

TO: MSPM Distribution
FROM: Karolyn Martin
DATE: October 18, 1968
SUBJECT: BX.8 and BY.2

Due to a recent redesign of the file system commands many sections of BX.8 and BY.2 are being re-issued. The major changes are in implementation in order to make the commands smaller and faster.

The main usage changes are covered in the overview to file system commands - BX.8.00. Points to look for are:

1. arrays as arguments are no longer allowed.
2. the commands take a fixed number of arguments and rely on the files procedure (BX.8.01) and the iteration feature of the shell (BX.2.00) to allow the user to act on a number of arguments, one after the other.
3. the star and equal conventions are restricted and simplified somewhat.
4. a new convention has been added in which directory entries are designated by a trailing ">".
5. users who have coded calls to procedures described in the old BY.2.01 should be warned that they are being replaced by those described in the new BY.2.01.

Sections BX.8 and BY.2 now consist of the following sections

<u>MSPM</u>	<u>Title</u>	<u>Replaces</u>
BX.8.00	Overview of File System Commands	BX.8.00, 8.00A
BX.8.01	List, files, status Commands	BX.8.01, 8.01A
BX.8.02	Acl Modification Commands	no change
BX.8.03	Listacl Command	no change
BX.8.04	Link Command	BX.8.04, 8.04A
BX.8.05	Unlink Command	new
BX.8.06	Branch Command	BX.8.05, 8.05A

<u>MSPM</u>	<u>Title</u>	<u>Replaces</u>
BX.8.07	Remove Command	BX.8.07, 8.07A
BX.8.08	Rename Command	BX.8.06, 8.06A
BX.8.09	Addname Command	BX.8.06, 8.06A
BX.8.10	Delname Command	BX.8.06, 8.06A
BX.8.11	Copy Command	new
BX.8.12	Movebranch Command	BX.8.09, BX.8.10
BX.8.13	Chasepath Command	new
BX.8.14, 8.14A	Working Directory Table Commands (section number change only)	BX.8.12, 8.12A
BX.8.15	Map_dir (section number change only)	BX.8.11
BY.2.00	Overview of File System Library Procedures	BY.2.00
BY.2.01	Command and File System Interface	BY.2.01, 2.01A
BY.2.02	File System Error Codes	no change
BY.2.03	File System Command Error Handling	new
BY.2.04	Setpath, entryarg	BY.2.04
BY.2.05	ACL Interface	no change
BY.2.06	Equals Convention Handler	BY.2.04
BY.2.07	Get_count, set_count	no change
BY.2.08	Star Convention Handler	BY.2.08
BY.2.09	Get_area	no change
BY.2.10	Entry_status	no change
BY.2.11	Working_segs	no change
BY.2.12	Deltree (section number change only)	BY.2.03
BY.2.13	Maplevel (section number change only)	BY.2.06

Published: 10/17/68
(Supersedes: BX.8.00, 01/27/67,
BX.8.00, 01/14/66)

Identification

Overview of File System Commands

C. A. Cushing, C. Garman, E. Q. Bjorkman, R. J. Feiertag

Purpose

The file system commands and procedures are designed as an interface between the Multics user and the Basic File System (Section BG); they provide frequently-used sequences of manipulations upon segments which reside in the file system hierarchy, as well as service functions such as the formatting and printing of segment status information.

Similarly, the file system procedures provide a more flexible interface between a user's programs and the "primitives" of the Basic File System, as well as error handling by means of the standard error procedures (BY.11). For further information see BY.2.00, Overview of File System Library Procedures.

Introduction

A file is a linear sequence of data elements; an element may be a machine word, an ASCII character, or a bit, or multiples of each of these, depending upon the context of a particular reference. In Multics, a file is generally found as a segment somewhere in the hierarchy of directories maintained, manipulated, and massaged by the Basic File System. A segment may be known to a user or his procedures by its segment name, which may include information on its location in the directory hierarchy, or by its (hardware) segment number, which provides a shorthand method for accessing the data of the segment in core memory.

A Multics user may create, modify, or delete segments only through the use of the Basic File System: directly, by programmed calls to user-accessible entries in the Basic File System and hardware segment addressing of data, and indirectly, by using the file system commands and subroutines. The first group is covered in MSPM section BG, The Basic File System, while the latter group is the subject of sections BX.8, File System Commands, and BY.2, File System Procedures.

Files, Branches and Entries

A directory is a special segment maintained by the Basic File System, which contains a list of entries. To the user, an entry appears to be a segment which is accessed in terms of one or more symbolic entry names. The names are composed of strings of ASCII characters, whose syntax is discussed below. An entry name need be unique only in the directory in which it occurs.

There are two types of entries each containing a "pointer":

If the entry is a branch, the pointer defines the location of the contents of the segment somewhere in the secondary storage hierarchy; this segment may or may not be a directory.

If the entry is a link, the pointer describes, symbolically, another entry in the same or another directory. A link may point to a link, etc., to a reasonable depth of nesting, but once a link points to a branch, the segment pointed to by that branch is accessible as if the original entry had been the branch pointing to the segment.

Each branch contains a description of the way the segment it points to may be used or referenced, which is termed its access control information; a link does not contain this information, but instead derives its access privileges from the branch to which it ultimately points.

File System Hierarchy

The Basic File System works with a basic tree hierarchy of segments across which links may be added to facilitate simple access to segments elsewhere in the hierarchy. With one exception each segment (e.g., directory) finds itself directly pointed to by a branch in exactly one directory; the exception is the root directory at the root of the tree, whose location is known to the Basic File System, but which does not appear in any directory.

A segment pointed to by a branch in some directory is immediately inferior to that directory and the directory is immediately superior to the segment. The master directory has level zero, and segments immediately inferior to it have level one. By extension, inferiority (or superiority) is defined for any number of levels of separation via a chain of immediately inferior (superior) segments.

Links are considered superimposed upon, but independent of, the tree hierarchy.

At any one time, a user is considered to be operating in some one directory, called his working directory (wdir). He may access a segment simply by specifying the entry name in his working directory which effectively points to that segment. An entry name is meaningful only with respect to the directory in which it occurs, and may or may not be unique outside that directory. In order to refer to an entry which is not in the working directory, it is necessary to have a symbolic name which unambiguously defines an entry in the hierarchy as a whole. Such a name is a path name. It consists of the chain of entries (branches or links) required to reach the desired entry from the root directory or from the working directory. A number of abbreviations may be used for the special directories at each Multics installation, e.g. the Multics Command and Subroutine Library, the Local Command and Subroutine Library, or the Process Directory (see BD.6 System Skeleton, and BD.4, the Search Module).

Names: Formation and Syntax

A file-system-name is a string of ASCII graphic characters which plays a role corresponding to that of an identifier in a programming language. Names are further classed as entry names, path names, and access-control names (user names).

Five punctuation characters (".", ">", "<", "*", and "=") are reserved and receive special interpretation when encountered in names in the context of the Basic File System and the file-system commands and subroutines. While the other punctuation characters and the ASCII control characters are not specifically excluded, most names will in fact consist of only the upper- and lower-case alphabetic characters, the digits, the underscore character "_", and appropriate usages of the five special characters named above.

A path name is written symbolically as a chain of entry names, each name separated by ">". If the first character of the path name is ">", the path name is an absolute path name; that is, it is fixed with respect to the root directory. If the first character is not ">", then the path name is relative to the working directory or one of the special system directories. For example, the path name of the directory corresponding to the box marked 1 in Figure 1 is >Z>H. If >Z>H is now the working directory, then a path name of the directory corresponding to the box marked 2 is L. The symbol '<' is the shorthand notation

for the path name of the directory which contains the link or branch that was used to access the current working directory. Using this notation and considering the path taken to directory 2 as >Z>H>L, a path name for directory 3 relative to directory 2 is <<A. Considering another path to directory 2 as >Z>A>B>C, a path name for directory 3 relative to directory 2 is <<.

An entry name is composed of one or more graphic components (not containing the reserved characters), separated by "."; the components may be referred to (in order, left-to-right) as the primary, secondary, etc., components of an entry name. Certain commands attach particular significance to various components of an entry name; thus the command

```
p11 my_seg
```

directs the PL/I compiler to attempt the compilation of the (ASCII) file,

```
my_seg.p11
```

(See also BB.5.01, Reserved Segment Name Suffixes.)

A successful compilation would produce the following new segments in the user's working directory:

```
my_seg          [note no secondary component]
my_seg.link     [linkage segment]
my_seg.symbol   [symbol table and binding information]
```

and if the "list" option was "on"

```
my_seg.list     [listing segment].
```

If an entry name is that of a directory branch it is followed by a ">". Thus the command:

```
branch newdir>
```

will cause the directory branch "newdir" to be created in the working directory whereas:

```
branch newdir
```

would create a non-directory branch.

A shorthand notation is available with certain commands. Special meanings are given to the character strings "*", "**", "=", and "==". The special command "files" generates a list of path names or entry names defined by the use of "*" and "**" in its argument. The list is generated by matching the argument of the "files" command with all entry names in the specified directory of the path name. If no directory is specified then the working directory is used.

If the entry name contains neither "*" nor "**" as a component, a character for character match is made against the entry names in the directory and if a match is found that name is returned as the value of files, with the specified directory to form the complete path name. If "*" is found as a component in the argument of "files" it will match any component appearing in the same position of an entry name in the directory. A list of the entry names matched will be returned. If "**" appears as the last component in the argument of "files" it will match any number of components (including 0) appearing at the end of an entry name in the directory. "**" appearing in other than the last component is treated as an ordinary component.

Examples:

my_seg will match only my_seg

.my_seg will match old.my_seg. PL/1
new.my_seg.ep1

but will not match my_seg.ep1.link
very.new.my_seg
my_seg
my_seg.ep1
new.my_seg

my_seg.** will match my_seg.ep1.link
my_seg
my_seg.symbol

but will not match ep1.link
new.my_seg

my_dir>** will match all entries in the directory "my_dir".

The command "files" should not be used as a command itself but only as an active command (see BX.1.00) in order to specify a list of entries to be acted upon. For example

```
link ([files >your_dir>xyz.**])
```

will create links in the user's working directory to all the entries in the directory "your_dir" whose names have first component "xyz".

Note: The command "list" uses this star convention implicitly and one should not issue the "files" command with the "list" command.

Some commands that require two arguments will give special meanings to the character strings "=" and "==" in the second argument. A "=" appearing as a component of the entry name in the second argument is equivalent to the corresponding component of the entry name in the first argument. If there is no corresponding component the "=" component is ignored. A "==" appearing as the last component in the entry name of the second argument is equivalent to that component and all components following it in the entry name of the first argument. A "==" appearing as other than the last component of the entry name of the second argument will cause all components following it to be ignored.

Examples:

```
rename my_seg.ep1 =.pl/1
```

is equivalent to

```
rename my_seg.ep1 my_seg.pl/1
```

```
rename my_seg.ep1.link your_seg.==
```

is equivalent to

```
rename my_seg.ep1.link your_seg.ep1.link
```

```
rename my_seg your_seg.=
```

and

```
rename my_seg your_seg.==
```

are both equivalent to

```
rename my_seg your_seg
```



```

rename my_seg.epl.link    your_seg.=.symbol
      is equivalent to
rename my_seg.epl.link    your_seg.epl.link

```

As an overall example consider;

```
addname ([files my_seg.*]) our_seg.=
```

This will add a name to all entries with two component names having first component "my_seg" in the working directory. The new names will have first component "our_seg" and the second component will be the same as the second component of the original name.

An access control name (or user name) is constructed similarly to an entry name; however, it always contains a fixed number of components (this number may vary from installation to installation) separated again by ".". The character strings "*" and "**" are defined as for entry-name arguments; since the number of components is fixed, however, missing components on the right are assumed to be "*". The single-character user name "*" denotes "all users of this Multics installation".

Access Control

In attempting to access a segment a user may or may not be successful depending upon his implicit intentions and his permissions with respect to the segment. The set of permissions with which a given user may access the segment pointed to by a particular branch is called the mode of the branch for that user. The permissions given the user of a particular branch are specified by an access control list for that branch. This list is a list of users (i.e., of sets of users) along with the corresponding mode associated with each user. The access control list is ordered according to the weights of the user names. The weight of a user name is equal to the sum of the weights of its components where the weight of the *i*th component (reading from right-to-left) is 2^{**i} . The components designated by "*" have weight 0. The ordering is from the highest to the lowest weighted user name. User names with the same weight have no ordering with respect to one another. For example, the following user names

Smith.proj_A.*

*

Jones.*

..aa

*.proj_A.aa

would be ordered

Smith.proj_A.*

Jones.*

*.proj_A.aa

..aa

* [equivalent to *.*.*]

The list is scanned from the top to discover the mode of a user. If all access control information required for the use of every segment in a particular directory is the same for certain users, this access control information may be put in the common access control list of the directory. If a user's name or class is not on the access control list of the branch pointing to the segment he wishes to use, then the common access control list is searched for this user's name or class, to determine his mode of access to the segment.

The mode consists of five attributes, named Trap, Read, Execute, Write, and Append, (sometimes abbreviated TREWA) each of which is either on or off. Collectively they define the apparent mode of the segment. The trap attribute is examined first. It has the power to override the other four attributes called usage attributes. The usage attributes indicate permission to perform the given activity only if the attribute is on.

TRAP ATTRIBUTE. When a branch has the Trap attribute on for a given user, a trap occurs whenever that user references that branch. That is, the procedure whose name is given as the trap procedure is called. A list of parameters may be defined with this procedure name and are passed as arguments to the called procedure.

The return from the trap procedure specifies the effective values of the four usage attributes, which may override the original values. A user can inhibit the trap mechanism, in which case all references to a branch by the user with the Trap attribute on will cause an error return to the calling procedure. (For further information see BX.8.02, Access Control Commands, BG.8, Directory Control, and BG.9, Access Control.)

USAGE ATTRIBUTES. Every operation on a given segment implies one of four intents, namely, read, execute, write or append. The interpretation of the intent depends upon whether the accessed branch points to a directory or a non-directory segment.

<u>Attribute</u>	<u>Directory</u>	<u>Non-directory branch</u>
Read	can "read" a directory to get information about any or all of the entries, including access control lists	can read the segment
Write	can delete or rename specifically named entries and change access control lists of specifically named entries.	can truncate or rewrite existing contents of the segment without adding to its length.
Append	can add entries without changing existing entries	can add to segment without changing original contents of file.
Execute	can search for specifically named entries in the directory in order to use them or to get information about them, excluding access control lists.	can execute the contents of the segment as a procedure.

EXAMPLE: Consider the request to delete all entries whose names contain the secondary component GAMMA in a particular directory. The user issuing this request must have the write attribute on in the access control list of the branch pointing to the directory in order to delete each of the entries in the set described. The user must also be able to Read the directory in order to find all the entries with this secondary component, GAMMA.

Organization of the File System Commands

The file system commands are described in the section which follow; the calling sequence for these commands (with examples of the kind of arguments expected) is generally "command entry args". (See also BX.1.00, The Command Language, and BX.2.00, The Shell.)

command is simply the name of a command

entry is an entry-name argument as discussed earlier, and defines the group of entries to be considered by command; entry may be a name defining a group of entries in the current working directory; or entry may be a path name with an entry name appended to it defining a set of entries in the directory pointed to by the preceding path, e.g.

<A	[entry A in the directory immediately superior to the working directory]
A>*.eplbsa	[all 2 component entries with secondary component "eplbsa" in directory A which is immediately inferior to the working directory]
>B>**	[all entries in directory B which is inferior to the root]
<<T>U>TEST.pl1	[This one is left as an exercise to the reader]

are all correct entry arguments.

args are arguments which vary from command to command. One type which is common to many file system command calling sequences is path.

path is a series of entry names separated by ">" or "<" which defines a directory different from the working directory. (See the previous discussion of path names in this section.)

Implementation

Many commands use certain routines in common. If the reader wishes to fully understand the implementation of the individual file system commands he should first be familiar with the following routines:

setpath	BY.2.04
entryarg	BY.2.04
command_error	BY.2.03
listfiles	BY.2.08
equalcomp	BY.2.06
ufo	BY.2.01

Note: The present convention is that entry names be 32 characters long. The file system command subroutines call for entry names to be 33 characters long. This extra character is to allow for a ">" following the name. However, any value less than 33 may also be used in declaring entry names to be given to the file system commands.

ROOT
DIRECTORY




Figure 1