
Subject: Re: Heavy Disk IO from a single VM can block the other VMs on the same host

Posted by [quantact-tim](#) on Fri, 02 Dec 2011 18:18:18 GMT

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

You can use `vzctl --ioprio` to set relative disk I/O priorities:
http://wiki.openvz.org/I/O_priorities_for_VE

-Tim

--

Timothy Doyle

CEO

Quantact Hosting Solutions, Inc.

tim@quantact.com

<http://www.quantact.com>

On 12/01/2011 09:27 AM, Kirill Korotaev wrote:

> That's most likely due to a single file system used for containers - journal becomes a bottleneck.
> fsync forces journal flushes and other workloads begin to wait for journal... In reality workload looks like this are typical for
> heavy loaded databases or mail systems only.

>

> How to improve:

> - increase journal size

> - split file systems, i.e. run each container from it's own file system

>

> Thanks,

> Kirill

>

>

> On Nov 29, 2011, at 20:13 , Hubert Krause wrote:

>

>> Hello,

>>

>> my environment is a Debian squeeze host with a few debian squeeze
>> guests. The private and root filesystems of the guest are located on
>> the same raid device (raid5) in an luksCrypt Container in an LVM
>> container on an ext4 partition with `nodelalloc` as mountoption. If I run
>> the tool stress:

>>

>> `stress --io 5 --hdd 5 --timeout 60s` (which means fork 5 threads doing
>> read/write access and 5 threads doing constantly `fsync`) the
>> responsiveness of the other VMs is very bad. That means, isolation for
>> IO operations is not given. I've tried to reduce the impact of the
>> VM with `'vzctl set VID --ioprio=0'`. There was only a
>> minor effect, my application on the other VM where still not

>> responsive.

>>

>> Any Idea how to prevent a single VM to disturb the other VMs regarding

>> diskIO?

>>

>> Greetings

>>

>> Hubert
