
Subject: system limits

Posted by [John Kelly](#) on Thu, 11 May 2006 01:47:04 GMT

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

For my needs, openvz appears to offer a better solution than xen. But I worry about system limits:

http://www.redhat.com/archives/rhl-devel-list/2006-April/msg_01356.html

With a large number of threads/pids, is it true? Can I exhaust 32 bit virtual address space?

Subject: Re: system limits

Posted by [dev](#) on Thu, 11 May 2006 08:27:37 GMT

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

the link you provided tells about per-process virtual address space limitations. It has nothing to do with OpenVZ, but rather with i686 arch. So what is your exact question, what do you mean? Whether this limitation will be global-wide and effect OpenVZ? The answer is NO then.

Subject: Re: system limits

Posted by [John Kelly](#) on Thu, 11 May 2006 10:35:23 GMT

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

dev wrote on Thu, 11 May 2006 04:27 the link you provided tells about per-process virtual address space limitations. It has nothing to do with OpenVZ, but rather with i686 arch. So what is your exact question, what do you mean? Whether this limitation will be global-wide and effect OpenVZ? The answer is NO then.

OK, thanks for clarification.

To use openvz for mass virtual hosting, with large number of VPS and PIDs, what is the practical upper limit for maximum number of PIDs, before scheduling overhead becomes a problem?

Subject: Re: system limits

Posted by [dev](#) on Thu, 11 May 2006 10:38:32 GMT

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

PIDs itself do not introduce scheduling overhead.

on scheduling only number of `_running_` processes influences.

Subject: Re: system limits

Posted by [John Kelly](#) on Thu, 11 May 2006 10:48:41 GMT

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

dev wrote on Thu, 11 May 2006 06:38PIDs itself do no introduce scheduling overhead.
on scheduling only number of `_running_` processes influences.

I admit I don't know much about linux kernel. But suppose I have a 40 to 1 ratio, total PIDs vs. running PIDs. And suppose I have 20,000 total PIDs, 500 running PIDs.

Do you think linux and the CFQ scheduler can handle that?

Subject: Re: system limits

Posted by [RapidVPS](#) on Thu, 11 May 2006 20:59:04 GMT

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

Hi, I really doubt you will support 20,000 pids or 500 running pids on normal server hardware (assuming you meant those numbers literally)- for example on a dual core dual opteron w/ 8GB mem. From my experience, you can have 4000-6000 total processes in such a production system and it will still be handled OK, that is probably 10-25 running procs average. 500 running processes would kill any system.

To answer your question, I have pushed vz to 10,000 pids and ~50 running average in a test case and the server will be slow but alive. All you really need to be concerned about is the number of running processes, I dont think total pids plays a factor in the 2.6 scheduler, they did in 2.4.

Subject: Re: system limits

Posted by [dev](#) on Fri, 12 May 2006 08:18:58 GMT

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

CFQ scheduler is about disk I/O, not pids and CPU.

RapidVPS provided you a correct and good info about running processes. If you have little running processes than even 200,000 of pids is not a problem.
