

To: Distribution

From: T. H. Van Vleck, W. S. Silver

Date: February 24, 1976

Subject: Interim version of mount and demount for disk volumes

Until the full scheme for user mounting and demounting of hierarchy volumes (storage system disk logical volumes) can be implemented, an interim scheme must be used. The full plan, described in MTB-229, will involve modifications to RCP to know about two new resource types, code in ring 1 to allocate disk drives to user requests analogous to that now used for tape, several new user commands, and new volume registration commands and operator commands. The full details of the eventual mechanism's user, librarian, and operator interfaces will be described in several forthcoming MTB's.

The interim mechanism is much simpler. From the operational point of view it appears to be a slight extension of the current operator commands for system startup. Since RCP is bypassed completely, the initial facility will lack any mechanism for allowing a user to await the mounting of a hierarchy volume; if a requested volume is not mounted, his call will fail immediately. When the user requests the mounting of a volume via telephone, or such mounting is scheduled, the operator must select a free disk drive and mount the pack, and then must type a command to the system indicating that the volume is mounted. He then tells the user, "try it now," and the user issues the mount command.

The "virtual mount" described in MTB-229 will be implemented. This change insures that correct access control discipline is obeyed by preventing a process from accessing segments on a volume unless the volume is public or the process has the volume id in its KST. The user's access to the hierarchy volume will be checked in ring 1 when the mount request is issued, and the virtual mount will be done from ring 1 if the user has access according to the volume registration data. Registration data will be checked for all volumes except the RPV; during a cold boot, the operator will register other volumes before using them. The registration data will not be the full data base which will eventually be implemented for RCP. One segment per hierarchy volume will be maintained in a directory under >system_control_1 (the root hierarchy volume's registration segment will reside in the root directory). These segments will contain enough information to allow ring 1 to check that the

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

hierarchy volume is correctly mounted: that is, public/private switch, AIM information, and a list of the physical volumes and their unique ID's. Every private volume must also have an Access Control Segment (ACS), which will be linked to from the system directory. This segment is a (possibly zero length) segment in a location specified by the volume owner. The ACL of the ACS is interpreted as specifying access to the hierarchy volume contents.

STEPS IN MOUNTING

To get a hierarchy volume virtually mounted to his process, a user first contacts operations and asks that the volume be mounted. This request may be conveyed via telephone or send_message, or it may be implicit in a schedule established by the system administrators.

Operator Preparation

When the operator decides to mount a volume as a result of such a request, he may choose a free disk drive if one is available, or he may use the following command to force a mounted hierarchy volume to be demounted:

```
demount_force <hvname>
```

This command will cause one or more physical volumes in use by the storage system to be shut down in an orderly fashion. All active segments on the volumes will be deactivated, and the label, VTOC, and free map for the volume will be updated. Users who attempt to use segments on a volume which has been demounted will encounter a seg_fault_error condition with the message "Volume not mounted." The supervisor will also print a message of the form

```
DEMOUNTED DSK7_04
```

giving the disk drive name.

When the operator has sufficient free disk drives to mount the requested hierarchy volume, he performs the physical mounting operation for each pack. Special interrupts generated by the disk units becoming ready will be ignored. The operator then types the following command for each physical volume he has mounted:

```
add_volume <pvname> <drive-name>
Example: add_volume pack32 dsk7_04
```

This command directs the initializer to call ring 1, where the following steps are taken:

- a) The combination <drive-name> is looked up in the disk_table, and that table entry is checked to make sure it is available for storage system use and currently free.
- b) The disk_table is also searched to insure that <pvname> is not mounted on some other drive.
- c) The registration information for <pvname> is located. Each per-hierarchy-volume registration segment has additional names of the form "pv.<pvname>" added to it for every physical volume in the hierarchy volume. If the physical volume is not registered it cannot be mounted.
- d) Ring 1 now calls the hardware to read and check the label of the pack on <drive-name>. The physical and hierarchy volume names and ID's and the AIM attributes are checked to make sure that the label matches the registration data. If everything matches the volume is accepted for paging and entered into the PVT.

When the operator has mounted all physical volumes he then issues the following command:

```
mount hvol <hvname>
Example: mount hvol student3
```

This command causes the initializer process to call ring 1 to cause the following steps to be performed:

- e) The registration data for hierarchy volume <hvname> is located.
- f) For each physical volume recorded in the registration of <hvname>, the disk_table is checked to insure that the physical volume is mounted. If the physical volume is recorded in the disk_table as assumed, steps c and d above are performed.
- g) The hierarchy volume <hvname> is entered in the hardware LVT by a call to initializer_gate_\$add_hv. If the registration data says the volume is public, any user process may then use it without further ado.

User Call

For a private volume, the user must now cause the virtual mounting of the hierarchy volume for his process. To do this, he may invoke the rcp-oriented command

```
mount hvol <hvname>
```

In a later RCP implementation, many other options and subcases of the command will be possible. But for the interim version, the

command simply calls the two entrypoints

```
rcp_$mount ("hvol", ifp, event, "", id, ec)
```

```
rcp_$check_mount (id, ifp, "", ix, ec)
```

in order. These entries will be called with their final standard calling sequence, but the rcp_gate will direct these calls to interim code which does not perform all the actions which the final version will perform. The following steps are taken:

- n) The registration information for <hvname> is located.
- l) The ACS for the hierarchy volume is located and the user's effective access to the hierarchy volume is derived. If the user does not have RW access to the hierarchy volume, an error is returned.
- j) Ring 1 calls the hardware to check that the hierarchy volume is in the LVT. If the volume is found this call will also enter hierarchy volume ID into the process's KST, unless the volume is a public volume.

The user process may then initiate segments on the hierarchy volume.

When the user has finished with a non-public volume, he may issue the command

```
demount hvol <hvname>
```

It will remove a user's KST item for a hierarchy volume and cause faults to be set in the SDW's of any active segments on the hierarchy volume for the user process. This operation can decrement a counter in the LVT which was counted up by the mount operation, so that the ring-1 programs can type

```
VOLUME STUDENT3 FREE
```

when the count becomes zero.

System Startup

The current temporary mechanism for system startup will be modified as follows:

- a) The DSKA command will be replaced by the add_volume and mount commands described above.
- b) The DSKG command will be eliminated.
- c) The automatic DSKG performed by certain commands such as startup, salv, and reload will be changed to be an automatic

"mount hvol root" command.

- d) If the special volume name "auto" is used in a mount command from the Initializer process, the disk_table will be scanned for volumes which are assumed to be in position but which have not yet been checked. Each hierarchy volume thus found will be mounted as described above. Installations which leave Multics running unattended can therefore place the command "mount hvol auto" in their system_start_up.ec to cause all volumes which were in use at the last crash to be rechecked and reaccepted automatically.
- e) Simple registration commands will be available that can be executed in the cold boot environment to register at least the volumes which are part of the RHV. These commands will be consistent with the commands used by the volume librarian, although the librarian commands may have more options.
- f) The initialize_disk command will be renamed initialize_volume. Only registered volumes can be initialized.
- g) When the system mounts a volume automatically because the disk_table shows that it was mounted at the time the system crashed, a registration file with default attributes will be generated if the volume appears to be unregistered. Thus, if the volume registration data is destroyed in a crash, it is reconstructed from the table of correctly mounted volumes if that data has survived.

The regular system startup procedure will thus differ from that used in 28-0 only by the detail that the system accepts hierarchy volumes other than the RHV from ring 4 rather than ring 1. More typing is required during a cold boot, since the volumes must be registered; and more typing is required after a disk reshuffle, since mount commands as well as add_volume commands must be typed.

DECLARATION OF THE VOLUME REGISTRATION FILES

The following PL/I declaration describes the structure of the interim volume registration segment.

```

dcl 1 volume_registration aligned,
    2 version fixed bin,
    2 hvid bit (36),
    2 hvname char (32),
    2 max_access_class bit (72),
    2 min_access_class bit (72),
    2 volume_owner char (32),
    2 flags,
    3 public bit (1) unal,

```

```

3 pad bit (35) unal,
2 npv fixed bin,
2 pv (0 refer npv),
3 pvid bit (36),
3 model fixed bin,
3 pvname char (32),
3 location char (32),
3 mfg_serial char (32),
3 date_registered fixed bin (71);

```

All volume registration segments except that for the RHV will reside in the directory >system_control_1>hvol. The RHV's registration segment will reside in the root directory, to insure that it is accessible while the system is coming up, before it has accepted the other volumes of the RHV.

Each volume registration segment will be named hv.<hvname> and will have additional names added to it of the forms hvid.<unique string> and pv.<pvname> and pvid.<unique string> for each physical volume. This is done so that master directory control can associate volume ID's with their volume names, and to insure the uniqueness of volume names.

The ACS for each hierarchy volume is a link in >sci>hvol with the name <hvname>.acs. The link target for all public volumes is a zero-length segment in the same directory, with an ACL of rw for *.*.*.

Similarly, the Master Directory Control File (MDCF) for the hierarchy volume will be pointed to by a link with the name <hvname>.mdcf.

NEW_OPERATOR_COMMANDS

This section gives brief descriptions of the new operator commands available in rings 1 and 4.

add_volume_registration (avr)

This command calls the rcp_llo_gate which is accessible to volume librarians and to the Initializer.

Format: avr pack <pvname> -hvol <hvname> -user <userid>

-hvol <hvname> This control argument is required. When avr is called from ring 1 the only legal <hvname> is "root".

-user <userid> This control argument is required. When avr is called from ring 1 the only legal <userid> is "system".

initialize_volume

This command writes a label and an empty volume map and VTOC onto a disk pack. It consults the volume registration for the volume to obtain the hierarchy volume information and the physical volume unique ID. It also checks the label of the volume and will refuse to label a volume if it appears to have a valid label for a registered volume.

Format: initialize_volume <pvname> <drive-name>

When initialize_volume is called from ring 1 only volumes of the RHV may be initialized.

-special If this control argument is specified, the system will ask the operator for request lines which may specify average segment length and partition definitions. The valid requests are:

part NAME low nrec	Define a partition on the volume at the low end.
part NAME high nrec	Define a partition on the volume at the high end.
avg fff.ff	Declare the average segment size to be ffff.ff records. (The default is 4.1.)
list	List the attributes of the volume.
quit	Exit without doing anything.
end	End of specifications; initialize the volume.

add_volume (addv)

This command is issued to inform the system that a volume is mounted and ready on a specified disk unit.

Format: addv <pvname> <drive-name>

No control arguments are allowed.

mount

This command is issued to inform the system that a hierarchy volume is completely mounted.

Format: mount hvol <hvname>

demount_force (dmf)

This command forces the demounting of a hierarchy volume.

Format: dmf <hvname>

One or more physical volumes will be demounted.

IMPLEMENTATION

This section gives a summary of the programs which must be modified or written for release 4.0.

1. Modifications to RCP.

- a) Fix rcp_device_info_ to accept device names of the form "diskX_01".
- b) Install new version of rcp_init_disk_sharing_ which respects the flag pvt.storage_system.
- c) Fix rcp_disk_ to read the label of IO disks and refuse to work on storage system packs except for privileged mount requests.

2. User commands.

mount
demount

3. User subroutines.

rcp_\$mount
rcp_\$check_mount
rcp_\$demount

4. Operator and Librarian commands. mount

demount_force
add_volume
del_volume
add_volume_registration
del_volume_registration
change_volume_registration
list_volume_registration

5. Operator and Librarian subroutines.

rcp_sys_\$demount_force
rcp_lib_\$set_volume_registration
rcp_lib_\$copy_volume_registration

system_startup_
rcp_vol_data_name
rcp_vol_data_uid
disk_table_
initializer_admin_

6. Hardcore

LVT manager
various checks that volume is mounted or public

To: Distribution
From: Robert S. Coren
Date: 02/25/76
Subject: New Strategy for Conversion of Terminal Input

INTRODUCTION

MTB 234 described a new method for processing terminal output in ring zero making extensive use of EIS. The design described has since been implemented in MIT system 27-6, and will be part of Multics Release 3.1. The new implementation shows approximately a threefold improvement in the efficiency of `tty_write`, measured in terms of the virtual CPU time spent in `tty_write` for each character sent to the 355; on MIT, about 1% of total time charged is now spent in `tty_write`, as compared to about 2.5% in pre-27-6 systems.

The current implementation of the ring-zero input-processing module, `tty_read`, has essentially the same problems as those described in MTB 234 for the old `tty_write`: characters are processed one at a time, even in "raw" mode; translation, canonicalization, and escape processing are handled simultaneously and driven by a single table; fixed tables in ring zero are used, pointers to which are constructed on every call. In addition, canonicalization is mishandled in some cases, as indicated in MTB 251; and the "prescan" function, which is intended to examine input for case-shift characters and to update the current column position for use by `tty_write`, is invoked at the wrong time and is therefore unreliable.

This MTB proposes a redesign of `tty_read` along the same lines as the recently-completed redesign of `tty_write`. Character-by-character processing is abandoned in favor of separate phases using PL/I builtin functions and ALM subroutines coded with EIS; canonicalization is reimplemented so as to conform to the rules set forth in MTB 251; the "prescan" function is removed from `tty_read` altogether, and its equivalent added to the 355 software (as described in a separate MTB). A version of `tty_read` implementing this design is intended for Multics Release

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

4.0.

One incompatible change that is being proposed is to discard all "invisible" characters (i. e., control characters that do not involve carriage or paper motion) whenever the channel is in "can" or "erkl" mode. The motivation for this proposal arises from these characters' invisibility: they do not show up on most terminals, and their retention violates the principle of canonicalization, that the contents of a line of input depend on its physical appearance. In other words, there is no way to distinguish visually between a#b and a<ETX>#b; what does the # erase? What column position does the ETX occupy?

The ability to input such characters directly (i. e., rather than by using octal escape sequences) seems to be of limited utility. The one exception might be the desire to use such a character as a kill or erase character; there are systems in existence which use CAN (octal 030, input by typing <CTL>x), as a kill character. User-replaceable kill and erase characters are planned for the future; it would not be too difficult to arrange not to throw away control characters which were being used for a Multics-defined purpose. For the present, a user employing special characters for erase and kill must process them in the user ring, and accordingly would not be in "can" or "erkl" mode. In addition, since the elimination of control characters would be a translation function (see below), user-substitutable translation tables (also a planned future improvement) would allow a user to admit selected control characters at will. In any case, all possible 9-bit patterns can be input as octal escapes.

One implication of this change is that the special meaning of the ESC character (octal 033) is eliminated for input. This character has been used primarily to insert ribbon-shift characters; this can be done by using the octal escape sequences \016 and \017.

PROPOSED_NEW_DESIGNOverview

The obligation of `tty_read`, when called through `hcs_`, is to return in a caller-supplied buffer either 1) as many characters as the caller specified; 2) all characters up to and including the first "break" character present in ring-zero buffers for the specified channel; or 3) all characters remaining in the buffers

for the specified channel, whichever is fewest. The "break" character is by default a newline character; there is currently no way to change this, but future modifications may permit it.

Certain transformations may be performed on the characters typed by the user, such as reduction to canonical form, removal of "erased" and "killed" characters, and the interpretation of escape sequences. The application of these transformations depends on both the modes associated with the channel and the contents of certain tables which are available to `tty_read`.

The functions of `tty_read` may be divided into the following phases:

1. Copying raw input data from `tty_buf`, and freeing the ring-zero buffers;
2. Translation to ASCII
3. Canonicalization of the contents of column positions
4. Erase_and_kill_processing
5. Escape_sequence_processing

Clearly, these five phases are not always necessary. Phases 3, 4, and 5 depend on "can", "erkl", and "esc" modes, respectively; in "rawi" mode, only phase 1 is required.

For convenience and to ensure consistency, conversion (the generic term used here for the relevant subset of phases 2 through 5) is done on all characters up to and including the first break character in the input buffers, whether or not the break character is found within the limit specified by the caller. This avoids the possibility of terminating conversion in the middle of an escape sequence or of a line that is subsequently killed, and also allows for the possible shrinkage of the input string (through the deletion of extraneous white space and the condensation of escape sequences, for example). "Extra" characters thus converted (i. e., those that cannot be returned because the caller has not provided sufficient space) are saved in reallocated buffers in `tty_buf`; these buffers are marked with a "converted" flag and chained to the head of the channel's input chain so that they can be picked up by the next call to `tty_read`. In two exceptional cases, conversion cannot proceed to the first break character: the first is, obviously, when no break character is present; the other is when the size of

tty_read's internal automatic buffers is exceeded. For reasons that will be explained later in this document, both these cases are expected to be very rare.

Reference is made in the course of this document to entries in the subroutine tty_util_, which is described in MTB 234. A new entry, tty_util_\$tct, has been added; it performs the same function as tty_util_\$find_char, except that it checks neither for characters with their high-order bits on nor for combinations of white-space characters.

The remainder of this document consists of the following:

1. A few remarks on the management of tty_read's internal buffer space;
2. A more detailed description of the five conversion phases mentioned above;
3. A description of the modifications required to the data structures described in MTB 234;
4. Module descriptions of the new column canonicalization routine, tty_canon (which replaces the old tty_con), and the new entry tty_util_\$tct.

Familiarity with the material in MTBs 234 and 251 is assumed throughout.

Space_Management

During conversion, intermediate forms of the input string result from each conversion phase; for the storage of these intermediate strings, two buffers are maintained in tty_read's automatic storage. Clearly this sets an upper limit on the allowable length of the input string. The normal limiting factor, of course, is the presence of a break character, and input lines longer than 100 characters are rare; a further limitation is imposed by the 355 software, which takes a channel out of receive mode if more than 600 characters are input without a break character. The input string can grow during canonicalization through the replacement of carriage returns by multiple backspaces, but this occurrence too is rare. All in all, a buffer size of 720 is very unlikely to be exceeded.

Consequently, no more than 720 characters are copied into the internal buffer from `tty_buf`. If the canonicalization phase attempts to increase the length of the string past 720, `tty_read` will start again from the beginning with a limit of 480 characters to be copied. This limit is entirely safe, since canonicalization cannot increase the length of the string by more than 50%. Because of the remote possibility that this restart may be necessary, buffers in `tty_buf` from which input characters have been copied cannot be freed until after the canonicalization phase is completed.

Since conversion is, if possible, carried out on all characters up to and including the first break character, the final converted string may be larger than the buffer provided by the caller. If this is the case, enough characters to fill the caller's buffer are returned; the remainder of the converted characters, as indicated above, are saved in buffers in `tty_buf` in each of which a "converted" flag is set. In addition, if one of these buffers contains a break character (the last one generally will), a "break" flag is set in that buffer. These buffers are added to the head of the chain of unconverted input buffers (the "read chain"), and the input pointer in the control block associated with the channel is set to point to the first "converted" buffer.

Copying

IN "rawi" MODE

The copying phase in "rawi" mode is very simple. Characters are copied from `tty_buf`, starting at the head of the read chain, directly into the caller's buffer, until either the caller's buffer is filled or the read chain is exhausted. Any buffer from which all the characters are thus copied is freed.

NOT IN "rawi" MODE

If there are any "converted" buffers at the head of the read chain, characters are copied from these buffers directly into the caller's buffer until either the caller's buffer is full, a break character has been copied, or the chain of converted buffers is exhausted. (In general, the last converted buffer contains a break character, and non-last converted buffers do not.) Any converted buffer from which all the characters are copied is

freed.

If there are no converted buffers, or the converted buffer chain is exhausted without encountering a break character or filling the caller's buffer, characters are copied from the unconverted read chain (if present) into the first of `tty_read`'s automatic buffers, until either a break character is encountered, the read chain is exhausted, or the internal buffer is filled. Buffers are not freed at this time, for the reason given above under "Space Management."

Because the 355 does not normally send input to the 6180 until a break character is typed, the read chain almost always ends with a break character. (Consequently, the converted chain usually does, too.) It might not if there was a quit on a channel not in "hndlquit" mode (in "hndlquit" mode the read chain is discarded on a quit), or if the channel exceeded the 355 software's 600-character limit.

If any characters were copied from unconverted buffers, conversion of the contents of `tty_read`'s automatic buffer begins.

Translation

If a translation table exists for the terminal type associated with the channel, it is used in a call to `tty_util_$mvt` to copy the characters from one internal buffer to the other, simultaneously translating it to ASCII. Translation is required for IBM-type terminals using either EBCDIC or Correspondence character codes; it is also used to translate capital letters to lowercase for uppercase-only terminals such as a Teletype Model 33. (Escaped letters will be changed back to uppercase by the escape-processing phase.)

The translation phase does not have to deal with case-shift characters. Under the new design, the 355 is responsible for recognizing case shifts, and for turning on the 100(8) bit in all uppercase characters (characters on shifting terminals are only six bits). All that is necessary on the 6180 side is a translation table that includes characters with the "100" bit on

and translates case-shift characters to ASCII NUL characters.

If the channel is in "can" or "erkl" mode, a further translation is done using a general table which translates "invisible" characters (see above) to NUL (all zero) characters. NUL characters are subsequently discarded by the canonicalization phase.

Canonicalization

Column-position canonicalization takes care of itself unless the input string contains leftward carriage motion, i. e., backspace and/or carriage return characters. In addition, backspaces and carriage returns at the left margin or immediately preceding a newline are discarded. In other cases, canonicalization must be performed in accordance with the rules given in MTB 251.

The canonicalization phase therefore begins by searching the internal buffer (using the PL/I "search" builtin) for a left-motion character (carriage return or backspace). If the first character is a left-motion character, the buffer pointer is advanced by one character, the string length is decremented by one, and the new string is searched as before. If a left-motion character is found, a verify builtin is used to discover if the rest of the line consists of white space (backspaces, carriage returns, spaces, horizontal tabs, or NULs) followed by a newline. If this is the case, the string length is reduced to the result of the search, and the newline is copied to the new end of the string. If a left-motion character is discovered in any other position, tty_canon is called to perform column canonicalization.

The subroutine tty_canon is a revised version of the old tty_con, and uses the same basic algorithm: store each printing graphic from the input string in an array along with its correct column position; sort the array by column position, and by character within each column position; restore the characters to the input string location in the resulting order, inserting backspaces and spaces as appropriate. Tabs must be treated as a slightly special case of printing graphic, so that tabs which are in no way overstruck are preserved but others are replaced by spaces.

A module description of `tty_canon` appears at the end of this document; the calling sequence has been modified so that the module could theoretically be called with an arbitrary string in other environments than that of the ring-zero typewriter DIM. The resulting calling sequence is still not ideal, as it contains arguments that are both input and output; this approach is retained for reasons of efficiency. Eventually, an essentially equivalent module can be implemented in the user ring.

The structure used for the elements of the sorting array makes the sort very easy, thus:

```
dcl 1 column_array (max_size) aligned,  
    2 column fixed bin (17) unaligned,  
    2 erase bit (1) unaligned,  
    2 kill bit (1) unaligned,  
    2 vertical bit (1) unaligned,  
    2 pad bit (5) unaligned,  
    2 not_tab bit (1) unaligned,  
    2 char char (1) unaligned;
```

The "erase" bit indicates an erase character; the "kill" bit indicates a kill character; the "vertical" bit indicates a non-newline character requiring vertical carriage motion (i. e., vertical tab or form-feed); the "not_tab" bit is on for any character except a horizontal tab. It can be seen that by treating each element of the array as a single value for the purpose of sorting, the characters automatically come out in column order and in character order in each column, except that: 1) an erase character will always be the last character in its column position; 2) a kill character will be last in its column position unless overstruck with an erase character; 3) a horizontal tab will always be the first character in its column position; and 4) a vertical-motion character will follow all characters other than an erase or kill character. Since during the initial scan, a vertical-motion character causes both the "current" column and the "starting" column to be set to the next highest multiple of 1000 (the "starting" column is the column assigned to the left margin, initially 0), a vertical-motion character cannot share a column position unless 1000 or more column positions are actually typed. A newline is assigned a column position of $2*17 - 1$ so that it will always sort to the end of the line.

Kill processing is not done by `tty_canon`; kill characters are sorted to the end of the column position to make things easier for the kill-processing phase of `tty_read`. Erase characters are only interesting to `tty_canon` if they are

overstruck; since an overstruck erase character sorts to the end of its column position, the rescan step, when it finds an erase character that is not first in its column position, deletes it and all preceding characters with the same column position.

Since a tab sorts to the beginning of its starting column position, it is sufficient to check whether the graphic following the tab has a column position less than the next tab stop; if it does, the tab is dropped, and spaces are inserted as they are whenever there is gap between two graphics. Otherwise the tab is inserted in the final string.

NUL characters are not stored in the column_array; thus tty_canon completes the elimination of "invisible" characters.

The maximum length of the input string is passed as an argument to tty_canon; if the final string exceeds this length, only max_length characters are returned, and a status code of error_table_\$long_record is returned.

Upon return from tty_canon, if the status code is zero, tty_read frees the ring-zero buffers from which characters were copied, as explained above; otherwise it resets its internal buffer size limit to 480 and starts again from the copying phase.

If the canonicalization phase completes without calling tty_canon, the string may still contain NUL characters; therefore if tty_canon has not been called, tty_read indexes the string for NUL characters, and copies the characters preceding and following each NUL into the other internal buffer, decrementing the string length by one for each NUL it finds.

Erase_and_Kill_Processing

Erase and kill processing is really done in two passes, kill and then erase. The string resulting from the canonicalization phase is indexed from the right for a kill character; if one is found, and the immediately preceding character is not a non-overstruck escape character, the pointer to the beginning of the string is incremented to point to the character following the kill character, and the length of the string is decremented accordingly. If the kill character is preceded by an escape character that is not preceded by a backspace, the pointer and the length are not changed, and the remainder of the string (if

any) is scanned for further kill characters.

The string resulting from the kill pass is now indexed for an erase character. If one is found anywhere but at the beginning of the string, the characters before and after the erased character(s) must be copied to the other internal buffer. The basic mechanism is to copy the characters to the left of the erased characters, decrement the count of total input characters by the number of erased characters plus one for the erase itself, and resume the scan starting with the character after the erase character. (If the erase character is preceded by an escape character not preceded by a backspace, the escape and erase characters are copied along with the preceding characters.) When the end of the string is reached, provided any copying has been done, all characters to the right of the last erase character are copied.

The number of characters to be erased (i. e., not copied) is determined as follows: if the character preceding the erase is "white space" (space or horizontal tab) the source string is searched backward for a non-white character, and all characters to the right of it are erased; if the character preceding the erase is a printing graphic, then the source string is searched backward until two non-backspace characters are found in succession, whereupon all characters from the one to the left of the leftmost backspace on are erased. Note that the character immediately preceding the erase character cannot be a backspace, since all overstruck erase characters are processed by `tty_canon`.

If the second or subsequent scan turns up an erase character as the first character in the string (as would happen if two erase characters were typed in succession), the determination of the number of erased characters is made in the same fashion as that described above, except that the characters at the end of the `target` string are examined; the erasing is carried out by decrementing the `target` pointer so that the erased characters will be overwritten, and decrementing the overall length accordingly.

Escape_Sequence_Processing

This phase, which is implemented in a similar manner to the formatting phase of `tty_write` (as described in MTB 234), actually deals not only with escape sequences, but with the elimination of white space before break characters and of characters designated as being "thrown away" for the current terminal type. It uses test character and translate (`tct`) instructions under control of

a table containing zero entries for ordinary characters, and indicators identifying four types of "interesting" characters: break character, escape character, form-feed, and "throw-away" character.

This phase uses `tty_util_$tct`, which scans for "interesting" characters and returns a tally of characters skipped over, the indicator value for the character stopped at, and an updated pointer to the character stopped at. If the tally is non-zero, `tty_read` copies the skipped characters into whichever internal buffer does not contain the source string; then it examines the indicator. For a break character, it scans the copied characters (if any) from the right for the last printing graphic; the break character is copied immediately to the right of it. If any intervening white space was found, the length of the final string is decremented by the number of white-space characters. Finally, a flag is set to indicate that a break was found.

If the scan finds a form-feed, and the terminal has a non-zero page length, the form-feed is thrown away, on the assumption that the user typed it for the purpose of starting a new page. Otherwise it is stored as a normal character. The interrupt handler, `dn355`, is responsible for adjusting the current line count on the page when a form-feed or newline is input.

If the indicator shows an escape character, `tty_read` must find out if it is in fact the start of an escape sequence. If the channel is not in "esc" mode, or if the character immediately preceding or either of the two characters immediately following the escape character is a backspace, the escape is copied as a normal character and the scan continues. (The backspace test is to ensure that neither the escape nor the column position to its immediate right is overstruck.) If the following character is an escape, erase, or kill character, it is copied to the target string; if it is an octal digit, the character whose value is represented by the one to three non-overstruck octal digits following the escape character is inserted in the target string; if the escape is followed by zero or more white-space characters followed by a newline, all characters from the break through the newline are skipped (the newline is not treated as a break in this case); otherwise the character following the escape is looked up in the `input_escapes` string in the appropriate `special_chars` structure (described under "Data Structures" later in this document). If it is found, the corresponding character from the `input_results` string is inserted in the target string. If the character is not found, then there is no escape sequence, and the escape character is copied as above. If an escape sequence is identified, the pointer used for the next call to

tty_util_\$tct is updated to point past the end of the escape sequence.

If the indicator shows that the character is to be thrown away, it is not counted in the length of the final string, and the scan continues starting with the following character. Note that "invisible" characters (see above) have already been thrown away by the time this phase is reached. The present default tables do not include any other characters to be thrown away; however, a user-supplied table might specify some other character which the user wishes the typewriter DIM to discard rather than returning it to the user ring.

If the first call to tty_util_\$tct returns an indicator of zero and uses up the entire source string, no characters at all are copied by this phase.

If the total number of characters in the now fully-converted string plus the number of previously-converted characters already copied into the caller's buffer is less than or equal to the number of characters requested by the caller, and the converted string ends in a break character, all the converted characters are copied into the caller's buffer, and tty_read returns. If the total number of converted characters exceeds the number requested by the caller, the caller's maximum is copied into the caller's buffer, and the remainder are placed in "converted" buffers in tty_buf as described above, to be picked up by a future call. If the total number of converted characters is less than the number requested by the caller, and the converted string does not end in a break character (either because a break character was escaped, or because the internal buffer size limit was reached), all available characters are copied to the caller's buffer and, if a read chain is still present, the next block of characters (up to the next break) is copied from the read chain and converted as above; any excess characters resulting from the latter conversion are saved in "converted" buffers as above.

DATA STRUCTURES

This section describes the modifications necessary to the data structures described in MTB 234 to make them useable for input conversion as well. Translation tables used by tty_util_\$mvt and tty_util_\$tct are similar to those used by tty_write, and, like them, are kept in ring zero by terminal type; future modifications will allow a user to specify his own

version of one or more of these tables.

Default_Table

The default table has been expanded and rearranged slightly, and the names of some of the items have been changed. The new format is shown below:

```
dcl 1 device_defaults aligned based,
  2 flags unal,
    3 shifter bit (1) unal,
    3 upper_case_only bit (1) unal,
    3 pad bit (7) unal,
  2 delay_char char (1) unal,
  2 upper_case char (1) unal,
  2 lower_case char (1) unal,
  2 delay_offset (4) fixed bin (18),
  2 output_tct_offset fixed bin (18),
  2 output_mvt_offset fixed bin (18),
  2 special_offset fixed bin (18),
  2 input_tct_offset fixed bin (18),
  2 input_mvt_offset fixed bin (18),
  2 break_char char (1) unal,
  2 pad bit (27) unal;
```

shifter	is "1"b if the terminal requires case shift characters.
upper_case_only	is "1"b if the terminal handles only capital letters.
delay_char	is the ASCII form of the character used for carriage movement delays.
upper_case	is the uppercase shift character.
lower_case	is the lowercase shift character.
delay_offset	is an array of offsets of the delay_tables (described in MTB 234) to be used for this terminal type at 110, 150, 300, and 1200 bps respectively.
output_tct_offset	is the relative offset (in tty_ctl) of the default table used by

tty_util_\$find_char for identifying "special" characters during output processing.

output_mvt_offset is the relative offset of the table used by tty_util_\$mvt for translation during output processing, or 0 if translation is not required for the particular terminal type.

special_offset is the relative offset of the default version of the special_chars table described below.

input_tct_offset is the relative offset of the default table used by tty_util_\$tct for identifying "special" characters during input processing.

input_mvt_offset is the relative offset of the table used by tty_util_\$mvt for translation during input processing, or 0 if translation is not required for the particular terminal type.

break_char is the break character for this device.

Special_Characters_Table

The special characters table is as described in MTB 234, except that the following items have been added at the end of the structure:

```

2 input_escape_length fixed bin,
2 input_escapes char (1 refer (input_escape_length):
                        unaligned,
2 input_results char (1 refer (input_escape_length):
                        unaligned;
```

input_escape_length is the number of characters in each of the strings input_escapes and input_results.

input_escapes is a string of characters each of which forms an escape sequence when preceded

by an escape character.

`input_results` is a string of characters each of which is to replace the escape sequence consisting of an escape character and the character occupying the corresponding position in `input_escapes` (above).

ADDITION TO MODULE DESCRIPTION OF `tty_util`

Entry: `tty_util_$tct`

This entry uses a `tct` (test character and translate) instruction to search a given string for "interesting" characters in the same manner as `tty_util_$find_char`.

Usage

```
declare tty_util_$tct entry (ptr);
```

```
call tty_util_$tct (argptr);
```

where `argptr` is a pointer to the structure described below. (Input)

```
dcl 1 tct_arg_structure based aligned,
    2 stringp ptr,
    2 stringl fixed bin,
    2 tally fixed bin,
    2 tablep ptr,
    2 indicator fixed bin,
    2 workspace (3) fixed bin;
```

All members of the structure have the same meaning as for `tty_util_$find_char`, except for the following:

`stringp` is a pointer to the string to be tested; it is updated to point to the first "interesting" character in the string.

(Input/Output)

indicator is the result of the search. It may have the following values: (Output)

0 -- no special characters

1 -- break character

2 -- escape character

3 -- character to be thrown away

MODULE_DESCRIPTION_OF tty_canon

Name: tty_canon

This subroutine is used to reduce a character string (which is expected to consist of one typed line image) to canonical form, i. e., sort the characters by column position and by ASCII value within each column position.

Usage

```
declare tty_canon entry (ptr, fixed bin (24), fixed bin (24),
                        char (1) aligned, char (1) aligned,
                        fixed bin (35));
```

```
call tty_canon (string_ptr, length, max_length,
               erase_char, kill_char, code);
```

string_ptr is a pointer to the string to be reduced; the result string replaces the input string. (Input)

length is the length of the string. It is adjusted to reflect the length of the result string. (Input/Output)

max_length is the maximum allowable length of the result string. (Input)

erase_char is the character which is to be interpreted as an erase character, or blank if no erase processing is to be done. (Input)

`kill_char` is the character which is to be interpreted as a kill character, or blank if no kill character is to be recognized. (Input)

`code` is a standard system status code. If the canonicalization of the string requires a result string whose length exceeds `max_length`, `code` is set to `error_table_$long_record`; otherwise it is set to zero. (Output)