Subject: Re: [PATCH 2.6.21-rc6] [netfilter] early_drop imrovement
Posted by Patrick McHardy on Fri, 06 Apr 2007 15:08:28 GMT
View Forum Message <> Reply to Message

Vasily Averin wrote:
> No, I've not investigated this scenario. It looks like you are right and my
> patch can leads to a long delays.
>
> In my experiments I've decreased ip_conntrack_max lower than number of hash
> buckets and got the "table full, dropping packet" errors in logs. I've looked on
> the conntrack list and found a huge number of conntracks that can be freed.
> However my hash bucket was empty and therefore I even did not have any chances
> to free something. That's why I would like to check the other hash buckets too.
>
> Ok, let's limit the number of conntracks that can be checked inside
> early_drop(). What do you prefer: some round number (for example 100) or
> fraction of ip_conntrack_max (for example 1%)?

A (small) fraction sounds better. We could even consider keeping track of the number of conntracks that can be evicted (not assured), so in a DOS situation we could save unnecessary table scans. Not sure if its worth the effort though.

Anyway, please base your patch on the net-2.6.22 tree, which doesn't include ip_conntrack anymore.