Subject: Re: [NETFILTER] early_drop() imrovement (v4)
Posted by Patrick McHardy on Tue, 03 Jul 2007 11:42:54 GMT
View Forum Message <> Reply to Message

Rusty Russell wrote:
> On Wed, 2007-06-27 at 15:54 +0200, Patrick McHardy wrote:
>
>>[NETFILTER]: nf_conntrack: early_drop improvement
>
> This looks good. The randomness in the hash means we no longer need the
> "hit the same hash bucket" heuristic to avoid hashbombing.
>
> I still wonder if we should batch up the drops a little while we're
> doing all this work? Should reduce stress under serious flood load.

Good point, I didn't think of that. Its a bit tricky though, we can't
destroy them while holding nf_conntrack_lock, so we'd either have to
release and re-grab it for every conntrack we want to kill or write
lock it from the beginning, clean them from the lists immediately and
put them on a temporary destruction queue. Or split death_by_timeout
so it can deal with callers already holding nf_conntrack_lock ..

I'll see if I can come up with a halfway decent looking patch :)