DragonFlyBSD - Bug #2137

vm_prefault: Warning, backing object race averted lobject

09/26/2011 04:56 AM - y0n3t4n1

Status:	New	**Start date:**	
Priority:	Normal	**Due date:**	
Assignee:		**% Done:**	0%
Category:		**Estimated time:**	0.00 hour
Target version:			

Description

Hello.
As I wrote in issue2118, after the fix (1f8fc82a) I started seeing
warning messages on the console under load:

vm_prefault: Warning, backing object race averted lobject 0xffffffe0749d3210

The system used to panic(backtrace at the end of this message) within
a day or two, but with the kernel built from more recent source(fa7cb8ce),
it no longer panics (well, at least for 4 days) running the same load
(pkgsrc bulk build with MAKE_JOBS=5 defined), just the warning messages
shown above all over the console.

#4 0xffffffff8039c50f in panic (
fmt=0xffffffff805edf90 "Bad link elm %p prev->next != elm")
at /usr/src/sys/kern/kern_shutdown.c:833
#5 0xffffffff80532b35 in vm_object_deallocate_locked (
object=0xffffffe0690b2d10) at /usr/src/sys/vm/vm_object.c:601
#6 0xffffffff80532c0f in vm_object_deallocate (object=0xffffffe06b9eed10)
at /usr/src/sys/vm/vm_object.c:451
#7 0xffffffff8052d5bd in vm_map_entry_delete (countp=<optimized out>,
entry=<optimized out>, map=<optimized out>)
at /usr/src/sys/vm/vm_map.c:2608
#8 vm_map_delete (map=0xffffffe05b0d9a00, start=<optimized out>,
end=140737488355328, countp=<optimized out>)
at /usr/src/sys/vm/vm_map.c:2753
#9 0xffffffff8052d661 in vm_map_remove (map=0xffffffe05b0d9a00, start=0,
end=<unavailable>) at /usr/src/sys/vm/vm_map.c:2774
#10 0xffffffff8052d6e9 in vmspace_terminate (vm=0xffffffe05b0d9a00)
at /usr/src/sys/vm/vm_map.c:310
#11 0xffffffff80370579 in _sysref_put (sr=0xffffffe05b0d9bd0)
at /usr/src/sys/kern/kern_sysref.c:325
#12 0xffffffff80388079 in sysref_put (sr=<optimized out>)
at /usr/src/sys/sys/sysref2.h:83
#13 exit1 (rv=<optimized out>) at /usr/src/sys/kern/kern_exit.c:417
#14 0xffffffff803884a4 in sys_exit (uap=<optimized out>)
at /usr/src/sys/kern/kern_exit.c:121
#15 0xffffffff805b0308 in syscall2 (frame=0xffffffe076d33c08)
at /usr/src/sys/platform/pc64/x86_64/trap.c:1188
#16 0xffffffff805a88df in Xfast_syscall ()
at /usr/src/sys/platform/pc64/x86_64/exception.S:320
#17 0x000000000000002b in ?? ()

Best Regards,
YONETANI Tomokazu

History

#1 - 10/04/2011 09:18 PM - y0n3t4n1

With heavier load using 6 jails, I could still trigger this same panic.
The kernel was built from source as of a5fc4.

On Mon, Sep 26, 2011 at 01:54:12PM +0900, YONETANI Tomokazu wrote:

> Hello.
> As I wrote in issue2118, after the fix (1f8fc82a) I started seeing
> warning messages on the console under load:
>
> vm_prefault: Warning, backing object race averted lobject 0xffffffe0749d3210
>
> The system used to panic(backtrace at the end of this message) within
> a day or two, but with the kernel built from more recent source(fa7cb8ce),
> it no longer panics (well, at least for 4 days) running the same load
> (pkgsrc bulk build with MAKE_JOBS=5 defined), just the warning messages
> shown above all over the console.
>
> #4 0xffffffff8039c50f in panic (
> fmt=0xffffffff805edf90 "Bad link elm %p prev->next != elm")
> at /usr/src/sys/kern/kern_shutdown.c:833
> #5 0xffffffff80532b35 in vm_object_deallocate_locked (
> object=0xffffffe0690b2d10) at /usr/src/sys/vm/vm_object.c:601
> #6 0xffffffff80532c0f in vm_object_deallocate (object=0xffffffe06b9eed10)
> at /usr/src/sys/vm/vm_object.c:451
> #7 0xffffffff8052d5bd in vm_map_entry_delete (countp=<optimized out>,
> entry=<optimized out>, map=<optimized out>)
> at /usr/src/sys/vm/vm_map.c:2608
> #8 vm_map_delete (map=0xffffffe05b0d9a00, start=<optimized out>,
> end=140737488355328, countp=<optimized out>)
> at /usr/src/sys/vm/vm_map.c:2753
> #9 0xffffffff8052d661 in vm_map_remove (map=0xffffffe05b0d9a00, start=0,
> end=<unavailable>) at /usr/src/sys/vm/vm_map.c:2774
> #10 0xffffffff8052d6e9 in vmspace_terminate (vm=0xffffffe05b0d9a00)
> at /usr/src/sys/vm/vm_map.c:310
> #11 0xffffffff80370579 in _sysref_put (sr=0xffffffe05b0d9bd0)
> at /usr/src/sys/kern/kern_sysref.c:325
> #12 0xffffffff80388079 in sysref_put (sr=<optimized out>)
> at /usr/src/sys/sys/sysref2.h:83
> #13 exit1 (rv=<optimized out>) at /usr/src/sys/kern/kern_exit.c:417
> #14 0xffffffff803884a4 in sys_exit (uap=<optimized out>)
> at /usr/src/sys/kern/kern_exit.c:121
> #15 0xffffffff805b0308 in syscall2 (frame=0xffffffe076d33c08)
> at /usr/src/sys/platform/pc64/x86_64/trap.c:1188
> #16 0xffffffff805a88df in Xfast_syscall ()
> at /usr/src/sys/platform/pc64/x86_64/exception.S:320
> #17 0x000000000000002b in ?? ()
>
>
> Best Regards,
> YONETANI Tomokazu

#2 - 10/13/2011 09:02 AM - marino

I also saw this kernel message this morning (vm_prefault: Warning, backing object
race averted lobject), and it was after I tested Dillon's patch last night.

No panic, just the message.

#3 - 10/18/2011 04:23 AM - y0n3t4n1

On Thu, Oct 13, 2011 at 09:02:22AM +0000, John Marino (via DragonFly issue tracker) wrote:
>
> I also saw this kernel message this morning (vm_prefault: Warning, backing object
> race averted lobject), and it was after I tested Dillon's patch last night.
>
> No panic, just the message.

I had to wait for several days before it panicked. But I guess
that the warning message itself suggests something bad is happening there,
even though it's not so easy to trigger the panic on other people's
machine?

#4 - 10/25/2011 10:36 AM - y0n3t4n1

On Tue, Oct 18, 2011 at 01:20:35PM +0900, YONETANI Tomokazu wrote:
> On Thu, Oct 13, 2011 at 09:02:22AM +0000, John Marino (via DragonFly issue tracker) wrote:
> >
> > I also saw this kernel message this morning (vm_prefault: Warning, backing object
> > race averted lobject), and it was after I tested Dillon's patch last night.
> >
> > No panic, just the message.
>
> I had to wait for several days before it panicked. But I guess
> that the warning message itself suggests something bad is happening there,
> even though it's not so easy to trigger the panic on other people's
> machine?

After upgrading to kernel built from the source as of e6b6e (I assume
this is the `Dillon's patch' mentioned above, right?), I haven't caught
panic for more than four days. However, I see the following message
popping up on the console:
Oct 25 01:24:18 atom64 kernel: refcount_wait objde2 long wait
Oct 25 01:26:19 atom64 last message repeated 2 times
Oct 25 01:36:19 atom64 last message repeated 10 times
Oct 25 01:46:19 atom64 last message repeated 10 times
Oct 25 01:56:19 atom64 last message repeated 10 times
Oct 25 02:06:19 atom64 last message repeated 10 times

And there's a few unkillable processes remaining, whose kernel backtrace
looks like this:

#2 0xffffffff803aa711 in _refcount_wait (countp=0xffffffe06a67a1b8,
wstr=0xffffffff80665b31 "objde2") at /usr/src/sys/kern/kern_refcount.c:82
#3 0xffffffff8053b954 in refcount_wait (wstr=<optimized out>,
countp=<optimized out>) at /usr/src/sys/sys/refcount.h:112
#4 vm_object_pip_wait (waitid=<optimized out>, object=<optimized out>)
at /usr/src/sys/vm/vm_object.h:261
#5 vm_object_deallocate_locked (object=0xffffffe06a67a160)
at /usr/src/sys/vm/vm_object.c:573
#6 0xffffffff8053b5a6 in vm_object_collapse (object=0xffffffe060e38370)
at /usr/src/sys/vm/vm_object.c:2003
#7 0xffffffff8053ba69 in vm_object_deallocate_locked (

object=0xffffffe060e38370) at /usr/src/sys/vm/vm_object.c:638
#8 0xffffffff8053beaa in vm_object_deallocate (object=0xffffffe061885210)
at /usr/src/sys/vm/vm_object.c:499
#9 0xffffffff80535a96 in vm_map_entry_delete (countp=<optimized out>,
entry=<optimized out>, map=<optimized out>)
at /usr/src/sys/vm/vm_map.c:2641
#10 vm_map_delete (map=0xffffffe05c518680, start=<optimized out>,
end=140737488355328, countp=<optimized out>)
at /usr/src/sys/vm/vm_map.c:2787
#11 0xffffffff80535b47 in vm_map_remove (map=0xffffffe05c518680, start=0,
end=<unavailable>) at /usr/src/sys/vm/vm_map.c:2809
#12 0xffffffff80388dc5 in exec_new_vmspace (imgp=0xffffffe05eec7918,
vmcopy=0x0) at /usr/src/sys/kern/kern_exec.c:785
#13 0xffffffff8036dbba in exec_elf64_imgact (imgp=0xffffffe05eec7918)
at /usr/src/sys/kern/imgact_elf.c:694
#14 0xffffffff8038944f in kern_execve (nd=0xffffffe05eec7a98,
args=<optimized out>) at /usr/src/sys/kern/kern_exec.c:304
#15 0xffffffff80389ef8 in sys_execve (uap=0xffffffe05eec7b58)
at /usr/src/sys/kern/kern_exec.c:607
#16 0xffffffff805bc282 in syscall2 (frame=0xffffffe05eec7c08)
at /usr/src/sys/platform/pc64/x86_64/trap.c:1188