

---

Subject: Re: D-state processes on i2o servers  
Posted by [vaverin](#) on Wed, 08 Nov 2006 09:13:02 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

Hello Gerlando,  
sorry for a long delay, but I do not have enough time for OpenVZ forum, I would like to recommend you to access our support.

Quote:

I noticed that the MTRR feature is not enabled anymore.  
Is this correct?

I've seen this behavior too, but I don't know the correct answer on your question. I've found in google that i2o developer Markus Lidel asked the question about MTRR, but I don't found any answers.

However I do not think that it may lead to the some problems. Drivers were taken from mainstream kernel, and plain mainstream kernel work by the same manner. We have tested new driver well and did not noticed any troubles. And we have a positive customers feedback.

Quote:

My general feeling is that write-caching on the disk is disabled, so when the I/O pressure is high the impact on machine responsiveness is noticeable.

How could I check if this is the case or not? I mean, how can I measure disk I/O throughput when there are several concurrent accesses?

As far as I understand write-caching on the disk should have a very low effect to the disk I/O throughput in case of heavy disk IO, just because of the cache will not used in this situation, new data will replace old cache content without any data-reusing.

As far as I know there is some tests for various filesystem operations (bonnie?), you can use it for measurements. However I would note that the IO performance depends vastly on where the data placed physically on the disk. Therefore it is very hard to analyze the test results.

Quote:

Also, could the overhead of vzfs be (at least partially) responsible?

As a side note, on the hardware node we are using an ext3 filesystem, and our VPSes make heavy use of LOTS of small files and LOTS of symlinks. Should we use a different filesystem to improve performance?

vzfs is a cow-like filesystem and it should not have a noticeable overhead.

Also I would note that ext3 performance is not so bad, but it's stability is much better than any other alternatives.

Therefore if you wish to improve filesystem performance i would like to recommend you think about disk IO-sysbsytem upgrade instead of the filesystem change.

Also you can store VPS data on the dedicated disks or disk partitions, it can decrease interference

between various VPSes.

Thank you,  
Vasily Averin

---