Subject: TEST IDEA - Scaling with openvz
Posted by kingneutron on Mon, 28 May 2007 06:21:44 GMT
View Forum Message <> Reply to Message

Say I have an external USB2 120GB drive. (Go on, say it!)

Now I would like to mount it for openvz use; no problem, I know how to do this.

I've made some basic mods to the standard Debian 4 openvz template and would like to use it as the "Base model" for testing, similar to "sysprep" builds in Windoze.

Can I just tar.gz up the 101 VE for Debian as a "master image", and vzctl create 102, etc based on that, if I copy the tarball to the Template dir? Or do I need "metadata"?

o Summary:

o Guest VE's will be batch-created to test scaling (10 at first, then 32 / 64 / 128 / 254 in order.)

o Host will have venet / veth of 172.16.0.1 // 255.255.0.0

o Master image will have IP of 172.16.1.1 // 255.255.0.0

oo Each guest VE will have IP like so:
VE 1002 == 172.16.1.2 ..
VE 1254 == 172.16.1.254 // 255.255.0.0

--I'm pretty sure I can do this programmatically with simple Bash scripting; but is there a way in-VE to have it determine what its number is on the host? (i.e. VE 1003 should "know" it's number is 1003 somehow without having networking enabled yet.)

--Once each VE gets powered on and sets up its IP, it will be netcat-listening on $port to make whatever local changes are needed after 1st-boot:

((' cd / ; time nc -l -p 32100 | tar xpf - '))

--We can then send a non-compressed tarball (what, you wanna overload your CPU? we may want to do this SIMULTANEOUSLY for stress-testing the disk subsystem) to each guest with whatever changes are needed on a localized basis, such as /etc/hostname, or the like.

--The basic plan is to easily setup an infrastructure for small ISPs, and network admins who want to deploy openvz on a semi-grand scale for customers or internal use.

--Any comments/suggestions appreciated... TIA