Subject: TUN causing instability.
Posted by KuJoe on Wed, 30 Nov 2011 08:42:55 GMT
View Forum Message <> Reply to Message

We recently migrated to new hardware nodes with the latest OpenVZ kernel and started to experience instability with TUN. At first, all live migrations with VPSs that have TUN enabled caused a kernel panic. Now when we enable TUN on a VPS and they run "tunctl -t tun0" they get the error ""enabling TUNSETPERSIST: Operation not permitted" and it causes the loads to spike over 300.0 which requires the node to be forcefully rebooted.

Is it really that easy to crash a whole OpenVZ node with a single command? Any ideas how to fix this?

Kernel: 2.6.18-274.7.1.el5.028stab095.1
vzctl version 3.0.29.3

Subject: Re: TUN causing instability.
Posted by KuJoe on Fri, 02 Dec 2011 05:48:59 GMT
View Forum Message <> Reply to Message

Any ideas of where to look? I've checked all logs, top, ps, and iotop, but there is no sign of what is causing the high loads. I moved the client to his own node with different hardware and the latest kernel but the problem continues.

Subject: Re: TUN causing instability.
Posted by KuJoe on Fri, 02 Dec 2011 06:56:07 GMT
View Forum Message <> Reply to Message

Resolved by using an older kernel.

Subject: Re: TUN causing instability.
Posted by Bryon on Sun, 15 Jan 2012 22:05:23 GMT
View Forum Message <> Reply to Message

We're experiencing the same issue with multiple servers running the latest OpenVZ kernel. Which older kernel did you switch to?

If a container (with TUN enabled) runs "tunctl" while the system is actively in use by customers, it begins to receive many hung / blocked task log messages. On a busy system, the load immediately spikes. Immediately after running tunctl no new connections to the system can be made (e.g. SSH) and tunctl never returns, it only outputs "enabling TUNSETPERSIST: Operation not permitted."

On a busy system the server will eventually crash with many lines similar to those below within

messages:

Jan 14 20:44:40 x22 kernel: INFO: task irqbalance:9026 blocked for more than 300 seconds.
Jan 14 20:44:40 x22 kernel: "echo 0 > /proc/sys/kernel/hung_task_timeout_secs" disables this message.
Jan 14 20:44:40 x22 kernel: irqbalance D ffff81043b8e2ba0 0 9026 1 9061 8885 (NOTLB)
Jan 14 20:44:40 x22 kernel: ffff81042c73bd78 0000000000000086 3933323036383031 00002b3d84fe2035
Jan 14 20:44:40 x22 kernel: ffff81043b8e2ba0 ffffffff8031bba0 0004e24c1dee8030 000bbd68d82c77c7
Jan 14 20:44:40 x22 kernel: ffff81043b8e2da8 0000000000030002 0000000000000000 ffffffff804a6280
Jan 14 20:44:40 x22 kernel: Call Trace:
Jan 14 20:44:40 x22 kernel: [<ffffffff8006520d>] __mutex_lock_slowpath+0x60/0x9b
Jan 14 20:44:40 x22 kernel: [<ffffffff8023215e>] dev_name_hash+0x1e/0x64
Jan 14 20:44:40 x22 kernel: [<ffffffff80065257>] .text.lock.mutex+0xf/0x14
Jan 14 20:44:40 x22 kernel: [<ffffffff80232598>] dev_load+0x18/0x46
Jan 14 20:44:40 x22 kernel: [<ffffffff80232cb0>] dev_ioctl+0x317/0x497
Jan 14 20:44:40 x22 kernel: [<ffffffff802277fa>] sock_ioctl+0x1d4/0x1e5
Jan 14 20:44:40 x22 kernel: [<ffffffff80043f2e>] do_ioctl+0x21/0x6b
Jan 14 20:44:40 x22 kernel: [<ffffffff8003154a>] vfs_ioctl+0x457/0x4b9
Jan 14 20:44:40 x22 kernel: [<ffffffff800c37ca>] audit_syscall_entry+0x1a8/0x1d3
Jan 14 20:44:40 x22 kernel: [<ffffffff8004ec3e>] sys_ioctl+0x3c/0x5c
Jan 14 20:44:40 x22 kernel: [<ffffffff800602dd>] tracesys+0xd5/0xe0
Jan 14 20:44:40 x22 kernel:
Jan 14 20:44:40 x22 kernel: INFO: task tunctl:22859 blocked for more than 300 seconds.
Jan 14 20:44:40 x22 kernel: "echo 0 > /proc/sys/kernel/hung_task_timeout_secs" disables this message.
Jan 14 20:44:40 x22 kernel: tunctl D ffff81037b0b61a0 0 22859 22784 (L-TLB)
Jan 14 20:44:40 x22 kernel: ffff8101c0635d68 0000000000000046 0000000100000000 ffff8102003cda08
Jan 14 20:44:40 x22 kernel: ffff81037b0b61a0 ffff81033bcf32e0 0004e2443ca8806b 000bbd55e7240efd
Jan 14 20:44:40 x22 kernel: ffff81037b0b63a8 0000000300000000 ffff8101da059980 ffff8102f59a0000
Jan 14 20:44:40 x22 kernel: Call Trace:
Jan 14 20:44:40 x22 kernel: [<ffffffff8006520d>] __mutex_lock_slowpath+0x60/0x9b
Jan 14 20:44:40 x22 kernel: [<ffffffff80065257>] .text.lock.mutex+0xf/0x14
Jan 14 20:44:40 x22 kernel: [<ffffffff88a1da20>] :tun:tun_chr_close+0x32/0x78
Jan 14 20:44:40 x22 kernel: [<ffffffff800128ef>] __fput+0xd3/0x1c2
Jan 14 20:44:40 x22 kernel: [<ffffffff800248f8>] filp_close+0x5c/0x64
Jan 14 20:44:40 x22 kernel: [<ffffffff8003a94e>] put_files_struct+0x63/0xae
Jan 14 20:44:40 x22 kernel: [<ffffffff80015ba0>] do_exit+0x74c/0xe2d
Jan 14 20:44:40 x22 kernel: [<ffffffff800c37ca>] audit_syscall_entry+0x1a8/0x1d3
Jan 14 20:44:40 x22 kernel: [<ffffffff8004b625>] cpuset_exit+0x0/0x88
Jan 14 20:44:40 x22 kernel: [<ffffffff800602dd>] tracesys+0xd5/0xe0

Jan 14 20:44:40 x22 kernel:

On a system with low activity nothing appears within any monitoring utilities (e.g. top, iostat) or within messages/dmesg.

---

Subject: Re: TUN causing instability.
Posted by Bryon on Sun, 15 Jan 2012 22:08:14 GMT
View Forum Message <> Reply to Message

I forgot to mention we're using kernel 2.6.18-274.7.1.el5.028stab095.1 - The same as KuJoe.

I can also say that we've experienced many "crashes" during live migrations. We're fairly certain the crashes occurred during the migration of a container with TUN enabled and most likely in use. We've been so far unable to retrieve a kernel panic because we do not have physical access to the systems.

---

Subject: Re: TUN causing instability.
Posted by KuJoe on Thu, 19 Jan 2012 04:10:33 GMT
View Forum Message <> Reply to Message

We reverted back to 2.6.18-238.19.1.el5.028stab092.2. This was the latest kernel we've found to not allow clients to crash the hardware node.

---

Subject: Re: TUN causing instability.
Posted by Bryon on Thu, 19 Jan 2012 04:14:09 GMT
View Forum Message <> Reply to Message

Thanks for letting me know. We've also downgraded to 2.6.18-238.19.1.el5.028stab092.2 on servers with clients that require TUN.

---

Subject: Re: TUN causing instability.
Posted by Bryon on Tue, 13 Mar 2012 17:29:22 GMT
View Forum Message <> Reply to Message

Looks like this may have been fixed in 028stab099.3 --

Lookup bugzilla #2207 - Can't include a link in my reply.

---

Subject: Re: TUN causing instability.
Posted by KuJoe on Thu, 29 Mar 2012 04:28:01 GMT
View Forum Message <> Reply to Message

My node running 99.3 just received a kernel panic from migrating a VPS using TUN. Next time it happens I'll take a screenshot. Back to 92.2 I go.

Subject: Re: TUN causing instability.
Posted by KuJoe on Thu, 29 Mar 2012 04:48:22 GMT
View Forum Message <> Reply to Message

Kernel panic #2. Screenshot attached.

File Attachments
1) 004.png, downloaded 320 times