Cpulimit and cpuunits are related but not dependent on each other. It took me a while to wrap my head around these concepts, I will explain them in detail..

cpulimit is an equivalent approximate MHz value that the VE can burst to. I say equivalent because its a number between 1 & 100*X where X is the number of cpus in your server. hyperthreading, dual core counts as another cpu in this case. cpulimit is easy to understand, if your server has two 3GHz processors, and you set the cpulimit to 50, your ve can burst to ~1500Mhz.

cpuunits is harder to understand because its specific to a shared resource architecture. Much like QOS on an IP network allows you to guarantee throughput to a class of traffic, cpuunits allows you to guarantee a certain slice of the cputime to a VE. Network QOS allows you to guarantee 7mbit/sec to WWW traffic on a 10mbit pipe. If you guarantee WWW 7mbit and FTP 5 mbit, when your traffic approaches 10mbit/sec, unpredictable things will happen to each traffic set. The same applies for cpuunits. cpuunits is a number based on the number of bogomips in your server. Its scaled up so its easier to divy out to many VEs. The number itself isn't important, it could be 100 or 1,000,000. A server with more cpuunits then another server doesn't necessarily make it faster- a P4 3GHz hyperthreaded will have more cpuunits then your dual opteron with 2Ghz one ach chip. what does that mean? nothing.

[root@neon ~]# vzcpucheck
Current CPU utilization: 454109
Power of the node: 478633

This server has 478633 cpunits to distribute. If you give half of them to a VE, that VE is guaranteed half the cputime in the server. You can say if this server has 6GHz of processing power and it has half the cpuunits allocated to it, its guaranteed around 3GHz of cpu power.

Hope this helps!
Rick Blundell