

---

Subject: A better answer

Posted by [dowdle](#) on Wed, 30 Jul 2008 14:54:26 GMT

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

---

Ok, the formula you gave (slightly modified to take into account swap) is a good formula if you do not want to overcommit resources. How does that ensure you don't overallocate resources? Well, if you don't give your containers access to more RAM than you have, whenever they try to use more than they have been given the OpenVZ kernel will cause a resource allocation failure and the User Beancounter (UBC) failcnt for the appropriate resource will be incremented.

You can get higher density by overcommitting resources... which works well for many people... but you really have to watch your UBCs and tweak your container's configs to avoid UBC failcnts.

The typical user that has a lot of all-the-same containers are hosting systems... and given the competitive nature of them... I haven't see much in the way of papers written that show best practices for density.

There are two theories on config tweaking. 1) Start low and tweak up. That certainly gives you the highest density... but will usually lead to any load spikes causing failcnts. 1) Starting high and tweaking down or just giving more resources than a container typically uses to better allow for load spikes not causing failcnts.

I'd have to say that the "make one config for a lot of containers all alike for maximum density approach" has potentially lead to the perception that VPSes, containers, or OS Virtualization doesn't work well. But what else are hosting providers going to do when the have to make configurations that match up to a small number service plans / products? As long as you give a container access to a reasonable amount of resources, you really shouldn't have a problem... so use the basic formula.

---