retriever, io, and offline segments. Each of these daemons require access to the particular part of the file system hierarchy upon which they are currently operating. This access is currently achieved by having these process's ids appear on all ACLs upon which they will have to operate. This current scheme forces users to be aware of the daemon accessing problems. This is unnecessary because the daemons are system processes and can be programmed so as to not reference segments unless they have a legitimate reason to do so.

Daemons can be split into two categories: those which perform operations at the explicit request of a user and those which perform operations without the knowledge of the user. The former category includes io, offline segment, and retriever which perform these specified operations only after receiving an explicit user request either by a command or by written form. In these cases the daemon only needs to verify that the user making the request has appropriate access to the segments in question. The daemon process itself can have access to all segments in the file system since it will, if coded properly only reference those segments implied by the explicit request and to which the requestor has appropriate access. The latter category includes backup and dumper. These daemons should, in theory, have access to all segments in the file system since they are supposed to reference all segments. If the Multics backup system were totally transparent and secure, these daemons should have free access to all segments, however, we will acknowledge that this is not the case by allowing users to inhibit the backing up of certain segments by turning off a backup attribute in the branch of that segment.

Wash

Essentially what is being proposed here is that all daemon processes potentially have full access to all segments and that the processes themselves have the responsibility for making sure they are legitimately referencing the segments. This has the advantage of isolating the facilities and services provided by these daemon processes from their implementation, i.e., that they are separate processes. The naive user should not have to be aware that the backup daemon needs access to his files or that he must grant some process called the IO Daemon access to his files in order to dprint them. The issuing of a dprint command or a retrieval request is all that should be necessary.

Daemon processes, for very practical reasons, should not have blanket new access to all segments in the file system. This could lead to chaos should there be bugs in the implementation of a daemon process. Daemons could simply be given the right to specify their access to specific segments or perhaps reference segments through special primitives. This would minimize the potential damage in case of bugs.

SPACE

Algorithm for determining which of two star names, x or y , should appear first on lists of initial ACL.

