

---

Subject: Unable to stop a container from blocking the host

Posted by [rsauvat](#) on Fri, 06 Aug 2010 08:46:10 GMT

[View Forum Message](#) <> [Reply to Message](#)

---

Hi,

On a production server I have setup 8 virtual machine. They share the 16Go of RAM on the host but only using 5Go when running. My problem is when I start an index creation in MySQL on a particular VM after 2-3 minutes the whole host stop responding correctly. I can see in top a lot of processes waiting for I/O and the load on the host goes up to 300. I have the ioprio parameter to 0 for this vm and 7 for other but the problem is still the same. The process take a really long time to respond to a kill -9, more than 10 minutes.

The mysql table is really big (29Go) so the index creation must be I/O consuming. But a iostat on the host show no more than a 1Mo/s on I/O.

I am running Proxmox 1.5 with debian 5 on the host.

```
root@HWnode:~# pveversion -v
pve-manager: 1.5-10 (pve-manager/1.5/4822)
running kernel: 2.6.24-11-pve
proxmox-ve-2.6.24: 1.5-23
pve-kernel-2.6.18-2-pve: 2.6.18-5
pve-kernel-2.6.24-11-pve: 2.6.24-23
pve-kernel-2.6.24-10-pve: 2.6.24-21
qemu-server: 1.1-16
pve-firmware: 1.0-5
libpve-storage-perl: 1.0-13
vncterm: 0.9-2
vzctl: 3.0.23-1pve11
vzdump: 1.2-5
vzprocps: 2.0.11-1dso2
vzquota: 3.0.11-1dso1
pve-qemu-kvm: 0.12.4-1
```

This problem is very annoying because I can't create any indexes in MySQL. And this defies the purpose of containers. How a single container can make the others unstable. I can understand a container slowing down the others but in this case even postfix daemon on an other container can't write mail to the user mailbox.

If anyone has an idea on what is causing this problem any help will be appreciated.

Best.

---

Subject: Re: Unable to stop a container from blocking the host

Posted by [rsauvat](#) on Mon, 09 Aug 2010 08:04:42 GMT

[View Forum Message](#) <> [Reply to Message](#)

---

I may have found the source of the problem. Due to the software RAID implementation by my reseller it seems that the number of fsync/second is very low compared to a standard disk. So this seems to be the problem since the mysql database might be using it very often so the other nodes don't have enough fsync for themselves.

So this leads me to a question. Why the number of fsync is not limited per VM? Or at least there should be a priority to let other containers do their fsync maybe based on ioprio.

---

Subject: Re: Unable to stop a container from blocking the host

Posted by [maratrus](#) on Fri, 13 Aug 2010 13:25:00 GMT

[View Forum Message](#) <> [Reply to Message](#)

---

Hello,

I do understand your problem. So I would like to recommend you the following way: if your server is located in a reliable place i.e. there is no electricity problem you can disable fsyncs inside the VE by writing 0 to /proc/sys/fs/fsync-enable on the HN. This means that all fsyncs inside VE would be very fast as fsync actually doesn't happen. Again, it wouldn't be a problem if your server never go down because of the power issues.

---