

---

Subject: TCP stack virtualization

Posted by [dagr](#) on Thu, 25 Jan 2007 09:02:26 GMT

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

---

AFAIK - sysctl parameters are same for all VEs including HN. I believe it could be great to virtualize some of them as well. For instance, recently i found out that when clients in apache (apache in vps) exceed ListenBacklog - they are not refused like i need, they just hang (just like when they exceed MaxClients but lower than ListenBacklog). The difference in these 2 cases is that in first - server socket is in state syn-recv , in second - in state established. But in both cases client side hangs. Moreover , server vps makes syn ack retries until net.ipv4.tcp\_synack\_retries exceeds, then socket dies. By default it takes about 3mins. What i need - is that it should just refuse connections over ListenBackLog parameter. I need this cause my VPS apaches are behind frontend reverse proxy. And if ANY inner apache exceeds MaxClients , proxy connections to it just hang and their number just grows like snow ball which affect all other apaches. I can get close to it if i set net.ipv4.tcp\_synack\_retries=1 (by default is 5), but id like this to be valid only for 1 VPS, not the whole HN. I thinks there are plenty other situations which will demand such functionality. Is such thing in TODO list ?

Also this can be fixed by iptables connlimit feature which is in patch-o-matic patch. And i really dont sure is it will be possible to apply it to openvz kernel, please tell me if its so.

---