

IEEE Standards Interpretation for IEEE Std 1003.1™-2001 IEEE Standard Standard for Information Technology -- Portable Operating System Interface (POSIX®)

Copyright © 2006 by the Institute of Electrical and Electronics Engineers, Inc. 3 Park Avenue New York, New York 10016-5997 USA All Rights Reserved.

Interpretations are issued to explain and clarify the intent of a standard and **do not** constitute an alteration to the original standard. In addition, interpretations are not intended to supply consulting information. Permission is hereby granted to download and print one copy of this document. Individuals seeking permission to reproduce and/or distribute this document in its entirety or portions of this document must contact the IEEE Standards Department for the appropriate license. Use of the information contained in this document is at your own risk.

IEEE Standards Department Copyrights and Permissions 445 Hoes Lane, Piscataway, New Jersey 08855-1331, USA

Interpretation Request #78

Topic: mmap page alignment conflict **Relevant Sections:** XSH mmap, mprotect, munmap, msync Page: 787 Line: 25836 Page 795 Line 26120 Page 827 Line 27064 Page 832 Line 27168

Previously conforming implementations of POSIX become invalid with newer versions. The previous specifications (POSIX 1003.1-1996 and earlier) do not require page alignment of `addr` and `offset`. They even have the wording "NOTE: It is expected that a later amendment of this standard will disallow the implementation from imposing the restriction on the alignment of the `off` and `addr` arguments." POSIX1003.1-2001 adopted the wording from SUSv2 which is conflicting, requiring the application to align the arguments. The 1996 and earlier wording allows both implementations to co-exist, but the 2001 and later wording does not. The 1003.1-2004 X3T11/X3P11 volume is still worded to match the earlier specification, page 120, line 4915, section B2.8.3: "The `mmap()` and `munmap()` functions allow a process to manipulate their address space by mapping portions of memory objects into it and removing them from it. When multiple processes map the same memory object, they can share access to the underlying data. Implementations may restrict the size and alignment of mappings to be on page-size boundaries." This alignment conflict applies to `mmap`, `munmap`, `mprotect` and `msync`.

Action:

Delete lines 25836-25841: The `off` argument is constrained to be aligned and sized according to the value returned by `sysconf()` when passed `_SC_PAGESIZE` or `_SC_PAGE_SIZE`. When `MAP_FIXED` is specified, the application shall ensure that the argument `addr` also meets these constraints. The implementation performs mapping operations over whole pages. Thus, while the argument `len` need not meet a size or alignment constraint, the implementation shall include, in any mapping operation, any partial page

specified by the range [pa,pa+len).

After line 25841 add: If MAP_FIXED is specified and addr is nonzero, it shall have the same remainder as the off parameter, modulo the page size as returned by sysconf() when passed _SC_PAGESIZE or _SC_PAGE_SIZE. The implementation may require that off is a multiple of the page size. If MAP_FIXED is specified, the implementation may require that addr is a multiple of the page size. The system performs mapping operations over whole pages. Thus, while the parameter len need not meet a size or alignment constraint, the system shall include, in any mapping operation, any partial page specified by the address range starting at pa and continuing for len bytes.

Delete lines 25882-25884: [EINVAL] The addr argument (if MAP_FIXED was specified) or off is not a multiple of the page size as returned by sysconf(), or is considered invalid by the implementation.

After line 25907 add: The mmap() function may fail if: [EINVAL] The addr argument (if MAP_FIXED was specified) or off is not a multiple of the page size as returned by sysconf(), or is considered invalid by the implementation.

Delete lines 26120-26121: The implementation shall require that addr be a multiple of the page size as returned by sysconf().

After line 26121 add: The implementation may require that addr be a multiple of the page size, {PAGESIZE}.

Delete line 26136: [EINVAL] The addr argument is not a multiple of the page size as returned by sysconf().

After line 26143 add: The mprotect() function may fail if: [EINVAL] The addr argument is not a multiple of the page size as returned by sysconf().

Delete lines 27064-27065: The implementation shall require that addr be a multiple of the page size as returned by sysconf().

After line 27065 add: The implementation may require that addr be a multiple of the page size as returned by sysconf().

Delete line 27099: [EINVAL] The value of addr is not a multiple of the page size {PAGE-SIZE}.

After line 27102 add: The msync() function may fail if: [EINVAL] The value of addr is not a multiple of the page size as returned by sysconf().

Delete line 27168: The implementation shall require that addr be a multiple of the page size {PAGESIZE}.

After line 27168 add: The implementation may require that `addr` be a multiple of the page size as returned by `sysconf()`.

Delete line 27192: [EINVAL] The `addr` argument is not a multiple of the page size as returned by `sysconf()`.

After line 27192 add: The `mmap()` function may fail if: [EINVAL] The `addr` argument is not a multiple of the page size as returned by `sysconf()`.

Interpretation Response #78

The standards states for page alignment for `mmap()` and related functions, and conforming implementations must conform to this. However, concerns have been raised about this which are being referred to the sponsor.

Rationale for Interpretation

None.