> I guess that what causes it, is having in guest more processes with runnable
> state then reduced virtual cpus available.
 Faster/simpler way:
 - go to guest, run 8x dd if=/dev/zero of=/dev/null
 - load should fairly quickly start creeping up to 8.0
 - on HN set cpus to 1
 - let it run, then stop all the dd's

 result:
Guest:
eyck@etchdev386:~/40p-ovz/work$ w
 16:13:44 up 5:21, 1 user, load average: 2.99, 2.94, 2.20
(and stays like this, 15-minute average even grows)
HN:
codev64:~# w
 16:14:05 up 5:25, 2 users, load average: 0.00, 1.76, 2.65

Playing with dd and set --cpus I also managed to cause load ~8 on HN while
guest reported only 2.0

.. I guess noone is doing things like this on production systems, and even
if so, running chkpnt/restore as precaution when performing such changes
is not out of the question.
--
Key fingerprint = 40D0 9FFB 9939 7320 8294 05E0 BCC7 02C4 75CC 50D9
 Total Existance Failure