DragonFlyBSD - Bug #2291

Panic in lwkt_remove_tdallq

01/26/2012 11:42 PM - rumcic

Status:	Closed	**Start date:**	01/26/2012
Priority:	Normal	**Due date:**	
Assignee:		**% Done:**	0%
Category:		**Estimated time:**	0.00 hour
Target version:			

Description

With a few days old master and under some I/O load I get the mentioned panic (the actual panic message is illegable, only the trace can be read) ... was unable to get a core dump (panic in ddb didn't produce a dump, only rebooted the machine). The panic seems to be easily repeatable (unfortunately this is the first time I actually saw a trace, because I had DDB_UNATTENDED switched on before + was in X, so could not see the message)

Part of the backtrace (only the function names, I managed to take pictures of the backtrace itself, attached to this report):
panic()
lwkt_remove_tdallq()
free_lock()
acquire_lock()
softdep_disk_io_initiation()
devfs_spec_strategy()
vop_strategy()
vn_strategy()
ufs_strategy()
ufs_vnoperate()
vop_strategy()
vn_strategy()
bawrite()
...
syncer_thread()
syncer_thread_start()
kthread_exit()

Related issues:

Related to Bug #2336: 3.0.3 catchall	**Resolved**	**03/26/2012**

History

#1 - 03/25/2012 04:33 PM - vsrinivas

Also seen on 3.0.2 by tuxillo: http://leaf.dragonflybsd.org/~tuxillo/archive/pics/2291/panic1.png

And on -master by marino and vsrinivas. Callchain can be rooted at kern_exit instead of syncer.

Perhaps some blame goes to softdep's locks. It uses mplock + critical section around softdep callbacks.

#2 - 03/26/2012 10:42 AM - vsrinivas

http://leaf.dragonflybsd.org/~marino/install_panic_1.jpg is marino's panic.

All of these panics are actually: 'td_critcount would go negative', from crit_panic. This appears to come from FREE_LOCK(), which exits a critical section "softupdates"; if one were to exit a critical section one more time than they entered it, that would cause this panic.

#3 - 03/27/2012 12:26 PM - vsrinivas

I believe

http://leaf.dragonflybsd.org/~vsrinivas/0001-kernel-ffs_softdep-Replace-softdep-MPLOCK-critical-s.patch

will correct the problem. Any testing highly appreciated.

Basically, softdep_disk_write_complete() was blocking, losing its lock, but leaving itself marked as the lock-holder. ACQUIRE_LOCK detected this and panic-ed. This patch switches to using lockmgr locks for softdep, which are hard locks and not lost on blocking conditions.

#4 - 03/27/2012 08:26 PM - vsrinivas

I take that back; my patch introduces a deadlock between BUFWAIT and getting the ffs_softdep lock. Working on it now.

#5 - 03/28/2012 11:27 AM - vsrinivas

Second try: http://leaf.dragonflybsd.org/~vsrinivas/ffs03.patch

Reviews/testing very much appreciated!

#6 - 03/29/2012 02:31 AM - vsrinivas

Commit 8e90f899fdf61479c5e76faa87e7ff716982ed08 *should* fix the problem.

#7 - 03/30/2012 08:14 AM - vsrinivas

There are remaining issues with this work, uncovered by fsstress:

1)

```
Fatal trap 12: page fault while in kernel mode
cpuid = 4
fault virtual address   = 0x391
fault code              = supervisor read, page not present
instruction pointer     = 0x2b:0x506bfc
stack pointer           = 0x10:0x1002c197538
frame pointer           = 0x10:0x1002c1978c0
processor eflags        = interrupt enabled, resume, IOPL = 0
current process         = Idle
current thread          = pri 12 (CRIT)
<- SMP: XXX
kernel: type 12 trap, code=0

CPU4 stopping CPUs: 0x00000000000000ef
stopped
Stopped at      0x506bfc:       testb   $0x20,0x391(%rdi)
db> where
bwrite() at 0x506bfc
softdep_fsync_mountdev() at 0x5f097e
buf_rb_tree_RB_SCAN() at 0x51b9aa
softdep_sync_metadata() at 0x5f0395
brelvp() at 0x51b171
vfsync() at 0x51bb9a
ffs_mountfs() at 0x5f6322
vop_fsync() at 0x52bae1
vinvalbuf() at 0x51bec4
vfs_getvfs() at 0x51ef57
mountlist_scan() at 0x51f8f6
mountlist_scan() at 0x51fb7c
db>
---
```

This panic appears to arise from us getting a NULL worklist buffer. This shouldn't be possible, it needs to be examined further.

2)

```
panic: handle_written_inodeblock: not started
cpuid = 0
Trace beginning at frame 0x10005d039e0
panic() at 0x4b1983
panic() at 0x4b1983
softdep_change_directoryentry_offset() at 0x5f1d31
bpdone() at 0x508429
biodone() at 0x5089ac
cluster_awrite() at 0x50f59c
biodone() at 0x508990
gptinit() at 0x671b7d
register_swi() at 0x48a3bd
Debugger("panic")
```

#8 - 04/02/2012 10:14 AM - vsrinivas

- Status changed from New to Resolved

Commit 24624a1562837ae797e4c1b05689f6f5b56006d9 prevents the latter two panics. I believe this issue is resolved.

#9 - 04/22/2012 12:49 PM - rumcic

- Status changed from Resolved to Closed

Unable to reproduce, I believe this issue can be closed

Files

DSC00104.JPG	611 KB	01/27/2012	rumcic
DSC00105.JPG	620 KB	01/27/2012	rumcic