

RFC: Sendfile patch (fwd)

Source: <http://unix.derkeiler.com/Mailing-Lists/FreeBSD/performance/2004-01/0001.html>

From: Mike Silbersack (*silby_at_silby.com*)

Date: 01/04/04

Date: Sat, 3 Jan 2004 21:08:38 -0600 (CST)

To: performance@frebsd.org

Eh, I guess this would be on topic for -performance as well. Especially the part about benchmarking. :)

----- Forwarded message -----

Date: Sat, 3 Jan 2004 21:06:39 -0600 (CST)

From: Mike Silbersack <silby@silby.com>

To: arch@frebsd.org

Subject: RFC: Sendfile patch

The attached patch (also available from the following url):

<http://www.silby.com/patches/raw/sendfile-wip4.patch>

modifies sendfile to ensure that the header is sent in the same packet as the data from the file. This greatly helps performance when dealing with http servers, as it allows small files to fit in one packet, instead of two.

In order to keep the implementation relatively simple, I have created two helper functions, `iov_to_uio`, and `m_uiotombuf`. You'll note that `iov_to_uio` contains the body of `readv/writev` almost exactly, and `m_uiotombuf` does much the same thing as `sosend`. (In the future, I may attempt to make `readv/writev` use `iov_to_uio`, and have `sosend` use `m_uiotombuf`, but that's a project for another day.)

I'm going to be gone for two weeks, so that should allow plenty of time for testing / review of the patch. The location(s) of `iov_to_uio` and `m_uiotombuf` in both source and header files sounds like an easy thing to bikeshed about, so please don't bring that up. Such matters can be discussed when I get back, because I commit it. What I am interested in are reviews of the actual code, specifically regarding any possible security holes which I may be introducing. <g> Also, I would appreciate it if someone could verify that I have the old sendfile compatibility part working correctly, as it is quite messy.

If you're not in the mood to review the patch, please feel free to test it. I have only tested with `apache2` and `thttpd` (with `alfred's sf` patches from the port enabled), there are undoubtedly many other programs which

freebsd-performance: RFC: Sendfile patch (fwd)

make use of sendfile and should be tested. If you have any sort of web benchmark which helps to show the performance increase given by this patch, please run it; it's always great to include statistics in a commit message.

Naturally, the patch is against -current. If you have any questions, please ask quickly, or prepare to wait a long time for an answer. :)

Mike "Silby" Silbersack

freebsd-performance@freebsd.org mailing list

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

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

- TEXT/PLAIN attachment: [sendfile-wip4.patch](#)