

## Re: sendfile(2) SF\_NOPUSH flag proposal

**Source:** <http://unix.derkeiler.com/Mailing-Lists/FreeBSD/arch/2003-05/0101.html>

---

**From:** Marc Slemko ([marcs\\_at\\_znep.com](mailto:marcs_at_znep.com))

**Date:** 05/27/03

Date: Tue, 27 May 2003 11:14:07 -0700 (PDT)

To: arch@freebsd.org

On Tue, 27 May 2003, Terry Lambert wrote:

- > *The overhead of toggling it would be costly. However, I really*
- > *don't understand why he isn't just not setting TCP\_NODELAY in*
- > *the first place, since it's an affirmative option, and then*
- > *leaaving the socket alone to act like it's supposed to act.*

Given the bug in FreeBSD's sendfile() that results in it sending the headers in a separate segment, then just leaving Nagle enabled will destroy your performance if you are implementing a protocol such as HTTP with persistent (pipelined or not) connections due to the interaction between Nagle and delayed ack.

However, that is somewhat of a moot point for HTTP anyway because if you support pipelined requests then for the most efficient implementation of HTTP on the network layer you need to ensure that the response headers and body from multiple requests can be coalesced into the same packet. When using sendfile(), this means you have to control TCP\_NOPUSH yourself, even if sendfile coalesced headers, data, and trailers.

sendfile() still should be fixed to properly coalesce them.

---

freebsd-arch@freebsd.org mailing list

<http://lists.freebsd.org/mailman/listinfo/freebsd-arch>

To unsubscribe, send any mail to "freebsd-arch-unsubscribe@freebsd.org"