Subject: Re: Hung Tasks on NFS (maybe not a OpenVZ Problem) - How to forcefully kill a container ?
Posted by Todd Lyons on Fri, 30 Mar 2012 12:40:14 GMT
View Forum Message <> Reply to Message

On Fri, Mar 30, 2012 at 4:03 AM, Sirk Johannsen
<s.johannsen@satzmedia.de> wrote:
> Hi everyone,
>
> I am running a lot of CTs with their roots located on an nfs share.
> Once in a while it happens that a process gets stuck which I fear has
> something to do with the nfs mount.
> See the dmesg out below.
> The problem now is that I can't kill this process anymore.
> This results into beeing unable to stop the CT running this process.
> vzctl stop <CTID> runs into a timeout.
> It is totally impossible to kill the process - The only solution is a
> reboot of the Host-System.

Yep, that's correct. Your only real option is to migrate all of the
other CT's to another host node, then reboot this host node, then
migrate the other CT's back.

> Is there a way to forcefully kill the CT ?

Nope, the kernel is hung waiting for IO which will wait until the cows
come home. Are you using TCP or UDP nfs mounts? Try switching from
one to the other and see if that affects your nfs timeout issue.

> In this case I don't care if the process remains running.
> I just want the rest of the CT to be stopped so I can start the CT again.

I don't think it can be done.

> Here is the dmes output:
> [194043.649945] INFO: task which:810615 blocked for more than 120 seconds.
> [194043.650077] "echo 0 > /proc/sys/kernel/hung_task_timeout_secs"
> disables this message.
> [194043.650274] which D ffff882f74146d50 0 810615 682640
> 125 0x00000084

So the "which" command was scanning the directories in the path and
that's what caused the nfs fault. Maybe you can tune your tcp
settings to compensate (assuming your using tcp nfs mounts) but I've
never tried to do anything like that, so I don't know if settings
could actually fix anything. Network congestion is likely your
biggest issue.

> [194043.650308] Call Trace:

Yeah, once you get a call trace, you're hosed.

Does slabtop show that your nfs slabs are using up extremely large
chunks of memory?

...Todd
--
Always code as if the guy who ends up maintaining your code will be a
violent psychopath who knows where you live. -- Martin Golding

---