Subject: Re: [NETFILTER] early_drop() imrovement (v4)
Posted by Patrick McHardy on Wed, 27 Jun 2007 13:35:49 GMT
View Forum Message <> Reply to Message

Patrick McHardy wrote:
> Vasily Averin wrote:
>
>>Patrick McHardy wrote:
>>
>>
>>>>+ 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;
>>
>>
>>It is incorrect: you should break nested loop here too.
>
>
>
> No, as I said, we want the last entry of the chain.

Ideally we should do something like this I think (please let it be
correct :)):

```
+       for (i = 0; i < nf_conntrack_htable_size; i++) {
+               entries = 0;
+               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;
+                       entries++;
+               }
+               if (ct)
+                       break;
+               if ((cnt -= entries) <= 0)
+                       break;
+               hash = (hash + 1) % nf_conntrack_htable_size;
        }
```

So we always walk chains up to the end and NF_CT_EVICTION_RANGE is
just a minimum. This ensures we will always get the last entry *and*
we won't scan less entries than currently if someone has a chain
longer than 8 entries.

What do you think?