
Subject: Re: OOM didn't save the machine
Posted by [lazy](#) on Mon, 30 Mar 2009 21:41:06 GMT
[View Forum Message](#) <> [Reply to Message](#)

thanks to sysreq i managed to reboot the machine without a crash, I didnt find anything interesting in proc beside filp beaing eaten in slab and this

CPU 1, VCPU 3000:0

Modules linked in: vzethdev(U) vznetdev(U) simfs(U) vzrst(U) ip_nat(U) vzcpt(U) ip_contrack(U) nfnetlink(U) ipip(U) tunnel4(U) tun(U) vzdquota(U) vzmon(U) vzdev(U) xt_tcpudp(U) xt_length(U) ipt_ttl(U) xt_tcpmss(U) ipt_TCPMSS(U) iptable_mangle(U) iptable_filter(U) xt_multiport(U) xt_limit(U) ipt_tos(U) ipt_REJE CT(U) ip_tables(U) x_tables(U) button(U) dm_snapshot(U) dm_mirror(U) dm_mod(U) mptctl(U) loop(U) sg(U) sr_mod(U) cdrom(U) sd_mod(U) ehci_hcd(U) ata_piix(U) libata(U) tg3(U) uhci_hcd(U) mptsas(U) mptscsih(U) mptbase(U) scsi_transport_sas(U) scsi_mod(U)

Pid: 10722, comm: httpd Tainted: P

^^^^^^^^ this is the unkillable 100% cpu eating monster

2.6.18-92.1.18.el5.028stab060.2 #1 028stab060

RIP: 0060:[<ffffff8004a9b7>] [<ffffff8004a9b7>] unix_stream_sendmsg+0x194/0x3d7

RSP: 0000:ffff8101f250fb88 EFLAGS: 00000203

RAX: 0000000000000000 RBX: ffff8101f250fee8 RCX: 000000000000003b8

RDX: ffffffffef0 RSI: ffff81011e35b4c0 RDI: ffff81012fbae000

RBP: 0000000000000227 R08: ffff8101c476cb80 R09: 0000000000000286

R10: 000053fe5fafdbc6 R11: ffff8101f250fb38 R12: 0000000000000000

R13: ffff8101f34b50c0 R14: 0000000000000000 R15: ffff8101b17cbc80

FS: 0000000000000000(0000) GS:ffff81022f494b40(0033) knlGS:00000000b7bd06c0

CS: 0060 DS: 007b ES: 007b CR0: 000000008005003b

CR2: 00000000b6d19030 CR3: 00000001f2634000 CR4: 000000000000006e0

Call Trace:

<NMI> <<EOE>> [<ffffff8001df4c>] __pollwait+0x0/0xe1

[<ffffff80055a09>] sock_sendmsg+0xd4/0xec

[<ffffff8003f92b>] memcpy_toiovec+0x36/0x66

[<ffffff80064edf>] _spin_lock_bh+0x9/0x14

[<ffffff800960cc>] autoremove_wake_function+0x0/0x2e

[<ffffff801de53f>] cmsghdr_from_user_compat_to_kern+0x180/0x20b

[<ffffff801ca06a>] sys_sendmsg+0x217/0x28a

[<ffffff8000c801>] do_sync_read+0xc7/0x104

[<ffffff80064edf>] _spin_lock_bh+0x9/0x14

[<ffffff800960cc>] autoremove_wake_function+0x0/0x2e

[<ffffff801ddd6c>] compat_sys_socketcall+0x159/0x172

[<ffffff800615fa>] ia32_sysret+0x0/0xa

this was the unkillable process, this apache process recives http requests and sends sockets to other http processes using sendmsg, there can be some error in this msghdr can be wrong or there might be some loop but it shouldn't kill the machine.

FD_SETSIZE is raised to 2048 if it makes any difference.

I think i might be able to reproduce some of the flow from this trace, to bad I did sysreq p only once

Basicly it should be sthing like poll(), accept a socket, sendmsg the socket to another process threw one of sockets created by socketpair(PF_UNIX, SOCK_STREAM, 0, socks)

Same code runs without problems on non openvz kernels 2.6.22 25 27 for months

Any pointers how to debug it ?
