
Subject: user_beancounters: failcnt increasing before resources reach limit

Posted by [david2389](#) on Thu, 28 Mar 2013 12:29:01 GMT

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

Hi,

I'm a customer of a VPS provider who uses OpenVZ. I don't have access to the OpenVZ configuration or kernel logs.

I'm finding those typical problems reported elsewhere (e.g. "fork: Cannot allocate memory"). I understand they happen because I'm using too much resources. However, the "maxheld" value in the /proc/user_beancounters file isn't near the "barrier" or "limit" ones, at which point I'd expect the "failcnt" to increase.

To illustrate my problem, consider this partial user_beancounters file:

Version: 2.5

uid resource	held	maxheld	barrier	limit	failcnt
67767067: kmemsize	26743189	26895615	31457280		34603008
9201127					
privvmpages	473723	485794	1048576	1153433	
30530					
numproc	87	89	128	128	635

Some minutes later, this is how it looks:

Version: 2.5

uid resource	held	maxheld	barrier	limit	failcnt
67767067: kmemsize	26960203	28481774	31457280		34603008
9227898					
privvmpages	431904	599424	1048576	1153433	
33525					
numproc	88	90	128	128	635

We can see that the failcnt for kmemsize and privvmpages has increased, however the maxheld field doesn't seem to have hit the limit for each field. I'd expect to find something like this:

uid resource	held	maxheld	barrier	limit	failcnt
67767067: kmemsize	26960203	34603008	31457280		34603008
9227898					

Is this normal behaviour? Is there something I can do to make a better usage of the resources I've been allocated?
