

Re: memory not released after using cdrecord/cdrdao (was: audio cd writing causes massive swap and crash)

Source: <http://linux.derkeiler.com/Mailing-Lists/Kernel/2004-07/4827.html>

From: Ed Sweetman (*safemode_at_comcast.net*)

Date: 07/26/04

Date: Sun, 25 Jul 2004 21:30:38 -0400

To: Eduard Bloch <blade@debian.org>

Indeed, i burned a smaller cd and got very similar results. I actually got to see the oom messages and all, of course nothing in userspace was hogging memory. The kernel is most definitely not releasing the memory it is using. I further would have to say that the kernel isn't being given the correct memory addresses that the data is actually at (and could be the reason why it's not freed) because every single audio cd i've burned partially or completely with 2.6.6+ has sounded garbled or been completely unreadable. This is the 20'th cd i've tested using different modes, settings, speeds, kernels and the results have been the same every time. I haven't tested kernels prior to 2.6.6 though. If there were any patches to the block layer or ide-cd code from 2.6.5 to 2.6.6 i'll go ahead and test 2.6.5... I'll go and check the changelog now but it's likely that a whole lot of code was changed between those versions if any was changed at all.

btw, the bug seems to be much more pronounced if you burn at 48x (some speed that you'd have to use burnfree or it would underrun)

My dmesg output.

Eduard Bloch wrote:

```
>#include <hallo.h>
>* Ed Sweetman [Sat, Jul 17 2004, 04:00:13PM]:
>
>
>>Both with 2.6.7-rc3 and 2.6.8-rc1-mm1 I get the same behavior when
>>writing an audio cd on my plextor px-712a. DMA is enabled and normal
>>data cds write as expected, but audio cds will cause (at any speed) the
>>box to start using insane amounts of swap (>150MB) and eventually cause
>>
>>
>
```

Re: memory not released after using cdrecord/cdrdao (was: audio cd writing causes massive swap and crash)

Linux-Kernel: Re: memory not released after using cdrecord/cdrdao (was: audio cd writing causes massive swap and crash)

>Just FYI: we have a similar bug description in the Debian BTS, where the
>user reports that kernel does not release memory assigned to userspace
>after cdrdao or cdrecord have used it (writing in DAO mode), though he
>could not find what allocated this memory. For details:
><http://bugs.debian.org/256871> (dump attached).

>
>Regards,
>Eduard.

>
>
>
>

>Debian Bug report logs – #256871
>cdrecord triggers memory leak in kernel space

>
>Package: cdrecord; Reported by: Jeff King <peff-debbug@peff.net>; Date: Tue, 29
>Jun 2004 15:48:02 UTC;
>Maintainer for cdrecord is Joerg Jaspert <joerg@debian.org>; Source for
>cdrecord is cdrtools.

>
>View this report as an mbox folder.

>
>

>-----
>Report forwarded to debian-bugs-dist@lists.debian.org, Joerg Jaspert
><joerg@debian.org>:
>Bug#256871; Package cdrecord. Full text available.

>
>

>-----
>Acknowledgement sent to Jeff King <peff-debbug@peff.net>:
>New Bug report received and forwarded. Copy sent to Joerg Jaspert
><joerg@debian.org>. Full text available.

>
>

>Message received at submit@bugs.debian.org:

>
>

>Received: (at submit) by bugs.debian.org; 29 Jun 2004 15:44:37 +0000

>>From peff-debbug@peff.net Tue Jun 29 08:44:37 2004

>Return-path: <peff-debbug@peff.net>

>Received: from 66-23-211-5.clients.speedfactory.net (peff.net) [66.23.211.5]

> by spohr.debian.org with esmtp (Exim 3.35 1 (Debian))

> id 1BfKmj-0003gN-00; Tue, 29 Jun 2004 08:44:37 -0700

>Received: (qmail 88814 invoked from network); 29 Jun 2004 15:44:40 -0000

>Received: from unknown (HELO coredump.intra.peff.net) (10.0.0.2)

> by peff.net with SMTP; 29 Jun 2004 15:44:40 -0000

>Received: by coredump.intra.peff.net (sSMTP sendmail emulation); Tue, 29 Jun 2004 11:44:36 -0400

>Content-Type: text/plain; charset="us-ascii"

>MIME-Version: 1.0

>Content-Transfer-Encoding: 7bit

>From: Jeff King <peff-debbug@peff.net>

>To: Debian Bug Tracking System <submit@bugs.debian.org>

>Subject: cdrecord triggers memory leak in kernel space

>Bcc: Jeff King <peff-debbug@peff.net>

Re: memory not released after using cdrecord/cdrdao (was: audio cd writing causes massive swap and crash)

Linux-Kernel: Re: memory not released after using cdrecord/cdrdao (was: audio cd writing causes massive swap and crash)

```
>X-Mailer: reportbug 2.62
>Date: Tue, 29 Jun 2004 11:44:36 -0400
>Message-Id: <E1BfKmj-0003gN-00@spohr.debian.org>
>Delivered-To: submit@bugs.debian.org
>X-Spam-Checker-Version: SpamAssassin 2.60-bugs.debian.org_2004_03_25
> (1.212-2003-09-23-exp) on spohr.debian.org
>X-Spam-Status: No, hits=-8.0 required=4.0 tests=BAYES_00,HAS_PACKAGE
> autolearn=no version=2.60-bugs.debian.org_2004_03_25
>X-Spam-Level:
>
>Package: cdrecord
>Version: 4:2.0+a30.pre1-1
>Severity: normal
>
>I am running cdrecord as root with the following options:
>cdrecord dev=/dev/hdc -v -dao -pad -audio *.wav
>
>One CD burns OK. When I try to burn another, about mid-way through the
>burn I notice random apps being killed. Sure enough, dmesg reports this:
> Out of Memory: Killed process 6700 (firefox-bin).
> Out of Memory: Killed process 6587 (xterm).
> Out of Memory: Killed process 2806 (netbiff).
> Out of Memory: Killed process 2861 (xterm).
>I'm fairly sure it's triggered by the cdrecord command, as I'm not doing
>anything else out of the ordinary and it happens consistently. Looking
>at the output of "free", something is using all of my memory:
>
> total used free shared buffers cached
>Mem: 1036972 1004756 32216 0 3000 72616
>-/+ buffers/cache: 929140 107832
>Swap: 0 0 0
>
>Running "top" or "ps" shows that userland processes aren't responsible
>for this:
> $ sum=0
> $ ps -AF | awk '{print $6}' | tail +2 | while read f
> sum=$(( $sum + $f ))
> echo $sum
> done
> ...
> 75288
>
>I've used cdrecord many times in the past and not run into this problem.
>Here are the things I'm doing differently recently:
> - move to 2.6.6 kernel; this is the first burn I've done with it, but I
> have done others with the 2.6 kernel series
> - use the /dev/hdc device syntax; I believe I have done this before
> and it worked OK, but I'm not sure
> - used -dao mode. This is DEFINITELY the first time I have tried using
> -dao mode.
>
```

Re: memory not released after using cdrecord/cdrdao (was: audio cd writing causes massive swap and crash)

Linux-Kernel: Re: memory not released after using cdrecord/cdrdao (was: audio cd writing causes massive swap and crash)

>Because of the massive amount of memory being gobbled, I can only guess
>that the buffer isn't being freed by the kernel for whatever reason. I
>suspect this is probably actually a kernel issue, but I thought I'd
>start here in case there have been other reports.

>
>-- System Information:
>Debian Release: testing/unstable
>APT prefers unstable
>APT policy: (500, 'unstable')
>Architecture: i386 (i686)
>Kernel: Linux 2.6.6-1-k7-smp
>Locale: LANG=C, LC_CTYPE=C

>
>Versions of packages cdrecord depends on:
>ii debconf 1.4.29 Debian configuration management sy
>ii libc6 2.3.2.ds1-13 GNU C Library: Shared libraries an
>ii makedev 2.3.1-70 Creates device files in /dev

>
>-- debconf information:
>* cdrecord/SUID_bit: false
> cdrecord/MAKEDEV:
> cdrecord/MAKEDEVNEW: true
> cdrecord/do_it_yourself:

>
>
>
>

>Information forwarded to debian-bugs-dist@lists.debian.org, Joerg Jaspert
><joerg@debian.org>:
>Bug#256871; Package cdrecord. Full text available.

>Acknowledgement sent to Andreas Metzler <ametzler@downhill.at.eu.org>:
>Extra info received and forwarded to list. Copy sent to Joerg Jaspert
><joerg@debian.org>. Full text available.

>
>

>Message received at 256871@bugs.debian.org:

>
>Received: (at 256871) by bugs.debian.org; 29 Jun 2004 17:11:28 +0000
>>From ametzler@downhill.at.eu.org Tue Jun 29 10:11:28 2004
>Return-path: <ametzler@downhill.at.eu.org>
>Received: from server.logic.univie.ac.at [131.130.190.41] ([fSfjV4L/RPRbY9sU7VqRJ9D2UhTd4nZy])
> by spohr.debian.org with esmtp (Exim 3.35 1 (Debian))
> id 1BfM8m-00017v-00; Tue, 29 Jun 2004 10:11:28 -0700
>Received: from m-134-246.adsl.univie.ac.at ([131.130.134.246])
> by server.logic.univie.ac.at with asmtmp (Exim 4.34)
> id 1BfM8j-0007qs-11; Tue, 29 Jun 2004 19:11:26 +0200
>Received: from ametzler by downhill.univie.ac.at with local (Exim 4.34)
> id 1BfM8j-0005Ya-RP; Tue, 29 Jun 2004 19:11:25 +0200
>Date: Tue, 29 Jun 2004 19:11:25 +0200
>From: Andreas Metzler <ametzler@downhill.at.eu.org>

Re: memory not released after using cdrecord/cdrdao (was: audio cd writing causes massive swap and crash)

Linux-Kernel: Re: memory not released after using cdrecord/cdrdao (was: audio cd writing causes massive swap and crash)

>To: Jeff King <peff-debbug@peff.net>, 256871@bugs.debian.org
>Subject: Re: Bug#256871: cdrecord triggers memory leak in kernel space
>Message-ID: <20040629171125.GA21354@downhill.at.eu.org>
>References: <E1BfKmj-0003gN-00@spohr.debian.org>
>Mime-Version: 1.0
>Content-Type: text/plain; charset=us-ascii
>Content-Disposition: inline
>In-Reply-To: <E1BfKmj-0003gN-00@spohr.debian.org>
>X-GPG-Fingerprint: BCF7 1345 BE42 B5B8 1A57 EE09 1D33 9C65 8B8D 7663
>User-Agent: Mutt/1.5.6+20040523i
>Delivered-To: 256871@bugs.debian.org
>X-Spam-Checker-Version: SpamAssassin 2.60-bugs.debian.org_2004_03_25
> (1.212-2003-09-23-exp) on spohr.debian.org
>X-Spam-Status: No, hits=-6.0 required=4.0 tests=BAYES_00,HAS_BUG_NUMBER
> autolearn=no version=2.60-bugs.debian.org_2004_03_25
>X-Spam-Level:
>
>On 2004-06-29 Jeff King <peff-debbug@peff.net> wrote:
>[...]
>
>
>> - use the /dev/hdc device syntax; I believe I have done this before
>> and it worked OK, but I'm not sure
>>
>>
>[...]
>
>Don't. Use dev=ATA:x,y,z.
> cu andreas
>
>
>
>-----
>Information forwarded to debian-bugs-dist@lists.debian.org, Joerg Jaspert
><joerg@debian.org>:
>Bug#256871; Package cdrecord. Full text available.
>-----
>Acknowledgement sent to Eduard Bloch <edi@gmx.de>:
>Extra info received and forwarded to list. Copy sent to Joerg Jaspert
><joerg@debian.org>. Full text available.
>-----
>
>Message received at 256871@bugs.debian.org:
>
>Received: (at 256871) by bugs.debian.org; 29 Jun 2004 17:26:25 +0000
>>From inet@zombie.inka.de Tue Jun 29 10:26:25 2004
>Return-path: <inet@zombie.inka.de>
>Received: from mailgate1.uni-kl.de [131.246.120.5]
> by spohr.debian.org with esmtp (Exim 3.35 1 (Debian))
> id 1BfMNE-0001sP-00; Tue, 29 Jun 2004 10:26:25 -0700
>Received: from aixs1.rhrk.uni-kl.de (aixs1.rhrk.uni-kl.de [131.246.137.3])

Re: memory not released after using cdrecord/cdrdao (was: audio cd writing causes massive swap and crash)

Linux-Kernel: Re: memory not released after using cdrecord/cdrdao (was: audio cd writing causes massive swap and crash)

> by mailgate1.uni-kl.de (8.12.10/8.12.10) with ESMTP id i5THQMJY006306;
> Tue, 29 Jun 2004 19:26:22 +0200
>Received: from rotes255.wohnheim.uni-kl.de ([131.246.178.65] helo=debian.zombie.inka.de)
> by aixs1.rhrk.uni-kl.de with esmtp (TLSv1:RC4-SHA:128)
> (Exim 4.20)
> id 1BfMNB-000Afw-RE; Tue, 29 Jun 2004 19:26:21 +0200
>Received: from inet by debian.zombie.inka.de with local (Exim 4.34)
> id 1BfMN9-0001o3-Iw; Tue, 29 Jun 2004 19:26:19 +0200
>Date: Tue, 29 Jun 2004 19:26:19 +0200
>From: Eduard Bloch <edi@gmx.de>
>To: Jeff King <peff-debbug@peff.net>, 256871@bugs.debian.org
>Subject: Re: Bug#256871: cdrecord triggers memory leak in kernel space
>Message-ID: <20040629172619.GA5713@zombie.inka.de>
>References: <E1BfKmj-0003gN-00@spohr.debian.org>
>Mime-Version: 1.0
>Content-Type: text/plain; charset=us-ascii
>Content-Disposition: inline
>In-Reply-To: <E1BfKmj-0003gN-00@spohr.debian.org>
>User-Agent: Mutt/1.5.6+20040523i
>Sender: <inet@zombie.inka.de>
>Delivered-To: 256871@bugs.debian.org
>X-Spam-Checker-Version: SpamAssassin 2.60-bugs.debian.org_2004_03_25
> (1.212-2003-09-23-exp) on spohr.debian.org
>X-Spam-Status: No, hits=-6.0 required=4.0 tests=BAYES_00,HAS_BUG_NUMBER
> autolearn=no version=2.60-bugs.debian.org_2004_03_25
>X-Spam-Level:
>
>#include <hallo.h>
>* Jeff King [Tue, Jun 29 2004, 11:44:36AM]:
>
>
>
>>I've used cdrecord many times in the past and not run into this problem.
>>Here are the things I'm doing differently recently:
>> - move to 2.6.6 kernel; this is the first burn I've done with it, but I
>> have done others with the 2.6 kernel series
>> - use the /dev/hdc device syntax; I believe I have done this before
>> and it worked OK, but I'm not sure
>> - used -dao mode. This is DEFINITELY the first time I have tried using
>> -dao mode.
>>
>>
>>
>
>Okay, this sounds more likely like a kernel problem than a cdrecord
>problem. If you have some time, please try with combinations of
>following factors to find the reason of the trouble:
>
> - enable swap (does not fight the root of the problems but gives you
> more time to experiment)
> - use a non-SMP kernel
> - try kernel 2.6.7

Re: memory not released after using cdrecord/cdrdao (was: audio cd writing causes massive swap and crash)

Linux-Kernel: Re: memory not released after using cdrecord/cdrdao (was: audio cd writing causes massive swap and crash)

> – run *cdrecord.shm* instead of *cdrecord* or *cdrecord.mmap* (note that
> there are some limits with *.shm*, see *README.Debian*)

>

>

>

>> *suspect this is probably actually a kernel issue, but I thought I'd*

>> *start here in case there have been other reports.*

>>

>>

>

> *AFAICS there were no such reports and that's what it is confusing. I do*

> *not count on much help from the upstream author there.*

>

> *Thanks for your efforts,*

> *Eduard.*

>

>

>

>

>-----
> *Information forwarded to debian-bugs-dist@lists.debian.org, Joerg Jaspert*

> *<joerg@debian.org>:*

> *Bug#256871; Package cdrecord. Full text available.*

>

>-----
> *Acknowledgement sent to Joerg Schilling <schilling@fokus.fraunhofer.de>:*

> *Extra info received and forwarded to list. Copy sent to Joerg Jaspert*

> *<joerg@debian.org>. Full text available.*

>

>

> *Message received at 256871@bugs.debian.org:*

>

> *Received: (at 256871) by bugs.debian.org; 4 Jul 2004 13:08:19 +0000*

>> *From schilling@fokus.fraunhofer.de Sun Jul 04 06:08:19 2004*

> *Return-path: <schilling@fokus.fraunhofer.de>*

> *Received: from mailhub.fokus.fraunhofer.de [193.174.154.14]*

> *by spohr.debian.org with esmtp (Exim 3.35 1 (Debian))*

> *id 1Bh6jC-0003Wl-00; Sun, 04 Jul 2004 06:08:19 -0700*

> *Received: from burner.fokus.fraunhofer.de (burner [193.175.133.116])*

> *by mailhub.fokus.fraunhofer.de (8.11.6p2/8.11.6) with ESMTP id i64D8FU15504;*

> *Sun, 4 Jul 2004 15:08:15 +0200 (MEST)*

> *Received: (from jes@localhost)*

> *by burner.fokus.fraunhofer.de (8.12.9+Sun/8.12.9/Submit) id i64D6fPT013802;*

> *Sun, 4 Jul 2004 15:06:41 +0200 (CEST)*

> *Date: Sun, 4 Jul 2004 15:06:41 +0200 (CEST)*

> *From: Joerg Schilling <schilling@fokus.fraunhofer.de>*

> *Message-Id: <200407041306.i64D6fPT013802@burner.fokus.fraunhofer.de>*

> *To: 256871@bugs.debian.org*

> *Cc: peff-debbug@peff.net*

> *Subject: Re: Bug#256871: cdrecord triggers memory leak in kernel space*

> *Delivered-To: 256871@bugs.debian.org*

> *X-Spam-Checker-Version: SpamAssassin 2.60-bugs.debian.org_2004_03_25*

> *(1.212-2003-09-23-exp) on spohr.debian.org*

Re: memory not released after using cdrecord/cdrdao (was: audio cd writing causes massive swap and crash)

Linux-Kernel: Re: memory not released after using cdrecord/cdrdao (was: audio cd writing causes massive swap and crash)

>X-Spam-Status: No, hits=-3.0 required=4.0 tests=HAS_BUG_NUMBER autolearn=no
> version=2.60-bugs.debian.org_2004_03_25
>X-Spam-Level:
>
>
>
>>AFAICS there were no such reports and that's what it is confusing. I do
>>not count on much help from the upstream author there.
>>
>>
>
>A sad truth :-(
>
>In the past, the Linux Kernel maintainers have not been very helpful with Linux
>Kernel bugs.
>
>
>I did stop sending bug reports to the Linux Kernel team two years ago for this
>reason. For me it makes not sense to waste my time with unwilling people.
>
>
>A possible reason for this Linux kernel bug is Kernel design bug that is known
>for a long time:
>
>Linux does not keep track of the size of the processes. If e.g. a process forks,
>Linux creates copy on write pages for the data segment of the fork. This idea
>has been taken from SunOS ideas in the mid 1980s. Unlike SunOS, Linux does not
>compute the needed virtual memory and overcommits the available size.
>
>When a process later modifies the parts of it's address space that is just a
>copy on write marked page, the kernel needs to create a second page before the
>modification may take place. The space for this page may not be available and
>the Linux kernel management is unable to tell which process is the cause for the
>missing virtual memory space. Linux for this reason kills random processes to
>regain virtual memory space. If you have a multi processor machine, it is even
>worse: Linux in this case becomes catatonic and may only recover by a "reset"
>induced reboot. This is caused by the way a dual or multi processor Linux
>works.
>
>BTW: As cdrecord touches all it's memory at process start, cdrecord is not one
>of the processes where Linux does not know the correct process size.
>
>
>
>J?rg
>
>--
> EMail:joerg@schily.isdn.cs.tu-berlin.de (home) J?rg Schilling D-13353 Berlin
> js@cs.tu-berlin.de (uni) If you don't have iso-8859-1
> schilling@fokus.fraunhofer.de (work) chars I am J"org Schilling
> URL: <http://www.fokus.fraunhofer.de/usr/schilling> <ftp://ftp.berlios.de/pub/schily>

Re: memory not released after using cdrecord/cdrdao (was: audio cd writing causes massive swap and crash)

Linux-Kernel: Re: memory not released after using cdrecord/cdrdao (was: audio cd writing causes massive swap and crash)

>
>
>

>-----
>Information forwarded to debian-bugs-dist@lists.debian.org, Joerg Jaspert

><joerg@debian.org>:

>Bug#256871; Package cdrecord. Full text available.

>-----

>Acknowledgement sent to Jeff King <peff-debbug@peff.net>:

>Extra info received and forwarded to list. Copy sent to Joerg Jaspert

><joerg@debian.org>. Full text available.

>-----

>

>Message received at 256871@bugs.debian.org:

>

>Received: (at 256871) by bugs.debian.org; 5 Jul 2004 05:56:13 +0000

>>From peff-debbug@peff.net Sun Jul 04 22:56:13 2004

>Return-path: <peff-debbug@peff.net>

>Received: from 66-23-211-5.clients.speedfactory.net (peff.net) [66.23.211.5]

> by spohr.debian.org with esmtp (Exim 3.35 1 (Debian))

> id 1BhMSb-00019j-00; Sun, 04 Jul 2004 22:56:13 -0700

>Received: (qmail 58673 invoked from network); 5 Jul 2004 05:56:16 -0000

>Received: from unknown (HELO localhost) (127.0.0.1)

> by peff.net with AES256-SHA encrypted SMTP; 5 Jul 2004 05:56:16 -0000

>Date: Mon, 5 Jul 2004 01:56:15 -0400 (EDT)

>From: Jeff King <peff-debbug@peff.net>

>To: 256871@bugs.debian.org

>Subject: Re: Bug#256871: cdrecord triggers memory leak in kernel space

>In-Reply-To: <200407041306.i64D6fPT013802@burner.fokus.fraunhofer.de>

>Message-ID: <Pine.LNX.4.60.0407050154250.19666@localhost.localdomain>

>References: <200407041306.i64D6fPT013802@burner.fokus.fraunhofer.de>

>MIME-Version: 1.0

>Content-Type: TEXT/PLAIN; charset=US-ASCII; format=flowed

>Delivered-To: 256871@bugs.debian.org

>X-Spam-Checker-Version: SpamAssassin 2.60-bugs.debian.org_2004_03_25

> (1.212-2003-09-23-exp) on spohr.debian.org

>X-Spam-Status: No, hits=-3.1 required=4.0 tests=BAYES_44,HAS_BUG_NUMBER

> autolearn=no version=2.60-bugs.debian.org_2004_03_25

>X-Spam-Level:

>

>On Sun, 4 Jul 2004, Joerg Schilling wrote:

>

>

>

>>missing virtual memory space. Linux for this reason kills random processes to

>>regain virtual memory space. If you have a multi processor machine, it is even

>>worse: Linux in this case becomes catatonic and may only recover by a "reset"

>>induced reboot. This is caused by the way a dual or multi processor Linux

>>works.

>>

>>

Re: memory not released after using cdrecord/cdrdao (was: audio cd writing causes massive swap and crash)

Linux-Kernel: Re: memory not released after using cdrecord/cdrdao (was: audio cd writing causes massive swap and crash)

>
>I'm not sure if this is my problem. I do get the OOM killer killing off
>random processes, but it's because the kernel thinks that it is legitimately
>out of memory. Using "free" reports that all memory is in use (and not by
>buffers/cache). I believe the real problem is a leak somewhere in the kernel.
>
>Sorry for the slowness of response on this...I've been out of town all week
>and will be for a few more days. I will try to do some tests and see if I can
>pin down the offending behavior (hopefully in a way that doesn't involve
>frying a CD-R each time!) and will report back.
>
>-Peff
>
>
>

>Information forwarded to debian-bugs-dist@lists.debian.org, Joerg Jaspert
><joerg@debian.org>:
>Bug#256871; Package cdrecord. Full text available.

>Acknowledgement sent to Joerg Schilling <schilling@fokus.fraunhofer.de>:
>Extra info received and forwarded to list. Copy sent to Joerg Jaspert
><joerg@debian.org>. Full text available.

>
>Message received at 256871@bugs.debian.org:
>
>Received: (at 256871) by bugs.debian.org; 6 Jul 2004 09:56:24 +0000
>>From schilling@fokus.fraunhofer.de Tue Jul 06 02:56:24 2004
>Return-path: <schilling@fokus.fraunhofer.de>
>Received: from mailhub.fokus.fraunhofer.de [193.174.154.14]
> by spohr.debian.org with esmtp (Exim 3.35 1 (Debian))
> id 1BhmfS-0008Pi-00; Tue, 06 Jul 2004 02:56:24 -0700
>Received: from burner.fokus.fraunhofer.de (burner [193.175.133.116])
> by mailhub.fokus.fraunhofer.de (8.11.6p2/8.11.6) with ESMTP id i669tCU15351
> for <256871@bugs.debian.org>; Tue, 6 Jul 2004 11:55:12 +0200 (MEST)
>Received: (from jes@localhost)
> by burner.fokus.fraunhofer.de (8.12.9+Sun/8.12.9/Submit) id i669qFwe024525;
> Tue, 6 Jul 2004 11:52:15 +0200 (CEST)
>Date: Tue, 6 Jul 2004 11:52:15 +0200 (CEST)
>From: Joerg Schilling <schilling@fokus.fraunhofer.de>
>Message-Id: <200407060952.i669qFwe024525@burner.fokus.fraunhofer.de>
>To: 256871@bugs.debian.org, peff-debbug@peff.net
>Subject: Re: Bug#256871: cdrecord triggers memory leak in kernel space
>Delivered-To: 256871@bugs.debian.org
>X-Spam-Checker-Version: SpamAssassin 2.60-bugs.debian.org_2004_03_25
> (1.212-2003-09-23-exp) on spohr.debian.org
>X-Spam-Status: No, hits=-3.1 required=4.0 tests=BAYES_44,HAS_BUG_NUMBER
> autolearn=no version=2.60-bugs.debian.org_2004_03_25
>X-Spam-Level:
>

Re: memory not released after using cdrecord/cdrdao (was: audio cd writing causes massive swap and crash)

Linux-Kernel: Re: memory not released after using cdrecord/cdrdao (was: audio cd writing causes massive swap and crash)

>Hi,

>

>I would not believe that there is a big memory leak in Linux.

>

>

>However, you seem to misunderstand me:

>

>The problem is not that Linux incorrectly believes that there is no memory.

>The problem is that Linux *_before_* believes that there is *_plenty_* of memory

>although there is none left over.

>

>Try this:

>

>char buf[16*1024*1024];

>main()

>{

> int i;

> int ret;

>

> for (i = 0; i < 1000; i++) {

> if (ret == 0)

> sleep(1000);

> if (ret < 0) {

> perror("fork");

> printf("i: %d\n", i);

> exit(1);

> }

> }

> }

>

>On Solaris, you will see something like:

>

>fork: Not enough space

>i: 53

>

>On Linux, I assume that you will see no error message.....

>

>If you modify the loop to:

>

> for (i = 0; i < 1000; i++) {

> if (ret == 0) {

> sleep(10);

> memset(buf, 0, sizeof(buf));

> sleep(1000);

> }

>

>...

>

>You will see the same problems as you did describe.

>

>

Re: memory not released after using cdrecord/cdrdao (was: audio cd writing causes massive swap and crash)

Linux-Kernel: Re: memory not released after using cdrecord/cdrdao (was: audio cd writing causes massive swap and crash)

>J?rg

>

>--

> EMail:joerg@schily.isdn.cs.tu-berlin.de (home) J?rg Schilling D-13353 Berlin

> js@cs.tu-berlin.de (uni) If you don't have iso-8859-1

> schilling@fokus.fraunhofer.de (work) chars I am J"org Schilling

> URL: <http://www.fokus.fraunhofer.de/usr/schilling> <ftp://ftp.berlios.de/pub/schily>

>

>

>

>

>-----
>Information forwarded to debian-bugs-dist@lists.debian.org, Joerg Jaspert

><joerg@debian.org>:

>Bug#256871; Package cdrecord. Full text available.

>

>Acknowledgement sent to Jeff King <peff-debbug@peff.net>:

>Extra info received and forwarded to list. Copy sent to Joerg Jaspert

><joerg@debian.org>. Full text available.

>

>

>-----
>Message received at 256871@bugs.debian.org:

>

>Received: (at 256871) by bugs.debian.org; 6 Jul 2004 21:02:11 +0000

>>From peff-debbug@peff.net Tue Jul 06 14:02:11 2004

>Return-path: <peff-debbug@peff.net>

>Received: from 66-23-211-5.clients.speedfactory.net (peff.net) [66.23.211.5]

> by spohr.debian.org with esmtp (Exim 3.35 1 (Debian))

> id 1Bhx4t-00023B-00; Tue, 06 Jul 2004 14:02:11 -0700

>Received: (qmail 1592 invoked from network); 6 Jul 2004 21:02:15 -0000

>Received: from unknown (HELO localhost) (127.0.0.1)

> by peff.net with AES256-SHA encrypted SMTP; 6 Jul 2004 21:02:15 -0000

>Date: Tue, 6 Jul 2004 17:02:15 -0400 (EDT)

>From: Jeff King <peff-debbug@peff.net>

>To: 256871@bugs.debian.org

>Subject: Re: Bug#256871: cdrecord triggers memory leak in kernel space

>In-Reply-To: <200407060952.i669qFwe024525@burner.fokus.fraunhofer.de>

>Message-ID: <Pine.LNX.4.60.0407061701540.22998@localhost.localdomain>

>References: <200407060952.i669qFwe024525@burner.fokus.fraunhofer.de>

>MIME-Version: 1.0

>Content-Type: TEXT/PLAIN; charset=US-ASCII; format=flowed

>Delivered-To: 256871@bugs.debian.org

>X-Spam-Checker-Version: SpamAssassin 2.60-bugs.debian.org_2004_03_25

> (1.212-2003-09-23-exp) on spohr.debian.org

>X-Spam-Status: No, hits=-3.1 required=4.0 tests=BAYES_44,HAS_BUG_NUMBER

> autolearn=no version=2.60-bugs.debian.org_2004_03_25

>X-Spam-Level:

>

>On Tue, 6 Jul 2004, Joerg Schilling wrote:

>

>

>

Re: memory not released after using cdrecord/cdrdao (was: audio cd writing causes massive swap and crash)

Linux-Kernel: Re: memory not released after using cdrecord/cdrdao (was: audio cd writing causes massive swap and crash)

>>The problem is not that Linux incorrectly believes that there is no memory.
>>The problem is that Linux *_before_* believes that there is *_plenty_* of memory
>>although there is none left over.
>>
>>
>
>I understand your complaints about memory overcommits. However, in my case,
>the OOM killing is not the problem, but rather a symptom that comes from not
>realizing that memory is available. The problem is that Linux incorrectly
>believes there is no memory. The sequence is something like this:
>
>1. I have 1G of memory. There is about 850M free (as reported by free, not
>counting buffers). Userland is using 100M or so of that (as reported by ps).
>
>2. I run cdrecord. It burns the CD correctly, then exits.
>
>3. There is now only 100M free. Userland is still using 100M of memory.
>Running cdrecord again causes OOM killing because of memory overcommits.
>
>Where did all of the memory go? If it's in use by the kernel, on whose
>behalf? There are no programs running after the burn that were not running
>before the burn. I believe that the kernel has marked some memory as used on
>behalf of cdrecord and failed to free it when the program exited.
>
>At this point it's just a theory. I will attempt to gather evidence later
>this week when I return home to either justify or disprove my theory.
>
>-Peff
>
>
>
>
>-----
>Information forwarded to debian-bugs-dist@lists.debian.org, Joerg Jaspert
><joerg@debian.org>:
>Bug#256871; Package cdrecord. Full text available.
>-----
>Acknowledgement sent to Jeff King <peff-debbug@peff.net>:
>Extra info received and forwarded to list. Copy sent to Joerg Jaspert
><joerg@debian.org>. Full text available.
>-----
>
>Message received at 256871@bugs.debian.org:
>
>Received: (at 256871) by bugs.debian.org; 9 Jul 2004 07:37:43 +0000
>>From peff-debbug@peff.net Fri Jul 09 00:37:43 2004
>Return-path: <peff-debbug@peff.net>
>Received: from 66-23-211-5.clients.speedfactory.net (peff.net) [66.23.211.5]
> by spohr.debian.org with esmtp (Exim 3.35 1 (Debian))
> id 1Bipx1-0001GI-00; Fri, 09 Jul 2004 00:37:43 -0700
>Received: (qmail 73572 invoked from network); 9 Jul 2004 07:37:48 -0000
>Received: from unknown (HELO coredump.intra.peff.net) (10.0.0.2)

Re: memory not released after using cdrecord/cdrdao (was: audio cd writing causes massive swap and crash)

Linux-Kernel: Re: memory not released after using cdrecord/cdrdao (was: audio cd writing causes massive swap and crash)

> by peff.net with AES256-SHA encrypted SMTP; 9 Jul 2004 07:37:48 -0000
>Date: Fri, 9 Jul 2004 03:37:42 -0400 (EDT)
>From: Jeff King <peff-debbug@peff.net>
>To: 256871@bugs.debian.org
>Subject: Re: Bug#256871: cdrecord triggers memory leak in kernel space
>In-Reply-To: <200407060952.i669qFwe024525@burner.fokus.fraunhofer.de>
>Message-ID: <Pine.LNX.4.58.0407090228540.4297@coredump.intra.peff.net>
>References: <200407060952.i669qFwe024525@burner.fokus.fraunhofer.de>
>MIME-Version: 1.0
>Content-Type: TEXT/PLAIN; charset=US-ASCII
>Delivered-To: 256871@bugs.debian.org
>X-Spam-Checker-Version: SpamAssassin 2.60-bugs.debian.org_2004_03_25
> (1.212-2003-09-23-exp) on spohr.debian.org
>X-Spam-Status: No, hits=-3.1 required=4.0 tests=BAYES_44,HAS_BUG_NUMBER
> autolearn=no version=2.60-bugs.debian.org_2004_03_25
>X-Spam-Level:
>
>OK, I tested a few different configurations. My procedure was more or less to
>run cdrecord -dummy and keep an eye on total memory free (using 'free'). I
>considered the problem to occur if memory usage (minus buffers) increased
>steadily as cdrecord ran, and then stayed high when the program exited. The
>rate of memory gobbling generally seemed to be several hundred kilobytes per
>second (one would guess if it is a leak of a data buffer, that it would be
>dropping about 176400 * 4 (the speed of the writer), which is close to what I
>saw). The machine was otherwise idle.
>
>Here's what I found:
> - cdrecord.mmap versus cdrecord.shm made no difference
> - problem occurred on kernel 2.6.7-1-k7-smp
> - problem also occurred on kernel 2.6.7-1-k7 (non-smp)
> - problem did not occur with -tao, but did occur with -dao; no other command
> line options seemed to effect whether it occurred
>
>Other fun facts:
> - I tried using cdrdao; problem also occurred there
> - all tests were done using ide-cd kernel module (dev=ATA:1,0,0)
>
>I think it's probably not related to the mmap'ing of the buffer memory
>(because it happens with cdrecord.shm), but rather to some problem within the
>kernel CD driver. Seeing what functional differences -dao and -tao have
>(which should hopefully lead us in the right direction) is kind of hard. The
>strace output is just a bunch of ioctls. :) I'm hoping somebody more familiar
>with the code can comment on how the behavior differs, and we can hopefully
>produce a very short program that triggers the kernel behavior.
>
>-Peff
>
>

To unsubscribe from this list: send the line "unsubscribe linux-kernel" in

Re: memory not released after using cdrecord/cdrdao (was: audio cd writing causes massive swap and crash)

Linux-Kernel: Re: memory not released after using cdrecord/cdrdao (was: audio cd writing causes massive swap and crash)

the body of a message to majordomo@vger.kernel.org

More majordomo info at <http://vger.kernel.org/majordomo-info.html>

Please read the FAQ at <http://www.tux.org/lkml/>

Re: memory not released after using cdrecord/cdrdao (was: audio cd writing causes massive swap and crash)