

crash on x86_64 – mm related?

Source: <http://linux.derkeiler.com/Mailing-Lists/Kernel/2005-11/9428.html>

From: Ryan Richter (ryan_at_tau.solarneutrino.net)

Date: 11/29/05

Date: Tue, 29 Nov 2005 10:44:09 -0500
To: linux-kernel@vger.kernel.org

Hi, I booted 2.6.14.2 with the MPT fusion performance fix patch about a week ago on my file server. The machine crashed last night while it was doing backups. You can see the voluminous kernel output below.

Someone else recently had seemingly the same thing happen, but didn't think it was a kernel problem. You can read about it here:
<http://bugs.debian.org/cgi-bin/bugreport.cgi?bug=338335>

I will reply later today with the kernel .config, right now I have to wait for someone to reboot the machine first.

Any help would be appreciated,
–ryan

Bad page state at free_hot_cold_page (in process 'taper', page ffff81000260b6f8)
flags:0x010000000000000c mapping:ffff8100355f1dd8 mapcount:2 count:0
Backtrace:

```
Call Trace:<ffffffffff80159f93>{bad_page+99} <ffffffffff8015a965>{free_hot_cold_page+101}
<ffffffffff80162007>{__page_cache_release+151} <ffffffffff802b8fe8>{sgl_unmap_user_pages+120}
<ffffffffff802b48fb>{release_buffering+27} <ffffffffff802b4fb1>{st_write+1697}
<ffffffffff8017af46>{vfs_write+198} <ffffffffff8017b0a3>{sys_write+83}
<ffffffffff8010db7a>{system_call+126}
```

Trying to fix it up, but a reboot is needed

Bad page state at free_hot_cold_page (in process 'taper', page ffff81000260b6f8)
flags:0x0100000000000081c mapping:ffff81005c0fc310 mapcount:0 count:0
Backtrace:

```
Call Trace:<ffffffffff80159f93>{bad_page+99} <ffffffffff8015a965>{free_hot_cold_page+101}
<ffffffffff80162007>{__page_cache_release+151} <ffffffffff802b8fe8>{sgl_unmap
_user_pages+120}
<ffffffffff802b48fb>{release_buffering+27} <ffffffffff802b4fb1>{st_write+1697}
<ffffffffff8017af46>{vfs_write+198} <ffffffffff8017b0a3>{sys_write+83}
<ffffffffff8010db7a>{system_call+126}
```

Trying to fix it up, but a reboot is needed

----- [cut here] ----- [please bite here] -----
Kernel BUG at include/linux/mm.h:341

Linux-Kernel: crash on x86_64 – mm related?

invalid operand: 0000 [1] SMP

CPU 1

Modules linked in: bonding

Pid: 2418, comm: taper Tainted: G B 2.6.14.2 #1

RIP: 0010:[<ffffffff802b8fcd>] <ffffffff802b8fcd>{sgl_unmap_user_pages+93}

RSP: 0018:ffff810035725e18 EFLAGS: 00010256

RAX: 0000000000000000 RBX: 0000000000000007 RCX: 000000000000000f

RDX: 00000000000000e0 RSI: 0000000000000001 RDI: ffff81000260b6f8

RBP: ffff810004852068 R08: 00000000ffffff R09: 0000000000000000

R10: 0000000000008000 R11: 0000000000000200 R12: 0000000000000008

R13: 0000000000000000 R14: 0000000000008000 R15: ffff810004949d10

FS: 00002aaaab53d880(0000) GS:ffff81004db880(0000) knlGS:0000000556b6920

CS: 0010 DS: 0000 ES: 0000 CR0: 000000008005003b

CR2: 00002aaaaaac0000 CR3: 0000000035691000 CR4: 00000000000006e0

Process taper (pid: 2418, threadinfo ffff810035724000, task ffff81017d680300)

Stack: ffff8101423f3600 ffff810004852000 0000000000000040 0000000000008000

ffff810004949c00 ffffffff802b48fb ffff810004852000 ffffffff802b4fb1

ffff810000000000 ffffffff00000001

Call Trace:<ffffffff802b48fb>{release_buffering+27} <ffffffff802b4fb1>{st_write+1697}

<ffffffff8017af46>{vfs_write+198} <ffffffff8017b0a3>{sys_write+83}

<ffffffff8010db7a>{system_call+126}

Code: 0f 0b 68 ba 12 3a 80 c2 55 01 f0 83 47 08 ff 0f 98 c0 84 c0

RIP <ffffffff802b8fcd>{sgl_unmap_user_pages+93} RSP <ffff810035725e18>

----- [cut here] ----- [please bite here] -----

Kernel BUG at mm/rmap.c:487

invalid operand: 0000 [2] SMP

CPU 1

Modules linked in: bonding

Pid: 2418, comm: taper Tainted: G B 2.6.14.2 #1

RIP: 0010:[<ffffffff8016f3f7>] <ffffffff8016f3f7>{page_remove_rmap+39}

RSP: 0018:ffff810035725ab0 EFLAGS: 00010286

RAX: 00000000ffffff RBX: ffff8100356976f8 RCX: ffff81000000f000

RDX: 0000000000000000 RSI: 8000000064c69067 RDI: ffff81000260b6f8

RBP: 00002aaaaadff000 R08: 0000000000000000 R09: ffff81000260b688

R10: 00000000ffffffa R11: 0000000000000000 R12: ffff810101c22380

R13: 8000000064c69067 R14: ffff81000260b6f8 R15: 0000000000000000

FS: 00002aaaab53d880(0000) GS:ffff81004db880(0000) knlGS:0000000556b6920

CS: 0010 DS: 0000 ES: 0000 CR0: 000000008005003b

CR2: 00002aaaaaac0000 CR3: 0000000035691000 CR4: 00000000000006e0

Process taper (pid: 2418, threadinfo ffff810035724000, task ffff81017d680300)

Stack: ffffffff80166ecd 00002aaaaab62000 ffff810035696aa8 00002aaaaab62000

00002aaaaab62000 00002aaaaab61fff ffff810035695550 00002aaaaab62000

ffff810167180 ffff810035725d68

Call Trace:<fffffff80166ecd>{zap_pte_range+477} <fffffff80167180>{unmap_page_range+496}

<fffffff801672e5>{unmap_vmas+293} <fffffff8016cfa2>{exit_mmap+162}

<fffffff80131ce1>{mmapput+49} <fffffff801371c6>{do_exit+438}

<fffffff8010f6f1>{die+81} <fffffff8010f9df>{do_invalid_op+159}

<fffffff802b8fcd>{sgl_unmap_user_pages+93} <fffffff80381f76>{thread_return+86}

<fffffff802a8662>{sym_setup_data_and_start+402} <fffffff8010e84d>{error_exit+0}

<fffffff802b8fcd>{sgl_unmap_user_pages+93} <fffffff802b8fe8>{sgl_unmap_user_pages+120}

crash on x86_64 – mm related?

Linux-Kernel: crash on x86_64 – mm related?

```
<ffffffff802b48fb>{release_buffering+27} <ffffffff802b4fb1>{st_write+1697}
<ffffffff8017af46>{vfs_write+198} <ffffffff8017b0a3>{sys_write+83}
<ffffffff8010db7a>{system_call+126}
```

Code: 0f 0b 68 9b 35 3a 80 c2 e7 01 48 c7 c6 ff ff ff ff bf 20 00

RIP <ffffffff8016f3f7>{page_remove_rmap+39} RSP <ffff810035725ab0>

<1>Fixing recursive fault but reboot is needed!

Unable to handle kernel NULL pointer dereference at 0000000000000000 RIP:

<ffffffff801b9c7b>{ext3_prepare_write+27}

PGD 355bc067 PUD 355c9067 PMD 0

Oops: 0000 [3] SMP

CPU 0

Modules linked in: bonding

Pid: 2416, comm: driver Tainted: G B 2.6.14.2 #1

RIP: 0010:[<ffffffff801b9c7b>] <ffffffff801b9c7b>{ext3_prepare_write+27}

RSP: 0018:ffff8100355e7b48 EFLAGS: 00010296

RAX: 0000000000000000 RBX: ffffffff8040f660 RCX: 000000000000017d

RDX: 0000000000000094 RSI: ffff81000260b6f8 RDI: ffff810035b09cc0

RBP: 000000000000000e R08: 00000000fffffffa R09: 00000000000000e9

R10: ffff81001190c818 R11: 0000000000000000 R12: ffff81000260b6f8

R13: ffff81000260b6f8 R14: 000000000000017d R15: 0000000000000094

FS: 00002aaaab53d8e0(0000) GS:ffff804db800(0000) knlGS:00000000555bc920

CS: 0010 DS: 0000 ES: 0000 CR0: 000000008005003b

CR2: 0000000000000000 CR3: 0000000035555000 CR4: 000000000000006e

Process driver (pid: 2416, threadinfo ffff8100355e6000, task ffff8100f43e8a80)

Stack: ffff81014e643310 ffffffff8040f660 000000000000000e ffff81000260b6f8

ffff81005c0fc310 0000000000000094 00000000000000e9 ffffffff80158247

0000000000000292 00002aaaaaac0000

Call Trace:<ffffffff80158247>{generic_file_buffered_write+551}

<ffffffff801c06bd>{__ext3_journal_stop+45} <ffffffff8019ec74>{__mark_inode_dirty+52}

<ffffffff801963ec>{inode_update_time+188}

<ffffffff80158a08>{__generic_file_aio_write_nolock+936}

<ffffffff80381f76>{thread_return+86} <ffffffff8013d349>{lock_timer_base+41}

<ffffffff80158cfe>{generic_file_aio_write+110} <ffffffff801b7783>{ext3_file_write+35}

<ffffffff8017ae43>{do_sync_write+211} <ffffffff8018ecc0>{__pollwait+0}

<ffffffff8014a2b0>{autoremove_wake_function+0} <ffffffff8018f681>{sys_select+1153}

<ffffffff8017af46>{vfs_write+198} <ffffffff8017b0a3>{sys_write+83}

<ffffffff8010db7a>{system_call+126}

Code: 48 8b 28 48 89 ef e8 aa 26 00 00 c7 44 24 04 00 00 00 89

RIP <ffffffff801b9c7b>{ext3_prepare_write+27} RSP <ffff8100355e7b48>

CR2: 0000000000000000

<0>Bad page state at prep_new_page (in process 'dumper', page ffff81000260b6f8)

flags:0x010000000000001d mapping:0000000000000000 mapcount:-1 count:1

Backtrace:

Call Trace:<ffffffff80159f93>{bad_page+99} <ffffffff8015a371>{prep_new_page+65}

<ffffffff8015ab2e>{buffered_rmqueue+302} <ffffffff8015ad85>{__alloc_pages+261}

<ffffffff801581bd>{generic_file_buffered_write+413}

<ffffffff80139509>{current_fs_time+105} <ffffffff8019636e>{inode_update_time+62}

<ffffffff80158a08>{__generic_file_aio_write_nolock+936}

crash on x86_64 – mm related?

Linux-Kernel: crash on x86_64 – mm related?

```
<ffffffff8031f4a4>{sock_common_recvmsg+52} <ffffffff8031bb30>{sock_aio_read+272}
<ffffffff80158cfe>{generic_file_aio_write+110} <ffffffff801b7783>{ext3_file_write+35}
<ffffffff8017ae43>{do_sync_write+211} <ffffffff8018ecc0>{__pollwait+0}
<ffffffff8014a2b0>{autoremove_wake_function+0} <ffffffff8018f681>{sys_select+1153}
<ffffffff8017af46>{vfs_write+198} <ffffffff8017b0a3>{sys_write+83}
<ffffffff8010db7a>{system_call+126}
```

Trying to fix it up, but a reboot is needed

Bad page state at prep_new_page (in process 'find', page ffff81000260b6f8)

flags:0x0100000000000064 mapping:ffff8100f3be9be9 mapcount:1 count:1

Backtrace:

```
Call Trace:<ffffffff80159f93>{bad_page+99} <ffffffff8015a371>{prep_new_page+65}
<ffffffff8015ab2e>{buffered_rmqueue+302} <ffffffff8015ad85>{__alloc_pages+261}
<ffffffff8015e7a3>{kmem_getpages+99} <ffffffff8015fbb0>{cache_grow+192}
<ffffffff8015fe3b>{cache_alloc_refill+459} <ffffffff80160226>{kmem_cache_alloc+54}
<ffffffff80193831>{d_alloc+33} <ffffffff80188fe9>{real_lookup+105}
<ffffffff801893c0>{do_lookup+112} <ffffffff80189e07>{__link_path_walk+2551}
<ffffffff8018a382>{link_path_walk+178} <ffffffff8018a8ce>{path_lookup+446}
<ffffffff8018aa9e>{__user_walk+62} <ffffffff801849b6>{vfs_lstat+38}
<ffffffff80184dff>{sys_newlstat+31} <ffffffff8010db7a>{system_call+126}
```

Trying to fix it up, but a reboot is needed

Unable to handle kernel paging request at 00002aaaab9c5b61 RIP:

<ffffffff8015fdb>{cache_alloc_refill+330}

PGD c2512067 PUD c2513067 PMD 0

Oops: 0002 [4] SMP

CPU 0

Modules linked in: bonding

Pid: 3011, comm: find Tainted: G B 2.6.14.2 #1

RIP: 0010:[<ffffffff8015fdb>] <ffffffff8015fdb>{cache_alloc_refill+330}

RSP: 0018:ffff810112f05c28 EFLAGS: 00010082

RAX: 00002aaaab9c5b59 RBX: 0000000000000010 RCX: 0000000000029ba6

RDX: 00002aaaab9c5bb3 RSI: ffff810064c69040 RDI: ffff81000c01a288

RBP: ffff8100f6fc4800 R08: ffff81000c01a250 R09: ffff81000c01a260

R10: 0000000000000000 R11: 0000000000000000 R12: ffff81000c01a240

R13: ffff8100f6fc3640 R14: ffff81000c01a288 R15: 00000000000000d0

FS: 00002aaaae00640(0000) GS:ffff804db800(0000) knlGS:00000000555bc920

CS: 0010 DS: 0000 ES: 0000 CR0: 000000008005003b

CR2: 00002aaaab9c5b61 CR3: 00000000c2bb3000 CR4: 00000000000006e0

Process find (pid: 3011, threadinfo ffff810112f04000, task ffff810102b46040)

Stack: ffff810112f05e68 ffff810179923cb8 ffffffff4 ffff810112f05d28

ffff810179923cb8 ffff810112f05d28 ffff810112f05e68 ffffffff80160226

0000000000000292 ffffffff80193831

```
Call Trace:<ffffffff80160226>{kmem_cache_alloc+54} <ffffffff80193831>{d_alloc+33}
```

```
<ffffffff80188fe9>{real_lookup+105} <ffffffff801893c0>{do_lookup+112}
```

```
<ffffffff80189e07>{__link_path_walk+2551} <ffffffff8018a382>{link_path_walk+178}
```

```
<ffffffff8018a8ce>{path_lookup+446} <ffffffff8018aa9e>{__user_walk+62}
```

```
<ffffffff801849b6>{vfs_lstat+38} <ffffffff80184dff>{sys_newlstat+31}
```

```
<ffffffff8010db7a>{system_call+126}
```

Linux-Kernel: crash on x86_64 – mm related?

```
Code: 48 89 50 08 48 89 02 48 c7 46 08 00 02 20 00 83 7e 24 ff 48
RIP <ffffffff8015fdbba>{cache_alloc_refill+330} RSP <ffff810112f05c28>
CR2: 00002aaaab9c5b61
NMI Watchdog detected LOCKUP on CPU 1
CPU 1
Modules linked in: bonding
Pid: 7, comm: events/1 Tainted: G B 2.6.14.2 #1
RIP: 0010:<ffffffff803837dd> <ffffffff803837dd>{.text.lock.spinlock+118}
RSP: 0018:ffff810004869dd0 EFLAGS: 00000086
RAX: ffff81000c01a240 RBX: ffff81000c01a288 RCX: ffff8100f6fc3640
RDX: 0000000000000003 RSI: 0000000000000003 RDI: ffff81000c01a288
RBP: ffff810100009dc0 R08: 0000000000000000 R09: 0000000000000000
R10: 00000000ffffffff R11: 0000000000000066 R12: 0000000000000000
R13: ffff810100009dd0 R14: 0000000000000292 R15: ffff810100009e40
FS: 00002aaaaae00640(0000) GS:ffffffff804db880(0000) knlGS:00000000556b6920
CS: 0010 DS: 0018 ES: 0018 CR0: 000000008005003b
CR2: 00002aaaaaf1df40 CR3: 000000017f448000 CR4: 000000000000006e0
Process events/1 (pid: 7, threadinfo ffff810004868000, task ffff8100f6fb6080)
Stack: ffffffff8015e35b ffff8100f6fc3640 ffff810100009f60 0000000000000001
      ffff810100009e40 ffff8100f6fc3640 ffff8100f6fc38e0 ffff810100009f88
      ffffffff80161414 ffff810004869e58
Call Trace:<ffffffff8015e35b>{drain_alien_cache+123} <ffffffff80161414>{cache_reap+164}
      <ffffffff80161370>{cache_reap+0} <ffffffff8014553c>{worker_thread+476}
      <ffffffff8012ed70>{default_wake_function+0} <ffffffff8012ed70>{default_wake_function+0}
      <ffffffff80145360>{worker_thread+0} <ffffffff80149c82>{kthread+146}
      <ffffffff8010ea02>{child_rip+8} <ffffffff80145360>{worker_thread+0}
      <ffffffff80149bf0>{kthread+0} <ffffffff8010e9fa>{child_rip+0}
```

```
Code: 80 3f 00 7e f9 e9 59 fe ff ff e8 58 41 e9 ff e9 6f fe ff ff
console shuts up ...
<0>Kernel panic – not syncing: Aiee, killing interrupt handler!
```

–

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