Subject: IO Limiting (Priorities)
Posted by CodyRo on Sun, 19 Apr 2009 18:46:22 GMT
View Forum Message <> Reply to Message

I've searched over the forums and don't think this is clarified anywhere, though I could be mistaken (and if I am, excuse me).

In the Wiki IO priority is described to work as the following:

Quote:
Priority range is 0-7. The more priority a container has, the more time for using block devices this container will obtain.

This seems fairly clear - the higher the priority given, the more IO the VPS will be allotted. The issue is later down in the Wiki entry (http://wiki.openvz.org/I/O_priorities_for_containers) is says this:

Quote:
The mapping from priority to time is the following: if 0 prio corresponds to time slice t, than 8 prio corresponds to time slice 2 * t. Default time slice value is HZ/2. The main criteria for fairness at the moment is that if you set I/O prio of CT 1 to p1 and I/O prio of CT 2 to p2, and p1 > p2 then CT 1 should do more I/O than CT 2.

.. which seems to contradict the above statement. If VE1 has IO Priority of 1, and VE2 has IO Priority of 2 .. VE1 should be doing more IO than VE2.

Am I missing something or completely misunderstanding the Wiki entry?

Thanks!

Subject: Re: IO Limiting (Priorities)
Posted by maratrus on Tue, 21 Apr 2009 06:09:54 GMT
View Forum Message <> Reply to Message

I cannot find any contradiction.
Could you please describe what do you have in mind?

The more priority a container has (p1 > p2), the more time for using block devices this container will obtain (CT 1 should do more I/O than CT 2).