

---

Subject: Re: [NETFILTER] early\_drop() imrovement (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 \_conntracks\_, 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, 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;

> + 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

---