
Subject: Re: [NETFILTER] early_drop() improvement (v4)

Posted by [vaverin](#) on Wed, 27 Jun 2007 13:02:24 GMT

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

Patrick McHardy wrote:

> Vasily Averin wrote:

>> it is incorrect,

>> We should count the number of checked `_contracks_`, but you count the number of
>> hash buckets. I.e "i" should be incremented/checked inside the nested loop.

>

>

> I misunderstood your patch then. This one should be better.

> +static int early_drop(unsigned int hash)

> {

> /* Use oldest entry, which is roughly LRU */

> struct nf_conntrack_tuple_hash *h;

> struct nf_conn *ct = NULL, *tmp;

> struct hlist_node *n;

> - int dropped = 0;

> + unsigned int i;

> + int dropped = 0, cnt = NF_CT_EVICTION_RANGE;

>

> read_lock_bh(&nf_conntrack_lock);

> - hlist_for_each_entry(h, n, chain, hnode) {

> - tmp = nf_ct_tuplehash_to_ctrack(h);

> - if (!test_bit(IPS_ASSURED_BIT, &tmp->status))

> - ct = tmp;

> + for (i = 0; i < nf_conntrack_htable_size; i++) {

> + hlist_for_each_entry(h, n, &nf_conntrack_hash[hash], hnode) {

> + tmp = nf_ct_tuplehash_to_ctrack(h);

> + if (!test_bit(IPS_ASSURED_BIT, &tmp->status))

> + ct = tmp;

> + if (--cnt <= 0)

> + break;

> + }

> + hash = (hash + 1) % nf_conntrack_htable_size;

> }

it is incorrect again: when cnt=0 you should break both cycles.

thank you,

Vasily Averin
