

Re: missing madvise functionality

Source: <http://linux.derkeiler.com/Mailing-Lists/Kernel/2007-04/msg02230.html>

- *From:* Ulrich Drepper <drepper@xxxxxxxxxx>
 - *Date:* Thu, 05 Apr 2007 19:52:09 -0700
-

Nick Piggin wrote:

Cool. According to my thinking, madvise(MADV_DONTNEED) even in today's kernels using down_write(mmap_sem) for MADV_DONTNEED is better than mmap/mprotect, which have more fundamental locking requirements, more overhead and no benefits (except debugging, I suppose).

It's a tiny bit faster, see

<http://people.redhat.com/drepper/dontneed.png>

I just ran it once so the graph is not smooth. This is on a UP dual core machine. Maybe tomorrow I'll turn on the big 4p machine.

I would have to see dramatically different results on the big machine to make me change the libc code. The reason is that there is a big drawback.

So far, when we allocate a new arena, we allocate address space with PROT_NONE and only when we need memory the protection is changed to PROT_READ|PROT_WRITE. This is the advantage of catching wild pointer accesses.

—

§ Ulrich Drepper § Red Hat, Inc. § 444 Castro St § Mountain View, CA V

Attachment: **signature.asc**

Description: OpenPGP digital signature