Subject: Re: Dropped packet, source wrong veid
Posted by ittec on Wed, 04 Jun 2008 08:34:44 GMT
View Forum Message <> Reply to Message

After that I examined properly all messages of syslog I found souspicious things with relation with my problem.

The symptom is that some VE cannot answer all TCP paquets. First I had troubles with kmemsize, after some minutes, into Ve bash I received this message:

Cannot to allocate memory

or

Unable to fork

I read openvz wiki for kmemsize parameter and "lowmemory". By default, using vzsplite the kmemsize value assigned is not enough. I read lowmemory is the first 832 MB in i386 system. I assigned 100 MB to kmemsize for the problematic VE. After that, Ve not failed more. But, wasn't able to serve all traffic. Then I noticed in syslog the next

TCP: time wait bucket table overflow

I read that this bug was fixed in 2.6.18

http://bugzilla.openvz.org/show_bug.cgi

Now im thinking of increase /proc/sys/net/ipv4/tcp_max_tw_buckets_ub but...based on ... what? Default value is 16384. How affects this on hardware node? Are you agree to increase max_tw_buckets is properly solution?

Thanks