

---

Subject: [PATCH net-2.6.25 0/10] Make fragments live in net namespaces

Posted by [Pavel Emelianov](#) on Tue, 22 Jan 2008 13:52:36 GMT

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

---

The overall design I propose is to keep the hash table global and tag `inet_frag_queue` with the net. Since the fragments hash is going to be re-sizable, this is OK to keep fragments from different namespace in one hash.

To speedup the evicting process LRU list is made per namespace.

As far as the CTL-tuned variables are concerned, the timeout and thresholds are made per namespace, since they have the per namespace sense, but the secret rebuild interval is read-only in sub-namespaces.

Since fragment management code is consolidated for ipv4 and ipv6 I make them all in one go. The `conntrack_reasm` netns-ization is not done - we have to make at least the core netfilter per namespace first, but this reasm code is patched to keep working in the initial namespace.

Signed-off-by: Pavel Emelyanov <[xemul@openvz.org](mailto:xemul@openvz.org)>

---