DragonFlyBSD - Bug #1580

Panic (Fatal trap 12: page fault while in kernel mode) while playing with pf and netif names

10/20/2009 06:52 AM - rumcic

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

Description

Not quite sure if the problem is in either of the two aforementioned acts or in
the combination, but while doing that (with netif names I'm thinking
of "ifconfig some_netif name new_name") I was able to panic the machine.

After getting "kernel: Non-unique normal route, mask not entered" in the dmesg
(immediatelly after loading my custom pf.conf - do note, that it does not
happen all the time, loaded the rules quite a few times to get this message),
I loaded my pf.conf again (pfctl -f /etc/pf.conf) and the machine fortunately
panicked for me (in the previous two instances of this problem, the machine
hanged and was completely unresponsive ... to ping, serial console, everything
was dead).

The dump is located at leaf:~rumko/crash/pf/ and the backtrace is:
#0 dumpsys () at ./machine/thread.h:83
#1 0xc0209e2d in boot (howto=260) at /usr/src/sys/kern/kern_shutdown.c:378
#2 0xc020a0f2 in panic (fmt=0xc0403d38 "%s")
at /usr/src/sys/kern/kern_shutdown.c:813
#3 0xc03d29f6 in trap_fatal (frame=0xddf26770, eva=<value optimized out>)
at /usr/src/sys/platform/pc32/i386/trap.c:1093
#4 0xc03d2b30 in trap_pfault (frame=0xddf26770, usermode=0, eva=305406860)
at /usr/src/sys/platform/pc32/i386/trap.c:994
#5 0xc03d314e in trap (frame=0xddf26770)
at /usr/src/sys/platform/pc32/i386/trap.c:674
#6 0xc03bd927 in calltrap ()
at /usr/src/sys/platform/pc32/i386/exception.s:785
#7 0xc02a3d21 in rn_walktree (h=0xc39f9080, f=0xc029d5db <pfr_walktree>,
w=0xddf267e0) at /usr/src/sys/net/radix.c:996
#8 0xc029bbe4 in pfr_mark_addrs (kt=0xdddf3c00)
at /usr/src/sys/net/pf/pf_table.c:723
#9 0xc029f1c1 in pfr_commit_ktable (kt=0xdddf3c00, tzero=1255984339)
at /usr/src/sys/net/pf/pf_table.c:1596
#10 0xc029f44a in pfr_ina_commit (trs=0xddf26a18, ticket=4, nadd=0x0,
nchange=0x0, flags=<value optimized out>)
at /usr/src/sys/net/pf/pf_table.c:1566
#11 0xc029792d in pfioctl (ap=0xddf26b8c)
at /usr/src/sys/net/pf/pf_ioctl.c:2669
#12 0xc01f05fe in dev_dioctl (dev=0xc3a24b80, cmd=3222029394,
data=0xddf26c18 "\r", fflag=3, cred=0xdd8f1478, msg=0xddf26cf0)
at /usr/src/sys/kern/kern_device.c:174
#13 0xc032e7ef in devfs_specf_ioctl (fp=0xdd3e5958, com=3222029394,
data=0xddf26c18 "\r", ucred=0xdd8f1478, msg=0xddf26cf0)
at /usr/src/sys/vfs/devfs/devfs_vnops.c:1354
#14 0xc022d92d in mapped_ioctl (fd=3, com=3222029394, uspc_data=0xbfbff584
<Address 0xbfbff584 out of bounds>, map=0x0, msg=0xddf26cf0)
at /usr/src/sys/sys/file2.h:88
#15 0xc022d9b6 in sys_ioctl (uap=0xddf26cf0)
at /usr/src/sys/kern/sys_generic.c:521
#16 0xc03d37a9 in syscall2 (frame=0xddf26d40)
at /usr/src/sys/platform/pc32/i386/trap.c:1339
#17 0xc03bd9d6 in Xint0x80_syscall ()
at /usr/src/sys/platform/pc32/i386/exception.s:876
#18 0x2811ac63 in ?? ()

Backtrace stopped: previous frame inner to this frame (corrupt stack?)
--
Regards,
Rumko

History

#1 - 09/10/2010 01:57 PM - rumcic

After some netif have been renamed (was not able to reproduce panic on a machine
where no netif have been renamed) I can easily panic the machine with "pfctl -s
all; pfctl -f /etc/pf.conf; pfctl -f /etc/pf.conf" ... after the first reload of
the pf.conf, I get "kernel: Non-unique normal route, mask not entered" and the
second reload panics the machine (still working on getting a dump):
Sep 10 15:49:18 zeus kernel: Non-unique normal route, mask not entered

Fatal trap 12: page fault while in kernel mode
mp_lock = 00000000; cpuid = 0; lapic.id = 00000000
fault virtual address = 0x12342378
fault code = supervisor read, page not present
instruction pointer = 0x8:0xc0297874
stack pointer = 0x10:0xd80ad9e0
frame pointer = 0x10:0xd80ada08
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 = Idle
current thread = pri 12
<- SMP: XXX
kernel: type 12 trap, code=0

CPU0 stopping CPUs: 0x00000002
stopped
Stopped at rn_match+0xf6: movl 0(%eax),%edi
db> trace
rn_match(c04d1c0c,d85e7100) at rn_match+0xf6
pfr_match_addr(db0e1000,db63aad0,2) at pfr_match_addr+0x4c
pf_match_translation(d80adbc4,db63aa00,14,2,daddade0) at pf_match_translation+0x1b2
pf_get_translation(d80adbc4,db63aa00,14,2,daddade0) at pf_get_translation+0x62
pf_test_rule(d80adc28,d80adc24,2,daddade0,db63aa00) at pf_test_rule+0x2b2
pf_test(2,d7be3000,d80adc7c,0,0) at pf_test+0x506
pf_check_out(0,d80adc7c,d7be3000,2,db63aa00) at pf_check_out+0x2e
pfil_run_hooks(c04f9b44,d80adcdc,d7be3000,2,d7be3000) at pfil_run_hooks+0x83
ip_output(db63aa00,0,dafec08c,10000,0) at ip_output+0x8c4
udp_send(c3c51888,0,db63aa00,0,0) at udp_send+0x271
netmsg_pru_send(db68eb80,0,c046cea0,d80add84,c02975c0) at netmsg_pru_send+0x1c
netmsg_service(db68eb80,1,0,c04fa2e0,ff800000) at netmsg_service+0xe0
netmsg_service_loop(c046cea0,0,0,0,0) at netmsg_service_loop+0x18
lwkt_exit() at lwkt_exit

#2 - 09/10/2010 07:22 PM - dillon

:all; pfctl -f /etc/pf.conf; pfctl -f /etc/pf.conf" ... after the first reload of
:the pf.conf, I get "kernel: Non-unique normal route, mask not entered" and the
:second reload panics the machine (still working on getting a dump):
:Sep 10 15:49:18 zeus kernel: Non-unique normal route, mask not entered

PF is going to be unstable until Jan tokenizes it. All of its entry
points are MP now and PF itself is not.

-Matt

#3 - 09/13/2010 11:40 AM - lentferj

Do you have the possibility to try out my working branch?
http://gitweb.dragonflybsd.org/~lentferj/dragonfly.git/shortlog/refs/heads/pf_mpsafe

I tried

ifconfig re0 name lan0
pfctl -s all
pfctl -f /etc/pf.conf
pfctl -f /etc/pf.conf
pfctl -f /etc/pf.conf

on a 4core VM with that branch and didn't experience any failure.

#4 - 10/25/2010 10:41 PM - rumcic

A bit more info...
it seems before I wasn't running a verbose boot 'cause with latest master I get
kernel: pfr_unroute_kentry: delete failed.
last message repeated 7 times
kernel: Non-unique normal route, mask not entered
and after that a panic ... so it seems it can't delete a few entries ...
rn_delete called from pfr_unroute_kentry() in sys/net/pf/pf_table.c returns null?
after that I guess that the same entries that couldn't be deleted are added through rn_addroute() and rn_addroute() at that time spits out "Non-unique normal route, mask not entered" ... then pf gets confused and panics? have not researched further than that yet

#5 - 12/21/2018 01:21 AM - martin1234

- Description updated