TO: Distribution

From: C. D. Tavares

Date: 12/01/77

Subject: Proposed Device and Volume Management Facility

INTRODUCTION

This document provides an overview of a proposed Device and Volume Management Facility. It is an attempt to describe the Device and Volume Management Facility in its ultimate form, and does not address either the questions of implementation detail or interim partial implementation of the Device and Volume Management Facility needed to satisfy product calendar requirements.

SCOPE

The proposed Device and Volume Management Facility provides a method for controlling, allocating, and accounting for:

1) Devices such as tape drives, disk drives, printers, ounches, card readers, MPC's, Network pseudo-channels, terminal channels, and major system modules such as CPU's.
2) Media and other volumes such as tape reels, disk packs, paper stock, print trains, Network hosts, and storage system logical volumes.

The proposed Device and Volume Management Facility provides the following features:

1) Registration facilities for devices/media
2) Access control for devices/media
3) Media verification facilities
4) Automatic runtime device/medium description facilities
5) Dynamic acquisition/release of devices/media
6) Device/media accounting facilities

Registration Facilities

o Preregistration of devices and volumes

All controllable devices and media which are to be available

Multics Project working documentation. Not to be reproduced or distributed outside the Multics Project.

to any user must be pre-registered on the system by a site administrator. Registration is simply the process of introducing (making known) the existence of a new device or volume to the system. The registration information will contain such specifications as the physical location of the device or volume, an AIM access class or access class bracket within which the use of the device or volume is allowed, and certain attributes of the device or volume (which are described more fully under "Device and Media Description.")

o PCP Registry

Registration information for all devices and volumes will be stored in a central database in ring 1, and will be accessible to normal users only through controlling procedures in ring 1. Typical items which are kept in this database include all the items supplied at registration time, the accounting owner, (1) the pathname of the ACS, a DTU and DTM, a volume UID for the volume dumper mechanism, a user comment field, and so on. Users will be allowed to alter a selected few of these items subject to these intermediary modules. Device and volume management will be performed in ring 1 by the user's process according to information stored in the Registry.

There will be a separate database containing volume management parameters. This database will be installed in the manner of other system databases (such as PDT's.) It will contain names for all device and volume types (e.g. "printer", "pack", "tape_drive") defined on the system, names for all the possible attributes for each device or volume type (See Device and Media Description), and other information about attributes needed to perform attribute matching as described in a later section.

Access Control

o Access to device or volume via ACS segment

The access control of all devices and volumes will be computed using both the access on an Access Control Segment (ACS) for that object and certain limiting information kept in the Registry.

o Methodology for Determining Access

The pathname of the ACS for any given device or volume will be contained in the Registry. Access to the device or volume

---

(1) The "accounting owner" is the person who gets charged a flat rate every month or so for the privilege of having the resource (e.g. one of the site's tape slots) for his exclusive use.

represented by the ACS will be computed partly from the ACL and ring brackets of the ACS. (If no ACS exists, the default ACL is "rew" to the accounting owner.) The allowable AIM class or range must be factored in also; however, because the AIM class of an ACS which is in the control of a user can be changed at will by deleting and re-creating the ACS, it cannot be trusted. Therefore, the Registry will contain the AIM class or range which is applicable to the device or volume.

Having read and write effective access (after ring brackets and AIM have been factored in) on the ACS will mean the obvious things with respect to the device or volume. There do not seem to be any reasons why a user should have to have write as well as read to a volume to mount it for read-only; however, if a device is in a less secure room, secure information can be transmitted to unauthorized onlookers via odd methods such as tape or head motion if we allow this generality for devices as well. Therefore, a user will still require "rw" effective access on a device to assign it to a process.

Executive ("e") access to the ACS confers on the possessor the power to alter certain selected attributes of the device or volume (more on attributes later.)

o AIM considerations

Because of the requirement that objects be able to carry a range of potential AIM access classes rather than a single access class, certain other restrictions are necessary to prevent compromise of data within any such bracket when using a volume as an intermediate repository for the information. When a user dynamically acquires (1) a resource, his current AIM class is noted in the registry and becomes the current AIM class of the resource for the duration of the acquisition. This attribute prevents a process from writing secure data onto a volume at the high end of the potential AIM bracket, then logging in at the low end and reading the data back in. The only case in which the current AIM class information on a volume will instead be allowed to be an AIM range is if the volume is a storage system physical volume, for which the storage system provides the proper AIM protection. Devices which are not dynamically acquirable (for instance, the accounting owner of most tape drives will be the Initializer for accounting purposes) the current AIM class information may also be a range.

o Protection for "discarded" data

---

(1) The word "acquisition" is used to distinguish between contracting with the system to become the accounting owner of a resource, and "assignment", which is a guarantee to the exclusive ability to physically use a resource for a block of time.

A notification will occur whenever a volume is released (the opposite of acquired) that will inform the operator of that fact, so that the volume may be degaussed. An option will be provided that will lock all released volumes in an unavailable state so that they may not be automatically acquired by any other user until the operator has explicitly indicated that the volumes are ready to re-enter the free pool.

Media Verification Facilities

o Label checking

All media will be checked for labels, including media which are "known" to be unlabeled. This label checking is for the purpose of detecting mounts of incorrect volumes. Label matching, verification, and authentication facilities will be very similar to those currently provided in MR6.0.

o Manual label authentication

The site should be able to specify that every label mismatch authentication must be explicitly typed (i.e. an authentication of "***" will not be permitted.)

o Mandatory label types

A site should be able to register volumes with mandatory label types. Although it may not be possible to prevent a determined user from creating arbitrary labels on I/O volumes, once he has done so the volume will become impossible to re-mount. For instance, a site may choose (for security purposes) not to register any volume with a potential attribute of "unlabeled". Also, a volume executive may set one of the potential label attributes as current and protected (as explained below), forcing all use of his volume to be performed in that mode.

Device and Media Description

o Potential attributes of a device or volume

The site administrator will be able to specify devices and volumes as possessing a set of potential attributes. These attributes may describe physical properties of a device or volume (such as certified to 6250 BPI, or mountable on an MSU450), logical properties (such as labeled in ANSI format, or storage system pack), or connectability (accepts removable print trains.) An item such as a charge-type for accounting purposes, although included in the Registry, is not considered an attribute. For instance, the administrator can specify that a certain block of

tapes can be recorded at 556, 800, and 1600 BPI, but not 6250 BPI. Most common and useful attributes will have system-defined names. However, sites can, if they choose, extend the attribute set in a tabular manner. For instance, a site could define both tapes and drives with the attribute "building_3", and require that any reel with that attribute be mounted on a drive also possessing that attribute (more on attribute matching, below.)

o Grouping of attributes

Certain mutually-exclusive potential attributes will be grouped, such that turning any one of the group on causes the others to be turned off (e.g. 556BPI, 800BPI, and 1600BPI.) The site administrator may define the groupings as desired. The groupings will then be largely transparent to the operation of the mechanism, except that they will remain exclusive.

o Syntax of attributes

These attributes will appear to the user very much like terminal "modes", and the syntax for printing or changing them will be similar. This will eliminate the current problem of having to register a new control argument for every special feature of any device which the user must be allowed to specify.

o Current attributes of a device or volume

Any user with "rw" effective access to a device or volume will have the ability to set any of these potential attributes as current attributes. In some cases, I/O modules may do this implicitly for the user; or the user can explicitly specify new current attributes. For example, the user may specify that he now wishes to use his tape as a seven-track tape.

o Protected attributes of a device or volume

The device or volume executive will have the ability to set any of the current attributes as protected attributes. If an attribute is not protected, I/O modules and other programs will be able to implicitly change that attribute. If it is protected, explicit user action (by a user possessing executive permission) must occur to change the state of that attribute. For instance, if a tape has the seven-track attribute current and an attempt is made to mount the tape on a nine-track drive, some (currently unspecified) intermediary procedure will attempt to set the tape's nine-track attribute. If the seven-track attribute was not protected, it will succeed. If it was protected, a diagnostic will be printed and the mount will fail.

o Attribute matching between devices and volumes

The administrator will also have the ability of specifying

which of the current attributes of a volume require a matching attribute in the description of a device before that volume may be mounted on that device. For instance, he can specify that any tape with the seven-track attribute currently active must be mounted on a seven-track tape drive. This not only catches device/volume mismatches, but for example allows the proposed reservation mechanism to realize that a user reserving a private disk logical volume consisting of two physical packs also requires two spindles of the right model which are allocated to storage system use (rather than I/O use.)

Dynamic Acquisition

Automated media acquisition should be possible. That is, rather than having a system administrator intervene whenever a user wishes to acquire a volume (or release one), a user should be able to request a volume directly from the system. The user may specify in which directory the ACS will be found (it will not be automatically created.) The user may additionally specify attributes which the volume chosen must possess (e.g., "Assign me a tape which is certified to 6250 BPI.") Sites should be able to enable or disable this facility as they choose, as well as to limit this facility to certain groups of users while denying others. The acquisition process may also be performed at registration time by the site administrator. This ensures that system tape drives cannot be dynamically acquired, and allows for arrangements between the site and a user group such that a flat rate per month may be paid for the privilege of guaranteed exclusive use of a device such as a disk drive.

Accounting Facilities

The Device and Volume Management Facility must be able to supply accounting information concerning all events which could potentially be viewed by a site as being auditable or billable. Examples would be volume acquisition, release, label mismatch, attempts to use inaccessible volumes, mounts, device assignments, and device assignment durations. Auditing and security information would be logged using current facilities. Accounting information would be transmitted to the Initializer as each billable event occurs, via the IPC mechanism. The Initializer will charge for these events according to site-defined parameters.

Also important is that the facility be able to handle, in an efficient manner, requests from users to list all their acquired volumes and devices, as well as requests from project administrators and real-time requests from system administrators and tuners, for the same type of information.