
Subject: Re: [NETFILTER] early_drop() improvement (v3)
Posted by [Patrick McHardy](#) on Mon, 25 Jun 2007 13:53:22 GMT
[View Forum Message](#) <> [Reply to Message](#)

Vasily Averin wrote:

```
> +static int early_drop(const struct nf_conntrack_tuple *orig)
> +{
> + unsigned int i, hash, cnt;
> + int ret = 0;
> +
> + hash = hash_conntrack(orig);
> + cnt = NF_CT_PER_BUCKET;
> +
> + for (i = 0;
> + !ret && cnt && i < nf_conntrack_htable_size;
> + ++i, hash = ++hash % nf_conntrack_htable_size)
> + ret = __early_drop(&nf_conntrack_hash[hash], &cnt);
```

Formatting is a bit ugly, looks much nicer as:

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

        ret = __early_drop(&nf_conntrack_hash[hash], &cnt);
        if (ret || !cnt)
            break;
        hash = ++hash % nf_conntrack_htable_size;
    }

> @@ -1226,7 +1243,7 @@ int __init nf_conntrack_init(void)
>   if (nf_conntrack_htable_size < 16)
>       nf_conntrack_htable_size = 16;
>   }
> - nf_conntrack_max = 8 * nf_conntrack_htable_size;
> + nf_conntrack_max = NF_CT_PER_BUCKET * nf_conntrack_htable_size;
```

I don't like the NF_CT_PER_BUCKET constant. First of all, each conntrack is hashed twice, so its really only 1/2 of the average conntracks per bucket. Secondly, its only a default and many people use `nf_conntrack_max = nf_conntrack_htable_size / 2`, so using this constant for `early_drop` seems wrong.

Perhaps make it `2 * nf_conntrack_max / nf_conntrack_htable_size` or even add a `nf_conntrack_eviction_range` sysctl.
