

---

Subject: High I/O in 028stab112 VZ kernels

Posted by [ccto](#) on Mon, 23 Jun 2014 08:29:39 GMT

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

---

Hello OpenVZ developer,

Under the kernel 2.6.18-371.6.1.el5.028stab112.3, when we transfer large files (the weekly full backup), the load average grow exponentially.

We used mdadm to form 2 RAID

/dev/md0 = /boot (consists of /dev/sda1 and /dev/sdb1)

/dev/md1 = LVM for / , /vz , /tmp , and swap (consists of /dev/sda2 and /dev/sdb2)

/dev/sdc1 = /backup

We run - tar czf /backup/<VEID>.tar.gz /vz/private/VEID/ - on weekly basis ,  
and use rsync over ssh to sync the tar.gz files to remote server.

During the rsync process, the load average grows from around 1 - 3 to over 200 (the OpenVZ server is very busy, not so responsive)

The transfer speed is around 400Mbps.

If we limit the rsync speed to around 100Mbps, the load average is OK.

In the past, it does not have such phenomenon.

It happens in both OpenVZ CentOS 5 x86 and x64 servers.

We already use a separate disk to store the backup files to prevent I/O congestion during tar or rsync process.

Is there any issue in the CPU scheduler or I/O priority?

It seems it does not happen in CentOS 6 series (e.g. 2.6.32-042stab090.3) kernel.

Please kindly have a look.

Thank you very much for your kind attention.

Regards  
George

---