

---

Subject: Tracking down a scheduling issue between 2 versions

Posted by [brama](#) on Mon, 09 Jan 2012 16:14:27 GMT

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

---

Hi,

I'm trying to track down an issue we have on our servers. All servers run the same software stack (other than host OS and OpenVZ version), and are pretty much identical. However, 2 newly added servers have an issue that I haven't been able to pin down.

The guest OS (and the only OpenVZ instance running) is Debian Lenny. This guest runs the software described below.

The servers spawn a lot of transcoding processes (ffmpeg) depending on demand. They are spawned by a python process that reads a queue over a remote DB connection, figures out which entries to process, and then launches the required processes.

With the new servers, I see that it appears that execution of the python process is much slower compared to the other servers when the load peaks (e.g. 20+, cpu-bound). The script is supposed to check a queue every 5 seconds, but on the new servers, this delay between checks fluctuates heavily, while the old servers are quite consistent regardless of the load. I debugged it, and could rule out the communication with the database. All that's left then is just processing logic. The new servers also currently run as many processes in parallel as the old ones.

Since the guest OS is an identical copy of the older servers, and the system uses the same architecture, I was wondering if this could be explained because the older servers do run an older redhat and ovenvz stack.

The servers we have:

Old Servers: RH5.4, tools: 3.0.25.2, kernel: 2.6.18-194.26.1.el5.028stab079.2PAE

New Servers: RH5.6, tools: 3.0.29.3, kernel: 2.6.18-274.7.1.el5.028stab095.1PAE

The newer servers should be faster both in Disk IO and total performance (identical server barebone, same CPU's running at the same speed, but more cores).

Could a change between these versions explain the odd behavior?

---