
Subject: Re: [NETFILTER] early_drop() improvement (v4)
Posted by [Patrick McHardy](#) on Wed, 27 Jun 2007 08:52:16 GMT
[View Forum Message](#) <> [Reply to Message](#)

Vasily Averin wrote:

> When the number of conntracks is reached nf_conntrack_max limit, early_drop()
> tries to free one of already used conntracks. If it does not find any conntracks
> that may be freed, it leads to transmission errors.
> In current implementation the conntracks are searched in one hash bucket only.
> It have some drawbacks: if used hash bucket is empty we have not any chances to
> find something. On the other hand the hash bucket can contain a huge number of
> conntracks and its check can last a long time.
> The proposed patch limits the number of checked conntracks and allows to search
> conntracks in other hash buckets. As result in any case the search will have the
> same chances to free one of the conntracks and the check will not lead to long
> delays.

Thanks Vasily. I have some patches queued to convert all conntrack hashes to hlists, which conflict with your patches. They need a bit more work, I'll integrate your changes on top of them once I'm done.

BTW, I played around with your last patch yesterday and it shows a big improvement when flooding the machine with new connections. Previously about 5% of the (valid) new connections would get dropped, with your patch not a single one :)
