

23 **I.3 Definitions**

24 ⇒ **2.2.2 General Terms** *Replace the contents of the subclause designated the*
 25 *definition of "character special file" with the following:*

26 **I.3.0.1 character special file:**

27 A file that refers to a device. One specific type of character special file is a termi-
 28 nal device file, whose access is defined in 7.1. Other character special files have no
 29 structure defined by this part of ISO/IEC 9945, but they can be accessed as defined
 30 in 21.

31 ⇒ **2.2.2 General Terms** *Add the following definition, in the right sorted order:*

32 **I.3.0.2 driver:**

33 A part of the implementation (possibly user supplied) that controls a device
 34 (2.2.2.x).

35 **I.4 Errors**

36 ⇒ **2.4 Error Numbers** *Add the following error value, in the right sorted order:*

37 [EBADCMD] Inappropriate I/O control operation. A control function was
 38 attempted for a file or a special file for which the operation
 39 was inappropriate. This error is synonymous with [ENOTTY],
 40 but shall be used when control operations are inappropriate
 41 for a non-TTY device.

42 **I.5 Functions**

43 ⇒ **22 Device Control** *Add a new chapter with the following new interface:*

44 The following function provides the device control capability:

45 `posix_devctl()`
 46 Control a Device

47 **I.5.1 Control a Device**

48 Function: *posix_devctl()*

49 **I.5.1.1 Synopsis**

```
50 #include <sys/types.h>
51 #include <unistd.h>
52 #include <devctl.h>
53 int posix_devctl(int fildes,
54                 int dcmd,
55                 void *dev_data_ptr,
56                 size_t nbyte,
57                 int *dev_info_ptr);
```

58 **I.5.1.2 Description**

59 If the Device Control option is supported:

B

60 The *posix_devctl()* function shall cause the device control command *dcmd* to
 61 be passed to the driver identified by *fildes*. Associated data shall be passed
 62 to and/or from the driver depending on direction information encoded in the
 63 *dcmd* argument, or as implied in the *dcmd* argument by the design and
 64 implementation of the driver.

65 The *dev_data_ptr* argument shall be a pointer to a buffer that contains data
 66 bytes to be passed to the driver and receives data bytes to be passed back
 67 from the driver or both.

68 If the data is to be passed to the driver, at least *nbyte* bytes of associated
 69 data shall be made available to the driver; if the data is to be passed from
 70 the driver, no more than *nbyte* bytes shall be passed.

71 If *nbyte* is zero, the amount of data passed to and/or from the driver is
 72 unspecified. This feature is obsolescent, and only provided for compatibility
 73 with existing device drivers.

74 The *dev_info_ptr* argument provides the opportunity to return an addi-
 75 tional device information word instead of just a success/failure indication.

76 The set of valid commands, the associated data interpretation, the returned
 77 device information word, and the effect of the command on the device are
 78 all defined by the driver identified by *fildes*.

79 Otherwise:

80 Either the implementation shall support the *posix_devctl()* function as
 81 described above or this function shall not be provided.

B

82 **I.5.1.3 Returns**

83 The *posix_devctl()* function shall return zero on success and the corresponding
84 status value on failure. The value returned via the *dev_info_ptr* argument is
85 driver dependent.

86 **I.5.1.4 Errors**

87 If any of the following conditions occur, the *posix_devctl()* function shall fail and
88 shall return the corresponding error number:

89 [EBADF] The *fdes* argument is not a valid open file descriptor.

90

91 If the following conditions are detected, the *posix_devctl()* function shall fail and
92 shall return the corresponding error number:

93 [EBADCMD] The *dcmd* argument is not valid for this device.

94 [EINTR] The *posix_devctl()* function was interrupted by a signal.

95 [EINVAL] The *nbyte* argument is negative, or exceeds an implementation
96 defined maximum, or is less than the minimum number of bytes
97 required for this command.

98 The argument *dev_info_ptr* is an invalid address, or the argu-
99 ment *dev_data_ptr* is an invalid address, or the *dev_data_ptr* +
100 *nbytes* - 1 is an invalid address.

101 [EPERM] The requesting process does not have sufficient privilege to
102 request the device to perform the specified command.

103 Driver code may detect other errors, but the error numbers returned are driver
104 dependent. See 21.4.9.

105 If the *posix_devctl()* function fails, the effect of this failed function on the device is
106 driver dependent. Corresponding data might be transferred, partially transferred,
107 or not transferred at all.

108 **I.5.1.5 Cross-References**

109 *close()*, 6.3.1; *dup()*, 6.2.1; *fstat()*, 5.6.2; *open()*, 5.3.1; *read()*, 6.4.1 *write()*, 6.4.2.

B

110 **I.6 Rationale Relating to Device Control**

111 An interface to be included in the POSIX standard should improve source code por-
112 tability of application programs. In existing UNIX practice, *ioctl()* is used to han-
113 dle special hardware. Therefore a general specification of its arguments cannot be
114 written. Based on this fact, many people claim that *ioctl()* or something close to it
115 has no place in POSIX.

116 Against this perception stands the widespread use of *ioctl()* to interface to all sorts
117 of drivers for a vast variety of hardware used in all areas of real time and embed-
118 ded computing, such as analog-digital converters, counters and video graphic dev-
119 ices. These devices provide a set of services that cannot be represented or used in
120 terms of read or write system calls.

121 The arguments in favor of *ioctl()* standardization can be summarized as follows:

122 Even if *ioctl()* addresses very different hardware, many of these devices are either
123 actually the same, interfaced to different computer systems with different imple-
124 mentations of operating systems, or belong to classes of devices with rather high
125 commonality in their functions, e.g. analog-digital converters or digital-analog con-
126 verters. Growing standardization of the Control and Status Register (CSR) space
127 of these devices allows or will allow exploitation of a growing similarity of control
128 codes and data for these. A general mechanism is needed to control these devices.

129 In all these cases a standardized interface from the application program to drivers
130 for these devices will improve source code portability.

131 Even if control codes and device data have to be changed when porting applica-
132 tions from one system to another, the definition of *ioctl()* largely improves reada-
133 bility of a program handling special devices. Changes are confined to more clearly
134 labeled places.

135 A driver for a specific device normally cannot be considered portable *per se*, but an
136 application that uses this driver can be made portable if all interfaces needed are
137 well defined and standardized. Users and integrators of real time systems often
138 add device drivers for specific devices and a standard interface simplifies this pro-
139 cess. Also, device drivers often follow their special hardware from system to sys-
140 tem.

141 **I.6.1 Existing Practice**

142 The *ioctl()* interface is widely used. It has provided the generality mentioned
143 above. This or a similar interface will build upon the current programming prac-
144 tice and existing code base, both at the application and device driver level.

145 Existing practice encodes into the second parameter information about data size
146 and direction in some systems. An example of such an encoding is BSD's use of
147 two bits of the command word as read/write bits. However, *ioctl()* has definite
148 problems with the way that its sometimes optional 3rd parameter can be inter-
149 preted.

150 This is similar to the existing POSIX.1 *fcntl()* function, in which the 3rd parameter
151 can be optional for `F_GETFD`, `F_GETFL`, an `int fildes` when used with the `F_DUPFD`,

Copyright © 1998 IEEE. All rights reserved.

This is an unapproved IEEE Standards Draft, subject to change.

152 F_SETFD, or F_SETFL commands or a struct *flock*, when used with the F_GETLK,
 153 F_SETLD or F_SETLKW commands. However, the *fcntl()* interface defines two dis-
 154 tinct and known data types as possible for the 3rd parameter. This is not the case
 155 in the *ioctl()* interface, where many device driver specific structures and com-
 156 mands are used.

157 **I.6.2 Relationship to *ioctl()* and the Perceived Needs for Improvement.**

158 POSIX.1 documents, in Annex B.7, the perceived deficiencies in existing implemen-
 159 tations of the *ioctl()* function. This discussion is in the context of those *ioctl()*
 160 commands used to implement terminal control. The POSIX.1 working group
 161 decided that, since the set of such control operations was fairly well defined, suit-
 162 able encapsulations such as *tcsetattr()*, *tcsendbreak()*, and *tcdrain()* could be
 163 standardized. These interfaces, while successfully standardizing portable termi-
 164 nal control operations, are not extensible to arbitrary user-supplied devices. The
 165 *posix_devctl()* interface replaces the various *ioctl()* implementations with a stan-
 166 dard interface which captures the extensibility of *ioctl()*, but avoids several of the
 167 deficiencies:

- 168 — The major problem with *ioctl()* is that the third argument is a generic
 169 pointer to a memory object which varies in both size and type according to
 170 the second *command* argument. It is not unprecedented in POSIX, or stan-
 171 dards in general, for a function to accept a generic pointer; consider the
 172 ANSI C library functions *fgets()* and *fread()*, or the POSIX functions *read()*
 173 and *mmap()*. However, in all such instances, the generic pointer must be
 174 accompanied by a user-specified size argument which specifies the size of
 175 the pointed-to object. Unlike the Ada language, it is, and has always been,
 176 the C programmer's responsibility to ensure that these two arguments form
 177 a consistent specification of the passed object. But traditional *ioctl()* imple-
 178 mentations do not allow the user to specify the size of the pointed-to object;
 179 that size is instead fixed implicitly by the specified command (passed as
 180 another argument). The *posix_devctl()* interface improves upon *ioctl()* in
 181 that it allows the user to specify the object size, thereby restoring the fami-
 182 liar C paradigm for passing a generic object by pointer/size pair.
- 183 — A secondary problem with *ioctl()* is that the third argument is sometimes
 184 permitted to be interpreted as an integer (int). This is non-portable to sys-
 185 tems where *sizeof(void *)* != *sizeof(int)*, not to mention a gross abuse of type
 186 casts. The *posix_devctl()* interface clearly requires the *dev_data_ptr* argu-
 187 ment to be a pointer.
- 188 — A related problem with *ioctl()* is that the direction(s) in which data are
 189 transferred to/from the pointed-to object is neither specified explicitly as an
 190 argument (as with *mmap()*), nor implied by the *ioctl()* function (as with
 191 *read()/write()*, *fread()/fwrite()*, or *fgets()/fputs()*). Instead, the direction is
 192 implied by the *command* argument. In traditional implementations, only
 193 the device driver knows the interpretation of the commands and whether
 194 data is to be transferred to or from the pointed-to object. But in networked
 195 implementations, generic portions of the operating system may need to
 196 know the direction to ensure that data are passed properly between a client

197 and a server, separately from device driver concerns. Two implementation-
198 specific solutions to this problem are: a) to always assume data needs to be
199 transferred in both directions; and b) to encode the implied direction into
200 the command word along with the fixed data size. The *posix_devctl()* inter-
201 face already provides the implementation with an explicit size parameter;
202 since the direction is already known implicitly to both the application and
203 the driver, and since workable methods exist for implementations to ascer-
204 tain that direction if required, this perceived problem is strictly an imple-
205 mentation issue, and solvable without further impact on the interface.

206 — Finally, *posix_devctl()* improves upon *ioctl()* by adopting the new style of
207 error return, avoiding all the problems *errno* brings to multi-threaded
208 applications. Because the driver-specific information carried by the non-
209 error return values of *ioctl()* still potentially needs to be passed to the appli-
210 cation, *posix_devctl()* adds the *dev_info_ptr* argument to specify where this
211 information should be stored.

212 **I.6.3 Which Changes to *ioctl()* Are Acceptable?**

213 Any change in the definition of *ioctl()* has to be perceived as a clear improvement
214 by the community of people touched by this change. We have to be aware that
215 drivers for *normal* peripherals are typically written by highly specialized profes-
216 sionals. Drivers for the *special devices* are very often written by the end-user or
217 by the hardware designer, sometimes with fairly limited software literacy. Any
218 interface definition which can be seen as overly complicated will simply not be
219 accepted.

220 Nevertheless, a few simple and useful improvements to *ioctl()* are possible, justify-
221 ing also the change of name from *ioctl()* to *devctl()*.

222 The major change is the addition of the size of the device data. For enhanced com-
223 patibility with existing *ioctl()* implementations, this size may be specified as zero;
224 in this case the amount of data passed is unspecified. (This allows a macro
225 definition of *ioctl()* which converts it into a *posix_devctl()* call.)

226 The method of indicating error return values differs from traditional *ioctl()* imple-
227 mentations, but it does not preclude the construction of *posix_devctl()* as a macro
228 built upon *ioctl()*.

229 **I.6.4 Rationale for the *dev_info_ptr***

230 The working group felt that it was important to preserve the current *ioctl()* func-
231 tionality of allowing a device driver to return some arbitrary piece of information
232 instead of just a success/failure indication. Such information might be, for exam-
233 ple, the number of bytes received, the number of bytes that would not fit into the
234 buffer pointed at by *dev_data_ptr*, the data type indication, or the device status.
235 Current practice for device drivers and *ioctl()* usage allows such a device depen-
236 dent return value. Thus the concept of an additional output argument,
237 *dev_info_ptr*, was born.

238 **I.6.5 Rationale for No *direction* Argument**

239 The initial specification for *posix_devctl()* contained an additional argument which
240 specified the direction of data flow — to the driver and/or from the driver. This
241 argument was later removed for the following reasons:

- 242 — The argument was redundant. Most (if not all) existing implementations
243 encode the direction data either explicitly or implicitly in the command
244 word.
- 245 — The argument increased the probability of programming errors, since it
246 must be made to agree with the direction information already encoded or
247 implied in the command word or an error would occur.
- 248 — The only real use of the argument would be if new drivers were written
249 which supported generic commands such as `TRANSFER_CONTROL_DATA`,
250 which was modified by the direction argument to indicate in which direction
251 the data should be transferred. This is contrary to current practice which
252 uses command pairs such as `GET_CONTROL_DATA`, and `PUT_CONTROL_DATA`.
- 253 — The primary purpose of the direction argument was to allow higher levels of
254 the system to identify the direction of data transfers, particularly in the
255 case of remote devices, without having to understand all the commands of
256 all the devices on the system. We believe that implementations which need
257 to ascertain the direction of data transfer from a command word will define
258 a consistent convention for encoding the direction into each command word,
259 and all device drivers supplied by the user must adhere to this convention.
260 A standard convention may be defined in the future when device driver
261 interface standardization is undertaken.

262 Thus the data direction argument was removed.

263 **I.6.6 Rationale for Not Defining the Direction Encoding in the *command*** 264 **Word**

265 Consideration was given to defining the direction encoding in the command word,
266 but was rejected. No particular benefit was seen to a pre-defined encoding, as long
267 as the encoding was used consistently across the entire implementation and was
268 well known to the implementation.

269 In addition, although only one encoding (BSD's) was known among the members of
270 the small working group, it could not be ruled out that other encodings already
271 exist, and no reason for precluding these encodings was seen.

272 Finally, system or architectural constraints might make a chosen standard encod-
273 ing difficult to use on a given implementation.

274 Thus, this standard does not define a direction encoding. Specifying a standard
275 encoding is actually a small part of a larger and more contentious objective, that of
276 specifying a complete set of interfaces for portable device drivers; if a future
277 amendment to this standard specifies such interfaces, the issue of device control
278 direction encoding will necessarily be addressed as part of that specification.

279 **I.6.7 Recommended Practice for Handling Data Size Errors**

280 In the event that the data from the device are too large to fit into the specified
281 buffer, as much data as will fit should be transferred, and the error posted. The
282 remaining data will aid in debugging.

283 **I.6.8 Recommended Practice for *nbyte* == 0**

284 The feature which permits an unspecified amount of control data to be transferred
285 if *nbyte* is zero is obsolescent, and exits only for compatibility with existing device
286 driver usage of *ioctl()*, i.e. the device driver always transfers an amount of data
287 implied by the command. Newly developed device drivers should always honor the
288 application's *nbyte* argument or return the error [EINVAL] if the argument is an
289 unacceptable value. Such a device driver should interpret a zero value of *nbyte* as
290 no data to be transferred.

291 **I.6.9 Recommended Practice for Driver Detected Errors**

292 If the driver detects the following error conditions, it is recommended that the
293 *posix_devctl()* function fail and return the corresponding error number:

294	[EAGAIN]	The control operation could not complete successfully because
295		the device was in use by another process or the driver was
296		unable to carry out the request due to an outstanding operation
297		in progress.
298	[EBADCMD]	The driver determined that the <i>dcmd</i> argument is not valid for
299		this device or subdevice.
300	[EINVAL]	The arguments <i>dev_dta_ptr</i> and <i>nbyte</i> define a buffer too small
301		to hold the data expected by or to be returned by this driver.
302	[EIO]	The control operation could not complete successfully because
303		the driver detected a hardware error.

Copyright © 1998 IEEE. All rights reserved.
This is an unapproved IEEE Standards Draft, subject to change.