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

Posted by [vaverin](#) on Wed, 27 Jun 2007 13:25:58 GMT

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

Patrick McHardy wrote:

> Vasily Averin wrote:

>> Patrick McHardy wrote:

```
> -static int early_drop(struct hlist_head *chain)
> +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, hnnode) {
> - 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], hnnode) {
> + tmp = nf_ct_tuplehash_to_ctrack(h);
> + if (!test_bit(IPS_ASSURED_BIT, &tmp->status))
> + ct = tmp;
```

It is incorrect: you should break nested loop here too.

```
> + if (--cnt <= 0)
> + goto stop;
> +
> + if (ct)
> + break;
> + hash = (hash + 1) & (nf_conntrack_htable_size - 1);
> }
> +stop:
> + if (ct)
> + atomic_inc(&ct->ct_general.use);
> + read_unlock_bh(&nf_conntrack_lock);
```
