

## Re: emulate an end-of-media

---

*Source:* <http://unix.derkeiler.com/Mailing-Lists/FreeBSD/hackers/2008-02/msg00454.html>

---

- *From:* Mike Meyer <mwm@xxxxxxxx>
  - *Date:* Wed, 27 Feb 2008 11:50:46 -0500
- 

On Wed, 27 Feb 2008 13:50:48 +0100 Joerg Sonnenberger <joerg@xxxxxxxxxxxxxxxxxxxx> wrote:

On Tue, Feb 26, 2008 at 04:00:00PM -0500, Mike Meyer wrote:

On Tue, 26 Feb 2008 21:28:53 +0100 Joerg Sonnenberger <joerg@xxxxxxxxxxxxxxxxxxxx> wrote:

On Tue, Feb 26, 2008 at 07:44:48PM +0100, Martin Laabs wrote:

I also made a comparison between gzip and bzip2 regarding the compression ratio on a dump of my home directory (3.2GB) bzip2 took about 74min to compress, gzip only 11minutes. And in terms of compression ratio bzip2 was only 3% better than gzip.

That's not a realistic test case. bzip2 normally takes trice the time and compresses 10% better. I can't comment on compress.

Considering we're talking about compression methods to use on dump output, that would seem to be the definition of a "realistic test case". Telling us what it "normally" does without defining what input is considered "normal" doesn't help much.

Source code in my case and various other documents. The test case above certainly was not normal.

So it sounds like your "normal" is mostly text documents of various kinds. I would expect such data to be a relatively small part of any dump data set, which, as you say, means that such data isn't

Re: emulate an end-of-media

normal. Given that the use case under discussion is abnormal, any tests using normal data are pretty much irrelevant.

<mike

--

Mike Meyer <mwm@xxxxxxxx> <http://www.mired.org/consulting.html>  
Independent Network/Unix/Perforce consultant, email for more information.

---

freebsd-hackers@xxxxxxxx mailing list

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

To unsubscribe, send any mail to "freebsd-hackers-unsubscribe@xxxxxxxx"