I think its worth noting--

As Dev pointed out, you do not have any cpulimits set on your VEs. This is generally bad especially if each VE has its own admin; ie you want to ensure fariness rather then optimal overall efficiency. When you add cpulimits to your VEs, the load will not decrease, and in fact it may increase. This is because a few of your VEs will reach their cpulimit and their runqueues will expand, which means more processes want to run as seen by your HN.

However, the end result is your cpuunits will kick into gear and each environment will have its amount of cpuunits cputime/equivalent MHz to count on. Each of your VEs will feel more interactive when cpuunits are well enforced, and cpulimit will protect the HN from reaching 0% idle time very often, which makes things even more interactive.

Be sure that your cpuunits are not oversubscribed. See my message #8620 on this page to read more about cpuunits/cpulimit and how they relate to each other.
http://forum.openvz.org/index.php?&t=msg&th=1551

The bottom line is that the load as seen by the HN is not always an accurate factor in determining if your server is overcommitted, underpowered, etc (you may have plenty of idle cpu time with a load of 100!). With regard to strictly your application load, it may and probably is overcommited; however if you are not the admin of your VEs, like in the case you are selling them as a service, you are fulfilling your responsibility in guaranteeing them a time slice of the cpu / equivalent CPU speed in MHz and the HN load is not something you can control.

Hope this isn't too wordy!