DragonFlyBSD - Bug #2099

page fault panic in vm system

07/05/2011 07:27 PM - pavalos

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

Description

I got a panic 2 nights in a row (both at around 2:55am which is weird).

```
DragonFly ylem.theshell.com 2.11-DEVELOPMENT DragonFly
v2.11.0.208.g14dc24-DEVELOPMENT #53: Tue May 17 12:52:46 HST 2011
root@ylem.theshell.com:/usr/obj/usr/src/sys/YLEM  i386

Fatal trap 12: page fault while in kernel mode
cpuid = 0; lapic.id = 00000000
fault virtual address   = 0x14
fault code              = supervisor read, page not present
instruction pointer     = 0x8:0xc02f0d84
stack pointer           = 0x10:0xe8007bd8
frame pointer           = 0x10:0xe8007be0
code segment            = base 0x0, limit 0xfffff, type 0x1b
= DPL 0, pres 1, def32 1, gran 1
processor eflags        = interrupt enabled, resume, IOPL = 0
current process         = 36365 (mysqld)
current thread          = pri 10
<- SMP: XXX
trap number             = 12
panic: page fault
cpuid = 0

(kgdb) bt
#0  _get_mycpu (di=0xc0441aa0) at ./machine/thread.h:79
#1  md_dumpsys (di=0xc0441aa0) at /usr/src/sys/platform/pc32/i386/dump_machdep.c:264
#2  0xc01a1388 in dumpsys () at /usr/src/sys/kern/kern_shutdown.c:927
#3  0xc01a199e in boot (howto=<value optimized out>) at /usr/src/sys/kern/kern_shutdown.c:389
#4  0xc01a1c01 in panic (fmt=0xc0394f68 "%s") at /usr/src/sys/kern/kern_shutdown.c:833
#5  0xc03618b7 in trap_fatal (frame=0xe8007b90, eva=<value optimized out>) at /usr/src/sys/platform/pc32/i386/trap.c:1118
#6  0xc03619dc in trap_pfault (frame=0xe8007b90, usermode=<value optimized out>, eva=<value optimized out>) at
/usr/src/sys/platform/pc32/i386/trap.c:1020
#7  0xc0361f63 in trap (frame=0xe8007b90) at /usr/src/sys/platform/pc32/i386/trap.c:707
#8  0xc033b367 in calltrap () at /usr/src/sys/platform/pc32/i386/exception.s:787
#9  0xc02f0d84 in vm_page_rb_tree_RB_LOOKUP (head=0x14, value=4) at /usr/src/sys/vm/vm_page.c:109
#10 0xc02f1a90 in vm_page_lookup (object=0x0, pindex=4) at /usr/src/sys/vm/vm_page.c:521
#11 0xc02e5826 in vm_prefault (map=0xedac66f8, vaddr=852598784, fault_type=<value optimized out>, fault_flags=<value
optimized out>)
at /usr/src/sys/vm/vm_fault.c:2050
#12 vm_fault (map=0xedac66f8, vaddr=852598784, fault_type=<value optimized out>, fault_flags=<value optimized out>) at
/usr/src/sys/vm/vm_fault.c:435
#13 0xc036196c in trap_pfault (frame=0xe8007d40, usermode=<value optimized out>, eva=<value optimized out>) at
/usr/src/sys/platform/pc32/i386/trap.c:1001
#14 0xc0361da4 in trap (frame=0xe8007d40) at /usr/src/sys/platform/pc32/i386/trap.c:589
#15 0xc033b367 in calltrap () at /usr/src/sys/platform/pc32/i386/exception.s:787
#16 0x0847adfc in ?? ()
Backtrace stopped: previous frame inner to this frame (corrupt stack?)
```

Files can be downloaded from:
http://www.theshell.com/~pavalos/crash/

It's crash28 and crash29.

--Peter

History

#1 - 07/09/2011 08:55 PM - pavalos

On Tue, Jul 05, 2011 at 08:47:46AM -1000, Peter Avalos wrote:
> I got a panic 2 nights in a row (both at around 2:55am which is weird).
>

I got another panic last night with mysqld as the current process. The
kernel for this one is:

DragonFly ylem.theshell.com 2.11-DEVELOPMENT DragonFly
v2.11.0.491.g96956-DEVELOPMENT #54: Tue Jul 5 23:12:44 HST 2011
root@ylem.theshell.com:/usr/obj/usr/src/sys/YLEM i386

I wasn't able to get a core, but this was on the console. Not sure if
much can be gained from this:

```
Fatal trap 12: page fault while in kernel mode
cpuid = 1; lapic.id = 02000000
fault virtual address   = 0x14
fault code              = supervisor read, page not present
instruction pointer     = 0x8:0xc02f2e20
stack pointer           = 0x10:0xecb89bd8
frame pointer           = 0x10:0xecb89be0
code segment            = base 0x0, limit 0xfffff, type 0x1b
= DPL 0, pres 1, def32 1, gran 1
processor eflags        = interrupt enabled, resume, IOPL = 0
current process         = 1017 (mysqld)
current thread          = pri 10
<- SMP: XXX
trap number             = 12
panic: page fault
cpuid = 1
Trace beginning at frame 0xecb89ab0
panic(ffffffff,1,c0397168,ecb89ae4,dd185060) at panic+0x198
panic(c0397168,c03a8d3e,d8abdd30,1,1) at panic+0x198
trap_fatal(dff60078,0,1,0,dff60078) at trap_fatal+0x37c
trap_pfault(ecb89b68,ff809000,3,ff81b000,ecb89b78) at trap_pfault+0x111
trap(ecb89b90) at trap+0x54b
calltrap() at calltrap+0xd
--- trap 0, eip = 0, esp = 0xecb89bd4, ebp = 0x4 ---
boot() called on cpu#1
Uptime: 3d6h38m40s
Physical memory: 3057 MB
Dumping 698 MB:
```

--Peter

#2 - 07/09/2011 11:02 PM - dillon

:On Tue, Jul 05, 2011 at 08:47:46AM -1000, Peter Avalos wrote:
:> I got a panic 2 nights in a row (both at around 2:55am which is weird).
:>=20
:
:I got another panic last night with mysqld as the current process. The
:kernel for this one is:
:
:DragonFly ylem.theshell.com 2.11-DEVELOPMENT DragonFly
:v2.11.0.491.g96956-DEVELOPMENT #54: Tue Jul 5 23:12:44 HST 2011
:root@ylem.theshell.com:/usr/obj/usr/src/sys/YLEM i386
:
:I wasn't able to get a core, but this was on the console. Not sure if
:much can be gained from this:
:

```
:Fatal trap 12: page fault while in kernel mode
:cpuid =3D 1; lapic.id =3D 02000000
:fault virtual address   =3D 0x14
:fault code              =3D supervisor read, page not present
:instruction pointer     =3D 0x8:0xc02f2e20
```

This is vm_page_rb_tree_RB_LOOKUP() .. the red-black lookup code
for a vm_page in an object. Called from vm_page_lookup().

Currently the global vm_token protects that field. I didn't see any
particular race there but the problem could be related to a more
general reference count problem with VM objects.

-Matt

#3 - 07/09/2011 11:16 PM - pavalos

On Sat, Jul 09, 2011 at 04:01:29PM -0700, Matthew Dillon wrote:
```
>
>    This is vm_page_rb_tree_RB_LOOKUP() .. the red-black lookup code
>    for a vm_page in an object.  Called from vm_page_lookup().
>
>    Currently the global vm_token protects that field.  I didn't see any
>    particular race there but the problem could be related to a more
>    general reference count problem with VM objects.
>
```

Ok, so it looks like it's the original panic then. Are the vmcores any
help to figure out what's going on?

--Peter

#4 - 07/10/2011 08:51 AM - pavalos

On Sat, Jul 09, 2011 at 01:15:21PM -1000, Peter Avalos wrote:
```
> On Sat, Jul 09, 2011 at 04:01:29PM -0700, Matthew Dillon wrote:
> >
> >    This is vm_page_rb_tree_RB_LOOKUP() .. the red-black lookup code
> >    for a vm_page in an object.  Called from vm_page_lookup().
> >
> >    Currently the global vm_token protects that field.  I didn't see any
> >    particular race there but the problem could be related to a more
> >    general reference count problem with VM objects.
> >
>
> Ok, so it looks like it's the original panic then.  Are the vmcores any
> help to figure out what's going on?
>
```

I just got the same panic on a more recent kernel, but the backtrace
looks the same. It's available at:

http://www.theshell.com/~pavalos/crash/crash31.tar.xz

--Peter