
Subject: Re: suspend/resume instead of stop/start from vz init script
Posted by [Vasily Averin](#) on Thu, 15 May 2008 06:34:56 GMT
[View Forum Message](#) <> [Reply to Message](#)

Hi Matt,

Matt Ayres wrote:

> Is there any logical reason why to not dump/restore a container when
> issuing a "service vz stop|start" command? It would make sense that it
> would be much quicker as no longer do we have hundreds of processes
> across the containers starting up. The other benefit is that the VPS's
> do not even see a reboot and maintain their "uptime" as the uptime
> command reports.
>
> Perhaps there could be an option in /etc/vz/vz.conf whether to do the
> dump/restore or stop/start variations.
>
> I'm open for arguments on why this is a bad idea.

1) suspend should write dump file to disk. This dump file can be a quite large and its write will be not faster than usual stop. Restore takes some time too.

IMHO we have measured these times -- and found that suspend/resume is not faster than stop/start. I would suggest you to re-check this experiment on your nodes, it would be very interesting for us to see some real numbers from real life.

2) suspend/resume is a bit more risky operation than usual stop/start, especially after kernel update. Also I would note that in this case some parameters (uname for example) will be changed and it can lead to troubles in some userspace applications.

On the other hand I tend to agree, this option can be useful for some scenarios, however I right now I'm not ready to contrive some understandable example. Do you probably able to describe the situation where this way will give us some significant benefits?

thank you,
Vasily Averin
